Date: Mon Mar 01, 1993 12:54 am PST
Subject: Jacobian code

[Avery Andrews 930301.1930]

For anyone who might be interested, I've hacked Bill Powers' 14df code so that it displays the components of the Jacobian associated with the first four df's (shoulder yaw & pitch, upper arm roll, & elbow flexion) (which I conjecture to constitute a hand-positioning system). IBM Compatible only, 286 / EGA or better recommended, tho not required. Send me a message if your're interested in a copy. Some of the Little Man 2 files are used, so you'll need them too.

Avery.Andrews@anu.edu.au

Date: Mon Mar 01, 1993 8:18 am PST
TO: Hortideas Publishing / MCI ID: 497-2767
Subject: Files

[From Dag Forssell (930301)]

Thanks for 2 disks. The file that triggered my call in the first place, 9011.zip got left out. No doubt because I confused the list by changing my mind about the others. That is a peculiar file. The 9011.zip I have shows 369,669 bytes with 87% compression. Then holds only 126 kb, being cut-off in mid sentence. Probably a zip problem rather than a Gary problem.

"shows" above means in the XTREE header under "open zip and arc".

I glanced at the Closed Loop. Am focusing tightly on my projects, and am barely downloading the net and glancing at the traffic. Sorry for lack of feedback. People respond ONLY when they have error signals, sez PCT. When YOU want feedback, that is discouraging.

Our 2-hour presentation went well. Now have a video, complete with picture inserts of all the slides. Have had to learn about editing videos too. And buying video equipment. How about making home videos of Durango this next time? Have a second presentation tomorrow (Tuesday) night. Am redoing all slides with a new graphic look. (Headline font, border). And putting the finishing touches on the hand out, which will also be sent with the video.

Thanks in advance for 9011.zip.   Best,   Dag

Date: Mon Mar 01, 1993 11:51 am PST
Subject: Misc. Comments

[From Rick Marken (930301.1000)]   Greg Williams (930228) --

I'm really glad to hear that Cam is recoving. Thank God (whatever THAT is).

>Did anyone like or not like the last CLOSED LOOP? (Hint, hint.)

I loved it; particularly that paper on "The Blind men and the Elephant". I'm currently negotiating options on the movie rights and your cut will be hugh -- about as big as my royalties for the LCS intro.

Ed Ford (930228:0920) --

>A close friend needs good references in the current literature
>(books and articles) on the best explanation of cognitive theory.
>I would appreciate anything you could offer.

I don't have easy access to the current literature on cognitive theory. And it would be hard for me to evaluate what might constitute the "best" explanation of this theory -- since I think they are all equally ridiculous. But I an currently reading a pop science book called "The Improbable Machine" by Jeremy Campbell. It's really about how neural networks are the cognitive theory of the future; but it does describe "current cognitive theory" as well.

Allan Randall (930226.1730) --

> You do not like information theory.

It's not a matter of like or dislike. It's a matter of "SO WHAT"? Information theory contributes zilch to our understanding of living control systems (though, I'm sure, Martin disagrees). But there is no need to argue; WITHOUT information theory, Bill Powers has been able to build a simulation of a system that can produce realistic complex behavior in a realistic environment; WITH information theory the life sciences (with thousands of bright researchers and decades of research) have been barely able build simulations of systems doing unrealistic, simple behavior in unrealistically simple environments.

Information theory may be cute and elegant and famous -- but it contributes nothing to our understanding of purposive systems (at least, I think so -- until I am proved wrong by Martin). I don't think PCTers "dislike" information theory any more than they dislike S-R theory. When it comes to the behavior of living systems, these theories are either irrelevant or wrong. No need for dislike. Of course, if you just like the theory and, for whatever reason, think that it just MUST be of some value -- well, feel free to think that way. In PCT we value demonstrations (via working models) of the value of theoretical ideas. I will believe that information theory is valuable when I see how it can improve models of purposive behavior.

Best Rick

Date: Mon Mar 01, 1993 12:06 pm PST
Subject: Ad; Closed Loop

[from Gary Cziko 930301.1540 GMT]

Greg Williams (930228) said:

>No problem with the ad, here. Let the new netters arrive -- I'm ready.

I'm not sure I am (I already get tons of returned messages every day from CSGnetters who become disconnected for one reason or another), but I'll go ahead with this anyway. How about someone out there composing a nice ad? I can add the technical info to it.

>But I hope that some more OLD netters will agree to let me use their
>posts in future CLOSED LOOPs (I need copyright permission). You (might)
>know who you are....

Greg, why don't you mention by name the people who contribute a lot to CSGnet but who have not given you clearance? Attach to that a copy of the form you need filled out and I bet you will get some results.

If you are too shy to go public with their names, send them to me (again with a copy of the clearance form) and I will contact each one privately.

>The doctor decided not to open up Cam's arm -- just stomped on it a
>bunch in the operating room. All going well now.

I hope that those were well-controlled stomps and not ballistic ones.

>Did anyone like or not like the last CLOSED LOOP? (Hint, hint.)

I sure did. It contained prettily formatted versions of two of my favorite previously underground papers. Both will be required reading for the graduate seminar I am now teaching which includes PCT.--Gary

P.S. You will soon get csg-l log9302d. Remember there is usually no log9302e in February so don't bug me about not receiving it (hint, hint).

Gary Cziko

Date: Mon Mar 01, 1993 3:29 pm PST
Subject: WORDS, VIEWPOINTS - RKC

[From Bob Clark (930301.1730 EST)] Bill Powers (930216), (930229.0800)
                                 Martin Taylor (930219.1900), (930221.1640)

Instead of making specific comments on each of the items distributed through February 24, I am offering items of a more general nature.

I. WORDS, MEANINGS, LANGUAGE
It seems to me that language arises from a desire to interact with other beings with the convenience provided by combinations of words. The one sending the message selects the words and sentence structures according to his (remembered) experience. The words he selects have no "intrinsic meaning." However, he assumes that they have "generally agreed" "meanings." Assigning "intrinsic meaning" to words is a shorthand convenience -- we need not take the time to review each "ordinary" word's meaning. In case of disagreement, the dictionary provides a convenient solution. If words are used with specialized, or other "unusual" meaning, specific definitions are necessary. There are many such special languages: mathematics, physics, English, Latin, French, etc etc. Bruce Nevin's impressive discussions describe ad illustrate the development and variations of language.

II. VIEWPOINTS: INTERNAL VS EXTERNAL -- Martin Taylor (930219.1900)

This is an important observation. You have focussed on a general concept: "viewpoint." In view of your remarks, I am trying to summarize my (present) orientation in the following. This turns out to be much more difficult than I expected -- and probably will change with additional review.

III. MY (RKC) VIEWPOINT
A. General View
Again quoting Powers (930216):
*This world, to the best of my knowledge, originates in signals
*emitted into the nervous system by sensory receptors....
*
*This means that the world we experience must consist of sensory
*signals and other signals derived from them. The "other signals
*derived from them" include the totality of what we can experience,
*from the taste of chocolate to Fermat's Last Theorem, as well as
*our experienced "interest" in that Theorem, if any, and any
*"thoughts" we may have about it. Nothing is exempt.
*
*When I say "it's all perception" this is what I mean. We live
*inside a nervous system and all we know is what goes on inside that
*nervous system.

As I noted in Clark (930219.1145 am EST), that is also my viewpoint.

B. Categories. When I investigate what I have available ("inside my nervous system"), I find several easily identified categories. Many other categories can be used as desired. I find the following four categories particularly convenient and useful.

a. "Decision Making Entity." "DME." "Center of Awareness."
This is the entity that "uses" viewpoints. "I" is not used because it tends to include too much. I proposed this in Clark (921205). This entity can direct its Attention to any of the neural signals entering the central nervous system. It can shift its attention rapidly from one signal (or group of signals) to another. It also can select which of the available signals has its attention at any given time. It responds to "built-in" reference levels by selectively "paying more attention" to some signals than to others.

b. "Recording Function." "Memory." "Conscious."
This is the entity that forms records of signals to which attention is directed. Attention can shift fast enough that it appears that all signals are recorded. Mere "exposure" to perceivable events seems to be insufficient for remembering. Conscious Attention, ie perception, appears necessary. Teachers, Parents, Supervisors, etc are invariably concerned that their students "pay attention."

c. "Perceptual Signals." "Attention."
These are the signals to which the DME's Attention may be directed. From time to time, the DME selects them from the available Signals. These form two goups:

1. "Sensory Signals" reporting the current condition of all physiological systems with neural connections to the central nervous system. They may form various combinations, resulting in production of additional, derived, Sensory Signals.

2. "Imaginary Signals" are recorded Sensory Signals and other recorded signals as selected by the DME. The Imaginary Signals include all perceptual signals derived

from recordings. Generally they are organized in some manner by the DME for convenience and accessibility. Such organization will distinguish between those coming from "External Sources" and those coming from "Internal Sources." When selected by the DME for examination, they resemble audio-visual-sensory recordings. They normally run from past time events toward the present and the DME can extrapolate them to future time. Likewise, memories can be combined in various ways, both sequentially and simultaneously. In this respect, they resemble editing of videotapes.

c."Output Signals."

These Signals are recorded in the memory together with the corresponding perceptual signals. After review, the DME determines the "desired" effects on the perceptual signals. The DME then applies the remembered perceptual signals to the corresponding Output Systems. They act as "Reference Signals" for the systems connected to them. Effects are determined by the nature of the systems to which they are connected. Although the DME cannot directly perceive these signals (they are not "incoming"), but their effects are determined by observing corresponding perceptual signals.

d. "Comparator Function."
The DME makes its selections on the basis of comparison of the "desired" effects with the anticipated results offered by alternative sets of Imaginary Signals in relation to current Sensory Signals (and their combinations).

C. VIEWPOINTS REGARDING THE HIERARCHY. Martin Taylor (930219.1900)
I have already been thinking about pointing out alternative views of the basic, feedback control system. However, see II above, you have focussed on a more general concept: "viewpoint." When I apply that concept to a minimal system, I find five (5) identifiable viewpoints. Perhaps others can be found. Different viewpoints may call for different classifications and definitions of the Hierarchical Levels/Orders.

a. The "USER'S" view. The User's DME selects the desired condition (activity, etc) of his own system as it relates to its surroundings, and applies the indicated reference signals. The User observes the resulting activity, etc for possible deviation from intended performance. If deviations are observed, corrections are applied as indicated. The corrections are selected from memory, including anticipation, analysis, and theory (as the User understands them). This process continues as long as results are acceptable. If the results are not within limits, changes may be needed in the remembered structures. Although the concept of a Hierarchy is not essential for the usual User, it can be very helpful when there is difficulty in finding adequate results.

b. The "ENGINEER'S" view. This view is "Objective" in that the Engineer treats the subject as external to himself, omitting the part(s) he plays in this activity. He studies the details of the various elements of the system(s) and their interconnections. Each element is evaluated in terms of the relation(s) between its input(s) and its output(s), expresses them in logical/mathematical terms, and analyzes the results. If this is unacceptable, modifications of one or more elements and/or interconnections are examined for possible alternatives. The Engineer supplies standards of performance selected from his memory by his DME. In this process, the Engineer's DME controls the activity. Although the concept of a Hierarchy is not essential for the usual Engineer (many are quite successful without it), it can be very helpful in more complex and multi-dimensional situations.

c. The "OUTSIDER'S" view. The Outsider, that is, his DME, is observing the activities of another "living-behaving" entity. His information about that entity is derived exclusively from his own input systems -- sensory, as modified and interpreted by his own established internal systems. He uses his knowledge to construct a description of the internal structure of the other entity. All this activity, together with the conclusions, is stored in his memory, and continues to be available for future application, modification, etc. These activities may include discussions, etc with other Outsiders. Although the concept of a Hierarchy is not essential for the usual Outsider/Observer, it can be very helpful in analysis and interpretation of results.

d. The "EXPERIMENTER'S" view. This view is also "Objective" in that the Experimenter treats the Suject as external to himself. He assumes that the subject's reference levels are determined by the Experimenter's instructions combined with the subject's pre-existing decisions. The Experimenter selects and applies some action to the subject's externally accessible inputs. The results are interpreted in terms of whatever behavioral theory he wants to apply. Although the concept of a Hierarchy is not essential for some experimental purposes, it can be very helpful both in experimental design and interpretation.

e. The "THEORIST'S" view. The Theorist pays attention to all the views listed above as well as any others that may be proposed. He resembles the Experimenter in searching for confirmation, or denial of proposed theoretical and/or analytical ideas. The User's and Outsider's views provide additional data for evaluation of proposals. The Engineer's view provides guidelines as to the logical and technical limitations that are intrinsic to the external surroundings. Although the concept of a Hierarchy is not essential for some theoretical purposes, it offers the most inclusive and effective theoretical framework I know of.

D. TWO VIEWS OF HIERARCHICAL LEVELS/ORDERS
These are both "Theorist's" views, as above.

a. Powers View, [Bill Powers (930218.0730)]
*My levels are intended to describe categories of experience
*that all people (and even animals) employ without any training
*or knowledge.
Bill is concerned with "categories of experience."

b. In my (RKC) approach, I have focussed on the perceptual signals as they combine to form the hierarchy. "Hierarchy" is defined in BCP, p 78:

*"This model consists of a hierarchical structure of
*feedback control organizations in which higher-order
*systems perceive and control an environment composed of
*lower-order systems; only first-order systems interact
*directly with the external world.
*
*"The entire hierarchy is organized around a single concept:
*control by means of adjusting referenced-signals for
*lower-order systems."

I am concerned with categories of perceptual signals as they combine to form a hierarchy of perceptual signals.

In this post and in my post, RKC (921205), I have attempted to summarize my present views. They are very highly condensed in these posts, and should be reworked.

Regards, Bob Clark

Date: Tue Mar 02, 1993 11:03 am PST
[from Joel Judd 930302 2:30CST (Caribbean Standard Time)]

I've been trying to figure out a good way to reenter the net--it seems stupid but I feel like I'm trying to contact a relative I haven't spoken to in months but whom I'd like to start talking to again. And we're talking about COMPUTERS here. So for the benefit of those who've signed on in the last eight months I'll just introduce myself briefly.

Last spring I finished five years of doctoral work at the U. of Illinois under the tutelage of Gary Cziko. My interests are in language and particularly second language learning and teaching, though I find most anything related to learning and behavior interesting. I am currently doing time on the island of Puerto Rico at the San German campus of the Interamerican University, one of the two large universities on the island (the other being the University of Puerto Rico). This campus just got hooked up to Internet last month, but unfortunately I must relearn to send and receive e-mail since everything here so far is IBM-based and I don't have the second-hand mail system I did in Illinois. Now I know how people like Rick have felt when doing this at home on their own. I used to laugh at the funny looking postings-- I don't anymore.

I am presently trying to work out a design to study one of the basic English courses here. These courses, which use translated English texts as the basis for the course (with accompanying vocabulary and comprehension question tests) assume that the students have some kind of literacy skills in Spanish which they may employ in the "learning" of English. Unfortunately, it is the sad experience of the professors that these "skills" whatever they may be, are lacking. I would like to show this, "scientifically," to the department to provide a reason to modify the basic English curriculum. Unfortunately, literacy is not an area I have dealt with, but which we will need to somehow define and measure as part of the evaluation process of these courses. So if anyone has some ideas about language literacy, please feel free to pass them on.

Well, time to go to work. Nos vemos.

Date: Tue Mar 02, 1993 12:19 pm PST
Subject: in a manner of speaking

[From: Bruce Nevin (Tue 93032 14:30:23)]

Too much to respond to, no hope of catching up, but you hooked me here, apparently.

[Bill Powers (930222.1300)] --

>At one level, we want to explain how people can come to speak in
>a certain way. That's not very hard to explain with PCT, I think.
>They make the sounds they hear match the sounds they intend to hear.

"Intend" has a peculiar duplicity here. Often, to judge from my actual utterances, I "intend" to hear myself saying things in a manner that is [socially defined as] "contrary" to the way that I am conscious of intending. No quarrel, or even quibble, just troubled by the depth of the "awareness" problem in this.

>But at another level, we have to ask "Why do they want to hear
>one set of sounds rather than another?" The answer to this
>question can't given on linguistic grounds, but on the basis of
>relationships with other people, self-image, and so forth.
> That's most of what I'm trying to get at. We have
>to model the perception and control of social interactions, with
>language as one facet of them but with many other considerations
>of equal importance.
> The norms we adopt pertain to all
>sorts of behavior, and our perception of "likeness" to those whom
>we want to be like include far more than the way they talk. The
>GENERAL problem we want to solve has to do with the perception of
>likeness and where people set their reference levels for it, not
>any particular example of it.

To this apply all the problematic aspects of mimicry that you noted in a later response to Martin. Only perhaps worse: that is, I can imagine seeing your glasses in your hand from your point of view, but unless I have considerable experience with recording equipment I cannot imagine the acoustic image of my voice as others hear it. The reference perceptions must be for the subjective (intracranial) acoustic image, not for the voice that others hear. Similarly for much of gesture, posture, and body language. So the reference perceptions must be adjusted and established with respect to factors in others' responses to us. Do they recognize us as the kind of person we intend to present ourselves as? Do they accord to us social roles and expectations that are acceptable or (better, of course) that we desire? Do we get a balance of our desire for independence vs. our desire for acceptance and intimacy that is acceptable to us? Oops! Do we even know how to talk about independence vs. intimacy in PCT terms?

>That what what I intended to communicate -- that the illusion of
>the glottal stop would appear only in speakers who had to prevent
>themselves from extending the vowel into a diphthong in order to
>imagine matching the heard vowel.

Now why would a speaker of one language *have* to end a word-final stressed syllable with a diphthong or a consonant? Clearly, there is no physiological or neurological constraint, which would have to apply to speakers of all languages. Rather, there is a socially instituted pattern on the basis of which speaker/hearer's expectations get set. There may be differences of detail between individuals, but we compensate for these very quickly (one imagines some analog of a lookup table or rules of conversion). A former professor, Dell Hymes, pronounced the "th" voiceless in "although" (like the th in thing rather than the th in all the other th words in this sentence); a CS professor at MIT whose name I forget (I saw him on a videotape of a lecture on user interface design) pronounces the medial consonant in "button" as a voiced flap (as in an informal, rapid pronunciation of latter, ladder, butter); my 6-year-old, Katrina, has resolved the neutralization of intervocalic t/d in favor of t, so that she actually says "datty" for "Daddy" when pronouncing carefully (just like a careful pronunciation of butter, latter). We compensate in the sense of knowing what was intended and what our own pronunciation of it (the recognized word) would be.

In conversation we may feel obligated either to maintain our contrary pronunciation or to conform to the other's pronunciation (out of friendliness, so as not to embarrass the other, etc.) Happens around here all the time around the pronunciation of "route" and "router" (is it like the nether appendage of a tree or like the flight of an army in defeat). We heard a lot of "hair-ess" vs. "her-ass" during the Anita defamations, er, Clarence Thomas hearings. We feel a need to be consistent with one another, and out of this does develop impressive consistency in social conventions or norms associated with subpopulations to which people identify themselves as belonging.

>When you eliminate the dialect itself as a means of identifying
>the subpopulation, what is left? Location of residence, income,
>sex, age, education, occupation -- all the marks that traditional
>psychology has tried to use as an "objective" way of identifying
>populations. It is notoriously difficult even to define a
>"population" in psychological experiments -- why should it be any
>easier for linguists? And why should these marks of membership in
>populations work for linguistics as predictors of verbal behavior
>any better than they do in psychology as predictors of any kind
>of behavior? I look askance at all purported "facts" that were
>established by these traditional means.

But we don't have to identify subpopulations "objectively." People--the members of the social groups--do it for us. Or rather they do it for themselves, and for each other. And social dialect is an excellent index when you consider natives of a place like NYC, for example. (Don't take my word for it, have a look at Labov's methods and results in the papers in _Language in the Inner City_ f'rinstance.) Sure, there are people that don't fit in, including obviously tourists and immigrants from other regions, but also effortful misfits of one sort or another. This is simply true: there are people that don't fit in the major social strata and groupings of any given community. Sometimes their seeming idiosyncrasies turn out to be negations or reversals of the norms of one or more existing groups, by which they deny membership. Sometimes they take on characteristics associated with one or more truly "foreign" groups (the guy from Topeka who talks with a put-on British accent; Bobby Zinnaman's knock-up of Ramblin' Jack Elliott in his emerging persona as Bob Dylan). And some people may truly bridge or alternate membership in more than one local group.

But the means that people use to identify themselves and others are not likely to be "location of residence, income, sex, age, education, occupation" and the like, but rather "subjective" factors of body language, gesture, and dialect, the stuff of creating and projecting a social face or persona. And people do this to elicit and sustain cooperation, it seems to me. Just because social psychology and sociology have often been inept handmaidens of government (and employers), trying to keep their hands aseptic and "objective", does not discredit social groupings as something that people identify and identify with or against.

Has anyone here read Deborah Tannen's work on manners of speaking? _You Just Don't Understand_ is the most popular one, on stereotypically male-female differences in conversational style and problems that follow from them.

Bruce   bn@bbn.com

Date: Tue Mar 02, 1993 1:41 pm PST
Subject: Re: in a manner of speaking

[Martin Taylor 930302 16:00]
(Bruce Nevin 930302 14:30)

I've been waiting for an opportunity to use this...

(even if I can't spell the dialect right)

>To this apply all the problematic aspects of mimicry that you
>noted in a later response to Martin. Only perhaps worse: that
>is, I can imagine seeing your glasses in your hand from your
>point of view, but unless I have considerable experience with
>recording equipment I cannot imagine the acoustic image of my
>voice as others hear it.
>...
> So the reference perceptions must be adjusted and
>established with respect to factors in others' responses to us.
>Do they recognize us as the kind of person we intend to present
>ourselves as?

Oh wad some Guid the giftie gie's
Tae see oursens as ithers see's.

(Anybody with a written version of Burns can correct this as they will).

T'would make mimicry much easier. We see our perceptions, imagined or based on sensory input, but we see other people's actions. Hard to mimic. But mimic actions we do. And I think much later do we find out what perceptions those actions serve to control, and only then can we integrate them into the hierarchy.

Martin

Date: Wed Mar 03, 1993 7:29 am PST
Subject: lateral specialization

[From: Bruce Nevin (Wed 93033 10:06:36)]

The following new book looks interesting. I don't recall any discussion of hemispheric specialization on CSG-L.

Wray, Alison   THE FOCUSING HYPOTHESIS
1992   xiv, 201 pp. + index
Cloth: 1 55619 389 0 $58.00
John Benjamins

Proposes that the left hemisphere teaches patterns to the right. The right then handles routine language processing, while the left deals with ideas and complicated or unexpected structures.

Date: Wed Mar 03, 1993 10:52 am PST
Subject: Language facts

[From Bill Powers (930303.0830)]   Bruce Nevin (930302.1430) --

>... I can imagine seeing your glasses in your hand from your
>point of view, but unless I have considerable experience with
>recording equipment I cannot imagine the acoustic image of my
>voice as others hear it. The reference perceptions must be for
>the subjective (intracranial) acoustic image, not for the voice
>that others hear.

That doesn't matter. EACH person speaks in such a way that the production is perceived as similar to (or deliberately different from) what is heard when others speak. The similarity is perceived only above a certain level (obviously it doesn't feel the same kinesthetically), and it's only at that level and above that there can be imitation. There can be great variation at the lower levels -- as you showed some time ago in remarking that the kinesthetic aspects of phoneme production vary widely from one person to another without altering the phoneme that is heard.

The reproduction of heard speech is more analogous to putting the fork "in" the water glass than to imitating the way the spectacles are held. The aspects of speech that are reproduced are those that are independent of point of view, while reproducing the orientation of the glasses specifically requires a transformation of the point of view.

>So the reference perceptions must be adjusted and established
>with respect to factors in others' responses to us.

They must be adjusted and established with respect to WHAT WE PERCEIVE OF factors in others' responses to us. There's no way for any individual to know what those factors "actually are."

>Do they recognize us as the kind of person we intend to present ourselves as?

Finding out how others recognize us is a matter of how we interpret our perceptions of their behavior. I have not noticed that people are particularly good at grasping the kind of person I perceive them as, nor am I very good at doing vice versa. Since effects on other people's perceptions are not actually observable, one can only guess what they are by comparing others' behavior with one's own: "If I acted that way, I would be perceiving the other person in a certain way." There is no way to verify this guess; one can only look for consistency in the evidence under one's model of the other person.

>Do they accord to us social roles and expectations that are
>acceptable or (better, of course) that we desire?

Nobody can "accord us a social role." Another person can behave as if we are playing a certain social role; that is quite independent of the social role we see ourselves playing. If these independent perceptions don't lead to actions and consequences compatible with both perceptions of social role, there will be conflict. If they do avoid conflict, there is still no reason to think that the two perceptions are the same. They are simply not in conflict.

>Do we get a balance of our desire for independence vs. our desire for
>acceptance and intimacy that is acceptable to us? Oops! Do we even
>know how to talk about independence vs. intimacy in PCT terms?

To "desire independence" is to misinterpret human nature according to PCT. The basic problem of human existence is that we ARE independent, whether we like it or not. Each person is free every waking movement of the day to set any goal whatsoever, including the goal of stepping off a cliff or stepping into traffic. All that prevents selecting such goals is a potential conflict with other goals -- nothing external has the slightest say in the matter.

Our actions have consequences outside us; we can't determine what those consequences are to be. But we can select any goals we please for those consequences: that they should occur, that they should have certain forms, that they should not occur.

Independence is not a problem, but intimacy is. All of us, I think, want to find evidence of a person like ourselves -- at least of the same basic nature -- behind the exterior appearances that others present to us. We do not want to be the only sentient being; we don't want to be alone. So we make great efforts to contact the person behind the facade, and we set goals that include doing what is necessary to achieve that contact -- if we can figure out what is necessary.

This is where a seeming problem with independence arises, because the other person has goals too, and conflict with them lessens the contact with the other person. We often find that in order to entice the other into revealing what we want to experience, we come into conflict with other goals of our own; to achieve this intimacy we must do things we don't want to do. But this is still an internal conflict, because nothing prevents us from giving up any goal, including the goal of knowing the other, or the goal of not doing certain things. The basic problem is achieving internal consistency among our goals.

>>That [is] what I intended to communicate -- that the illusion
>>of the glottal stop would appear only in speakers who had to
>>prevent themselves from extending the vowel into a diphthong
>>in order to imagine matching the heard vowel.

>Now why would a speaker of one language *have* to end a
>word-final stressed syllable with a diphthong or a consonant?
>Clearly, there is no physiological or neurological constraint,
>which would have to apply to speakers of all languages.

Of course not. I said they "had to"

>>prevent themselves from extending the vowel into a diphthong
>>in order to imagine matching the heard vowel.

The "in order to" supplies the reason they have to prevent themselves. The goal is to match the heard vowel. But the vowel they are used to hearing is extended into a dipthong, so as speakers of a given language that is their normal reference perception for the vowel. When they imagine saying "no" they imagine "no-ooo", but to say "no" in the Spanish way they must prevent the "ooo" from appearing, which creates an error in the normal control system that they have learned. At a higher level they experience this as cutting the vowel short, as in a glottal stop. People whose reference signal is the Spanish "no" would have no experience of the illusion. (My own illusion is that the "o" is nasalized -- sort of the opposite of a glottal stop).

>Rather, there is a socially instituted pattern on the basis of
>which speaker/hearer's expectations get set.

This is not an explanation, but a description. There is a social pattern, which means only that most of the people with whom the test person interacts speak in a similar manner. The existence of this pattern can't, by itself, have any effect on anyone. If a person's speech changes when the person interacts extensively with this group, that is only because selecting the group's way of speaking suits the person's purposes. There are plenty of cases to show that there is nothing necessary about the social influence on pronunciation -- just consider all the British receptionists who were hired for the neat accent, and make sure they don't lose it because their job depends on it. Or the movie and TV actors. Or the taxi drivers from Brooklyn. Or certain ex- Europeans who consider American phonemes ugly and refuse to adopt them. You can't base a serious law of speech acquisition on a generalization with so many exceptions.

Expectations are set, according to HPCT, by higher-level systems, not by perceptions. All that perceptions can do is provide raw material in memory, from which higher systems select the instances that fit their goals. Most people can imitate accents other than the ones they normally use, at least well enough to be speaking in an obviously different way. It is not the set of recorded perceptions of speech that determine the accent we choose, but our goals for our relationships to those around us. Some people think (almost always mistakenly) that the way to entertain people around them is to tell jokes in dialect. They may not achieve an entertained state in those who hear the joke, but that is certainly the goal, and the correct explanation for the way the joke-teller is pronouncing words at the moment. I have observed that people who tell such jokes also perceive that the others are entertained, which explains why they persist. They perceive themselves in a social role that matches what they want it to be; the others perceive something else, but their self- images prevent them from being frank.

>We feel a need to be consistent with one another, and out of this does
>develop impressive consistency in social conventions or norms associated
>with subpopulations to which people identify themselves as belonging.

That I can agree with. Particularly if you add that most people would identify themselves as belonging to a number of different subpopulations (e.g., a sports-loving homosexual stockbroker living with her Vietnamese husband in Bedford-Stuyvesant).

>But we don't have to identify subpopulations "objectively."
>People--the members of the social groups--do it for us. Or
>rather they do it for themselves, and for each other.

Do you mean that they walk up to you and tell you what subpopulation they belong to, so you don't have to do any interpreting of your own?

>And social dialect is an excellent index when you consider natives
>of a place like NYC, for example. (Don't take my word for it,
>have a look at Labov's methods and results in the papers in
>_Language in the Inner City_ f'rinstance.)

You can't run that logic both ways at once. If dialect is used to identify the subpopulation, then you can't prove that there is any shared characteristic other than the dialect, in which case all you've shown is that some people speak with that dialect and some don't. Either you pick the dialect as the discriminator, or some other set of characteristics like living in the inner city. To prove an association you'd have to show that not using the dialect means you don't live in the inner city, or that living in the inner city means that you have to speak with that dialect, and the two converse cases. If ALL the cases don't prove out every time, then all you've got is one of those statistical generalizations that's true of a population but not of any particular person. You certainly don't get any explanation for why one person speaks in a particular dialect on a particular occasion, and another doesn't.

>Sure, there are people that don't fit in, including obviously
>tourists and immigrants from other regions, but also effortful
>misfits of one sort or another.

So everyone is the same except for those who are different. This is, if you will excuse me, the usual kind of excuse that is made for statistical generalizations. The longer you examine the actual individuals involved, the more excuses are needed. This is why I have no use for "facts" about individual behavior that are based on statistics. They're facts, except when they're not.

>This is simply true: there are people that don't fit in the
>major social strata and groupings of any given community.
>Sometimes their seeming idiosyncrasies turn out to be negations
>or reversals of the norms of one or more existing groups, by
>which they deny membership. Sometimes they take on
>characteristics associated with one or more truly "foreign"
>groups (the guy from Topeka who talks with a put-on British
>accent; Bobby Zinnaman's knock-up of Ramblin' Jack Elliott in
>his emerging persona as Bob Dylan). And some people may truly
>bridge or alternate membership in more than one local group.

More of the same. Excuses, excuses. Let's face it: this way of understanding people just doesn't work. So-called social strata and groupings and norms exist mostly in the imagination of the statistician. They are perceptual prejudices. If you pick any individual at random and try to apply your generalization, you will find that you need a dozen excuses to explain why this person doesn't fit the category or behave like others whom you put in the same category. When you finish with the excuses (none of which is in fact investigated to see if it actually fits), all you have done is to discard from your sample all the individuals who don't fit the generalization. If you go through the population person by person in this way, you may end up with nobody left.

>Just because social psychology and sociology have often been
>inept handmaidens of government (and employers), trying to keep
>their hands aseptic and "objective", does not discredit social
>groupings as something that people identify and identify with or against.

I have not said that ONLY statisticians commit this sort of foolishness. People who informally identify with social groupings don't know any more about what is actually a social grouping than statisticians do -- if anything, their concepts are even looser and less consistent, and would lead to confusion even sooner under any kind of detailed questioning. Everyone imagines living in a certain kind of society, which doesn't actually exist and isn't like the society that the next person believes in. When any individual says "us" you haven't the least idea what he means, and neither, come to think of it, does he.

Society, in my view, is a continuum in which individuals vary along a hundred dimensions in ways independent of each other. Yet the members of the society like to look at superficial secondary characteristics like where one lives, how one dresses, how one speaks, what one does for a living, what one says aloud about beliefs, and so forth, and forms perceptions of groupings that have almost nothing to do with what makes each person that person and no other. Worst of all, the individuals then begin to deal with others whom they see as sharing these apparent patterns of secondary characteristics as if they were all right in the middle, and all alike. The use of formal statistics to define such groupings simply makes the prejudice harder to recognize for what it is.

>Has anyone here read Deborah Tannen's work on manners of
>speaking? _You Just Don't Understand_ is the most popular one,
>on stereotypically male-female differences in conversational
>style and problems that follow from them.

It's really a book about what's wrong with stereotyping -- that is, with speaking in terms of averages of populations -- isn't it?

Best,   Bill P.

Date: Wed Mar 03, 1993 12:46 pm PST
Subject: Lit. Survey

[From Allan Randall (930303.1415)]

Hi.
Just a note to say that if anyone is interested in seeing a draft copy of my literature survey, I'd be happy to mail them a copy. This is the first step in a new contract with Martin Taylor to do PCT modelling. Anyone who has an interest in the recent discussion on Ashby might like to see it, as it has a summary of Ashby's information theory development and applies it to PCT. It also has an overview of the genetic algorithm literature as applied to neural nets and my own (initial) attempt to draw a parallel between PCT and the general organization of the cerebral cortex. I'd be quite pleased to get some feedback from those more expert in these areas.

There are several pieces, available in different formats, so please specify the options you want when ordering, and don't forget your cheque or money order for $0.00: US funds ONLY please - do not send cash via email).

Survey + References: troff/refer, ASCII text or Postscript.
 (the paper and the references can be in different formats)

Bibliography: a larger set of references for a broader range of related areas w/ some annotations (refer/bib or ASCII text).

Powers Citation Index: all citations of W.T. Powers from the
Science Citations Index 1955-Aug.1992 (refer/bib or ASCII).

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Wed Mar 03, 1993 10:02 pm PST
MBX: randall@dciem.dciem.dnd.ca
Subject: Yes, Please

[From Dag Forssell (930303)

Dear Allan:

>Survey + References: troff/refer, ASCII text or Postscript.
> (the paper and the references can be in different formats)
>
>Bibliography: a larger set of references for a broader range of
> related areas w/ some annotations (refer/bib or ASCII text).
>
>Powers Citation Index: all citations of W.T. Powers from the
> Science Citations Index 1955-Aug.1992 (refer/bib or ASCII).

I am very much interested. Will appreciate a complete set.

I work with a postscript printer for all my fancy graphics work and have
a number of windows word processors and graphics programs. Have never
received any postscript files as E-mail, but would like to try. Why
postscript? Are there illustrations? What is refer/bib? If that is some
program name, I do not have it. Obviously have not heard of it. Please
use your own judgement.

Thanks!

Dag Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195   Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Thu Mar 04, 1993 10:10 am PST
Subject: Ad; Human Error

[from Gary Cziko 930304.1115 GMT]

This note from Rick Marken apparently made it to me but never made it to CSGnet
(for some unfathomable reason). So I am forwarding it for me. Hopefully Rick will
be normally connected again soon.

[from Rick Marken (930302.0800)]   Gary Cziko (930301.1540 GMT) -

>How about someone out there composing a nice ad?
>I can add the technical info to it.

Well, I made my little contribution. I would like to see an ad from Avery (if he has the time now) since he seems to be the one who is most aware of the goals of cognitive psychologists. He is also most unlikely to include in the ad claims about PCT being the most remarkable intellectual achievement since Eve took a bite out of the serpent's apple.

To whom it may concern (but particularly to Martin Taylor):

Does anyone have some very recent ('91 -'93)references to articles on "human error"? I was looking over some old attempts of mine to write a paper on "human error" (being a Human Factors engineer and all -- some would say that "reducing human error" is what Human Factors is all about) and I liked them so much that I've decided to present on this topic at the Durango meeting. The most recent work I'm familiar with is Norman's "The psychology of everyday things". Are there any more recent works on the subject that anyone knows about?

If anyone wants to start a net discussion on this subject, that would be great. I guess I would start with the question "how is it possible to deal with the notion of human error AT ALL without understanding the nature of control?"

Best Rick

Date: Thu Mar 04, 1993 2:10 pm PST
Subject: CLOSED LOOP business

From Greg Williams (930304) I sent the following last month:

>From Greg Williams (930217)

>1. Most CSG members will have received the Winter '93 issue of
>CLOSED LOOP by now. It inaugurates inclusion of "Research Reports"
>with two papers. I'd like to encourage submission of articles for
>possible inclusion in future issues (send manuscripts to me at
>460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A. or via e-mail),
>AND I'd like to have some volunteers to read/evaluate/critique
>such submissions. Perhaps we could set up an ongoing panel of
>reviewers? Publishing semi-peer-reviewed Research Reports in
>CLOSED LOOP is another step toward a JOURNAL OF LIVING CONTROL SYSTEMS.

I received only one reply, from Martin Taylor, who basically said that he thought having reviewers would be nice, but didn't volunteer for the job.

Yesterday I sent the following to Gary Cziko, privately:

>Martin T. posted privately to me that maybe CSG can get reviewers for
>CL Research Papers. That was the only reply I got to my suggestion that
>maybe we could do better than have just a horticulturist/journalist/
>mechanical engineer/roofbuilder as the sole gatekeeper for CL! How can
>we get folks to talk this up -- and, more important, volunteer???

Gary replied as follows:

>>We need someone willing to coordinate this, but it shouldn't be too
>>much of a bother if all or most could be done electronically. Anyone
>>with a paper to submit could send the abstract to the coordinator.
>>The coordinator would send the abstract out to CSGnet (blind review?)
>>with a request for reviewers to contact the coordinator who would then
>>forward to them the paper (hopefully electronically; figures are a
>>pain). Reviewers would write reviews, submit to coordinator who would
>>submit back to (principal) author. If opinion is in favor of
>>publication, article is sent to you. If not, but author wants to
>>publish anyway, reviewers are given a chance to polish up reviews for
>>publication (anonymously or otherwise) with paper.

>>I'm busy enough keeping the net floating and trying to finish my book,
>>so I would hope that someone else would volunteer as publication
>>coordinator.

Comments, suggestions, volunteers?

With help from Gary (thanks!), I'm actively seeking copyright clearance for using CSGnet posts in future issues of CLOSED LOOP. If you haven't yet sent me your permission, please consider doing it soon. In a private reply to my copyright permission inquiry, Oded Maler said:

>I will send you my permission (it's good to have alternative routes to
>eternity). I will only ask if possible to send me instead of royalties,
>a copy of the issue where my posts will appear.

Thank you, Oded. I hope others appreciate the access to an alternative (Gregory Bateson thought the ONLY) route, also. And, henceforth, I decree that all those whose posts are included in CLOSED LOOP and are not CSG members will receive free copies of the issues in which their posts are printed. (Ed, I'll work this out with you for each mailing.)

As ever,   Greg

Date: Thu Mar 04, 1993 8:40 pm PST
Subject: Journal proposals

[From Bill Powers (930304.2000)]   Greg Williams (930304) --

RE: the journal

I think that something like Gary Cziko's suggestion would work -- sending abstracts to CSGnet with a request for reviewers. Come to think of it, the job of a potential coordinator would be eased even more if those who wanted to submit articles to the journal simply posted abstracts to the net with a request for reviews. Anyone who wanted to do a review could then ask for a direct email copy, review it, and send the review to the coordinator.

Elaborating along the same lines, reviewers could send their reviews directly back to the person writing the article, raising objections, offering suggestions, etc..

In this way the reviewers would have to take responsibility for their opinions, and would interact directly with the author(s) at least once and possibly several times without an coordinator standing in the way or needing to relay things back and forth.

To keep this process from dragging out interminably, I suggest that when an abstract is posted, it be marked with a date. Reviewers would then be required to notify the author that they are taking on the review within a week of that date, and all interactions would cease (1,2,3?) months from the submission date. At that point authors could submit a final version to the coordinator and the reviewers, and the reviewers could send their final commentary to the coordinator. In this way the coordinator would receive a paper that has already been criticized and revised to the extent that the author is willing, and the reviewers would have thrashed out minor points with the author that the coordinator should not have to bother with. Really bad articles would have been strongly discouraged by the reviewers, so perhaps the authors would have the sense to withdraw them.

The coordinator, of course, would then edit the reviews for brevity and they would be published along with the article, signed.

There remains the question of how to choose what goes into the journal, even after this process. Perhaps we can create guidelines that will save us from committing the same crimes of which many of us have been victims. Some of them:

Crime 1: rejecting a paper because it disagrees with the official point of view or the reviewer's own opinions.

Crime 2: rejecting a paper because the reviewer did not understand it. Of course if NO reviewers understand it, one must ask why.

Crime 3: rejecting a paper because it challenges the findings of other papers that have been published in the same journal (the in-group syndrome).

Crime 4: rejecting a paper without reference to its substance, on grounds that have nothing to do with its scientific soundness.

Crime 4, of course, sort of wraps up the other three.

My own view is that we should not accept papers that are not principally about PCT but simply use PCT terminology as a way of talking about other subjects. All articles should teach something or demonstrate something about PCT. I'm not even sure that the "worlds" article by Tom Bourbon and me would qualify under this heading, or that Rick Marken's article would fit (although both articles did develop some relationships between PCT and specific aspects of conventional psychology -- we'll have to think about that).

What I'm trying to say is that this journal should be the basic research journal for the study of living control systems, and not attempt also to evangelize. The articles should enhance our understanding of living control systems or improve on the theory we use for that understanding. The journal should be the place for people who are already committed to the exploration of PCT to publish their ideas and results, and to find out what others have thought and found. It is not the place to convince psychologists or anyone else that they ought to pay attention to PCT -- that, I fear, we must continue to try to do through submissions to conventional journals.

I feel that all propositions about PCT presented in the journal should be accompanied by replicable demonstrations, and should adhere to the statistical concepts we have worked out in several years on this net. No use of group statistics to derive laws of individual behavior. All propositions to be tied to a specific testable model. Replicability of data to be some negotiable figure corresponding to correlations in the 90s (and confidence levels of $p < 0.00001$ or so). Exceptions to be very carefully scrutinized.

Suppose that someone wants to publish an article showing that behavior is NOT the control of perception. As I interpret my own words above, it would qualify for publication if it were soundly argued and accompanied by clear demonstrations of high reliability. We must never lose sight of the fact that the basic thing we want to know about living systems is whether they are in fact control systems; this question must remain forever open if we want to claim the name of science.

But suppose that this article is based on a philosophical por theoretical argument, or a mathematical theorem, or some fundamental principle of physics. I would vote against publishing it. Without a clear demonstration of the applicability of the argument to some specific set of real observations, such an article would be an exercise in pure reason, and that is the opposite of what I would like to see in THIS journal.

I will argue strongly against excluding articles that are written at a low level of sophistication. There are many facets to PCT, not all of them mathematical or amenable to computer demonstrations. There are formal uses of PCT and informal ones. People from many fields of interest and with quite different levels of technical education are contributing to PCT. If you say "Every time I contradict a person's statement about his own personality, that person makes an objection intended to counteract my comment," that is data about PCT and the nature of a living control system. It is replicable with high reliability. Simply by being tried, that experiment tests the generality of the claims of PCT. I would recommend publishing it even if the study is very simple and very simply reported.

I think that the review process I described above would encourage articles from people with good ideas but little experience with rigorous theorizing. Those who submit articles would benefit from submitting abstracts for reviews while they are still designing their studies; reviewers with more experience can help them see how to convert a vague idea into something actually testable, and can help them select relevant (if narrower) tests and write a good article on the results. There's no reason for reviewers simply to sit back and snipe. A person with some experience with PCT can pick out of a too-general concept many specific concepts that could be tested; if the author is willing to do the work, an interesting and useful contribution to PCT could result. Reviewers can be mentors as well as critics.

Finally, one of my intentions (which I invite others on the net to help support) is to make sure that there is no bread-and- butter way of getting published just by turning a crank. The conventional journals are choked with articles that are simply another study of the effect of A on B, the author having got lucky and found a significant, but trivial, effect. If we want high quality results to be reported, we must set the standards so that only high quality results are accepted. This means that the first run-through of an experiment will probably not yield results publishable here. Only by testing and refining hypotheses, then testing and refining again, is it possible to get results with very small errors, very high replicability. We may have a slim journal if we enforce such standards, but by God every word in it will be worth reading -- when published, and ten years later.

Best to all, Bill P.

Date: Fri Mar 05, 1993 7:52 am PST
Subject: journal volunteers

(from Joel Judd 930305.11:30)

I would like to volunteer to contribute to the organization of the journal. I'd be willing to be the coordinator if others feel I have the wherewithal to do it, and I'm not too far out of the e-mail, and snail mail, mainstream to make coordinating a logistical pain.

With respect to reviewING, would it be feasible to periodically publish a list of netters and their declared interests/specialties? That way abstracts could be addressed to them (or to them through the net). Or is something like this already available?

Date: Fri Mar 05, 1993 8:20 am PST
Subject: Allan Randall's review and bibliography

[From Bill Powers (930305.0900)]

Allan Randall (930304) --

I have received Allan Randall's review of PCT and annotated bibliography. It is a major effort and of great value. The bibliography is actually available, with keyword search, via telnet! One warning: the entire package is about 360 K characters long.

Allan, thanks for your contribution to the future of PCT.

Bill P.

Date: Fri Mar 05, 1993 9:53 am PST
Subject: Re: Journal proposals

[Martin Taylor 930305 12:15]
(Bill Powers 930304.2000)

I suppose I shouldn't argue with Bill on the topic of his baby growing up, but there are a few features of his commentary that I find disturbing.

If there were a Journal of Living Systems, based around PCT as the foundational idea, then it should incorporate all aspects of development of PCT. That includes not only experiments, but theory and explanation.

I really don't like Bill's proposal:

>I feel that all propositions about PCT presented in the journal
>should be accompanied by replicable demonstrations, and should
>adhere to the statistical concepts we have worked out in several
>years on this net. No use of group statistics to derive laws of
>individual behavior. All propositions to be tied to a specific
>testable model. Replicability of data to be some negotiable
>figure corresponding to correlations in the 90s (and confidence
>levels of $p < 0.00001$ or so). Exceptions to be very carefully
>scrutinized.

How do you make a "reliable demonstration" of how a proposed Decision Making Entity interacts with one or a set of ECSs? Or of whether a program-type ECS can be segregated from the sequency-type ECSs that presumably support it? How do you make a reliable demonstration that PCT-based methods of teaching a second language work? As Joel Judd once remarked, you can't even replicate such things, even if you can experiment with them.

The theoretical discussion of what could be, must be, and cannot be, is as important to PCT as demonstrations of one or two level models that fit real data. Those are just demonstrations. They even mislead many people as to where the power of PCT lies.

And I think Bill violates his own principles (committing Crime 2 or Crime 1) by his dogmatic rejections of group statistics. If maxim 1 of PCT is that you can't do anything about information you can't get, maxim 2 must be that you can use information you can get. It IS valuable to the educator to have 70% of the children rather than 30% understand say Pythagoras' Theorem. A change in teaching method that achieves this is a good thing. Can you (have you) achieved 99% accuracy in modelling in a situation where a 4-level hierarchy is necessary? Do you expect to?

>What I'm trying to say is that this journal should be the basic
>research journal for the study of living control systems, and not
>attempt also to evangelize. The articles should enhance our
>understanding of living control systems or improve on the theory
>we use for that understanding.

If the theory is any good, all the reports will evangelize. Review articles serve to bring together results, and again if they are good, the article evangelized. Tutorial articles do the same. The articles by you and Rick enhance people's understanding. I would object to dishonest attempts to evangelize, but "the truth shall set you free."

The proposal that articles might improve on the theory seems to contradict the suggestion that all propositions should be accompanied by reliable demonstrations. As does:

>But suppose that this article is based on a philosophical por
>theoretical argument, or a mathematical theorem, or some
>fundamental principle of physics. I would vote against publishing
>it. Without a clear demonstration of the applicability of the
>argument to some specific set of real observations, such an
>article would be an exercise in pure reason, and that is the
>opposite of what I would like to see in THIS journal.

So you don't want theoretical articles. I do. But it's your baby, so you get the casting vote. Fundamental principles of mathematics or physics won't go away, and need no demonstration. The demonstration they need is that they apply to the situation at hand.

On reviews: in principle I like the electronic review idea, but particularly in PCT diagrams are very important. Many of the ongoing arguments could be resolved if one high-quality picture were available for both discussants to use. You couldn't visualize my 2-level 2-ECS structure for talking about conflict, and that was very trivial.

As I just posted to Greg, I see no need for volunteer reviewers. I receive a stream of articles for review, sometimes from journals I never heard of. Never have I volunteered as a reviewer, but if time permits, I usually accept the job. I think most scientists are the same. Why ask for volunteers?

Martin

Date: Fri Mar 05, 1993 10:54 am PST
Subject: Journal proposals - Taylor and Powers et al

****************** FROM CHUCK TUCKER 930305 *******************

I have long been a critic of how the journals present us with what we ought to know about a subjectmatter. From my experiences in all aspects of the processes of "journal work" I came to the comclusion that what we should have is what we have on this net; anyone who is on the list can write anything they care to write and everyone else has the opportunity to comment as they see fit. There are no official gatekeepers to say you can or can't publish this or that. But as we note even those with the best of intentions, highest of standards, greatest of concern for mankind, and most intelligent get into the mode of "specifying what should be in or out" when the question arises "What would be the best way to operate this journal?" I CAN RECOGNIZE THIS BECAUSE I HAVE DONE IT AND SEE IT IN MYSELF [see the recent exchange between Bruce and Bill on speaking for an elaboration on this action]. So my suggestion would be not a compromise of the proposals by Taylor and Powers but something more like we are doing at the moment on the net with a bit more specification (I think Greg actually comes close to what I as proposing but the formatting and style of CLOSED LOOP do not clearly indicate it to all but the most careful reader and those who did the writing).

Put in CL those pieces which involve a major statement along with the reviews, comments on reviews and a final revise and resubmit by the initial author(s). There are several journals that actually use this style now (I can't think of their names at this moment). This would allow for the exchange and education to go on as both Martin and Bill want (as do all of us) and yet not set up a situation where we have a set of approved ideas (I don't believe that Bill wants this type of situation either although one reading of his comments has this appearance)

Under these specifications, the major choice that would be made would be "Which of these many discussions would be in CL?" Hopefully that choice would be as easy as it is now and made on the same basis. The other way to decide would be that if no one was interested in commenting on a statement then it would not be considered; it dies of lack of being seconded. To guard against a piece being "missed" those who wanted their statement considered could put it on the fileserver so we could all retrieve it by GOPHER (or another means or get it by smail from the author who would remind us of it). If there was enough interest to have a series of comments and a revision (or restatement) by the author then it would be considered for inclusion. I think, finally, the author should agree to the request that the exchange be "published" in CL knowing full well how important such a publication will be to his or her future on the planet {{ :-D & :-J & ;-D }}

I do NOT give permission that the above statement be published in CL!!

Best regards to all,   Chuck

Date: Fri Mar 05, 1993 2:51 pm PST
Subject: Re: journal

[From Bill Powers (930305.1330)]   Joel Judd (930305.1130) --

Nice to know we have at least one volunteer for coordinator. I think we're not quite ready for that yet -- unless we are. Why not start by keeping a record of all the suggestions?

Martin Taylor (930305.1215) --

Well, I didn't expect everyone to like my proposal. And by the way, I dissasociate myself from your statement

>But it's your baby, so you get the casting vote.

If that's the case, I won't even vote. What I know about control theory doesn't make me an authority on journals, and even if it did, I would expect all good PCTers to insist on controlling for what they want, not for what I want. I have noticed that this is generally what they do anyway, whether I give them permission or not.

>If there were a Journal of Living Systems, based around PCT as
>the foundational idea, then it should incorporate all aspects
>of development of PCT. That includes not only experiments, but
>theory and explanation.

I can't disagree with that. Experiments, theory, and explanation are the toolkit of science. But I don't think you can leave experiments out and still have science. My opinion.

>How do you make a "reliable demonstration" of how a proposed
>Decision Making Entity interacts with one or a set of ECSs?

So far, you don't. That's why we toss such ideas and many others around on the net. There's no point in taking them seriously until we can think of a way to show that no other explanation appears possible. I wouldn't submit an article on the DME, pro or con, in the present state of the idea.

>Or of whether a program-type ECS can be segregated from the
>sequency-type ECSs that presumably support it?

Ditto. Until someone thinks of an experiment that will distinguish between these concepts, they're just concepts to me. Worth thinking about, but not worth publishing for posterity.

>How do you make a reliable demonstration that PCT-based methods
>of teaching a second language work? As Joel Judd once
>remarked, you can't even replicate such things, even if you can
>experiment with them.

Then I guess we don't yet have anything publishable on that subject. If we don't have anything to say, should we publish anyway?

>The theoretical discussion of what could be, must be, and
>cannot be, is as important to PCT as demonstrations of one or
>two level models that fit real data.

I totally, heartily, and vehemently disagree. Nature pays no attention to what we think could be, must be, and cannot be. All such musings are based on incomplete premises and inadequate knowledge. Maybe somewhere in the realm of such ideas are things that ARE. But without a demonstration, a test against what is observed, there is no basis for either accepting or rejecting such notions.

>Those are just demonstrations. They even mislead many people
>as to where the power of PCT lies.

Insofar as PCT has any power, its power lies in its ability to explain phenomena. If you don't compare what PCT says against phenomena, it has no power at all.

>And I think Bill violates his own principles (committing Crime
>2 or Crime 1) by his dogmatic rejections of group statistics.
>If maxim 1 of PCT is that you can't do anything about
>information you can't get, maxim 2 must be that you can use
>information you can get. It IS valuable to the educator to
>have 70% of the children rather than 30% understand say
>Pythagoras' Theorem. A change in teaching method that achieves
>this is a good thing.

It's a good thing for the teacher at evaluation time, and for 7 out of 10 of the students, although you can't say in advance WHICH 7. Insurance companies, casinos, doctors, teachers, and generals benefit greatly from knowing the odds. But knowledge of the odds tells us nothing about why 30 per cent of the students didn't learn the theorem. Population measures have their uses, but helping to build a model of individual human behavior isn't one of them.

>Can you (have you) achieved 99% accuracy in modelling in
>a situation where a 4-level hierarchy is necessary?

No.

>Do you expect to?

I expect someone to. When they do, I hope to read about it in the journal.

>The proposal that articles might improve on the theory seems to
>contradict the suggestion that all propositions should be

>accompanied by reliable demonstrations.

I don't see the contradiction. If an article proposes a change in the theory, and demonstrates that it is necessary in order to account for phenomena in an experiment, surely the theory has been improved by this.

>So you don't want theoretical articles. I do.

I don't want articles that are JUST theoretical. That doesn't mean they won't appear in the journal. It just means that I won't be much interested in them. I'm of the school that says the way to find out if men have more teeth than women is to open some mouths and count.

>Fundamental principles of mathematics or physics won't go away,
>and need no demonstration. The demonstration they need is that
>they apply to the situation at hand.

Fine, then demonstrate that they DO apply to the situation at hand. Just claiming that they do, or ought to, isn't enough for me. The way in which fundamental principles are applied always depends on premises and assumptions; that is the great loophole in all pure theory.

By the way, the only reason that we have reliable principles of physics is that someone DID demonstrate them.

>On reviews: in principle I like the electronic review idea,
>but particularly in PCT diagrams are very important.

That's a touchy point. I don't see how we can avoid snail mail until somebody finds a simple universally usable way to transmit GOOD diagrams over the net.

>As I just posted to Greg, I see no need for volunteer reviewers. I receive
>a stream of articles for review, sometimes from journals I never heard of.

I was trying to accomplish two things: first, to reduce the load on the coordinator; second, to remove from the coordinator the power to decide who gets reviewed and who does not. Every little reduction in power helps. By setting up the review process as a direct relationship between authors and reviewers, we also make the connection more personal and the reviewers more answerable for what they say.

Chuck Tucker (930305) --

The idea of publishing everything submitted, as we do on the net, is appealing. But I don't think we can afford it. The CSG's annual membership fee of $45, plus the student fee of $5, covers the cost of printing and distributing Closed Loop four times per year -- but it wouldn't support a much larger effort. And we can't go on forever expecting Greg Williams to provide his superb skills and valuable time for next to nothing.

Since we don't have the resources to publish everything submitted, we have to choose somehow. My mechanism for choosing is to raise the standards and get rid of the sort of fluff that fills the conventional journals. CSGnet already publishes everything that anyone submits, so we don't have to give that up.

After we've batted these ideas around for a while longer, I'd like to see a group effort to establish criteria for publishability that don't violate our sense of scientific honesty and fairness. I think we can think of guidelines that emphasize good science and eliminate politics. But we'll see.

Mary has had an excellent idea, which is to set very high standards even if this means only one or two articles per issue, or none. The remainder will be the usual fare from the net. The title of the journal, Mary suggests, would be Closed Loop: a journal of living control systems (a nice ambiguity). Just add the subtitle and we're there.

Best to all, Bill P.

Date: Fri Mar 05, 1993 8:56 pm PST
Subject: Starting Over?

[From Rick Marken (930304.1400)]

There are two kinds of psychologists who embrace Perceptual Control Theory (PCT). One kind (including myself) believes that we should start psychology over from scratch based on the principles of PCT. The other kind (probably the majority, unfortunately) believes that PCT should build on the psychology that has already been done. I think it is clear that many non-PCT psychologists are turned-off by the revolutionary implications of PCT; so PCTers of the first kind are not likely to win many recruits from the ranks of conventional psychology. On the other hand, it is also clear that many non-revolutionary PCTers (like Carver and Scheier), by not abandoning convention, have had to abandon some of the fundemental tenets of PCT; so PCTers of the second kind are likely to recruit conventional psychologists to a model with the right name (PCT) and the wrong organization.

PCTers of the first kind argue that MOST conventional psychological research can be ignored because research carried out in the framework of the wrong model is not likely to provide the kind of information that is needed by a PCT modeller. PCTers of the second kind argue that data is data; if the data can be handled by a conventional model then PCT should be able to handle it too (if PCT is really worth its salt); moreover, by ignoring conventional data PCTers of the first kind give the impression of being evasive.

So here is a chance for the two kinds of PCTer to slug it out again. The most recent issue of American Scientist (Jan/Feb 1993, v.81) contains an article by Harold Pashler (a psychologist at UC San Diego) on "Doing two things at the same time". It is straight cognitive psychology; no mixer, no ice. It is also apparently "state-of-the-art" because American Scientist (I think) tries to present the latest, the best and the most interesting science in various fields to a general scientific (or intelligent lay) audience. Thus, I submit that it might be to discuss this article from the points of view of the two kinds of PCTer.

The experiments described are VERY simple (from a conventional perspective); the subject is presented with a higher or low tone (S1) followed, after a variable delay, by one of two possible letters (S2). The subject is to say, as quickly as possible, whether the tone was high or low and press a button to indicate whether the letter was an A or B. So the subject must do two things -- respond to S1 (tone) and S2 (letter). As the interval between S1 and S2 increases, the time to respond to S2 decreases. This result (and the results of other experiments) is consistent with the idea that there is a "bottleneck" in the "processing" of the two stimuli; the "processor" must finish processing S1 (deciding which response to make, say the auther) before it can START processing S2. The model of these results treats the stimuli as "clients" that must be delt with one at a time by the processor.

Now, how would PCT explain these results?

Best Rick

Date: Sat Mar 06, 1993 3:45 am PST
Subject: More CLOSED LOOP business

Actually, from Greg Williams (930306)

I tried to post the following to Ray Jackson directly, but it bounced back. I hope the whole net won't mind seeing it. Ray, how do ATT Mail addresses work?

From Greg Williams (930205 - direct)

>As you're looking for resources to help out on Closed Loop issues, I'd
>like to offer my services as an editor. I'm an ex-English teacher, so
>proofing comes fairly easily for me; I'd also be glad to make readibility
>comments if you'd like.
>I'm quite short on time until the end of April (finishing my Masters),
>but I'll have more time after that.
>Let me know if this is something that could help you out.

Thanks very much for the offer, Ray. I do have a project in mind which you could help with if you want. Last year there was an extended discussion of interpersonal manipulation and PCT on the net, mainly between Bill Powers and myself. By the time the debate wound down, several netters had asked that it be edited and printed in CL. I've been putting that off, mainly because I don't like the idea of editing my (sometimes heated) arguments involving me as a main participant. Would you consider editing the material for CL? I know it needs a lot of cutting; there was significant repetition, not to mention several blind alleys. A likely issue for publication would be the one due out ca. July 15th (which will probably go out ca. August 1, if history is any guide).

Let me know what you think. As ever, Greg

P.S. I hope you will be sending copyright clearance for your posts.

Date: Sat Mar 06, 1993 8:16 am PST
Subject: A challenge to information theorists

[From Bill Powers (930306.0700)]

It has been said that information theory contains the real meat of control theory. If this is true, then as Rick Marken has said information theory ought to tell us how to improve our models of control systems. It would also follow that information theory should lead to correct predictions about control processes, or at least not contradict what is observed in simple experiments. I believe that I have an example of a control situation in which information theory will have difficulty doing that.

The situation is simply an implementation of the general diagrams W. Ross Ashby used to describe disturbance-driven and error-driven regulation. These two situations, and combinations of them, are easy to set up on a computer with a human subject to provide quantitative data. From quantitative data, the information theorist should be able to calculate the amount of information (or variety) represented by all the variables, and from the principles of information theory (or the Law of Requisite Variety) show that the observed degree of regulation follows from the theory.

I am issuing a formal challenge to information theorists. I will program this experimental situation and offer the program free to any information theorist who wants to use it to run the experiment, or will run experiments on a human subject (myself) and provide the raw data on disk or on the net in ASCII-numerical form, for analysis. I don't want to waste my time preparing this experiment if there are no takers, so before I do it I want to know who is accepting the challenge, if anyone. I believe that people on the net have seen enough of my programming output to know that I can produce a program that will do what I claim. But I will allow challengers to write their own programs and run their own experiments, as long as they can convince me that the program satisfies the conditions to be described here.

The experiment:

The basic experiment involves a disturbance that affects an essential variable through a transmission channel of specified properties, and a regulator that acts through the same transmission channel on the same essential variable.

The experiment involves three conditions:

1. The regulating person R has continuously available direct information about the state of the essential variable AND the state of the disturbing variable, and acts to keep the variations in the essential variable as small as possible.

2. The information directly available from the state of the essential variable is denied to the regulating person.

3. The information directly available from the state of the disturbing variable is denied to the regulating person.

Condition 2 corresponds to Ashby's diagram of a disturbance-driven regulator, which I call a compensator:

```
D ------> T -------> E
|         |
|         R
|        /
 ---->-
```

Condition 3 corresponds to Ashby's diagram of an error-driven regulator, which I call a control system:

```
          D ------> T -------> E
                    |          |
                    R          |
                     \         |
                       --<-----
```

Condition 1 combines these situations:

```
          D ------> T -------> E
          |         |          |
          |         R          |
          |       /  \         |
           --->--      --<-----
```

As I understand the information-theoretic approach, information passes from D to E via the transmission channel T. To the extent that E is regulated, the actions of R via T are such that some information is kept from being transmitted to E, thus reducing the information content of the variations in E. This is equivalent to regulating E.

In condition 2, the Regulator receives information directly from D alone, and so in principle could produce outputs affecting T that completely block the flow of information, thus permitting the information in E to be reduced to zero and achieving perfect regulation.

In condition 3, the regulator receives information directly from E alone. The better the regulation, the less information is available to R from E, because the action of R via T is diminishing the information flow from D to E. As a result, perfect regulation is not possible because perfect regulation would reduce the information content in variations of E to zero, preventing any information from passing from E to R.

In condition 1, the regulator receives information from both D and E. In principle, perfect regulation should be possible because of the information received from D. The information received from E is redundant.

Experimental details

The disturbance D is the output of a pseudo-random-number generator passed through a three-stage low-pass filter, each stage being a simple time constant of 0.3 second. The transmission channel T is a simple noise-free adder which adds D to the output of the regulator R, passing the result to E. The essential variable E is a visual display, a moveable object on the screen, the vertical position of which relative to stationary reference marks is proportional to the output of T. A second moveable object of the same kind, located adjacent to the path of E, also moves vertically in a way proportional to the magnitude of D, so that D is represented visually in the same way that E is represented. D is scaled so its position accurately represents its effects on E, with the zero point corresponding to the position of the reference marks and to zero effect on E.

The time-course of both variables -- D and E -- is sampled 50 to 80 times per second (depending on the display characteristics) and stored in arrays sufficiently long to save the data from a 2- minute experimental run (3000 to 4800 data points per table). In addition, the output of the regulating person is saved in a third table, containing a record of the mouse positions during the run and thus the person's contribution to the state of E. The tables are saved to disk after a run, in ASCII format with triples of decimal numbers separated by spaces and terminated by a carriage- return-line-feed (\x0d\x0a).

The task of the participant is to use a mouse to maintain the object E exactly even with the reference marks. The subject is allowed to practice on condition 1 as long as necessary to reach and maintain a minimum in the RMS variations of E averaged over 1 minute.

Then condition 2 is established by turning off the display of the state of D, and the participant is allowed to continue practicing until a minimum in the RMS variations of E is achieved.

Finally, condition 3 is established by turning the display of E back on, and turning the display of D off. Once again, practice is allowed until a minimum in the RMS variations of E is achieved.

The runs for each condition are saved in separate files: cond.1, cond.2, and cond.3.

The predictions.

Because absolute information content is hard to specify, the challenge issued here is concerned only with relative measures. The problem is to form a theoretical ranking of the goodness of regulation for cases 2 and 3 above, on the basis of either information theory or PCT.

The PCT analysis predicts that condition 3 will provide the best regulation, condition 1 either the same degree of regulation or slightly worse, and condition 2 a degree of regulation that is worst of all by a large margin. In other words, between conditions 2 and 3, PCT predicts that regulation will be unequivocally BEST when the participant gets the LEAST information about the actual state of the disturbance D -- condition 3.

I believe that information theory will make the opposite prediction: that condition 2 will provide better regulation than condition 3. But I will leave it up to information theorists to derive and explain the authoritative prediction. I expect them to make their prediction, as I do, before the experiment is run.

The experimental data should then settle the question of the relative power of PCT and information theory.

Best to all, Bill P.

Date: Sat Mar 06, 1993 1:26 pm PST
Subject: conference report

[Avery.Andrews 930306]

For general information. Note the explicit swipe at control theory (probably just cybernetics, not bp in particular)

--

Workshop on Computational Theories of Interaction and Agency

Philip E. Agre
Department of Communication
University of California, San Diego
La Jolla, California 92093-0503

phone: (619) 534-6328
fax: (619) 534-7315
internet: pagre@weber.ucsd.edu

About fifty researchers with a broad range of interests gathered at the University of Chicago on February 20th and 21st 1993 for the Workshop on Computational Theories of Interaction and Agency. This meeting brought together the authors who submitted papers to the AI Journal special issue of the same name that Phil Agre and Stan Rosenschein are currently editing, along with several of their students and a few other of the usual suspects. Its purpose was to discuss the submitted papers, both to help the authors improve them and to assess the current situation and future prospects for research in this area. Its format called for an hour to be spent in focused discussion of each paper, in two tracks. The participation of graduate students was a particularly important part of this process, and a grant from AAAI allowed several students to attend. The workshop was also supported by Philips Research Laboratories New York.

The general subject of the special issue is the development of principled characterizations of agent-environment interactions, and the use of these characterizations in explaining existing agents and designing new ones. The idea might be illustrated through the example of classical control theory, which models interactions between a controller and a plant in terms of a differential equation whose properties are the subject of considerable analysis. Such equations may not be useful for analyzing qualitatively complex agents and environments, but a remarkable variety of computational research has recently been focused on the broader idea of describing and analyzing the interactions between particular categories of agents and environments. It would be impossible to summarize the intense and wide-ranging discussions provoked by the papers at the workshop. At the risk of slighting the majority of authors through omission, I will report on a small sample of representative points.

Randy Beer's paper, "A dynamical systems perspective on autonomous agents", provides a framework, based on dynamical systems theory, for analyzing the interactions between robots and their environments. The general idea is that the robot and environment, understood as a point in a large-dimensional parameter space, trace a certain complex path as they interact. Dynamical systems theory provides a vocabulary for talking about these trajectories. Beer focuses on a case drawn from his own work, in which a robotic insect is programmed to walk by a genetic algorithm that adjusts the parameters on a neural net that control its legs and sensors. Discussion focused on the power of a dynamical systems framework to analyze broader classes of agents and environments.

Bruce Donald's paper, "On information invariants in robots", provides a formal analysis of the information complexity of robots' interactions with one another and with their environment, given particular assumptions about their sensors. For example, if one robot is supposed to follow another, the designer is faced with a trade-off between the amount of communication between the robots, the length of the path they follow (for example when going around blind corners), and the power of their sensors. The formalism captures the intuition that something is conserved as the designers shifts from one design to another. This "something" can be measured in bit-seconds, that is, the amount of information that is communicated among the various devices during the performance of the task. The theory draws heavily on concepts from robotics. Discussion focused on the ways in which the theory might be generalized so that its concepts of information, for example, might be applied to a wider category of tasks in which the geometric arrangement of devices is not a central issue.

Kris Hammond, Tim Converse, and Josh Grass's paper, "The stabilization of environments", concerns the ways in which people actively modify their environments in order to simplify the cognitive tasks involved in living in them. This topic is part of their larger concern with what they call long-term activity, as opposed to the one-shot stretches of action envisioned by classical planning systems. Designing agents for long-term activity requires, among other things, an understanding of the long-term relationships maintained between an agent and its customary environment, as well as an understanding of how an architecture might organize and deploy the wide range of knowledge necessary for such activity. Hammond, Converse, and Grass's algorithm is based on case-based planning. They observe that a case-based system works best when the agent tends to encounter the same situations over and over, and this observation provides an engineering justification for the active stabilization of the environment. They describe an architecture that modifies its plans in response to various indications that some kind of active stabilization might be desirable. Discussion focused on the relationship between this kind of active stabilization and the more emergent kinds of stability that emerge when one's activities in a given environment settle into a routine.

Yoav Shoham and Moshe Tennenholtz's paper, "On social laws for artificial agent societies: Off-line design", concerns the ways in which multiple agents sharing an environment can avoid interfering with one another by sticking to a predetermined set of "social laws" constraining their behavior. The design of these laws faces a trade-off: the more freedom is left to the agents, the more they can step on each other's toes. As a case study, they provide a set of traffic laws that allow a set of agents to coexist in a rectilinear grid through various conventions about when and where they move. They are able to prove that these conventions do indeed prevent the agents from colliding while simultaneously allowing close to optimal behavior. They also provide a general framework for deriving these social laws and prove some complexity results about the process of doing so. Discussion focused on the relationship between Shoham and Tennenholtz's proposals and the other ways in which multiple agents can coordinate their activities, for example through local negotiation or centralized control.

The workshop discussed a total of twenty-five papers, of which four have been mentioned above. It was altogether remarkable how thoroughly and usefully the workshop participants were able to discuss the papers, given that their backgrounds included fields as diverse as natural language processing, control theory, planning, logic, vision and robotics, neural networks, and philosophy. This is an excellent portent for the future of research in this area. The special issue itself, which is expected to appear toward the end of this year, will provide more details.

Date: Sat Mar 06, 1993 6:33 pm PST
Subject: Conference report

[From Rick Marken (maybe)]

Avery Andrews

The report on the conference was a hoot. Obviously, these people know all the right things to say. Oof da.

Best   Rick

Date: Sat Mar 06, 1993 7:12 pm PST
Subject: ERRORS & PCT - RKC

[From Bob Clark (930306.1930 EST)]

"Errors" ("Mistakes" etc) occur. "Correction of Errors" is very common and is an important aspect of behavior.

Errors are of many kinds and occur for many reasons. Rather than a catalog of types and reasons, the following is an analysis of the general relations between "Error" and Perceptual Control Theory.

Recent discussions on the Net have considered suggestions that "Feedback Systems" consist of "Error Correcting Systems." In criticizing these suggestions, comments have emphasized that feedback control systems neither "Control Errors" nor "Correct Errors." Rather, they "Oppose Disturbances" of a "Controlled Variable," maintaining its condition as specified by their Reference Signals. The detection and correction of errors by Feedback Control Systems were not included in those discussion.

Both of these views define "Error" as the discrepancy between current and intended conditions of a perceptual signal. For an Error to be perceived as an Error, both conditions of the signal must be available somewhere at the same time. "Control of Error" consists of reducing the size of the discrepancy.

In the discussions noted above, both the "Error Correcting System" and the "Feedback Control System" assume a single control system, rather than a combination of control systems. The suggested "Error Correcting System" is not defined in terms of the origins and nature of the variables involved. The details are left to imagination, but limited to a single system, with "Compensatory" elements suggested. In contrast, a PCT System includes specification of perceptual variables, reference levels, environmental variables, error signals and their interacting components. Since the PCT System operates as a unit, no "Error" can be perceived until the system has completed its action. Thus a single PCT System cannot even perceive, much less "correct," its own "error." Therefore a single feedback control system cannot act as an "Error Correcting System."

However, as above, Error Correction is a very common and important form of behavior. These activities can be included by the addition of higher levels, as in the Hierarchical Control Theory. Higher levels can provide for both the perception and control of errors.

THE FOLLOWING VIEWPOINTS ARE SUGGESTED IN MY POST: Bob Clark (930301.1730)

The "USER'S" view.
This is the viewpoint of higher levels as they use lower levels to satisfy their own reference levels. With at least two levels, the higher level perceives the current condition of the variable(s) controlled by the lower level. The higher level also perceives the memory that provides the reference signals to the lower level. The higher level has a longer time-scale, so that it observes the situation as the lower level's action is completed, and compares it to the condition that was intended. The "intended condition" would have come from a still higher level. The User may use that level, or still higher levels, for "corrective action."

The ENGINEER'S" view.
This view lacks direct access to events inside the human system. However a "hardware" system having conceptually similar components can be used for analysis. Thus it is possible, in concept and in "hardware," to examine the operation of the system. Thus events can be observed during the reduction of the discrepancy between the current and required conditions of the controlled variable.

These events may look like "correcting the error." But it is merely the ordinary operation of a feedback control system. An Error does not exist until the system has completed its control action. In addition, the "discrepancy" cannot be perceived by the system itself -- the only perceptual variable existing within the system is the feedback signal. Other "signals" (e.g., the output signal) within the system are not available outside the system. They are not "Perceptual Variables."

The Engineer can analyze the effects of changing the system's parameters.

The Engineers's view can easily include the interconnections and other relationships between levels, as in Hierarchical Perceptual Control Theory. Thus it can also be used to analyze the User's view.

Since these events have the attention (of the Engineer's DME), they are being recorded as they are occurring.

The "OUTSIDER'S" or "Observer's" view.
If an error, or any "corrective action," occurs, that event may be observed -- that is, get the Observer's attention. That is, the Observer (his DME) pays attention to the events. The observations are interpreted in terms of those theories available to him. These events are also recorded in the Observer's memory because they have the attention his DME.

The "EXPERIMENTER'S" view.
The Experimenter also has the view of an Outsider, but, as an Experimenter, he may make some changes in the situation. He may change the environmental conditions. He may give information or instructions to the Subject. For interpretation ("Understanding") of the subsequent events, he also uses the theories available from his memory.

The "THEORIST'S" view.
The Theorist recognizes the validity of each of these views and relates them to his own experiences and explanations. He examines proposed theoretical descriptions and explanations for adequacy, relevance and consistency with the available data. Some theories may be useful in special situations, but are seriously limited in describing or explaining important aspects of the situation.

Other theories range from being irrelevant to down-right wrong. Perceptual Control Theory is the only theory that comes anywhere near accounting for all aspects of the observed events.

I am sure that the topic of "Errors" and Perceptual Control Theory can be treated much more effectively than I have done. I offer this as a beginning.

Regards, Bob Clark

Date: Sat Mar 06, 1993 10:24 pm PST
Subject: Re: A challenge to information theorists

[Martin Taylor 930307 00:50]   Bill Powers 930306.0700)

Bill issues a challenge to compare PCT versus information theory. I don't understand this. If by "information theorist" Bill means me and Allan (since I have had much to say on this topic I suppose he does), then there is hardly likely to be a situation in which PCT and information theory could make opposite predictions. Information theory as I understand it leads directly to PCT as Bill understands it, or so I believe and have tried to explain. Since before Xmas, I have been trying to get time to write this in a serious paper, but I haven't got as far as I would have hoped with it.

Somebody's Law: Things take longer.

(Sorry if this is full of spurious characters. I see them in my echo, and I try to edit them, but I may not catch them all)

The difference I see between PCT with information theory and PCT without is that with IT it should be possible to make models that need not have as much arbitrary fitting of parameters in order to account for real data. If you know that the acquisition rate of information from say a cursor position for one condition is 60 bps, and for another condition is 20 bps (perhaps because of viewing distance or brightness contrast or something) then you should be able to derive the predictions for the second condition from those for the first. Without IT, a PCT modeller would just have to find a new best-fitting parameter set--perhaps a different slowing factor.

There's no conflict in my mind between the information-theoretic approach and "straight" PCT. I've said this over and over. The models are the same. The ways of finding the parameters are different. At least I have no reason so far to think there is a difference in othr respects.

If you want to compare Ashby's diagrams as skeletons for models to contrast with PCT models, that's a completely separate issue, an issue that does not concern me.

In respect of the specific experiment Bill proposes, the information theorist has to worry about the information transmission between R and T as well as T and E. If the task is to keep E small, information about the future of D might be useful, but E is where the action is. Given a rigid robot so that there is no loss of information between R and T and between T and E, information about D just might permit keeping E within reasonable tolerance, but the time lags woud have to be very small compared to the bandwidth of the disturbance for condition 2 to work at all well even under such unrealistic conditions.

Martin

Date: Sun Mar 07, 1993 8:06 am PST
Subject: For Ray Jackson
Actually from Greg Williams (930307 - via the net)

Sorry folks, I've tried to send this to Ray three different ways directly, but none worked. Does ANYONE know how to address e-mail to folks on ATT Mail from MCI Mail? I'd hate to have to call the MCI Mail customer assistance line and stay on hold for hours....

From Greg Williams (930306 - direct)

Ray, I hope this makes it to you!

>By the way, you must really run up some serious monthly
>MCI bills (I know I do on ATT)...

I get 40 5KB messages (outgoing) per month for a flat $10 monthly fee, with unlimited incoming messages.

>>I do have a project in mind which you could help with if you want....
>>Would you consider editing the material for CL?
>Sure, I'd be glad (honored, in fact) to help out on that project (in May).
>That is, if I don't totally lose my mind before I graduate.

Great! I have faith that your sanity will still be intact in May. Some time before then I'll gather all of the relevant posts in chronological order and send them to you. The easiest way for me would be on IBM disks in ASCII format, but if you're 100% Mac (or God knows what else!?!?), I could send disks to somebody (like Gary Cziko) with an academic e-mail bankroll for e-mailing to you. Or maybe Gary ("Mr. Mac") could convert from IBM format to Mac disks and send you the disks. Which would be best for you?

>>I hope you will be sending copyright clearance for your posts.
>I'm sorry, can you refresh me on what I need to do for that?

I've been sending the following to netters:

I ask that IF YOU HAVE NO OBJECTIONS TO MY USING EXCERPTS FROM YOUR POSTS in CLOSED LOOP, please fill in the form below and postal-mail a copy to me at the address given at the bottom of the form. When I say "excerpts," I mean parts edited (to the extent I am able) ONLY for brevity, NOT for content. If you want to place different and/or additional requirements on my use of your posts, simply state all the requirements over your signature and send me a copy. If you want to haggle, phone me at 606-332-7606.

IF I DON'T RECEIVE PERMISSION FROM YOU, I WILL USE NO PARTS OF YOUR POSTS IN "CLOSED LOOP." IF I DO RECEIVE PERMISSION FROM YOU, THERE IS NO GUARANTEE THAT EXCERPTS FROM YOUR POSTS WILL APPEAR IN "CLOSED LOOP."

Many thanks to all for the colorful threads on the NET, whether or not you decide to be immortalized (???) in hardcopy.
--Greg Williams

===

TO GREG WILLIAMS:

YOU HAVE MY PERMISSION TO USE EXCERPTS FROM MY POSTS ON CSGNET IN "CLOSED LOOP." I RETAIN ALL COPYRIGHTS TO MY POSTS, AND YOU WILL INDICATE THAT FACT BY INCLUDING A LEGAL COPYRIGHT NOTICE IN "CLOSED LOOP" FOR EACH EXCERPT FROM MY POSTS. I MAY CANCEL PERMISSION (NON-RETROACTIVELY) WITH REGARD TO ANY PORTION OF MY POSTS BY GIVING YOU SIX WEEKS' NOTICE.

SIGNED ____________________________________

DATE ______________________________________

NAME ______________________________________

ADDRESS ___________________________________

____

____

Send to: Greg Williams, 460 Black Lick Road, Gravel Switch, KY 40328 USA

Date: Sun Mar 07, 1993 10:01 am PST
Subject: Re: A challenge to information theorists

[Martin Taylor 930307 12:40]   Bill Powers 930306.0700

I suspect my response early this morning might have been a little incoherent. To it, I would like to add a quote that more or less agrees with what I want to say:

(Bill Powers 911119.1000)

>I suppose that
>in principle one could construct an entire control hierarchy out of
>probabilistic calculations. Such a model might actually come closer to
>the basic mode of operation of a nervous system. But a model that ignores
>statistical noise and treats signals as smoothly varying frequencies is
>far simpler to express, and its behavior is much easier to calculate. Is
>the added rigor of a probabilistic treatment needed, considering the
>level at which we can measure and characterize behavior?

The heart of my feeling about using information theory is in that word "needed" in the second-last line. It is not needed if all you want to do is to describe the results of specific experiments by fitting parameters to correctly constructed control circuit diagrams. It is needed if you want an explanation of why those diagrams are correct and how the parameters may change under different circumstances of task, environment, or perceptual (or I guess motor) skill.

PCT is often touted as explanation in contrast to description. Looking from the other side, it can seem to lack explanatory power. See my Occam's razor discussion, posted a few weeks ago. There really isn't any difference between explanation and description except in succinctness of expression and range of validity. Incorporate information theory within PCT and you improve both (since nothing need be added to IT already described when it is used within a PCT framework, and it should reduce the specific fitting of model parameters to experimental data).

The child asks "why is the sky blue." The adult explains that the different wavelengths of light are differentially scattered and obsorbed by the particles in the air, which is an explanation. The child asks "why does that happen." The adult answers with lots of quantum-mechanical equations, which provide an explanation. The child asks "why do those equations work." The adult has no answer. They are just a description of the way the world seems to be. So it is with all explanation.

Let me try to get the PCT-IT paper written (God knows when), and we can reopen this discussion with a hope of doing what you want--providing numerical predictions for experiments. If I don't get it done by July, we can at least talk about it around a blackboard in Durango.

Martin
Date: Sun Mar 07, 1993 10:44 am PST
Subject: Brain Evolution

[from Gary Cziko 930307.0450 GMT]

Would anybody out there in CSGnetland know of some good references concerning the evolution of the brain as a progressively hierarchical perceptual control system?

Of course, I don't expect there to be much out there which discusses brain evolution in this way. But even a good source of information comparing nervous systems of extant species form jellyfish to worm to insect to vertebrates, etc. would be very helpful.--Gary

Date: Sun Mar 07, 1993 4:30 pm PST
Subject: Re: error and mistakes; challenge

[From Bill Powers (930307.0930)] Bob Clark (930306.1930) --

Very nice, Bob. The "error" in a control system exists in but is not perceived by the control system itself. Ergo, the opinion that an error exists must be a perception in some other system.

There's a related problem with language here. People have often interpreted the target position in pursuit tracking as the reference signal, because the person doing the tracking seems to be reacting to the target-to-cursor distance as an "error." Behind this there is an unconscious assumption, which is that a distance of zero between target and cursor is the only possible reference condition. To show that this assumption is incorrect, I usually just ask the tracker to maintain the cursor at some fixed distance from the target for a while. The apparent reference condition suddenly disappears. Now there is only the target- cursor distance to be seen, with no reference distance visible. No observer can now see what the reference condition is, except by inferring it from the way the person changes the cursor position when the target position changes. The reference condition has no physical existence in the environment.

Still, there is a problem yet to be solved here. When a control system fails to operate properly, for example because something has changed in the external part of the loop or because there is an overwhelming disturbance, the person as a whole certainly knows that something has gone wrong -- that there is an error larger than errors should be when control is successful. Somehow something in the whole system knows that a subsystem has failed to bring its perceptual signal to the reference level requested. I don't have any neat solution to that problem. The idea of model-based control might lead to a solution, but so far I haven't been able to work out the details that would make such a system work.

Martin Taylor (930307.0050) --

>Bill issues a challenge to compare PCT versus information
>theory. I don't understand this. If by "information theorist"
>Bill means me and Allan (since I have had much to say on this
>topic I suppose he does), then there is hardly likely to be a
>situation in which PCT and information theory could make
>opposite predictions.

All I'm asking is that you show me that this is true in a specific experiment.

>Information theory as I understand it leads directly to PCT as
>Bill understands it, or so I believe and have tried to explain.

OK, I have heard your words. I am asking that you take the specific experiment I described, analyze it in terms of information theory, and tell me what information theory says about the expected result. If information theory leads directly to PCT, then obviously information theory will lead to the right prediction. Just show me how it does, in this specific case.

>The difference I see between PCT with information theory and
>PCT without is that with IT it should be possible to make
>models that need not have as much arbitrary fitting of
>parameters in order to account for real data. If you know that
>the acquisition rate of information from say a cursor position
>for one condition is 60 bps, and for another condition is 20
>bps (perhaps because of viewing distance or brightness contrast
>or something) then you should be able to derive the predictions
>for the second condition from those for the first.

Fine. That would be very helpful. The problem I have set is much easier, though, because you only have to analyze two conditions. What you must do to meet the challenge, I presume, is to lay out the way you would determine the number of bits per second in the various data sets, how you would compute the degree of regulation of the essential variable in terms of information theory, how you would treat the person doing the Regulating in information terms, and finally how you would calculate the predicted result, as defined in the challenge. If PCT follows from information theory, and PCT can make a general prediction before the fact, then information theory should also be able to make a prediction before the fact if it is fundamental to PCT. When the actual data become available for the experimental run, you should be able to plug them into your quantitative formulas and show that they match the predictions. I expect to be able to do that with the PCT model.

I'm a bit puzzled at the claim that information theory would allow fitting the model to the data with fewer arbitrary parameters. In the model we use most often, we obtain an excellent fit for an individual by changing just one parameter, the integration factor of the output function. It is hard to imagine fitting the data with fewer parameters.

If you're offering an alternative test experiment, then lay it out for me and I will try to meet the counter-challenge. In the meantime, how about meeting my challenge?

>In respect of the specific experiment Bill proposes, the
>information theorist has to worry about the information
>transmission between R and T as well as T and E. If the task
>is to keep E small, information about the future of D might be
>useful, but E is where the action is. Given a rigid robot so
>that there is no loss of information between R and T and
>between T and E, information about D just might permit keeping
>E within reasonable tolerance, but the time lags would have to
>be very small compared to the bandwidth of the disturbance for
>condition 2 to work at all well even under such unrealistic
>conditions.

What the information theorist has to worry about is the information theorist's problem. It strikes me that you're talking about determining rather a large number of parameters from the data, probably more than one. How many parameters are you allowed to adjust in an information-theoretic model order to reduce the number of parameters the PCT model must adjust?

We are not given a rigid robot, but a person doing a real behavior in a real experiment. It is up to you to assume what you must in order to apply information theory to your own satisfaction. I have laid out the data that will be available and how they will be obtained. PCT can arrive at a prediction on the basis of those data. If information theory is fundamental to PCT, it should be able to do the same thing with the same data and come out with the right answer. The conditions are not unrealistic: I have spelled out exactly how they will be accomplished, and it will be perfectly possible to do an experimental run under those conditions.

Look, I'm offering to do as much of the actual detailed work for you as you wish, including writing the program, running the experiment, and providing you with the raw data. You can do any part or all of this for yourself to determine that the experiment was properly done and reported. You can even send me the formulae for the analysis of the raw data, and I will write the required program and apply them for you.

You say that information theory is capable of offering explanations more fundamental than what PCT offers. Good. But I want to see it actually done, not just talked about. Just to up the ante a bit: until my challenge is met, I will feel free to claim that information theory is incapable of predicting ANY behavior that PCT can predict.

Challengingly, Bill P.

Date: Sun Mar 07, 1993 5:35 pm PST

Subject: Ray's address; hardnosed experimentalist

[From Bill Powers (930307.1730)] Greg Williams (930307)

I have Ray's address listed as RLJACKSON@ATTMAIL.COM

And that works for me. Maybe MCI just refuses to send mail to a rival's network.

In general:

The reason for my modeling challenge and a generally hard-nosed attitude toward experimentation is not specifically aimed at information theory. I think that we tend to get so speculative on this net that we confuse plausible explanations with testable (and tested) ones.

Speculative theories have one great advantage: they don't risk being disproven by an experiment. As long as we stay at the "in principle" level, we can safely propose any idea that seems to hang together logically, mathematically, or philosophically. If the idea is wrong or irrelevant, nobody will ever know.

But that doesn't really get us closer to founding a real science of living control systems. If we want to build that science, we have to stick our necks out and risk making specific predictions that can fail the test of experiment. The biggest problem with the conventional psychological sciences (and others) is that the people in them don't want to risk being wrong. So instead of making specific predictions from their theories, deliberately challenging them with difficult experimental tests, and judging the results in terms of rigorous standards, they try too hard to be right the first time. And they relax the standards so that even just a mere suggestion of agreement of experiment with theory is enough to get by, and get published. In even larger numbers, other people don't even think of testing their ideas experimentally -- all they really want is to prove some theorems or develop some plausible statements and establish the internal consistency of their ideas. They may argue about whose theorem to apply, but it doesn't seem to occur to them to ask whether real systems behave consistently with ANYONE's theorem.

It's not all that hard to make a prediction from a theoretical idea. All you really have to do is ask "What does this idea predict about what will happen the next time I look?" Then, of course, you must actually MAKE a prediction, and either perform an experiment or watch carefully what happens by itself. When the prediction fails, you can then ask what made it fail. The way in which it fails gives you information about how to correct the model, so it won't fail quite so badly the next time. If you keep doing this, you will eventually have a model that doesn't fail -- ever. That's how physics and chemistry got the way they are. The model had to behave correctly within the limits of measurement. If it didn't, it was wrong.

I say that we should expect no less of PCT and HPCT. But we're just going to end up as a dilettantish debating society unless we do what is necessary to get there. To me, that means making specific predictions about specific experiments and testing them in a way that really tests them, without excuses.

I'm probably asking too much (even of myself) in saying we should publish only ironclad experimental results. But if we do publish anything else, it should be made plain that nobody is satisfied with it, and that the report is a progress report, not a conclusion to be used in the fabric of a scientific system. We have to talk with each other on the way to finding reliable models, but let's not confuse that sort of communication with the product of a true science. When such products do turn up in our journal, the difference from the other kind of report will be glaringly obvious.

Best to all, Bill P.

Date: Sun Mar 07, 1993 6:16 pm PST
Subject: error, information theory

[From Rick Marken (930307.1500)]

Well, I think I can post to the net again; the secret seems to be using the correct addess.

Bob Clark (930306.1930 EST) --

Thanks for the comments on human error.

I think your distinction between the "user's" and the "engineer's" view of error is very important; and it was one of the points I was planning to make in my discussion of error (I was going to call it anthropocentric vs egocentric). But I think my approach to this distinction is a bit different than yours. So rather than respond to your post directly, let me first broach my ideas and then see if we can converge (if we want to try).

I would begin my attempts to understand human error by pointing out that perceptions are just perceptions; in themselves they are neither right nor wrong; neither correct nor errors. They are just what they are -- varying intensities, sensations, transitions, configurations, sequences, relationships, programs, categories, principles, system concepts. When I see a glass of wine (configuration) knocked off of a table (relationship, transition) onto an expensive (category) white rug (configuration), that's what I see. Error implies deviation from a comaprative reference (specifying the way things should be). If the perception of spilled wine seems like an error to an observer it must be because that perception deviates from some specification IN THE OBSERVER of what should be perceived. This seemingly obvious fact about perception is completely missed in all discussions of human error that I have read. In these discussios, one get's the impression that an error is a perception itself -- corresponding to something "out there" in boss reality. But we know that very often a perception that is an error to one person is not an error at all to another. The spilled wine, for example, might be just what someone wanted to see -- that's why they knocked over the glass. So PCT enters the picture before it is even invited; the fact that a person can see an event as an error implies that they have a reference for how they want that perception to be; they might not be actively trying to control that perception, but they have a reference for it nonetheless. An engineer who studies human error must realize that a perception may be an error from his or her perspective but not necessarily from the perspective of "the user" -- the person who is involved in contributing to that perception (like the person who knocked over the glass of wine). So my version of Bob Clark's "engineers" perspective on human error requires that the engineer know that the "human error" he or she is talking about exists because the engineer him or herself IS a control system.

The engineer must then realize that his or her perceptions (whether they are considered errors are not) may be the side effects of the efforts of other control systems to keep their perceptions matching their own specifications for these perceptions. The engineer must realize that the "user's" perspective is not anything like his or her own perspective. The engineer who deals with human error must decide whether his or her goal is to eliminate his or her own perceptual error (created as a side effect of the user's actions) or the perceptual error of the user (if there is any). If the engineers goal is the latter (which the humane and really the only possibly achievable goal from a PCT perspective) then he or she must learn 1) what perceptual variables the user is trying to control and 2) what values of these perceptual variables are considered "right" by the user. Eliminating human error from the user's perspective is then largely a matter of figuring out how to design the feedback function (from user output to input) so that the user has better CONTROL. Some human factors engineers have successfully designed systems that help user's control better -- but, since these engineers don't understand PCT, they are not able to go about the process in a systematic manner (using the test for the controlled variable, for example).

Martin Taylor (930307 00:50) --

>There's no conflict in my mind between the information-theoretic approach
>and "straight" PCT. I've said this over and over. The models are the same.

I believe that there is no conflict in your mind; I just think there should be -- and a BIG one. Information theory has been part of psychology since the start of the "cognitive revolution". If it is really the same as "straight" PCT then why aren't ANY information processing type psychologists aware of some of the fundemental facts about living systems from a control system perspective; that they control perceptual variables; that input-output transfer functions depend on characteristics of the environment, not the internal "processing" capabilties of the organism; that the use of statistics is an unnecessary consequence of the failure to test for controlled variables; etc etc. In other words, information theory based psychology should already be where PCT was in the 1960s. In fact, psychology, with the benefit of information theory, demonstrably has NO CLUE about the nature of the phenomenon of control (purposive behavior) or how to study it.

I can see that you want to cling to this information theory thing; apparently it's very important to you. And I am honestly willing to be convinced of it's value (answering Powers' recent challenge successfully would go along way toward convincing me). But as it sits, it looks to me like you want to cling to infomation theory the way other psychologists want to cling to their favorite theory -- even while embracing PCT. I'm afraid that it just can't be done (and at the same time get PCT right).

Wanting to stick with a grand old theory is a very common phenomenon -- and it's the reason why 1) most psychologists don't get into PCT and 2) if they do, they don't get PCT right. It's why we say PCT is revolutionary. I know that it seems impossible that all the old, revered theories in psychology are invalidated by the work of a nice engineer from Chicago who doesn't even have a PhD in psychology -- but that's the fact Jack. I know that psychologists in particular are used to clinging to some remant of the past while moving on to new verbalizations (theories). But PCT is a whole new enchilada -- it is not like anything else that has been dreamed of before in your philosophies (except possibly by James who did say that purposefully produced results were intended SENSORY CONSEQUENCES of action -- but he had NO IDEA how this worked or what it meant for the study of behavior).

If you don't want to let go of the old stuff, that's OK. I understand. But if info theory really has something important to contribute to the PCT model then SHOW ME WHAT IT IS. I don't want to hear that it just DOES. Heck, I can go to Agre's conference and hear about how important it is for me to have a "principled" (I hate that word) understanding of the interaction between agents and their environments. I don't want philosophical BS -- I want to see precisely how info theory fits into my models.

Best Rick

Date: Sun Mar 07, 1993 10:18 pm PST
Subject: Randall's bibs

I'd like to add my voice to the approbation of Allan Randall's bibliographies, etc -- these are extremely useful, and just the sort of thing that PCT needs.

Avery.Andrews@anu.edu.au

Date: Mon Mar 08, 1993 5:41 am PST
Subject: Re: Language facts

[From: Bruce Nevin (Fri 93035 14:20:00)]

[I wrote this in my last hour here Friday, covering the easier half of my notes scattered through Bill's (930303.0830). When I came in today 3/8 I found it poised ready to send--but I had never pressed the carriage return! I'll resist temptation to review it before sending (that would drag through another week). So, warts and all.]

(Bill Powers (930303.0830)) --

This must be terse, I hope not too telegraphic.

I think my (930302.1430) was a disturbance because you believe I am saying that social groups, strata, classes, etc. control individual people. I am not, I did not, and I don't expect to.

If you now reread what I have said with this in mind, I think that many of your objections and comments will have no occasion.

It is very easy to mix discussion of matters of different logical type. To pick a set of distinctions that have often bedevilled our discussions, are we talking about language learning in the child, established reference perceptions for language in the adult, the evolutionary origins of language, the immediate historical origins of some characteristics of a particular language such as English? It makes a difference. Or (in this post) are we talking about acoustic and kinesthetic perception of phonetic differences between dialects, or about phonemic contrasts in which these phonetic differences don't make any difference? Or about the perceptual control by means of which we arrange that they don't make any difference?

Some individual points:

>kinesthetic aspects of phoneme production vary widely from one
>person to another without altering the phoneme that is heard.

Some kinesthetic aspects of some phonemes. In certain regions of the mouth, differences in articulation of the tongue make little or no acoustic difference. These (acoustic and articulatory) regions that are relatively disturbance-proof tend to be favored as "targets" for phonemes in various languages. However, there are exceptions (usually languages in process of change, I think), and there are other sorts of phonemes that are not so nicely favored by the acoustic properties of the vocal tract. I would be very cautious about overgeneralizing here.

>The reproduction of heard speech is more analogous to putting the
>fork "in" the water glass than to imitating the way the
>spectacles are held. The aspects of speech that are reproduced
>are those that are independent of point of view, while
>reproducing the orientation of the glasses specifically requires
>a transformation of the point of view.

Repetition in my dialect of what I heard you say in yours, yes. It doesn't matter that I'm using a table fork and a tumbler where you had a pickle fork and a goblet. Imitation of manner gets into mimicking detailed differences that don't make any difference at the higher level of phonemic contrast: my way of saying the diphthong in "light" vs. yours. For repetition, it doesn't matter that I realize the contrast (between that diphthong and, say, the "a" of "barn") at a different point in the space of acoustic possibilities available to us both. I do maintain the contrast and each of us is able to map my way onto your way and vice versa, so we agree that it is the same contrast, and so indeed it is. But if I am to reproduce your speaking, I must learn to control the contrast in the same manner that you do--to speak your dialect.

As a child learning the language, I imitate until I learn the contrasts. Then I repeat. As an adult, I no longer imitate (or rarely), I repeat in my own manner. If I were mirroring your dance movements, comparable differences would be errors; in repeating what you say, they are not errors, they are differences that don't make any difference, below the level of phonemic contrast.

Adapting your words, the aspects of speech that are REPEATED are those that are independent of point of view, while the aspects of speech that are IMITATED specifically require a transformation of the point of view. The terms of repetition (phonemic contrasts, morphemes, words, etc.) are independent of point of view because they are socially instituted.

>I said they "had to"
>
>>>prevent themselves from extending the vowel into a diphthong
>>>in order to imagine matching the heard vowel.
>
>The "in order do" supplies the reason they have to prevent
>themselves. The goal is to match the heard vowel. But the vowel
>they are used to hearing is extended into a dipthong, so as
>speakers of a given language that is their normal reference
>perception for the vowel.

They "had to" prevent themselves from extending the vowel into a diphthong in order to imagine matching the vowel they IMAGINED that they heard. They substitute an imagined perceptual signal, the reference perception for controlling perception (at a higher level) that amounts to the patterning in their language. Before Sapir's students learned about the glottal stop, they substituted an imagined diphthong, because they could imagine no consonant there. After they learned about the glottal stop, they substituted an imagined consonant (the other possibility in the English pattern for stressed final syllables). This imagined consonant was the new one they had just learned to write.

(Incidentally, the first part of the o is nasalized in both English and Spanish "no", but in Spanish the second part, which is denasalized, is shorter than in the English diphthong.)

>>Rather, there is a socially instituted pattern on the basis of
>>which speaker/hearer's expectations get set.
>
>This is not an explanation, but a description. There is a social
>pattern, which means only that most of the people with whom the
>test person interacts speak in a similar manner. The existence of
>this pattern can't, by itself, have any effect on anyone.

The pattern is socially instituted in the individual. The individual tries very hard to learn and participate in the same patterns as others, and the child is typically assisted in this by adults. Look again at Bruner, _Child's Talk_.

>If a person's speech changes when the person interacts extensively
>with this group, that is only because selecting the group's way
>of speaking suits the person's purposes.

Of course. And that is how and why (a) the individual institutes the prevalent social patterns in his or her own reference perceptions and (b) by participating in them contributes to their ongoing institution as social facts that children learn, etc.

>There are plenty of
>cases to show that there is nothing necessary about the social
>influence on pronunciation -- [...]
>You can't base a serious law of speech acquisition on a
>generalization with so many exceptions.

I never proposed such a simplictic "law of speech acquisition." Nor did I say that social patterns, etc., control people. Take the Martha's Vineyard example. Kids are exposed to a MV dialect and various mainland dialects (mainly Boston/New England and New York) all their lives. At puberty, they pick one. THEY do the picking. The dialect comes with the package when they choose an image of self for themselves. But they PICK from a range of possibilities that is socially available. They don't invent something entirely new. (Even a nonconformist creates a persona by patching together available materials that are meaningful to others in intended ways. We all know about the nonconformist uniform, reinvented with each generation. An individual may create some unexpected combination or incorporate a new element in an existing pattern, and is very likely thereby to start a fashion.) In this sense individuals are constrained in their choices (even when choosing "against" prevailing norms) by the social patterns that prevail at the time. But this is not to say that the social norms control them.

The examples you give as exceptions to the supposed "law of speech acquisition" are not exceptions, they are each a case in point. They learned their manner of speaking in childhood and established an adult persona in their teens which subsequently they resisted changing.

Character actors, on the other hand, cultivate an ability at mimicking dialect, body language, etc. to project a persona with socially identified characteristics. Remember the Oklahoma study I cited quite a while back? Four grad students in the speech department recorded multiple readings of a short text, varying things like speed, pitch variation, nasality, orotundity (not sure how that was defined), etc. Audiences listening to the tapes perceived these as many different people, and their evaluations of the probable personality characteristics of these different people were virtually unanimous. These correlations of speech characteristics with perceived personality traits must be known to the child choosing what to include in a persona, it seems to me; they apparently were known to the college students participating in this study, and are known to the audiences of storytellers and character actors.

>Most people can imitate accents
>other than the ones they normally use, at least well enough to be
>speaking in an obviously different way. It is not the set of
>recorded perceptions of speech that determine the accent we
>choose, but our goals for our relationships to those around us.

When most people imitate accents, they do so laughably, and would not be mistaken by anyone as a native speaker of the dialect they are imitating. But more important here is the question: how, then, do our goals for relationships with those around us have any connection with our pronunciations of words? I have proposed a basis for such a connection. What do you propose as a basis? Or is my proposal more acceptable to you, now that you know I am not saying that social structures control individuals?

>To "desire independence" is to misinterpret human nature
>according to PCT. The basic problem of human existence is that we
>ARE independent, whether we like it or not. Each person is free
>every waking movement of the day to set any goal whatsoever,
>including the goal of stepping off a cliff or stepping into
>traffic. All that prevents selecting such goals is a potential
>conflict with other goals -- nothing external has the slightest
>say in the matter.

So PCT tells us. So we may come to perceive. But so most people most of the time in fact do NOT perceive. That is one reason PCT is revolutionary. But many people do perceive their independence threatened by the actions of others. That is their perceptual reality.

>Do you mean that they walk up to you and tell you what subpopulation
>they belong to, so you don't have to do any interpreting of your own?

Yes. But not by name. Ostensively. If you want to name their social affiliations or talk about them, you have to do some interpreting. You may not recognize all that they are "telling" you (or would be telling, if you had learned the patterns they are participating in but you are not).

>Everyone imagines living
>in a certain kind of society, which doesn't actually exist and
>isn't like the society that the next person believes in. When any
>individual says "us" you haven't the least idea what he means,
>and neither, come to think of it, does he.

(A) How do you know? (B) These perceptions, real or imagined, constitute an important part of his real world.

[Tannen's books]:
>It's really a book about what's wrong with stereotyping -- that
>is, with speaking in terms of averages of populations -- isn't it?

Nope. What goes wrong in communication because of being members of groups whose norms of "conversational style" differ.

Got to run. Snow and blow still in progress. Train to catch. Until Monday, then,

        Bruce Nevin   bn@bbn.com

Date: Mon Mar 08, 1993 10:10 am PST
Subject: Re: error, information theory

[Martin Taylor 930308 12:00]   (Bill Powers 930307.0930)

I now understand what Bill Powers wants, which is different from what I thought. I will try to do it, though it will lack the background I am trying to develop in the paper. See the end of this posting.

(Rick Marken 930307.1500)

>>There's no conflict in my mind between the information-theoretic approach
>>and "straight" PCT. I've said this over and over. The models are the same.
>
>I believe that there is no conflict in your mind; I just think there
>should be -- and a BIG one. Information theory has been part of
>psychology since the start of the "cognitive revolution". If it is
>really the same as "straight" PCT then why aren't ANY information
>processing type psychologists aware of some of the fundemental facts
>about living systems from a control system perspective; that they
>control perceptual variables; that input-output transfer functions
>depend on characteristics of the environment, not the internal "processing"
>capabilties of the organism; that the use of statistics is an unnecessary
>consequence of the failure to test for controlled variables; etc etc.
>In other words, information theory based psychology should already be where
>PCT was in the 1960s. In fact, psychology, with the benefit of information
>theory, demonstrably has NO CLUE about the nature of the phenomenon of
>control (purposive behavior) or how to study it.

You are right on target here, except, I think, in that line about "should be." (And, I think though I won't argue the point, in rejecting statistics entire).

As I understand the problem, the failure of information theorists has been that of most other theorists, that they take a Newtonian view rather than an Einsteinian view. They take "probability" as being some kind of a limit after an infinite number of replications of the frequency with which one event happens rather than another. This leads to the idea of information as something that is transmitted through a channel whose size can be exactly determined. The reason I am taking so long about the information theory PCT paper is that I am trying to go to fundamental principles to show how wrong and unworkable that idea is. In essence, the error is exactly the same as the error of a cognitive planning system that works only if the environment and the effects of the actuators are precisely known.

Probability is something relating only to the information available at the point where that information is used. It relates to individual possible events that might be detected at that point, not to frequency. Frequencies of past events that happened when conditions at the point of interest were "for all practical purposes" identical may well be used in assessing the probability of some future event, but they are not its probability.

>I can see that you want to cling to this information theory thing;
>apparently it's very important to you. And I am honestly willing to
>be convinced of it's value (answering Powers' recent challenge
>successfully would go along way toward convincing me). But as it sits,
>it looks to me like you want to cling to infomation theory the way
>other psychologists want to cling to their favorite theory -- even
>while embracing PCT. I'm afraid that it just can't be done (and at
>the same time get PCT right).

Well, you may be right. I don't think so. I find that every aspect of PCT is illuminated and makes sense if I think from an informational point of view. I have had my disagreements with Bill Powers on technical aspects of PCT, but those have usually been based on a (possibly incorrect) belief that Bill either is asssuming that information is available somewhere it isn't, or that he is ignoring a source of information that could be used. Sometimes we come to an agreement when I see that I am wrong in my assumption, sometimes (as in the reorganization style issue) Bill changes. But the issue is almost always "where is information available, and is it used."

Bob Clark's Engineer's viewpoint is fine, but in using it the engineer must try to empathize with the many viewpoints that occur at all places within the system. If point A is a perceptual signal that has as part of its input a sensory signal B, the engineer cannot assume that every variation in B is reflected exactly in A. The question must be "what does A see of the variation in B" before the engineer can properly assess what will happen at A. None of Bob Clark's viewpoints seem to me to be of the class that I might call "internal." They are all "external," even the User's. An information theoretic approach must take the internal view, even at second-hand, as an empathetic engineer.

>In other words, information theory based psychology should already be where
>PCT was in the 1960s. In fact, psychology, with the benefit of information
>theory, demonstrably has NO CLUE about the nature of the phenomenon of
>control (purposive behavior) or how to study it.
>....
>Wanting to stick with a grand old theory is a very common phenomenon -- and

>it's the reason why 1) most psychologists don't get into PCT and 2) if
>they do, they don't get PCT right. It's why we say PCT is revolutionary.

Quite right. I don't know if I get PCT right. I can't know that, ever. What I can know is that I get what seems like a considerable error reduction, or subjectively, the first satisfying feeling, after 30 years as a professional psychologist, about a theory that fits what I believe about the world.

Some 10 years ago I had a conversation with someone who was in graduate school with me, in which we bemoaned the fact that there seemed to have been nothing new learned in psychology since we were at school. All that seemed to have happened was a cyclic change in fashions. We looked for the Newton of psychology to show us how all these little fashionable micro-theories could be superseded by a unifying viewpoint. Then 2 years ago I discovered Bill Powers, and although it took me a year to really feel that I had a good enough handle on what was going on that I could contribute to the discussion in ways other than simply asking questions, it took much less time than that to see that the unifying viewpoint was at hand.

It is still the case that (like poor Bruce Nevin) every moment I put into PCT is a moment stolen. But I have been trying to shift the constructs of human-computer interaction theory into a PCT framework, which turns out to be relatively easy, given that my existing Layered Protocol Theory can be seen as PCT applied to two interacting hierarchies, if I fudge a little on the requirement for scalar perceptual variables in classical PCT. That work is not on stolen time.

But you are right that I have a reference for sticking with information theory. It has served me very well for nearly 40 years in understanding economics, aesthetics (from my bachelor's thesis, in which I see little need for change even now), perceptual psychology, and now PCT. Maybe I would be better if I dropped it. But I see no evidence pointing in that direction other than arguments that sound quite analogous to those who say "Believe in Christ and You Shall Be Saved."

>If you don't want to let go of the old stuff, that's OK. I understand.
>But if info theory really has something important to contribute to
>the PCT model then SHOW ME WHAT IT IS.

I tried, many times, last year. The result was that some people thought I was just trying to puff up my ego. So I swore off trying to show you until I could do a more thorough job, which is why I'm trying to find the odd moments to work on the information->PCT paper, and make it something that works from principles everyone can agree on, rather than from higher level constructs that different people read in different ways. I have to redevelop information theory using subjective probability, in such a way that I do not have to put up with claims such as "Shannon information has been shown not to relate to the everyday notion of information or meaning."

>I don't want philosophical BS -- I want to see precisely how info
>theory fits into my models.

That's what I want to give you.

Back to Powers:

>I am asking that you take the
>specific experiment I described, analyze it in terms of
>information theory, and tell me what information theory says
>about the expected result. If information theory leads directly
>to PCT, then obviously information theory will lead to the right
>prediction. Just show me how it does, in this specific case.

I'll try to do that, ASAP.

>I'm a bit puzzled at the claim that information theory would
>allow fitting the model to the data with fewer arbitrary
>parameters. In the model we use most often, we obtain an
>excellent fit for an individual by changing just one parameter,
>the integration factor of the output function. It is hard to
>imagine fitting the data with fewer parameters.

If you change, say, the contrast of a tracked target (or the tracking cursor), or the bandwidth of the disturbance, or put a grid on the visual surface...is the integration factor the only thing that changes?

I once had an information-based theory of the displacement of the figural aftereffect. It fitted many different aftereffects in different perceptual dimensions with only two parameters. Later, a study was published in which the contrast (a parameter not included in the earlier studies) was changed. The theroy fitted it with ZERO parameter estimation from the data. The effect of contrast on information rate was determined from studies that looked as if they were looking at something entirely different. That's the sort of thing I mean by extending the range of prediction.

>Just to up the ante a bit: until my challenge is met, I will feel free to
>claim that information theory is incapable of predicting ANY
>behavior that PCT can predict.

I'll use your circuits, then. And if PCT predicts behaviour that seems to demand information that cannot be obtained at the point it is used (the perceptual signal compared to the reference signal), and nevertheless makes good predictions, then I'll allow your claim.

Look, in any one experiment, you will get as good prediction as you can achieve with any circuit by varying the circuit parameters. You say that for all experiments only the integration factor has to be changed. Fine. If that's true, the challenge is really for me to show whether information theory could predict the change of integration factor across experiments, taking into account the conditions of the experiment. I don't know how I would do that. The integration factor is not a construct that I have so far analyzed in my thinking. I would assume it relates to motor information rates as seen by the environment. But I don't know.

In making an analysis for a specific experiment, I would use other experiments to obtain data--perceptual discrimination experiments, for instance, or control experiments that use presumed lower-level control systems common to the specified experiment.

Enough for now. I have a book to edit. Analysis this evening if I can. At the weekend if I get the #%$@! book review out of my hair that's been taking up every weekend for the last goodness knows how many. (Rick: that book exemplifies what you are talking about, but I don't know how to review it in a way that the audience that believes in its basic principles will understand. And I find it very difficult to read more than 20 pages at a sitting, without stopping to mull over why they seem so horribly wrong and yet so erudite).

Martin

Date: Mon Mar 08, 1993 10:14 am PST
Subject: MCI to ATT

[from Gary Cziko 930308.0525 GMT]   Bill Powers (930307.1730) said that:

>I have Ray's address listed as   RLJACKSON@ATTMAIL.COM
>
>And that works for me. Maybe MCI just refuses to send mail to a
>rival's network.

Greg, if that is the case, then you might be able to fool MCI by using the following address for Ray:

RLJACKSON%ATTMAIL.COM@VMD.CSO.UIUC.EDU

This should send the message first to "my" machine and then change the % to @, chop off VMD part, and then send it on to ATTMAIL.--Gary

Date: Mon Mar 08, 1993 11:16 am PST
Subject: Re: Conference report

[From Oded Maler 930308-ET]

* [From Rick Marken (maybe)]   Avery Andrews
*
* The report on the conference was a hoot. Obviously, these people know
* all the right things to say. Oof da.

I wonder how this has escaped Gary's censoreship - maybe you should augment your mailer with some time lags..

One of the main reasons for me to go to Paris last week was to here Agre's talk about the very same topics. My overall initial impression was positive. He said a lot of things with which I agree, especially about embedding computation within geometry. His analysis of the transition functions of the three basic elements in the kitchen (tools, materials and containers) and their interacations was also interesting. When he talked about other people's work it was clear that there things he undersatnds better than others, but it is understandable when you want to make a synthesis.

What I didn't like was some of the repeating cliches (probably from the jargon of anthropologists/sociologists/phenomenologists): "establish a vocabulary", "leave this topic here and take it as we go along", "unpack" etc. But this is a personal taste - I can see how it can have a positive effect on groupies.

>From the local (PCT) point of view, I would not classify him as good or bad but rather as an "orthogonal" guy, working on higher levels of the hierarchy (symbols, plans) and having interesting observations which were not self-evident (at least in the AI community at the time).

About one of the works he mentions (Shoham & Tenenholtz) - I know this work, and I think it is a very good and important work. It shows (mathematically and by simulation) how "social" laws that all the agents obey locally (e.g., "take the right lane"), allows each of them to achieve its own goals, with a moderate degradation of performance (compared to the situation where he would be alone with no possible conflicts). There are some similarities with the "crowd" program, but the emphasis is different.

--Oded

Date: Mon Mar 08, 1993 11:31 am PST
Subject: Re: A challenge to information theorists

Bill Powers (930306.0700) writes:

> ...PCT predicts that regulation will be
> unequivocally BEST when the participant gets the LEAST
> information about the actual state of the disturbance D --
> condition 3.
>
> I believe that information theory will make the opposite
> prediction: that condition 2 will provide better regulation than
> condition 3. But I will leave it up to information theorists to
> derive and explain the authoritative prediction. I expect them to
> make their prediction, as I do, before the experiment is run.
>
> The experimental data should then settle the question of the
> relative power of PCT and information theory.

This whole challenge is based on the idea that information theory makes opposite predictions about human control situations than PCT does. In order for this to be interesting, I think you need to do one of two things:

1) Explain in a little more detail why you think information theory predicts that the compensatory system will perform better. If you do not have the time to do a formal proof, even a rough outline or sketch might give us some idea what you're getting at. Personally, I wouldn't expect information theory, all by itself and without PCT, to make ANY prediction about which will do better (unless it is possible, as I think Martin is claiming, that PCT can be derived from information theory, but in that case there's hardly any conflict).

2) Find an information theorist that actually does make this claim. I am NOT claiming that there are no such information theorists, but I am not one. Martin is not one. Ashby is not one. For a challenge to have any clout, you need to have someone to challenge. However, if you can provide #1 above, then that's fine too, since you would be able to show that information theoretic PCTers like Martin and I are being inconsistent.

Allan Randall, randall@dciem.dciem.dnd.ca

NTT Systems, Inc.   Toronto, ON

Date: Mon Mar 08, 1993 11:44 am PST
Subject: Boating and PCT

[Allan Randall (930308.1410)]

Here's a quote from Daniel Dennett's "Consciousness Explained," which I am currently reading (Little, Brown & Co., 1991). I thought this might be interesting to consider in PCT terms:

"Pleasure-boaters sailing along a tricky coast usually make sure they stay out of harm's way by steering for a mark. They find some visible but distant buoy in roughly the direction they want to go, check the chart to make sure there are no hidden obstacles on the straight line between the mark and where they are, and then head straight for it. For maybe an hour or more the skipper's goal is to aim directly at the mark, correcting all errors. Every so often, however, skippers get so lulled by this project that they forget to veer off at the last minute and actually hit the buoy head on!"

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc.   Toronto, ON

Date: Mon Mar 08, 1993 1:11 pm PST
Subject: Re: CLOSED LOOP business

[From Dick Robertson] (930308)
Greg, it sounds like a good idea. I'd be glad to do some reviewing in my field.
Bouquets to you for starting another promising enterprise.
Best, Dick.

Date: Mon Mar 08, 1993 1:17 pm PST
Subject: Re: Journal proposals

[From Dick Robertson] (930308)
Bill Powers' suggestions about the Journal sound really good to me. I would suggest though, that there be a section of "brief Reports" for those first round articles that he said would usually not be good enough to publish. I think it would be good to brief report them so that others working on similar studies could know what's in the works.
Best, Dick.

Date: Mon Mar 08, 1993 2:25 pm PST
Subject: CSGnet Protocol

[from Gary Cziko 930308.2140 GMT]

Over the last few days Greg Williams had difficulty getting direct messages to Ray Jackson and so sent some to him through CSGnet. In checking messages returned over the weekend to the CSGnet listserver today, I found that CSGnet messages have also been bouncing back from Ray Jackson. So apparently the problem was/is at Ray's end and not Greg's.

There is lesson to be learned from this. If you have difficulty communicating directly with someone but you have no problems with other addresses, chances are that the problem is at the other end and that CSGnet is also having difficulty communicating with this person. Therefore, it makes little sense to send personal messages to the individual in question via CSGnet hoping to get around the problem. Sending direct you will find out if your message is bounced back. You will get feedback (closed-loop, if delayed). Sending via CSGnet is open loop (bad).

Therefore, the simple rule that makes sense to me and perhaps to others is NEVER to send PURELY personal messages via CSnet (such as, "Joe, send me a copy, too"). The only exception to this is if you have a very short note that can be tacked onto a CSGnet post of general interest (as Bill Powers sometimes does as in "Greg, the ms. is in the mail") and you know that the other person religiously (or irreligiously) reads EVERYTHING posted to CSGnet.

On the other hand, posting to CSGnet is the way to go if even though your message is directed to one person you suspect that it might be of interest to others (even if only a few) on CSGnet.

If you wish to find the e-mail address of anyone (and everyone) on CSGnet, just send the following command as the text of a message to LISTSERV@VMD.CSO.UIUC.EDU

rev csg-l (countries --Gary

Date: Mon Mar 08, 1993 3:48 pm PST
Subject: advert working

[Avery.Andrews 930309.932]

Here is some tentative advert wording. The stuff in brackets should maybe be omitted as taking up too much space - I don't have time for psycholloquy, so I don't know the format of such adverts in it. Otherwise my inclination is to keep it pretty straight & administrative in nature.

It is generally thought that feedback control is well-understood and thoroughly integrated into the modern behavioral sciences. But there are some indications that it is not in fact quite as well-understood as one might expect. [For example, with the revival of interest in `Central Pattern Generators' in the 1970s, one possibility that does not seem to have occurred to people is that a CPG might function by delivering time-varying set-points to feedback mechanisms, thereby reducing the need to pre-compute parameters for motor-programs. In fact, even today, it is not clear that this rather simple possibility is clearly formulated in the standard literature, and adequately investigated in cases where it has some degree of prima facie plausibility.]

CSGNet is a group of people interested in aggressively investigating the application of closed-loop models to all areas of psychology, broadly construed, following the general lead of the work of William T. Powers. Its current discussants include psychologists, linguists, sociologists, control-system engineers, ... [not sure how to go on here].

Its postings appear in the Usenet newsgroup bit.... [not sure what here], and it may be subcribed to by sending a message to g-cziko....

Date: Mon Mar 08, 1993 4:03 pm PST
Subject: Defining information; challenge

[From Bill Powers (930308.1520)] Martin Taylor (930308.1200) --

What I like about you, Martin, is the civilized way you respond when people attack your life's work with bludgeons.

I'm looking forward to seeing your paper on IT and PCT. It is going to be different from the sort of IT work up with which I have the most difficulty puting -- that seems clear from your discussions, e.g.:

>Probability is something relating only to the information
>available at the point where that information is used. It
>relates to individual possible events that might be detected at
>that point, not to frequency.

I wonder whether you have considered this point yet:

In your terminology, the perceptual signal in an ECS, an elementary control system, stands for the state of a CEV -- a complex environmental variable. As we all seem to agree now, that CEV is a construct created primarily by the forms of all the perceptual functions lying between the primary sensory interface and the place where the perceptual signal finally appears.

When we speak loosely (for coherence and convenience), we say that the perceptual signal is an analog of some aspect of the environment. This sets up the picture of the environment and the aspect of it that exists outside the organism, and a perceptual signal that represents that aspect. It is natural, then, to say that the perceptual signal contains information ABOUT that aspect of the environment.

From there, we can go on to ask how well the perceptual signal represents that aspect. This is where the difficulty starts. In order to answer that question, we must have an independent measure of the aspect of the environment in question, one that tells us its true state so we can compare that state with the representation in the form of the magnitude of the perceptual signal (or any other form, come to think of it).

But this contradicts the previous statements, which say that the CEV is a construction by the nervous system. From the internal point of view, the perceptual signal is always a PERFECT representation of the CEV -- in fact, that signal plus all the perceptual functions leading to its production defines the CEV.

What, then, if the perceptual signal is noisy? We could say that this defines a noisy CEV, or we could say that there is a noiseless CEV out there, and that the perceptual signal is being derived through a noisy channel -- or any combination of these effects.

Consider a television set tuned to a very distant station. The screen shows a picture with a lot of noise dots dancing over its surface. Inside the person looking at the TV set, presumably, is a set of signals representing the state of the TV screen. This set of signals would show an amplitude envelope that would be varying at a high frequency, like some amount of noise superimposed on a average picture. These noise signals don't reflect the state of affairs we would measure with optical instruments at the face of the TV screen -- for one thing, they don't show the frame rate, and the highest-frequency noise dots are smoothed out. So the human perceptions of the TV screen aren't quite as noisy as the screen itself would look to fast- responding physical instruments.

What the human being experiences is a noisy set of signals. Some part of this noise, theoretically, is channel noise in the perceptual system. So the CEV is defined as being noisy in exactly the way the perceptual signals are noisy. The implies that the noisy perceptual signal is a noiseless representation of the hypothetical CEV.

Of course the person is trying to ignore the noise in the perception -- meaning that some higher perceptual system must be constructing something closer to a noiseless CEV, extracting the average picture, which presumably is closer to what the TV station is transmitting.

How, then, do we define information in this situation?

Some of our disputes seem to depend on the inherent noisiness or ambiguity of the information being transmitted; others seem to come down to the nature of the transmitting channel itself (the discrete-impulse nature of neural signals).

In the first case, the brain itself could be completely free of noise and ambiguity, and still have a problem with deciding what the message is. The brain may have to rely on (noise-free) estimations of probabilities to resolve such problems, but we're talking about learned algorithms now, not fundamental problems of information theory or probability in the brain's own operation.

In the second case, the importance of the channel noise depends on the signal amplitudes (frequencies) involved in the brain. Here the problem is dynamic range -- how small an average signal can be detected in the presence of the irreducible channel noise. In this context, information theory is simply a version of statistical analysis cast in terms of logarithms, isn't it? Now the problem can be handled in many ways. Electronically, we would handle it in terms of noise power spectra and filters designed to favor the signal spectrum over the noise spectrum. In analyzing a system operating in this region, we could use a statistical analysis or a Bode diagram or probably many other methods -- the results are equivalent. The statistical approach would call the unpredictable component of the signal "uncertainty" while an electronics analysis would call it "noise" -- but the phenomenon is the same.

In this second case I doubt that we have any important divergences of understanding.

RE: the challenge experiment

>If you change, say, the contrast of a tracked target (or the
>tracking cursor), or the bandwidth of the disturbance, or put a

>grid on the visual surface...is the integration factor the only
>thing that changes?

Another way to ask this question is to ask about the conditions under which such changes would make a difference. We can try these things later, if you wish. In previous posts, the question of disturbance bandwidth has already been raised, and the answer is that the best-fit integration factor varies with bandwidth. This could be treated as a problem with nonlinearity in the system, or as a problem of information theory. I would take the approach of trying to define a nonlinear function -- perceptual or output -- as a way of making the same model work over a range of disturbance bandwidths. The least nonlinearity possible would be the addition of a square term in the output function, and that is what I would try first.

As to varying contrast or background, I suspect that these factors would not make any difference until some extreme was reached -- very low contrast, or very strong interfering background patterns. And there, information theory might surprise me with some useful predictions. This, however, would bring us into the realm of behavior near the limits of perception, which isn't a consideration in most real behaviors.

>... the challenge is really for me to show whether information
>theory could predict the change of integration factor across
>experiments, taking into account the conditions of the
>experiment. I don't know how I would do that. The integration
>factor is not a construct that I have so far analyzed in my thinking.

The use of an integration factor is motivated by PCT. You may find an equivalent motivation in IT, but the way you analyze the situation for your purposes is independent of PCT. If the integration factor isn't a natural part of your analysis, there's no rule that say you have to use one.

>In making an analysis for a specific experiment, I would use
>other experiments to obtain data--perceptual discrimination
>experiments, for instance, or control experiments that use
>presumed lower-level control systems common to the specified experiment.

You have two conditions, with data, to use as you wish -- one without feedback and one with (conditions 2 and 3, and also 1 if you can make any use of it). The PCT analysis does not need any other experiments with perceptual discrimination -- that is, it uses simple assumptions about perceptual discrimination on which its performance depends. You can make any reasonable assumptions you like. You can, for example, assume that given an output signal, the handle will come instantly to a position proportional to it. The PCT model assumes that.

I suspect you will find some difficulties at the point where the disturbance effects and the regulator's output effects converge to produce a net effect on the essential variable.

I'm glad you're going to do it. Don't feel pressed to hurry: I'll take your IOU.

Allan Randall (930308) --
RE: challenge

>In order for this to be interesting, I think you need to do one of two things:

>1) Explain in a little more detail why you think information
>theory predicts that the compensatory system will perform better.

I'm just guessing and could be wrong. The object is for an information theorist to show that I am wrong by showing exactly how IT would be used to arrive at a correct prediction.

My reason for suspecting that IT will produce the wrong result is Ashby's analysis on p. 224 of Intro to Cybernetics. On reading this, I realized that the kinds of measures involved in information theory or "variety" calculations are such that they can't predict the equilibrium state of a system with a negative feedback loop in it. As I just said to Martin above, there is a problem when two information channels come together to produce an output signal that affects the essential variable. The nature of these statistical calculations is that they can say nothing about the coherence of the two converging signal channels. If the variables are treated only in terms of information content or probabilities, they will add in quadrature, so the variance of the signal from T to E will be larger than the variance of either signal entering T. It is hard, in fact, to see how even the disturbance-driven regulator could be predicted to reduce the variance of E.

>2) Find an information theorist that actually does make this claim.

Ashby made it: the disturbance-driven regulator was said to be potentially perfect, while the error-driven regulator was said to be inherently imperfect. But I'm simply voicing a strong suspicion -- that if an information theorist were to analyze this as a problem in information flow, the wrong answer would be found. Anybody can prove my suspicion wrong right here, in public.

Best to all, Bill P.

Date: Mon Mar 08, 1993 5:04 pm PST
Subject: Greg, Allan, PCT video

[From Dag Forssell (930308-1)]

Greg:
Ray Jacksons address: This is an internet address, just like mine shown below. Just send your message to internet. It will be passed to ATTmail.

Thanks for 9011. My CSGnet files are now complete. 9011 C,D &E were rich. I am glad I caught the gap.

Allan Randalls bibliography: Thanks Allan. I shall study later.

Progress report:

On Feb 4, Christine and I presented an introduction to PCT to a Deming Users Group. This included a presentation of:

Results
Profound knowledge
Theory, Revolutions

Role play, before
Self-directing system / Control (cooperation, conflict)
(Demo's came out fine)
Role play, after
Systems thinking
Seven tools - Perception
The essence of TQM
Dr. Deming's 14 points - Integration

We have just finished editing this into a two hour video for promotional purposes.

Saturday, I mailed courtesy copies to my active advisors: Bill, Rick, Ed, Hank, Ray, Jim Soldani, Toto Grandes and of course the archive, Greg.

We look forward to comments on the net. We have already made some changes which will be incorporated in the next live talk (this Thursday). Judge by the comments as they come.

I will mail this video and supplemental booklet to anyone who asks, but cannot afford to do it for free. (We are living off savings at this stage of development). I'll ask for $10 for the tape and booklet with additional for postage as follows:

Postage to	Surface 4th class	Air
USA	$2.-	$3.-
Canada	$3.-	$4.-
Europe	$4.-	$9.-
Pacific	$4.-	$11.-

I will honor E-mail, fax or letter request and trust that you add up and send U.S. funds by snail mail.

Please note: The video is 119 minutes, 1/2 inch VHS, NTSC (U.S. video signal). Overhead slides are shown in closeups.

Dag & Christine Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808
Phone (805) 254-1195 Fax (805) 254-7956
Internet: 0004742580@MCIMAIL.COM

Date: Mon Mar 08, 1993 8:58 pm PST
Subject: Saving Trees

[From Gary Cziko 930309.0335 GMT] Dick Robertson] (930308) says:

>Bill Powers' suggestions about the Journal sound really good to me. I would
>suggest though, that there be a section of "brief Reports" for those first
>round articles that he said would usually not be good enough to publish.

>I think it would be good to brief report them so that others working on simi-
>lar studies could know what's in the works.

People can submit "brief reports" anytime they want to CSGnet. I say save the journal paper for articles worth cutting down trees for.--Gary

Date: Tue Mar 09, 1993 1:18 am PST
Subject: jacobian transpose blues

[Avery Andrews 920309.1900]

I've gotten far enough with my variant on Bill Power's 14df (for me, currently 2dfs), to begin to perceive problems with control via transpose Jacobians in Cartesian space. These arise because the Jacobians (partial derivatives of end-point movement w.r.t. joint movement) frequently don't give very good guidance as to what to do to correct the error. For example, if the arm is extended, and the target position is, say, at the present location of the elbow, the Jacobians for both the shoulder and the elbow point straight up, and so are orthogonal to the error, & the controller does nothing.

There is an easy solution to this particular problem, tho I don't know about the general case. Maybe using several different coordinate systems at once is the way to go. But I also wonder if the true solution isn't just to train a neural net to accept a set of joint angles & an error vector, & spit out a bunch of joint forces.

Avery.Andrews@anu.edu.au

Date: Tue Mar 09, 1993 5:38 am PST
From: CHARLES W. TUCKER
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Yes, please send video and booklet

Dear Dag,

See subject. Check in the mail. Regards, Chuck

Charles W. Tucker (Chuck)
Department of Sociology
University of South Carolina
Columbia SC 29208
O (803) 777-3123 OR 777-6730 FAX (803) 777-5251
H (803) 254-0136 OR 237-9210
BITNET: N050024 AT UNIVSCVM (0=ZERO)
INTERNET: N050024 AT UNIVSCVM.CSD.SCAROLINA.EDU (0=ZERO)

Date: Tue Mar 09, 1993 5:46 am PST
Subject: MCI/ATT Mails Finis; ANN. REV. PSYCHOL. atricle

From Greg Williams (930309)

Thanks to all who offered help on my MCI Mail -> ATT Mail problems. The "forwarding" (using %) trick worked fine! Thanks to Bruce Nevin especially for being first to suggest it. Gary C., you'll have to be quicker next time (:>>). By the way, Gary, the problem was not with ATT -- just arcane addressing and MCI Mail refusing to send to ATT Mail via INTERNET (except with the % trick). So now we all know....

New reference to PCT and related work: Paul Karoly, "Mechanisms of Self-Regulation: A Systems View," ANNUAL REVIEW OF PSYCHOLOGY 44, 1993, 23-52. Even cites LIVING CONTROL SYSTEMS! Wow!! And claims that it was edited by R.S. Marken!!! Ouch!!!! Oh, well, I suppose Rick could use some academic brownie points. I surely don't need any. The bottom line on Karoly: lots of devil's bibliography quotes. For example (p. 30):

> Based upon a century-old insight attributed to William James (cf Powers 1989), that humans are 'unique' in nature because they can produce consistent ends by variable means, a number of contemporary (post-1960) models of dynamic self-regulation have been developed under the imprimatur of cognitive theory, control/systems science, cognitive social learning, or European action theory. All presume that on-line regulation is a dynamic process, continuous and holistic rather than linear, built upon the operation of feedback [now hang on] (knowledge of results) and feedforward (stand-produced disequilibrium) [say what?] , sensitivity to action-produced environmental changes, the accessibility of goal representations [?], and a capacity for the selective mobilization of energy [damn straight, if they're going to move around!], attention, and relational judgment.

And so it goes....

As ever (even as Pat and the kids STILL have the flu -- remarkable!),

Greg

Date: Tue Mar 09, 1993 6:44 am PST

Please put me on the mailing list for the control systems group

thanks and regards bruce digney digney@dvinci.usask.ca

Date: Tue Mar 09, 1993 8:54 am PST
Subject: misc. replies, comments, remarks, ideas

[Hans Blom, 930309]

I'm really jealous of all you people who can spend so much time on the net. Thank you all for your replies, especially Bill Powers, Rick Marken and Martin Taylor. I just don't have the time to reply to everything that y'all said, so just a few remarks.

Rick Marken (930217.1100):

>I give. I thought you might have some EVIDENCE for your point of view.

Sorry to be so irritating, but 'evidence' is, in my opinion, a word very much like 'fact'. What is evidence to one person may not convince another. As Popper noted, evidence is an ephemeral thing. There is never enough of it to prove your point. Like Popper, I think that counter-evidence (debunking) is much more important. One counterexample may prove a theory wrong. But then again, one needs a common understanding of what would be a counter-example. In other words, there might be an infinite regression in a search for common ground. If the languages -- or the world views -- are too different, no agreement whatsoever will be possible ('three men and an elephant'). My personal solution to this predicament is to try to temporarily adopt the other's world view, however foreign it might appear, and try to make sense of what the other means. Contradictions often disappear and prove to be just different perspectives. This process has its own difficulties, of course. But sometimes this process appears to succeed and it suddenly seems as if you have another perspective on the same old 'reality' which you suddenly realize you saw only partly before. You know all this. My point is, that in engineering, as in life, we soon discover that a point of view or a solution is good iff it works for you. But here the same regression threatens: when does it work? What goals does it satisfy? It helps to have a common (scientific, engineering, natural) language. It may not be the best language we have, but without SOME common language we cannot cooperate.

Bill Powers (930217.1030):

> When we study human
>behavior, we aren't comparing it with some "optimal" or "best"
>way of controlling. We're just trying to understand what people
>are actually controlling under various circumstances. In some
>regards, people control things very well indeed, by clever means
>that surpass what any engineer knows how to build. In other ways,
>people control stupidly and poorly, and suffer the consequences.

That is not my impression. In my opinion, in the billions of years of experimentation through evolution people (and organisms in general) have found superb ways to realize their goals. If we think that they are stupid, then we are in error: we just have not properly identified their (many!) goals. This is in line with your remark that

> Much of the apparently chaotic nature of behavior
>becomes more understandable when we ask about higher-level goals.

In my world view, an organism's behavior is perfectly in line with its top-level goals. Reaching idiosyncratic goals may, of course, be hindered by the laws of nature and of society. Every organism is always at its own local optimum. Of course, we may not agree with its definition of optimum and think that it is just plain stupid. We may even have convinced the organism of that 'fact'.

I realize that this is a personal world view that can in no way be proven. Nevertheless, it is one of my basic life rules, until a better-working one appears. By the way, your use of 'suffer the consequences' applies in any case. Behavior has unforeseeable short and long range side-effects, always. Our perception is limited, although training may improve things slightly.

>The highest-level goal is to win the contest, not to jump as high as possible.

How come you know? The rules of the game are usually considered to be: when I invent a hypothetical situation, I know what goes on in that situation, because I invented it. You go against the rules here. I say, in effect, 'assume that X', and you reply 'no, I cannot assume X, I assume Y'. You do not play according to what I think the rules are.

When I think of a reason, I can only come up with the suggestion that high-jumping looks different to you than to me. Your high-jumper wants to win the contest. My high-jumper really wants to jump as high as possible; he is not interested in winning the contest since he already knows that he is by far the best of those he meets today. No, he is setting his sights much higher: he is training for the next Olympics. He has to compete not with his direct competitors this day, he has to compete with the figures in the World Records book that he studies every day. But not even that is enough. He knows that a world record holds only for six years on average. He wants to do better than that and hold the record for many years to come. He will just give this jump his very best effort.

Are these extra perceptions helpful in seeing the situation differently? You could have been right. Your understanding might have explained somebody else's behavior. But in different persons identically looking actions may result from completely different motives. A few lines later you do seem to take that position:

>You can't tell what a person is doing just by looking at what the person
>is doing.

And later again:

>My point is that pure reason isn't going to identify the actual
>variable under control by a given person in a given circumstance.
>A guess about what someone is controlling for could be quite
>right, or quite wrong.

Yes, this is the whole discussion of 'facts' versus hypotheses.

>I suspect that this is another of those myths about control, this
>time about spinal control systems. For a long time, it was
>thought that the tendon reflex had the purpose of "limiting"
>muscle tension to prevent damage.

I was not talking about the tendon reflex but about pain as a protective mechanism. I cannot agree with the following statement:

>"Pain" is not an either-or sensation; it begins at zero and rises
>from there, with some level being considered "too much" and
>calling for action to reduce it. Most "pain", I suspect, is
>really just an ordinary sensation, like the sensation of having a
>fold of skin squeezed.

You seem to have a weird conception of what pain is. Your "pain" sensors seem to be my pressure/deformation sensors. That does not correspond with my experience. When my dentist gives me a shot of painkiller, extraction of a tooth does not cause pain, but it does cause a massive sensation of "pull". For me, these are different things. In my opinion, pain is a stimulus generated by our body when attention needs to be drawn to more or less massive ongoing or threatening distruction of bodily tissues. Pain, in this view, is not normally present. Pain, moreover, is a strong motivator to get away from the situation that brings it forth if still possible and, probably even more important, a very strong motivator to avoid similar situations in the future. This is slightly paradoxical: we want to keep away from something that we do not feel. Isn't such a construct, control relative to an 'imagined' reference, possible in PCT?

About my formula

> x (t + T) = a * x (t) + b * u (t) + e (t)

you say:

>Why not the position of the car relative to a point 1 foot to the
>right of the middle of the road? You're sneaking a reference
>condition into this argument without mentioning it. ...
>In a steady crosswind, e certainly does not have an average value
>of zero. Assuming disturbances with an average value of zero
>conceals the real control problem -- such as standing up in a
>gravitational field.

You are right, of course. However, in control systems a CONSTANT never provides a problem. Control of CONSTANT disturbances is trivial; one might not even call a constant disturbance a disturbance. Add a constant to the formula:

x (t + T) = a * x (t) + b * u (t) + c + e (t)

The constant c can be given three different equivalent interpretations:

1. it is an offset for x, say 1 foot to the right of the middle of the road;
2. it is a constant disturbance, say a steady crosswind or a gravitational field;
3. it is an offset for u, say a base rate of its metabolism to keep the engine running.

The third interpretation shows why we have a non-problem.

>So this control law will leave the car weaving back and forth
>from one side of the road to the other under each gust of wind,
>with the driver attempting to steer only when a limit gets too
>close. If this is how you drive, I'm not sure I would like to be
>a passenger! Nor do I think that this behavior would look much
>like the way a real human driver steers a car.

When you work out the control law, it shows up as 'behavior' that does not attach much weight to BEING AT at the middle of the road, but a great deal to GOING TOWARDS the middle of the road. Mathematically, those are very different things.

>   The simplest kind is a reference setting of zero.
>If you set your reference level for the perception of a loose
>tiger to zero, then any perception of a loose tiger constitutes
>an error, and you will act to reduce the perception of the tiger
>to zero by moving it away or yourself away from it.

My problem: if you "set your reference level for the perception ... to zero", then, if you succeed, you see nothing. How can seeing nothing tell you where you are? "No", you say, "you still use any perception ... by moving away". You seem to have to balance on the edge between light and darkness. My problem with your solution is that you are stuck there.

I have encountered the same problem in engineering. We are currently designing a muscle relaxation control system for use in the operating room. Its goal is, of course, to abolish motion reflexes that might cause the surgeon's knife to cut where it should not. Muscle relaxation is estimated from the EMG (obtained from a muscle in the hand) that is evoked by applying a supramaximal stimulus (actually a train of four short pulses) to the nerve going to the muscle. Measurements have shown that the sensitivity of patients to the drug can vary a great deal. Some patients require little of it, some a lot. Giving a massive dose of the drug is certain to provide relaxation in all cases but this is regarded a primitive method: overdosing is not nice practice. Moreover, it may be harmful to the patient if he is extremely sensitive or allergic to the drug. One would therefore like to just barely abolish the response. This approach is unacceptable, however: going from 0% to 100% relaxation by infusing the drug _in a controlled way_ takes upto 15 minutes. Current manual practice is to give an initial bolus dose that is barely enough for the least sensitive patients but overdoses the more sensitive ones. This takes about 3 minutes in the worst case. The difference of 12 minutes at the start of surgery cost thousands of dollars; surgeons do not come cheap. A control system must therefore adopt the current practice method, even if it has to temporarily relinquish control. A consequence is, that for tens of minutes or even several hours the controlling system does not know "where the patient is" and therefore cannot answer questions like "when we stop the infusion now, how long will it take before the patient is able to breathe again", which are of great practical importance.

>   human beings
>roam free through an undisciplined environment that is far more
>complex than any of them can understand. That environment is also
>full of disturbances that can't be predicted (weather, for
>example) or even be sensed before they occur. Most of our
>"predictions" are statistical in nature; sometimes they work and
>sometimes they don't. So there's no way that living systems could
>evolve to anticipate every circumstance or act correctly every time.

That is not my point. My point is that the human perceptual + conceptual systems are so beautifully designed that they even extract information from very 'noisy' perceptions. It is extremely common to 'filter' stochastic observations in such a way that the information contained in them is preserved whereas the noise is discarded. Kalman filtering is one of the engineering methods to do this. Originally designed for satellite tracking and position control, it is a technique to extract the maximum possible information from few and imprecise measurements. Mathematically, it is similar to 'real-time' ANOVA. What is does is to 'pack' a great number of observations into just a few numbers through a least squares averaging process.

>My job is actually easier than yours. I'm not trying to optimize
>anything -- just to match the behavior of a model with that of a
>real human subject.

Have to be precise here: our jobs are very similar. You ARE trying to optimize something: you are trying to find an optimal match between a model and a real human subject.

>Of course real control engineers know a lot more than I do about
>the design of complex control systems ...

Maybe, maybe not. Anyway, that extra knowledge may not account for much when it comes down to designing good control systems. After all, there is not much good theory around to travel by. 'Feeling' and 'intuition' are required as substitutes for knowledge. I don't think you lack those. I have to agree with Avery.Andrews (930220.1130):

>On the topic of `real' control engineers...
> There may simply not be much in the way of theorems that
>help with understanding how complex living control systems work.

Martin Taylor (930218 10:40):
> I always thought
>psychology was essentially a problem in engineering, which is why I seemed
>to switch fields (according to society--I never thought I switched).

Great! Agreed! But how about this: once in a while I sit back and take the opposite perspective and consider engineering a problem in psychology. Why, for instance, do we think that our current scientific approaches are so great? Why do we believe in grand unifying theories? I frequently see chance piled upon chance when I ponder questions like: why are there humans? Where do I come from, biologically and mentally? Where do our theories come from? Are other approaches feasible? Attribution theory is one of the psychological tools that 'explains' why similar perceptions lead to different conceptions (or higher level perceptions) in different persons or in the same person under different conditions. What is pure sensory input and what are the personal 'illusions' that I mix in. Infinite regres- sion again, yet a process akin to vacuum-cleaning my mind.

>Where there is no feedback, CSG-L tends to use terms such as "affect,"
>"influence," "linkage," and the like.

The question of 'control' versus 'affect' seems to have to do with either intended versus unintended or full versus partial correlation. In either case, it has to do with our limited predictive powers. The first raises the question what it means to 'intend' or to have 'goals'. The second raises the problem that actions will always have effects in addition to those 'intended'. Control must always be limited; the world is just too complex for our three pounds of brains to model it and our fifty pounds or so of muscles to subdue it.

>>I would call this a selection between different possibilities of action.

>That's exactly what I would call "decision." Do we have another source of
>confusion based solely on a different dictionary? What do you mean by
>"decision?"

For me, decision has the connotation of 'willed' or 'conscious'. I wanted to avoid that. Why? Well, from modelling theory we know that, whatever we want to model any part of the world, an infinite variety of models is possible. We also know that the simpler the model, i.e. the fewer the number of degrees of freedom that it contains, the faster it converges. But simple models of complex realities are of course less accurate. In every model built up from real world information there is a compromise between detailedness and accuracy. In particular, in a changing world a very detailed model is impossible; convergence will not be reached, i.e. the variance of the parameter estimates will remain large. In information theoretic terms you might say that the (limited) available information is (often about equally) distributed amongst the degrees of freedom of the model.

That is why frequently different models are employed. A very detailed one can be obtained and used if the world that it is about proves to be stationary (a posteriori; if you know this a priori than you need only one model). When the world changes rapidly, only a coarse model can be obtained. How selection amongst the models takes place I do not know; somehow the most appropriate one must be selected. Child psychologists do indeed find a developmental sequence of models which become more and more detailed the more experience the child gathers. Sometimes, when the current model does not work, there is a 'regression' to an earlier model. This regression is no tragedy, at least not for the individual that shows it; it provides the opportunity to backtrack and start a different, hopefully better, new model.

Dick Robertson (930219) points to the selection mechanism:

>... Daniel Dennett ... proposes that our thinking consists of "multiple
>drafts" in which different control systems (my terminology, not his) in
>the brain compete for consciousness...and therefore we often don't know
>what we think until we hear ourselves say something. I thought of it in
>terms of different versions of a program competing to satisfy some error
>in a principle-level system.

>>Luckily, humans are wired in such a way that they can sense their
>>outputs; this is called the "body image".

>No. They sense inputs from many sensors, some of which are detecting the
>conditions inside the body. But it is quite possible (and discussed in
>BCP) that the anticipated (imagined) effects of outputs can be used as
>inputs through what you call models.

A model basically consolidates observed correlations. If I want to know what the world is like, I'd better correlate my actions on that world with the way that world reacts to my actions. It is therefore important to me to have an accurate sensory picture of what I do to the world. And I do seem to have the sensors for that, in particular muscle spindles, Golgi tendon organs, skin pressure and temperature sensors. These provide a sensory image as far from the body and as close to the world as is possible. That is what I meant.

Bill Powers again:

>When I think of the "output" of a system, I mean the physical
>effect on the environment that is due to the actions of the
>behaving system ALONE. In the human system, this would mean
>muscle tensions, because that's that last place in the chain of
>outgoing effects where environmental disturbances can't get into
>the process and alter the consequences.

Very much in line with what I said above.

>So this is more a matter of labeling than ideology. I'm sure you
>would agree that a servomechanism doesn't control the torque
>applied to the armature of its motor, but only some consequence

>of that torque measured farther downstream in the causal chain.

and

>I think that my way of defining output and control is the least ambiguous.

This IS a matter of labelling. In engineering, we take great liberty in defining inputs, outputs and systems. I can take for an input anything that I can manipulate and for an output anything that I can measure. A system is anything in between. Ambiguity does not appear as long as we clearly state which is which. One person's choice may differ from another one's, but in your example that appears hardly relevant as long as the relations between different output choices are rigid (mathematically: can be one-to-one transformed into each other).

>The only aspect of a control loop that is under reliable control,
>therefore, is the sensor signal.

In practice, this is true only if the sensor signal is ever-present and noise-free. If the sensor signal contains (much) noise, as is often true in engineering applications, some means is required to separate the signal from the noise. You often call such means an input function. There are two general approaches to getting rid of the noise: averaging over multiple sensors (redundancy) or averaging over time (filtering).

If the sensor signal may be absent for shorter or longer periods, a model (in the sense of a built-in or acquired approximation of the object to be controlled) is required that temporarily provides an alternative means of pseudo-feedback. Such models are usually not highly accurate and therefore drift will occur away from the optimal operating point.

An easy-to-do experiment to demonstrate both phenomena:

The next time you go for a walk, try the following in a well-known, unsurprising environment.

First, close your eyes while you walk. Your visual perception, though not completely absent, will be so disturbed that you have to fall back on the feedback that other senses provide. You will feel a strong urge to slow down and you will undoubtedly do so: you need more time to do the same thing (averaging over time). [Try the test of averaging over multiple sensors when you can find a number of companions: link arms, eyes closed. Your common walking speed will tend to go up.]

Second, force yourself to walk in your natural rhythm. Open your eyes briefly whenever you feel like it, but try to do so as infrequently as possible without feeling uncomfortable. You will probably find that you will have to open your eyes only very briefly (like an inverse 'blink') every three or four steps or less. An internal model (memory) provides the required additional pseudo-feedback. A similar notion of yours is the 'imagination mode'.

Martin Taylor (930223 14:20) hints at the same things in his answer to Bill Powers (930223.0800):
>When a control system is trying to keep a percept near a reference, it
>can do so only to the extent that it can imagine or analyze the incoming
>sensory data.

>The idea is that reference signals ARE played-back recordings of
>perceptual signals, in organisms. This will remain only an idea
>until someone does the implied experiments, to see if reference
>signals are ever set to values that have never been experienced.
>I make no predictions one way or the other.

In adaptive control systems, we can discriminate between two types of reference signals: those that are hardwired (analogues to hunger, thirst and bodily integrity) and those that are acquired (analogues to social conventions). In A.I., the former are sometimes called primary goals, the latter can be called secondary, derived or contingent goals. The former must always be met, regardless of outside conditions, the latter depend on external circumstances to a much greater degree. Hints of nature and nurture...

Best to all, Hans Blom

Date: Tue Mar 09, 1993 9:42 am PST
Subject: Correcting many errors

[From Rick Marken (930309.0800)]

Well, now that I'm back on the net I seem to be canonizing some unintended targets. Here follow some mea culpas.

Oded Maler (930308-ET) -- re: my brief coment on Agre conference.

>I wonder how this has escaped Gary's censoreship - maybe you should
>augment your mailer with some time lags.

Gary is not to blame; I can now get directly to the net, circumventing Gary's wise editorial hand. My comment on the Agre conference was a cheap shot; probably not the best way to test my ability to access the net (that's what I was doing). But, as Avery mentioned, Agre's description of the conference did include a cheap shot at control theory. I know we (PCTers) should be more gracious than the "opposition" -- but I'm afraid I just lose it occasionaly. What do you expect from a loose canon?

I said to Martin Taylor:

>Wanting to stick with a grand old theory is a very common phenomenon -- and
>it's the reason why 1) most psychologists don't get into PCT and 2) if
>they do, they don't get PCT right. It's why we say PCT is revolutionary.

Martin Taylor (930308 12:00) replies;

>Quite right. I don't know if I get PCT right.

I am really sorry. I did not mean to imply that you (Martin) do not get PCT right. In fact, I think you DO get it right; and you have made some excellent suggestions about extending the model. I meant that there are people (NOT YOU) who don't get the fundementals of PCT right (control of perception, S-R as disturbance resistance, test for controlled variables) due to their desire to integrate PCT into the framework of their favorite "grand old theory" that just MUST be right. If anything, you are working from the opposite direction -- trying to integrate IT into PCT. I still don't see the value of that effort -- but I'm willing to be convinced. But I think that you DO understand the fundementals of PCT very well; and I did not mean to imply that your interest in IT diminishes your understanding in any way.

Which leads me to:

Bill Powers (930308.1520) --

>What I like about you, Martin, is the civilized way you respond
>when people attack your life's work with bludgeons.

Apparently my comments to you were perceived by others as bludgeons as well. I certainly did not intend them as such. I am very interested in seeing what IT can contribute to PCT (skeptically interested, but interested nonetheless). But I don't mean to be attacking Martin's life work -- which would be real silly because it seems to me that Martin's life work is the same as mine; promoting and extending the PCT model of human nature.

Greg Williams (930309) --

>New reference to PCT and related work: Paul Karoly, "Mechanisms of
>Self-Regulation: A Systems View," ANNUAL REVIEW OF PSYCHOLOGY 44,
>1993, 23-52. Even cites LIVING CONTROL SYSTEMS! Wow!! And claims that
>it was edited by R.S. Marken!!! Ouch!!!!

That is really weird. I don't have easy access to the ANNUAL REVIEW but I wonder, did Karoly also refer to the Americal Behavioral Scientist PCT issue that I actually did edit? It seems like this must be the source of the confusion. Whether or not he referred to ABS in the REVIEW article, the confusion at least suggests that Karoly was poking through some of the literature of PCT. I guess I'll have to try to get a copy of the article. That may take some time; in the meantime, could you (Greg) give a brief review of Karoly's REVIEW?

Best Rick

Date: Tue Mar 09, 1993 11:15 am PST
Subject: Just TRY to correct THESE errors!?!?

From Greg Williams (930309 - 2) From Rick Marken (930309.0800)

>That is really weird. I don't have easy access to the ANNUAL REVIEW
>but I wonder, did Karoly also refer to the Americal Behavioral Scientist
>PCT issue that I actually did edit?

No.

>It seems like this must be the source of the confusion.

I doubt it!

>Whether or not he referred to ABS in the REVIEW article, the confusion at
>least suggests that Karoly was poking through some of the literature of PCT.

"Poking," indeed.

>I guess I'll have to try to get a copy of the article. That may take some
>time; in the meantime, could you (Greg) give a brief review of Karoly's
>REVIEW?

I thought I already did. What a glutton for punishment... suffice it to say that you aren't going to like it. The main sections are "Paradigm Variations and Basic Models," "Proximal Volitional Regulators," "Metaskills and Boundary Conditions," and "Self-Regulatory Failure: Concepts and Dynamics." Karoly's "bottom line," as opposed to mine: "As a concept akin to 'getting one's life together,' self-regulation has not achieved a simple or uniform paradigmatic embodiment, nor should we expect this in the foreseeable future." (p. 45) Gee, I can't imagine why not.... :-!

As ever,   Greg

Date: Tue Mar 09, 1993 12:22 pm PST
Subject: Re: Defining information; challenge

[Martin Taylor 930309 11:30]   (Bill Powers 930308.1520)

>>Probability is something relating only to the information
>>available at the point where that information is used. It
>>relates to individual possible events that might be detected at
>>that point, not to frequency.
>
>I wonder whether you have considered this point yet:
>
>In your terminology, the perceptual signal in an ECS, an
>elementary control system, stands for the state of a CEV -- a
>complex environmental variable. As we all seem to agree now, that
>CEV is a construct created primarily by the forms of all the
>perceptual functions lying between the primary sensory interface
>and the place where the perceptual signal finally appears.

And Bill goes on to say that:

>the CEV is a construction by the nervous system. From the internal
>point of view, the perceptual signal is always a PERFECT
>representation of the CEV -- in fact, that signal plus all the
>perceptual functions leading to its production defines the CEV.

This is correct.

Between these two statements, Bill brings in an apparently contradictory one that also seems correct, so there is an issue:

>It is natural, then, to say that the perceptual signal contains
>information ABOUT that aspect of the environment.
>
>From there, we can go on to ask how well the perceptual signal
>represents that aspect. This is where the difficulty starts. In
>order to answer that question, we must have an independent

>measure of the aspect of the environment in question, one that
>tells us its true state so we can compare that state with the
>representation in the form of the magnitude of the perceptual
>signal (or any other form, come to think of it).

The issue is, as usual, one of viewpoint. From the outside view, there is a complex in the world that seems to be what the "subject" is controlling. It is the experimenter's view of the putative CEV. The theorist outsider can also "see" the subject's perceptual signal that is the actual controlled variable. As far as the subject is concerned, that signal IS the CEV. It is all that the ECS in question can know about the state of the world.

There are various kinds of "outsiders" as Bob Clark has pointed out. One of them is the DME, which views all sorts of signals in the hierarchy. All outsiders use their own perceptions rather than the one actually being controlled by the observed ECS. It is from the outsider's viewpoint that we can see a dichotomy between the CEV in the world and the perceptual signal. The subject cannot see it.

The outsider, who may be using very precise measuring instruments, can see that there are discrepancies between state of the putative CEV and the state of the derived perceptual signal, even if the total perceptual input function is correctly interpreted. These discrepancies have to do with the resolution of the perceptual system. The subject may not be able to detect that any individual discrepancy exists, but may be able to detect the possibility that discrepancy exists, by virtue of the success of control. (This is much the same in principle as the way astronomers judge the numbers of meteor craters on the moon that are smaller than they can see, or the way ecologists judge the number of species never yet identified).

The perceptual signal, in this way of looking at things, does not define the CEV. It defines the operations on the sensed world that create the CEV, but the CEV is a structure in the world, not in the mind. It is a conceptual structure that mirrors the mind, and it may not be detectable to anyone else than the mind that created it, but nevertheless, it is in the world not in the mind. For example, a CEV may be "the distance between my fingertip and my nose." Forgetting the irregularities of skin and the like, there is a perceived value for that CEV--the perceptual signal that corresponds to it. If I hold up my finger, I may perceive that distance as stable (or nearly so, with a slow drift), but I know from other information that if I could only see it there is a rapid oscillation in the distance. Someone with a laser interferometer could probably measure fluctuations that are not in my perceptual signal. But I would say that they are in the CEV that the perceptual input function determines. So, the CEV is not defined by the perceptual signal; it is represented by the perceptual signal. It is defined by the perceptual input function.

There's a hidden issue here, one that relates to reorganization. There is no CEV that corresponds to the function that causes the actions of the subject to control the intrinsic variable. Reorganization controls the control operations, but it does not work on any perceptual signal in the usual sense; a perceptual signal based on a function of sensory input variables. Reorganization works, but it works only because the behaviour of the world (unperceived) is factually stable over periods longer than the time it takes to reorganize. That factual stability can be inferred from the success of the reorganization. It cannot be perceived (I'm tempted to say "in principle" but I don't know if I could argue that). An outsider with a perceptual function that operated over a long time scale (I include memory here) could perceive the stability that permits reorganization to happen. Likewise, with a normal perceptual signal and its corresponding CEV, an outsider could perceive discrepancies between the CEV and the perceptual signal that represents it, even though the user of the perceptual signal cannot. But as with reorganization, the user of the perceptual signal might possibly infer that there is a factual discrepancy.

I realize that the word "factual" in the above paragraph raises its own issues about Boss Reality and the like. I assume that all such issues are resolved against the solipsist position.

Martin

Date: Tue Mar 09, 1993 3:32 pm PST
Subject: Re: Correcting many errors

[Martin Taylor 930309 18:00]   (Rick Marken 930309.0800)

I didn't see you as using bludgeons. In fact, I took Bill's comment to refer more to his own posting(s) (perhaps in context of yours). I was quite serious when I said that I don't know if I get PCT right. It's theoretically impossible for me to know that. I don't even know whether you or Bill get it right, and I think both of you have said from time to time that you yourselves don't know. If you did, you could make ex cathedra statements about how it works, rather than performing your experiments. I happen to appreciate experiments and experimental results, but am not good at proposing and designing experimental tests of theories. I'm much better at seeing the patterns in other people's results and in developing theories that fit them into a wider field of enquiry.

It's that width of validity that attracts me to PCT, at least as much as (and probably more than) the precision of fitting experimental results. Adding information theory may or may not enhance PCT by helping improve experimental data fitting within its own range of applicability, but since IT applies anywhere, well beyond the claims of PCT, then if PCT can be integrated within it (or vice-versa), then both are enhanced. It always helps a theory if its foundation is set from outside its range of claim, independently of the observations on which it is based. In Occam's razor terms, the foundation then comes for free, leaving only the domain-specific additions as contributions to the length of the hypothesis. That can be a great enhancement to the credibility of the theory. That's the heart of what I want to do in IT->PCT, but if improved data fitting also happens, then that's quite a bonus.

Martin

Date: Tue Mar 09, 1993 3:40 pm PST
Subject: The heart of the matter

[From Rick Marken (930309.1500)]   Bill Powers said:

>The only aspect of a control loop that is under reliable control,
>therefore, is the sensor signal.

Hans Blom (930309) answers --

>In practice, this is true only if the sensor signal is ever-present and
>noise-free. If the sensor signal contains (much) noise, as is often true in
>engineering applications, some means is required to separate the signal
>from the noise. You often call such means an input function. There are two
>general approaches to getting rid of the noise: averaging over multiple
>sensors (redundancy) or averaging over time (filtering).

But then it is this filtered sensory signal that is controlled, right?

>If the sensor signal may be absent for shorter or longer periods, a model
>(in the sense of a built-in or acquired approximation of the object to be
>controlled) is required that temporarily provides an alternative means of
>pseudo-feedback. Such models are usually not highly accurate and therefore
>drift will occur away from the optimal operating point.

And again, it is these model produced sensor signals that are controlled.

I don't get you point here, Hans. When a system is controlling (holding some aspect of the environment in a fixed or variable reference state while resisting the effects of disturbances) then it is the sensory (perceptual) signal (whatever its cause) is really controlled. If there is lots of noise or intermittancy in the system that there will be NO CONTROL. In this case the sensory signal is, indeed, NOT CONTROLLED because NONE OF THE VARIABLES IN THE LOOP ARE CONTROLLED.

The central insight of PCT is that WHEN THERE IS CONTROL THEN IT IS THE FUNCTIONAL EQUIVALENT OF THE SENSORY SIGNAL (PERCEPTION) THAT IS CONTROLLED. This is not just true of living systems; it is true of ALL NEGATIVE FEEDBACK CONTROL ORGANIZATIONS when they are IN CONTROL (that is, when their loop gain and dynamics are set appropriately so that the system is STABLE and keeping a variable at a fixed or varying reference against disturbance). The behavior of every single control system IS THE CONTROL OF PERCEPTION. The evidence for this fact about control systems is

1) mathematical; solving the simultaneous equations for a negative feedback control system we find that

p = r

(the perceptual signal, AND ONLY THE PERCEPTUAL SIGNAL, depends on the setting of the reference signal -- or setpoint for the controlled variable).

2) experimental; varying the sensory representation of the controlled variable changes the environmental level at which it is maintained. For example, wearing prism glasses changes the arm position that is counted as pointing at a target.

This fact about control system operation must be of little interest to a control engineer. But it is of FUNDEMENTAL significance a student of the behavior of living systems.

Your statement above gives the impression that you don't completely agree with the idea that control ALWAYS and ONLY involves the control of a sensory (perceptual) variables. It sounds like you are suggesting that other variables, besides sensory variables, can be controlled in a control loop. I would really like to know what these variables are because, if they exist, it would take an awful lot of the wind out of the PCT sails.

Best Rick

Date: Tue Mar 09, 1993 4:05 pm PST
Subject: Re: Just TRY to correct THESE errors!?!?

[From Rick Marken (930309.1530)]

Greg Williams (930309 - 2) on the Karoly REVIEW --

>What a glutton for punishment... suffice it to
>say that you aren't going to like it.

The chapter headings look pretty grim. Did he put down PCT in particular? If he gave a real embarassing review of PCT then maybe you dodged a bullet when he attributed editorship of LCS I to me instead of you.

Love Rick

Date: Tue Mar 09, 1993 6:46 pm PST
Subject: Replies to Hans Blom

[From Bill Powers (930309.1800)] Hans Blom (930309) --

RE: Evidence

Yes, people do disagree on what constitutes evidence. I don't think that this is any reason to abandon the idea of experimental test. Common sense has to come in here somewhere -- if I try to sell you the Statue of Liberty, and you have any interest in buying it, you'll want some evidence that I have the right to sell it. You won't let me off just because I start a philosophical argument about the nature of evidence.

Or will you? As a matter of fact, I happen to be selling shares in some nice real estate in South Central Florida ....

RE: Optimal control systems

I don't think many evolutionists would agree with your statement

>... in the billions of years of experimentation through
>evolution people (and organisms in general) have found superb
>ways to realize their goals.

Evolution doesn't optimize anything; it just weeds out unworkable organisms. What's left is just barely good enough to survive -- for a while.

I would have to agree with your implication that organisms control as well as they can. That's a matter of definition. But in looking at the state of our world, I am not greatly impressed with the way people control for social harmony, economic viability, or maintenance of an environment fit to live in.

>In my world view, an organism's behavior is perfectly in line
>with its top-level goals.

I think you're defining top-level goals from outside the organism. When I speak of goal-seeking I'm not normally thinking of "goals" like maintaining the life-support system and combating invasive microorganisms, or even "surviving," the unlearned goals that I assume to drive reorganization. I'm thinking more in terms of the learned goals, things like being a good person, making a decent living, and so forth. I don't think that people are particularly adept at constructing systems of goals that hang together, are consistent with each other. Most of the people in the world live in poverty, hunger, and illness. I don't see how you can claim that they are optimal control systems.

In offering alternatives to the highest-level goals you suggested (jumping as high as possible), I wasn't denying that some people might actually have the goal of jumping as high as possible. I was only pointing out that other goals are equally plausible, and in my experience more common (particularly when you ask what the _immediate_ goal is). In explaining to me that in different persons identical actions may come from different motives, you're simply echoing my point.

>You seem to have a weird conception of what pain is. Your
>"pain" sensors seem to be my pressure/deformation sensors.

Let someone put a fold of your skin in a pair of pliers, and then start increasing the squeeze. At low levels of squeeze, you will simply detect the amount of squeeze. But I predict that at some level of that sensation, you will say "Ouch!"

>In my opinion, pain is a stimulus generated by our body when
>attention needs to be drawn to more or less massive ongoing or
>threatening distruction of bodily tissues. Pain, in this view,
>is not normally present. Pain, moreover, is a strong motivator
>to get away from the situation that brings it forth if still
>possible and, probably even more important, a very strong
>motivator to avoid similar situations in the future.

If a sensation rises above a certain reference level, you will take steps to reduce it. Isn't that how we define pain? For some sensations, like stimulation of a nerve in a tooth, that reference level is set quite low. For others, like pinching the skin, it is set fairly high (although the reference level varies according to the part of the body involved).

It isn't the sensation that serves as a motivator (perceptions don't motivate anything). The reference level determines what level of the sensation will be enough to produce efforts to reduce it.

RE: zero reference levels

I don't see any particular problem here. As you say, setting a reference level for a particular perception to zero is not very practical, because you'll end up with the perception oscillating above and below zero. Better to set it at a very low but nonzero level, so you can maintain control of it at a harmless level.

In your anthesthesis problem, you don't want any reaction from the patient at all, so you have to try to guess how much anthesthesia is required without administering too much. And you don't want to use the patient's reaction as a gauge, because you don't want any reaction at all. Obviously, tight control is impossible; you just do the best you can.

>This is slightly paradoxical: we want to keep away from
>something that we do not feel. Isn't such a construct, control
>relative to an 'imagined' reference, possible in PCT?

I think you must mean controlling an imagined perception relative to a reference level. Reference signals are just reference signals, generated by the next level up, not by the environment.

Controlling an imagined perception is possible, but doesn't give you any control over a real one. What you can do is interpret signs from the environment as an indication that an unwanted perception is likely to occur, and correct the environment to a state in which you predict that it will not occur. Of course if you continue to do this, the most likely outcome is superstitious behavior. You have to check now and then to make sure that if you don't take the preventive step, the unwanted perception will in fact occur. The most persistent forms of superstitious behavior arise when the unwanted perception implies great danger or possible death. Then you're stuck with the superstition because you don't dare test it.

To break out of that sort of supersitious behavior, all you can do is improve your understanding of nature. You wouldn't try experimentally stepping on a crack to see if it really would break your mother's back as a way of testing that supersition. But if you understood how the world works a little better, you might decide that there's no basis for the rule, and simply cancel it. The less you understand, of course, the harder it is to give up any superstition.

RE: your control formula

>... in control systems a CONSTANT never provides a problem.
>Control of CONSTANT disturbances is trivial; one might
>not even call a constant disturbance a disturbance.

But in your design, you have to assume that the average value of the disturbance is zero. If you don't assume that, will your formula work?

This is a prime candidate for a working model. I claim that your design for a model of steering will end up with the driver oscillating between the edges of the road, whereas a PCT-type design will stabilize the car on some path well away from the edges. The two become equivalent, by the way, if there is a region on the road where both edge-avoiding systems are simultaneously active. Then you get a conflict system with a virtual reference level in the middle.

RE: optimization again

>You ARE trying to optimize something: you are trying to find an
>optimal match between a model and a real human subject.

You're a pretty slippery customer. What you say is true: I'm controlling for the best fit between the model and the real behavior. Achieving this requires the same sort of trial and error that tuning a radio or focusing a lens requires, because the amount of error doesn't tell you which way to move, and there's no a priori

way to specify the magnitude of the effect at the minimum (or maximum). This sort of control does happen. It's not very common. And it's not very tight.

Same subject:

>My point is that the human perceptual + conceptual systems are
>so beautifully designed that they even extract information from
>very 'noisy' perceptions.

They do that only as well as the statistics and the accuracy-time tradeoff permits. I don't worry much about extracting signal from noise; most of the behaviors we observe work at signal levels where noise can be neglected.

>Control must always be limited; the world is just too complex
>for our three pounds of brains to model it and our fifty pounds
>or so of muscles to subdue it.

Well, I won't be nasty and remind you of how wonderful our evolved control systems are supposed to be.

What's really wrong with your statement is the implication that it's hard to find instances of good control. Control is, to be sure, limited -- but it's hard to find examples of behavior in which control isn't pretty good by anyone's standards. "Limited" is one of those qualitative terms; the importance of the limits depends on quantitative definitions. Human motor behavior works with a bandwidth of only about 2.5 hz -- certainly too limited to enable us to balance one end a stick one inch long. On the other hand, this bandwidth seems to be just sufficient to handle most of the disturbances that actually occur on scales that matter to us. On those scales, the limitations are irrelevant.

RE: input and output

>In engineering, we take great liberty in defining inputs, outputs and systems.

I think this is one of the reasons that engineers failed to come up with PCT. When you're focussed on producing some outcome in the environment, there's no organizing principle for laying out the control system. You can put your stabilizing filters in the input function, or add little loops anywhere you like that will do the job. The result is that there are no real principles of design in control engineering (that I know of). There are plenty of principles, but none having to do with how to design the functions of a control system in some systematic way. Basically, you kludge up a design that looks as if it will work, and then buckle down to analyzing what you designed.

The PCT approach is to define the problem in terms of sensed variables: it is the sensed variable that will ultimately be controlled, so it should represent something specific in the environment to be controlled. The engineer can violate this principle, because the engineer knows what is to be controlled. But if the control system is in an organism, its perceptions have to be useful in a variety of higher-level systems, and can't have haphazard relationships to the outside world. This forces the modeler to propose a consistent set of definitions of input, output, system, and environment.

I think that a little more systematicity would also help control engineers, but that's their business.

Best,   Bill P.

Date: Tue Mar 09, 1993 9:12 pm PST
Subject: Psycholoquy Ad

[from Gary Cziko 930310.0235 GMT]

Avery.Andrews 930309.932 offered the following as a start for the Psycholoquy ad to which I added some stuff and turned his parenthetical observation into an asterisked footnote.

==
ANNOUNCING THE CONTROL SYSTEMS GROUP NETWORK (CSGnet)

It is generally thought that feedback control is well-understood and thoroughly integrated into the modern behavioral sciences. But there are some indications that it is not in fact quite as well-understood as one might expect.*

CSGnet is an electronic forum for the investigation and application of closed-loop models to all areas of psychology, broadly construed, following the general lead of the work of William T. Powers and his associates. CSGnet links together over 120 individuals in 17 countries and its participants include experimental and clinical psychologists, counsellors, educators, linguists, sociologists, control-system engineers, cognitive scientists, and roboticists, among others.

CSGnet exists as an unmoderated LISTSERV group as well as a NetNews (Usenet) group. The LISTERV address for CSGnet is LISTSERV@VMD.CSO.UIUC.EDU or LISTSERV@UIUCVMD.BITNET and the Usenet group is BIT.LISTSERV.CSG-L. More information about CSGnet can be obtained from Gary Cziko <g-cziko@uiuc.edu>.

*For example, with the revival of interest in 'Central Pattern Generators' (CPGs) in the 1970s, one possibility that does not seem to have occurred to people is that a CPG might function by delivering time-varying set-points to feedback mechanisms, thereby reducing the need to pre-compute parameters for motor-programs. In fact, even today, it is not clear that this rather simple possibility is clearly formulated in the standard literature, and adequately investigated in cases where it has some degree of prima facie plausibility.
===

So should I try to get this posted on Psycholoquy? Any suggestions for changes? If we can get this on Psycholoquy once, we may be able to get it on periodically in which case we can try various approaches and tones from the soft-sell one here to a more radical ones.--Gary

Date: Tue Mar 09, 1993 9:12 pm PST
Subject: Bounced Messages from CSGnet

[from Gary Cziko 930310.0250 GMT]

One of the joys of "owning" CSGnet is that I am priviledged to receive all CSGnet messages that never reach their final destination. Since it is not uncommon for five or so addresses to be out of commission at any one time and since it is not uncommon for 10 or messages to be posted to CSGnet in a day, it is not uncommon for me to receive 50 or more undelivered messages in any given day. This amount can only grow as CSGnet grows, so there is some incentive for me to eliminate addresses from CSGnet which are inoperative for any substantial period of time.

Usually the listserver software will try to deliver returned messages for a period of three days so if your system is off-line for three days or less it is likely that you will not miss any CSGnet traffic.

If you are unconnected for longer than three days or if you suspect that you have missed some messages which you wish to recover, you should request the appropriate log from the listserver. For example, sending the following command as the text of a message to LISTSERV@VMD.CSO.UIUC.EDU will result in your being sent all messages disseminated via CSGnet during the first week of March 1993 (i.e., March 1 through 7).

GET CSG-L LOG9303A

Substituting LOG9303B will procure the second week of March 1993. Remember, send these commands to LISTSERV, NOT to CSGnet (the commands are ineffective no matter how many humans read them). The log is continually updated so you don't have to wait until the end of the week to request the current log file. Note that there is a limit as to how much info the listserver will send you over time, so that if you request a log file at noon you may have to wait until 10 pm or so before your request for another one will be honored.

You may also be able to recover missed message to CSGnet if you have access to NetNews (Usenet). But not all NetNews systems receive bit.listserv.csg-l and those who do keep posts for different durations (my system keeps them for 10 days).

If you are unconnected for more the one week I may remove you from CSGnet. It will then be up to you to either resubscribe directly to the CSGnet listserver or send me a request to do it for you.

I hope that this mode of operation seem reasonable. My goal is to have individual subscribers take on the responsibility of recovering missed messages and dealing with service outages so that I can lead a close to normal life. But still don't hesitate to call on me if you are in a bind. CSGnet is a friendly place and I want to be a friendly listowner.

Your faithful CSGnet servant, Gary

Date: Wed Mar 10, 1993 2:54 am PST
Subject: Re: Psycholoquy Ad

[Avery.Andrews 930310.2036]

It might be an idea to wait a bit till Bill's 14df project is more consolidated - I think there are a lot of practical lessons in the application of PCT to complicated situations there (e.g. why Bill's design works so much better than `naive cartesian transpose Jacobianism'), and it might make a better impression on any new recruits to have a reasonably extensive & coherent writeup of this.

Avery.Andrews@anu.edu.au

Date: Wed Mar 10, 1993 8:34 am PST

[From Rick Marken (930310.0800)]

Avery.Andrews (930310.2036) --

>It might be an idea to wait a bit till Bill's 14df project is more
>consolidated - I think there are a lot of practical lessons in the
>application of PCT to complicated situations there (e.g. why Bill's
>design works so much better than `naive cartesian transpose
>>Jacobianism'), and it might make a better impression on any new recruits
to have a reasonably extensive & coherent writeup of this.

I heartily disagree; my experience is that you just can't tell what will make an impression on new "recruits". It is very likely that when the 14df project is finished, the nay sayers will simply yawn and explain how it "could be done" with the appropriate output generation model. Look at what happened to Bill's ARM demo; you may think this is because there were not enough df in that model; but when there are plenty of df in the model the "ho hums" will be based on something else. I've gone through this for years with my attempts to produce convincing demos and experimental tests of PCT; I try to anticipate every complaint and misunderstanding -- and then when I submit the paper or show the demo there is always some completely unanticipated kvetch. The "number of df" complaint about the ARM demo is just a roose -- anyone who understands the principles of the HPCT model would know that it works with an arbitrarily large number of relatively orthogonal df (see my spreadsheet model, for example) in a highly nonlinear (even slightly non-monotonic) enviroment.

So the 14 df model may mean a lot to you (and me) -- and I think it's a wonderful project and demonstration. But if we (PCT) had to wait for every demo that we thought would "clear up" things for the uninitiated (and usually unconvincable) we would wait forever to publicize PCT. Believe me, supporters of other theories -- with FAR LESS (how about virtually zero) experimental and working model support -- have absolutely no hesitation about touting the virtues of their brilliant models. PCT does not have to be so shy -- especially given its remarkably strong experimental and modeling base.

I believe that the ad still could be improved a bit (I'll try to suggest some changes if I get a chance today or this evening). I suggest waiting until the end of the week (Sunday) for all input on the ad and then shoot for posting the final version of the ad sometime next week (3/15).

Best Rick

Date: Wed Mar 10, 1993 9:21 am PST
Subject: Re: Defining information; challenge

[Allan Randall (930309.1320)] Bill Powers (930308.1520) writes:
>
> Allan Randall (930308) --

> >1) Explain in a little more detail why you think information
> >theory predicts that the compensatory system will perform better.
>
> I'm just guessing and could be wrong. The object is for an
> information theorist to show that I am wrong by showing exactly
> how IT would be used to arrive at a correct prediction.

Turning away from information theory and PCT for a moment, lets travel back to the time of Isaac Newton, to an alternate universe, where Isaac Newton became an apple farmer and it was his brother Phil who came up with the laws of motion. Unfortunately, Phil wasn't bright enough to also invent the differential calculus (but then he could bake a really nice apple pie). Imagine that someone has come up with a competing theory to Newton's. They've written a lovely book full of ambiguous, vague language, making it impossible to even formulate what exactly their theory is, let alone put it to experimental test. This book is filled with many "proofs" using Liebniz's differential calculus. Let's say a trend develops and there is a whole spate of similarly motivated physical theories, all in direct competition with Newton's theory. Newton mounts an opposition against this new movement, and rightly so. Says Newton:

"Calculus has absolutely nothing useful to say about physical systems. It can't even make a valid prediction about what will happen in a real physical experiment. It is obviously bogus and invalid."

This is exactly how the anti-information theory arguments sound from my end. Information theory is a tool. There is no reason it must be able to make correct predictions about real world control systems in order to be valid. Repeatedly, I have seen criticisms of information theory in this group that treat it as if it MUST be some kind of competing theory to PCT.

Information theory is NOT a theory of cognition or of living systems. Shannon's theory defines information in a way that can be separated from issues of perception and control in living systems (it does not have to be, but it can be). Entropy is a mathematical measure, based on probabilities. Like calculus, it can be used to better understand many things, such as temperature, work and heat flow. Ashby applies it to control systems. His analysis jives with everything I have learned about PCT. Perhaps I am wrong - I have much to learn. Perhaps, while it is a valuable tool for some things, it adds no explanatory power to PCT. Nonetheless, it is still not a competing theory. Of course, if Martin can show that PCT follows from information theory, then I will be most happy. But this would be a fundamental new discovery. PCT is not currently a subset of information theory, and predictions consistent with PCT should not be expected from it, any more than differential calculus can be expected to predict the orbit of Mercury.

> ...As I just said to Martin above, there is a
> problem when two information channels come together to produce an
> output signal that affects the essential variable. The nature of
> these statistical calculations is that they can say nothing about
> the coherence of the two converging signal channels. If the
> variables are treated only in terms of information content or
> probabilities, they will add in quadrature, so the variance of
> the signal from T to E will be larger than the variance of either
> signal entering T.

I guess I just don't understand this. I assume by "variance" you mean the information content? When two information channels converge, the result of course will be something of higher entropy. This is the Second Law of Thermodynamics. However, as long as the process gives off heat, the resulting information channel may or may not have higher entropy than the sum of the original two. In other words, the total information in the resulting channel is greater than or equal to the sum of the information in the individual channels only if there is no information loss at the point of convergence. It sounds like you are assuming all computations are reversible processes. This is not so.

The signal to E has very low entropy if the system is controlling. This is pretty straightforward, no? Now surely you will agree that if it were impossible to encode the disturbance D into a number of bits that could be handled by the output channel, then the system could not control. Thus, for real-world complex systems with limited bandwidth output channels, you need the perceptual functions to compress the inputs into a single scalar value (for comparison) that STILL RETAINS THE ESSENTIAL CHARACTERISTICS OF THE DISTURBANCE. If no information about the disturbance can be extracted from this data, then there is no way the system can translate the error from this signal into an action on the world that will counter the disturbance. Is this or is this not true? If you agree, then you are agreeing with an information theoretic analysis. Minimal encodings are what information theory is all about. So I ask you: where above did my reasoning go astray?

Note that to talk about minimal encodings (i.e. information content), you need an encoding scheme. This is a basic principle of information theory (the encoding scheme is what Martin calls the subjective probability distribution). In HPCT, the encoding scheme is the structure of the hierarchy. The act of structuring the world into CEVs can be understood in these terms. Maybe it can be understood in other ways too, but that does not invalidate the information theoretic perspective.

Note that if we are dealing with a very simple system, and the disturbance entropy already matches the output capacity (with some trivially simple encoding scheme) then the Law of Requisite Information does indeed become rather redundant. But then so does the hierarchy.

> >2) Find an information theorist that actually does make this
> >claim.
>
> Ashby made it: the disturbance-driven regulator was said to be
> potentially perfect, while the error-driven regulator was said to
> be inherently imperfect.

Well, I guess you know how I'm going to respond, but I'll say it anyway. Yes, Ashby does make this claim, but this is NOT the claim that you were attributing to information theorists. You were suggesting that information theory would predict better control with condition 2 (compensation) in a real world situation than with condition 3 (error-control). Ashby's claim is quite different You have stated it correctly above, and it is quite true. If I can actually have complete knowledge of the disturbance D, it is theoretically possible for me to respond appropriately before it has had an effect on the controlled variable. For instance, the thermostat in the bath sees someone coming with cold water. This thermostat also has incredibly powerful sensors that can detect all the relevant positions/momentums of the particles in the room necessary to precompute the EXACT location, amount and timing of hot water it must add to counteract the cold water.

Of course, it also has unlimited (but not infinite) time to compute the results it needs. This thermostat could actually achieve perfect regulation of the water temperature within the desired range! An error control system could never do this. It is impossible in principle, since by definition the error-control thermostat cannot act until an imperfection is introduced. The fact that the compensatory system will NEVER work in the real world does not change the fact that I can build a simple toy world in which it does (such as the compensatory systems Ashby discusses).

Note that in the case of this perfect compensatory thermostat, information theory still applies. The system has a high bandwidth for output, unbelievably humungous bandwidth for input, and equally ridiculous processing power. But it still has the problem of all control systems in complex environments (unlike Ashby's toy systems): its input bandwidth is much much bigger than its output. In this case, the input is compressed by the algorithms that pre-compute the necessary outputs to correct for the coming cold water.

As for your challenge, I would only respond if we could nail down what the challenge really is. If all you want is a quantitative information-based analysis of the data that is consistent with a PCT analysis, then this is a perfectly reasonable request. However, you seemed to be asking for a *prediction* about which (condition 2 or 3) will be better in a real-world situation. This seems to be beyond the scope of information theory as it stands now. Information theory is not a theory of living systems, and as I mentioned, the Law of Requisite Information applies equally to the (unrealistic) perfect compensatory system as it does to the (realistic) error control system. So it will NOT provide the prediction you are looking for.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Wed Mar 10, 1993 11:14 am PST
Subject: Re: Psycholoquy Ad

[from Gary Cziko 930310.1426 GMT]

Avery.Andrews 930310.2036 says about the proposed Psycoloquy advertisement:

>It might be an idea to wait a bit till Bill's 14df project is more
>consolidated - I think there are a lot of practical lessons in the
>application of PCT to complicated situations there (e.g. why Bill's
>design works so much better than `naive cartesian transpose
>Jacobianism'), and it might make a better impression on any new recruits
>to have a reasonably extensive & coherent writeup of this.

I don't think most people on Psycholoquy (or elsewhere) would even know or care what 14 df refers to (other than perhaps having a sample of size 15). We an always change our ad later as new stuff becomes available. I'd like to send this as it is now if there are no serious objections or suggested revisions.

By the way, Avery, how about a short, sweet paragraph for nonspecialists on what "naive cartesian transpose Jacobianism" is. It can't be all THAT complicated since it only has to do with moving stuff like limbs around in space and if a linguist can understand it there must be hope for other non-physicists as well.--Gary

Date: Wed Mar 10, 1993 11:32 am PST
From: Robert K. Clark / MCI ID: 491-2499
TO: * Dag Forssell / MCI ID: 474-2580
Subject: PCT APPLICATIONS - RKC

SUBJECT:PCT APPLICATIONS - RKC

[From Bob Clark (930310.1300 EST)]
When talking to Greg Williams last night, I mentioned that I may be getting involved with a small organization providing management training to middle and lower level management. I don't yet know exactly what will be involved, but I expect to find applications for PCT theory. Greg agrees that it is important to find and develop applications of PCT theory if it is ever to be generally accepted.
Greg suggested that some of your materials might be helpful in this activity. Please send a sample or two to:

Bob Clark
834 Holyoke Drive
Cincinnati, Ohio 45240   Thanks,   Bob Clark

Date: Wed Mar 10, 1993 11:42 am PST
Subject: intimacy

from Ed Ford (930310:1230)

Got back from Green River, Wyoming (first day it was -15 degrees as I left my hotel - I learned what you do up there is go out, start your car, turn the heater on, then go back into the hotel and read the paper. When you're finished, your car is warmed up and you're ready to go).

It was an extremely successful trip. I was able to train 24 people in four days in reality therapy well enough that they could not only do it, but they had the tools to teach others. The only other program offers 18 months just to learn it, another 18 months to become a certified teacher and the cost is around $6,000.

Far more exciting was how easily they took to PCT. Any time you teach people PCT and tie the concept directly to very practical things they can use, they've bought into PCT. Ultimately, they'll learn to teach it and eventually understand it, but the buy-in comes from the tie-in to practical techniques for working with or teaching others. I guess, Rick and others, the problem is you are trying to convince others of an idea who are not going to perceive any benefit to their lives from learning it so your are blowing into the wind every time you try as you've mentioned on the net over and over again.

(Bill Powers (930303.0830)

>Independence is not a problem, but intimacy is. All of us, I
>think, want to find evidence of a person like ourselves -- at
>least of the same basic nature -- behind the exterior appearances

>that others present to us. We do not want to be the only
>sentient being; we don't want to be alone. So we make great
>efforts to contact the person behind the facade, and we set goals
>that include doing what is necessary to achieve that contact --
>if we can figure out what is necessary.

Bill, I'd be most anxious to learn more of your thoughts on intimacy. From a PCT point of view, how is this idea pulled out of the model, or isn't it? Is there anything in PCT that tells us we "don't want to be the only sentient being."   You only spoke two paragraphs worth but they were interesting and beautiful thoughts. Would be interested in more of your thoughts in this area.

To all..

Off tomorrow to Tucson for two days of R&R then Sunday I go to Toppenish, Washington for another four day program with a school district. Nothing like starting a new career at the tender age of 66. The only thing I dread is trying to catch up on my E-mail when I return. Life goes on........

Best, Ed.

Date: Wed Mar 10, 1993 3:15 pm PST
Subject: Information -- the challenge

[From Rick Marken (930310.1400)]   Allan Randall (930309.1320) --

First, thanks for the paper; it looks like it will, indeed, be an excellent addition to the PCT library. As to your post -- I'm sure Bill P. will have all kinds of helpful things to say. I'll just comment on a couple things that caught my attention.

First, you say that Information Theory (IT) is to PCT as calculus is to Newton's laws; IT is a tool like calculus. But calculus helps us make detailed predictions from the basic model (in my business we do this all the time -- it helps to know the future position and velocity of a satellite). But at the end of your post you say:

>you seemed to be asking for a *prediction* about which
>(condition 2 or 3) will be better in a real-world situation.
>This seems to be beyond the scope of information theory as it stands now.

Well, IT isn't much of a tool if it can't help us predict things; looks like PCT WITH IT is no better off than Isaac's brother Phil WITHOUT calculus.

At another point in your post you say:

> If no information about the
>disturbance can be extracted from this data, then there is
>no way the system can translate the error from this signal into
>an action on the world that will counter the disturbance. Is this
>or is this not true?

This is NOT TRUE. Surprise!

>If you agree, then you are agreeing with an

>information theoretic analysis.

So I guess I disagree with an IT analysis -- and if I'm right (which I am) then the IT analysis is wrong, right? That's really what Bill's challenge was about -- demonstrating this fact experimentally.

It is this fact about control systems that nails everyone to the wall -- and proving it to myself is what turned me into a PCT "fanatic" because it kabashes the entire ediface on which the behavioral sciences are built; the input - output model of behavior. In a high gain, negative feedback control loop, the output DOES NOT depend on the sensory input; rather, SENSORY INPUT IS CONTROLLED BY OUTPUT. What you do in a tracking task is NOT caused by what you see; there is a LOOP so that what you see is both a cause AND A RESULT of your output. The nearly perfect relationship between output and disturbance does NOT exist because the system has access to information (from the sensory input) about the disturbance. When control is good, there IS NO information about the disturbance in the stimulus -- NONE, ZILCH, NADA. The output mirrors the disturbance because this is what the output MUST DO in order to keep the input IN THE REFERENCE STATE; this is the magic of closed loop control. We're talking about purposeful behavior, here, and it is not the result of "feedback guidance", "stimulus control", or "programmed output"; we're not talking about behavior that is the control of perception. If you haven't already done so, I highly recommend the first paper in Chapter 3 of "Mind Readings". This is an experimental "proof" that sensory input IS NOT, never was and never will be the cause of control movements in a closed loop tracking task.

Best Rick

Date: Wed Mar 10, 1993 3:39 pm PST
Subject: oops, let those people go

[From Rick Marken (930310.1430)]

Oops. At the end of my last post I said:

>we're not talking about behavior that is the control of perception.

Of course, I meant to say that we ARE talking about behavior that is the control of perception.

Gary Cziko --

In the spirit of passover (a holiday that threatens to disturb my peace sometime soon because my mother already called and TOLD US when we would be coming over for the seder) why not let all the CSG-L captives go -- I'm referring, in particular, to the two people who recently posted to CSG-L pleading to get off. After re-reading some of my latest posts I can understand how they feel. I strongly advise against chasing after them if they head towards the Red Sea.

Best Rick

Date: Wed Mar 10, 1993 4:33 pm PST
Subject: naive jacobian etc.

[Avery Andrews 931103.0928] [Gary Cziko 931003.????]

```
                           ^
                           |  (effect of elbow flexion))
   X ----------0-----------
               T

  X = shoulder, 0 = elbow
```

For each joint, one can calculate the partial derivative of tip movement w.r.t. flexion at that joint. The resulting collection of vectors represent the immediate effects of flexion at each joint, and can be represented as a matrix, which is called the Jacobian, whose columns are the various partial derivative (various aspects of the representation are arbitrary, such as the order of the columns representing the partial derivations). So if we have three joints moving the tip in three dimensions, we get a Jacobian that looks like this:

```
     D1,1   D1,2   D1,3

     D2,1   D2,2   D2,3
```

Each column represents a vector sticking out from the tip in some direction, representing the direction and velocity with which the tip will start moving if the corresponding joint is flexed a little bit.

One thing you can do with the Jacobian is calculate an approximation to how the tip will respond to a combination of joint movements: If we (pre) multiply the Jacobian by a column vector of joint movements:

```
  F1       D1,1   D1,2   D1,3        T1

  F2   *   D2,1   D2,2   D2,3    =   T2

  F3
```

We get an approximation to how the tip will move.

Another thing we can do with the Jacobian is find out a bit about which joint movements will be helpful in reducing the distance between the perceived and reference locations of the tip. Suppose we have an `error vector', pointing from the tip to the reference location. If the Jacobian component for a joint is more or less in the same direction as the error-vector, then flexion of the joint will reduce the size of the error vector, while if it is in the opposite direction, flexion will have this effect, while if the two are at right angles, neither flexion or extension will have much effect. An operation called the `dot product' measures the extent to which two vectors are pointing in the same direction: it consists of just multiplying corresponding components & adding the results (it gives the cosine of the angle between them, which seems magical to me, but is nonetheless true):

```
      A1          B1
            .           =  A1*B1 + A2*B2
```

```
      A2          B2
```

So if we take the dot product of each Jacobian component with the error-vector, we get an indication of what *immediate* contribution movement at each joint will make to reducing the error. But this is equivalent to pre-multiplying the error-vector by the transpose of the Jacobian (baby matrix algebra), whence `transpose Jacobianism'. It's cartesian from the use of Cartesian coordinates, and naive if you just use the components of the result of this multiplication (scaled or limited, perhaps) to tell you how to move each joint.

And, it doesn't work very well, at least in the naive form, tho there may well be ways of fixing it up. The problem is that the Jacobian just tells you what effect a tiny change in joint angle will have, whereas in many cases you need to know more. In the elbow case, for example, if the target is located at the elbow, the system will lock, because both components of the Jacoboian point straight up:

```
                          ^
                          |^
                          ||  (effect of movement at shoulder, elbow)
   X ----------0-----------
              T

  X = shoulder, 0 = elbow
```

And the shoulder component will be larger than the elbow one, due to mechanical advantage. So if the target is a bit above the elbow, the shoulder angle will increase at first, even though to get a sensible result at the end it should actually decrease.

The Little-Man arm controller avoids this kind of problem by using appropriately chosen coordinate systems, & a bit of hardwired `foreknowledge'. Suppose we use shoulder-centered polar coordinates rather than Cartesian ones. Furthermore, being smart, we know that the only way to correct a distance error is by changing the elbow angle - flex to decrease the distance, extend to increase. We can think of ourselves as *pretending* that the Jacobian for elbow flexion points toward the shoulder, even though this is literally false almost everywhere (though it gets truer as the elbow flexes). The Jacobian for shoulder angle change on the other hand really will be a pure change in perceived angle, without any effect on the distance coordinates. If we represent the error-vector as a (distance, angle) pair, then the second component is guaranteed to be in line with one of our Jacobian components, so that good control is trivially available. For the distance component, flexing the elbow introduces an angle error as well as changing the distance, so our Jacobian is not fully in line with the error, but since the shoulder angle control is so easy, a shoulder-angle control system can make up for this, & we can get away with pretending that the Jacobian for shoulder-angle point shoulder-ward.

So we might regard the Little-Man style controller (and Bill Powers' 14df, if I understand it more or less correctly) as a kind of clever transpose Jacobian scheme, the cleverness residing in a good choice of coordinate systems (really, perceptual dimensions), plus foreknowledge of the fact that when the arm is fully extended, the only way to correct a distance error is to flex, in spite of the fact that the Jacobian component for elbow flexion in that position is (0,0) (a bit more vector calculus there).

I think there is a moral lesson here, which is this. Many people seem to think that PCT is a sort of blind faith that feedback systems can solve all problems, with no thought required, & I wouldn't be surprised to see the Naive Cartesian Jacobian Transpose Controller produced as an argument that this faith is false, which it certainly is. But the contrast in workability between NCJTC and Little Man shows vividly that the performance of these systems depends crucially on how there perceptual systems are set up - one of the many aspects of PCT that the critics tend not to notice.

It remains for us to find out whether we can get a useful story about serious manipulator control out of all this - the fact that I was able to invent a horrible 2df controller for 2 dimensions simply shows that you can be an OK linguist without being much of a roboticist.

Avery.Andrews@anu.edu.au

Date: Wed Mar 10, 1993 5:35 pm PST
From: marken MBX: marken@aero.org
Subject: Thanks

Hi Dag (and Christine);

I got the tape and the supporting materials. I'll try to watch the tape over the next couple of days when I get home. I've seen the beginning and it looks quite good (a couple little technical flaws -- back of Christine's head in one shot; out of focus in another -- but overall it seemed quite professional). And the supporting material was EXCELLENT. Very nice Introduction brochure.

I'll try to give you some feedback withing the next couple of weeks.

Best Rick

Date: Wed Mar 10, 1993 8:27 pm PST
Subject: robot comment

[Avery Andrews 930311:1407] (John Gardner (930125:0030))

A rather belated response to part of John Gardner's posting, in the light of my recent experiences:

>Here is the heart of our argument. Implicit in the above
>statement (and, I assume, in the approach) is that it is
>somehow easy to 'set up the system so that it can have
>relatively independent influences on the degrees of
>freedom'. This is where we differ and this is also the
>heart of the more degrees of freedom argument.

I would say that the heart of the approach is not the idea that it is easy to set up the system in this way, but it is *possible*, at least for human arms, and that this possibility is part of the reason why human arms can work as well as they do. For example, the eye-centered `polaroid' coordinates (distance, & xy-location on retina) used by Little Man have much nicer properties w.r.t. what the joints do than xyz coordinates have, as do real polar coordinates, wherein an increase in shoulder pitch always produces an increase in hand elevation angle (where the natural maximum is with the hand just behind the shoulder-blade, natural minumum just under it).

Getting the right analysis isn't necessarily easy, but what the circuits are actually doing might be easier than expected.

The motor control literature & robotics books I've been browsing in don't seem to say a whole lot about setting up coordinate systems so as to make control problems easier, so maybe there's something here. On the other hand, maybe this approach is just well-known to be hopeless for industrial robots.

Avery.Andrews@anu.edu.au

Date: Wed Mar 10, 1993 9:12 pm PST
Subject: Mostly the challenge

[From Bill Powers (930310.1845 MST)] Gary Cziko (930310.0235 GMT) --

The advert sounds find to me -- go for it. I don't think we need to wait for the 14 df to be finished (as per Avery Andrews, 930310.2036). It still needs vetting by a real control engineer, and there's a way to go to make it doing anything useful (- looking). By the way (to Avery), I think the net would be interested in your evaluation of the Jacobian approach and I'd be interested in responses from engineers.

Allan Randall (930309.1320) --

I get your point about Phil Newton -- it's possible to apply a perfectly good technique to a fuzzy problem, and come up with fuzzy results. The good technique shouldn't be rejected because of its misuse.

To make the parallel work better, however, you should think of something other than the calculus, which in fact CAN be applied to specific problems in mechanics, so Ike Newton would never have rejected it.

>This is exactly how the anti-information theory arguments sound
>from my end. Information theory is a tool. There is no reason
>it must be able to make correct predictions about real world
>control systems in order to be valid. Repeatedly, I have seen
>criticisms of information theory in this group that treat it as
>if it MUST be some kind of competing theory to PCT.

I don't think of it as a competing theory. I'm just trying to point out, through my challenge, that an analysis of a control system in information-theoretic terms might be able to come up with true statements, but they're not useful for any of the purposes that interest me, like designing a PCT model that will actually do something like what real systems do.

I'm sure that AFTER I have designed a control system to behave in a certain way, an information theorist could estimate the information flows in the system, the entropies, and all that lot. If I were trying to model behavior that takes place under difficult conditions, this analysis might offer more of interest by way of

predicting limits of performance. But as I see it now, information theory can't tell me anything that would help in designing a model of ordinary human behavior.

RE: design details: when two signals converge.

Just to keep the diagram in mind:

```
         D ---> T ---> E
                |      |
                R --<--
```

Signals from D and from R converge on T. The result is -- that is, should be -- that the signal going from T to E contains less "variety" than it would if R were not acting. At least this is how regulation has been characterized by Ashby, and you have said that his characterization is derived from information theory.

My point was that when you characterize signals in terms of information flow rather then in terms of amplitude and phase, it no longer is possible to predict the result of the above convergence. If ordinary statistical measures like variance are used, we would NOT in general expect the variance of E to be less than that of D or of R.

>I assume by "variance" you mean the information content? When
>two information channels converge, the result of course will be
>something of higher entropy. This is the Second Law of
>Thermodynamics. However, as long as the process gives off heat,
>the resulting information channel may or may not have higher
>entropy than the sum of the original two.

So you appear to agree that we can't predict whether regulation will actually occur in the above arrangement on the basis of information theory alone. Let's look further into that. What information is missing in the specifications, that you would need in order to make a prediction?

Suppose I supply more specific information. The disturbing waveform is a sine wave of frequency omega and amplitude A: it is D = A*sin(omega*t) where t is elapsed time. T is an adder that adds the disturbance to the output of the regulator. The regulator senses the state of E, which without the regulator would be another sine wave of frequency omega. The regulator subtracts the state of E from a constant value, E*, and outputs the time-integral of the difference to T: the output is just k*integral(E* - E). So the adder receives D and k*integral(E* - E) and as a result, the state of E is

E = A*sin(omega*t) + k*integral[(E* - E)dt].

This is now a completely-defined system, as far as PCT is concerned: by plugging in values for k,A,omega, and E* you can calculate the behavior of E, and from that you can decide whether amplitude regulation is taking place, and how much. Omega and A are specifications for the disturbance, and E* is the state in which the subject is asked to maintain E. These values are known before the experiment. The only parameter to be determined from the data is k. In my actual challenge experiment, the disturbance is a smoothed random wave, but the same principles will hold. Note, by the way, that E appears on both sides of the above equation, a typical closed-loop form.

Where I predict a problem for an information theorist in trying to meet my challenge is in explaining, on IT grounds alone, why the amplitude of E becomes less than the amplitude of D. According to my vast knowledge of information theory, shown here life-sized (.), information theory does not deal with continuous physical variables on a moment-by moment basis: there is no way to say in terms of information that signal A is "equal to and opposed to" signal B. The idea that the information in one channel is "opposite to" the information in another would seem to me to be nonsense, or to get into semantic information. Yet to design the above system, we have to rely on quantitative subtraction of one signal from another. I don't see that Shannon information theory in any way predicts this subtraction, in either Ike's hands or Phil's.

The phenomenon of control depends on the oppositeness of the signals, not on their respective information contents. The two signals could have the same information content by any common measure, yet their amplitudes would not, just because of that, have to cancel each other quantitatively. We could imagine a sine-wave (or any other waveform) coming out of R that had the same amplitude and frequency as the sine-wave disturbance, yet the net effect on E could be anything between completely cancelling variations in E to doubling them. What makes the difference is not the information carried in the two converging channels, but the quantitative phase relation between the waveforms. This, by the way, is just as true for condition 3 (compensation) as for condition 2 shown above.

I sort of object to using a physical term like entropy in dealing with information in living systems. If you think that informational entropy is connected to physical entropy, you can't separate the formal entropy in a neural signal from the real physical entropy involved.

I would guess that the entropy increase involved in simply transmitting a neural spike from one neuron to the next would be hundreds of times greater than the hypothetical decrease involved in the actual synaptic event - the transfer of a bit of information. The propagation of a neural signal expends energy all the way along the fiber, energy that must be replaced from metabolic sources. Neural signals don't CARRY energy -- they EXPEND energy. Shannon and others used to claim that information entering the nervous system amounted to negative entropy, but if they intended to mean some mysterious connection with dQ/T, they had it backward: getting neural signals to their destination occurs at the cost of a great INCREASE in physical entropy, which is compensated only by increasing the chemical entropy of the surrounding fluids, which is compensated only by eating and breathing. The actual direction of information flow in the nervous system is opposite to the direction of energy flow. So there can't be any connection between entropy as physicists use the term and entropy as it appears in information theory.

>The signal to E has very low entropy if the system is controlling.

Now why would you say that? When the system is not controlling, E is varying exactly as D is, and if D is a signal within a certain bandwidth, then E is also a signal with that bandwidth. All the information in D is being transmitted to E.

When the system is controlling, the information formerly reaching E is now very nearly cancelled; E now contains very little information. Would not this loss of negative entropy amount to a great INCREASE in entropy?

>Now surely you will agree that if it were impossible to encode
>the disturbance D into a number of bits that could be handled

>by the output channel, then the system could not control.

If that number of bits could not be handled by the output channel (by which I presume you mean the path from T to E), then the only the bits that survived passage through the channel would be able to disturb E. All that the regulator R would have to do would be to provide bits to T that subtract, bit for bit, from those that survive the passage through T, and regulation would be perfect, or at least no worse than about 1/G of the original variations in E, where G is the loop gain.

>Thus, for real-world complex systems with limited bandwidth
>output channels, you need the perceptual functions to compress
>the inputs into a single scalar value (for comparison) that
>STILL RETAINS THE ESSENTIAL CHARACTERISTICS OF THE DISTURBANCE.

No! It is not the disturbance that has to be represented, but the state of the controlled variable itself. The output only has to act on that variable with sufficient speed to keep the perceptual representation of that variable matching the reference signal. The disturbance itself (D above) could contain megahertz variations; most of those would disappear because E can't respond significantly to them. E, however, will still contain frequency components that are not represented in the perceptual signal. Those frequency components will be uncontrolled (from the viewpoint of an external observer). The perceptual signal itself, however, will be controlled.

>If no information about the disturbance can be extracted from
>this data, then there is no way the system can translate the
>error from this signal into an action on the world that will
>counter the disturbance.

You have to make up your mind whether you're talking about the variable responsible for the disturbances (D), or the effect of that variable (variations in E). The control system only has to sense the state of E; it needs no direct information about D or the channel between D and E.

>So I ask you: where above did my reasoning go astray?

In assuming that the control system needs information about the state of the disturbing variable D, and in assuming that control is Èn2terms of the external variable E rather than the perceptual representation of it. The first is the most important error. There can be any number of independent disturbing variables acting through environmental links on E, all at the same time and uncoordinated with each other. The control system needs no information about them, singly or collectively. All it needs to sense is the state of E. The reasons for the variations in E are irrelevant to the process of control. The control system senses the state of E, compares that with a desired state E*, and acts DIRECTLY ON E to oppose any perceived deviations of E from E*. If the control system is so organized that for each possible kind of deviation of E from E* the system can act on E in the opposite direction, nothing more is required.

>If I can actually have complete knowledge of the disturbance D,
>it is theoretically possible for me to respond appropriately
>before it has had an effect on the controlled variable.

What you actually should have said, to be precise, was "If I can actually have complete knowledge of the disturbance D, and if I can use this knowledge to produce actions having a precisely calculable effect on E, and if the actions actually produced affect E exactly as calculated so as to cancel the effects due to variations in D, then the actions will affect E so as to cancel the effect of variations in D."

By the time you have specified ALL the assumed conditions on which the conclusion depends, the conclusion has become redundant: it is stated in the conditions. So the entire statement becomes merely an assertion that these conditions hold. All I have to do is deny the assertion, which I do. No theory is involved; only a tautology.

In your example of the thermostat, you stated some of the conditions and abilities necessary for the bath to be regulated by compensation. If you had stated all of them, you would merely have stated that the hypothetical thermostat is able to do what you say. Having said that, what do you add by saying "therefore such a thermostat can do what I say it can do?"

If a cow could jump 187,000 miles, then theoretically a cow could jump over the moon.

>As for your challenge, I would only respond if we could nail
>down what the challenge really is. If all you want is a
>quantitative information-based analysis of the data that is
>consistent with a PCT analysis, then this is a perfectly
>reasonable request.

Fine. Which condition will result in the best regulation, defining regulation in IT terms however you please: condition 2 or condition 3?

>However, you seemed to be asking for a *prediction* about which
>(condition 2 or 3) will be better in a real-world situation.
>This seems to be beyond the scope of information theory as it
>stands now...So it will NOT provide the prediction you are looking for.

Well, isn't this sort of the point? If information theory is fundamental to PCT, and if PCT can make such a prediction, why can't IT make it, too? If IT can't do that, I would think that IT has only a peripheral relationship to PCT -- otherwise, how could PCT make correct predictions without benefit of information theory, while information theory can't make any predictions without first knowing what PCT has to say?

So I have exactly one taker for my challenge, Martin Taylor. He is looking braver to me by the moment.

That is more than enough for one post.

Best to all, Bill P.

Date: Thu Mar 11, 1993 9:11 am PST
Subject: even more naive jacobian

[From Rick Marken (930311.0800) Avery Andrews (931103.0928) --

Thank you for that very helpful explanation of control using the jacobian transform matrix. I would just like to add some even more naive comments about it just in case anyone is actually interested -- because I think it a very important topic for PCTers.

First, I would like to emphsize that the Jacobian approach you describe is CLOSED LOOP. So it can be thought of as one approach to the control of perception. You mentioned this in your post when you said:

>Another thing we can do with the Jacobian is find out a bit about which
>joint movements will be helpful in reducing the distance between the
>perceived and reference locations of the tip. Suppose we have an
>`error vector', pointing from the tip to the reference location.
>If the Jacobian component for a joint is more or less in the same
>direction as the error-vector, then flexion of the joint will reduce
>the size of the error vector,

but it might have slipped by. Let me try to paraphrase what you are saying; feel free to correct me if I have it wrong. The system you describe is controlling the perceived x,y,z coordinates of the finger tip. There is a reference position -- x',y',z' -- and the current perceived position ---x,y,z. The difference between reference and perception is the "error vector" that you mention above -- x'-x, y'-y,z'-z. Now the problem is to turn the errors in this vector into outputs (changes in limb positions) that will reduce the error. That's what the Jacobian is for; it is a means of transforming error into output based on the facts of arm kinematics (as defined by the jacobian matrix). So there is an "output computation" based on the jacobian -- but it produces just a slight change in the outputs that affect the arm, resulting in a slight change in the perception of finger tip position and a slight change in the error matrix which is used to compute the next change in output via the jacobian.

This is a legitimate "control of perception" approach to building a robot. The jacobian is actually acting as a complex "output function" that transforms error DIRECTLY into environmental effects that alter the perception that is being controlled. This is NOT a pure output generation approach -- and it WILL PRODUCE CONTROL -- ie. resist disturbances to the controlled variable; for example, if you "push" on the finger tip or any of the arm components that influence the perceived position of the fingertip this system will compensate.

The problems that Avery is running into with this approach are a bit esoteric for me -- but I think it has to do with the fact that certain fingertip targets (reference settings) result in outputs of the jacobian that specifiy physically incorrect or mathematically impossible states of the arm components.

Avery points to one advantage of the PCT approach:

>So we might regard the Little-Man style controller (and Bill Powers'
>14df, if I understand it more or less correctly) as a kind of
>clever transpose Jacobian scheme, the cleverness residing in a good
>choice of coordinate systems (really, perceptual dimensions),

I think another way of saying this is that the success of the Little Man stems, in part, from incorporating appropriate perceptions to be controlled. But an important implication of this is that the PCT model is HIERARCHICAL; the "error vector" in PCT specifies references for LOWER LEVEL PERCEPTIONS -- not for environmetal outputs, as in the jacobian approach. So it is not just selection of clever perceptual dimensions that characterizes the PCT approach -- it is the fact that errors ARE specifications for lower level perceptions (except at the lowest level, of course, where errors "command" physical effects like muscle fiber contractions or glandular secretions). In the jacobian approach there is only one perception being controlled, finger tip position in Cartesean coordinates. You could change the jacobian approach by having it control the fingertip in polar coordinates, for example. But that would not change the structure of the model in any important way (though it MIGHT improve performance in some situations); the jacobean would still transform error (now indicating deviation from polar references) into outputs.

So I think another moral of the jacobean experiments is that there is virtue in having higher level control systems control by setting specifications for lower level perceptions. Bill Powers has mentioned several principles from which this moral can also be derived (convergent rather than divergent control functions, experiential dependence of "higher order" perceptions on existence of "lower order" perceptions [can't see configurations without sensations], etc) -- I mention some of these in my "Hierarchical behavior of perception" paper. But the jacobean experiments seem like a very good practical demonstration of this moral.

Great work, Avery -- especially for a linguist!

Best   Rick

Date: Thu Mar 11, 1993 9:43 am PST
Subject: new summary re language

[From: Bruce Nevin (Tue 93039 12:17:00, Thu 930311 12:08:14)]

In our discussions of language we have had the proposal (from Bill) that a word-perception be one of the inputs satisfying the input function for a category-level comparator.

To the extent that (most) words are to be identified with category perceptions, analysis of structure in discourses should tell us something about the organization of the category level. The advantage of this is that words and word dependencies are much easier to manage and analyze than nonverbal perceptions, and it is much easier to reach agreement about results. This is partly because words are by "design" public, and nonverbal perceptions are ineluctably private. I believe it is also because dependencies among words are much more well defined, but a comparison with the aim of demonstrating this is virtually impossible so long as we have only the words with which the nonverbal perceptions are associated as our means for reaching agreement about those nonverbal perceptions. To me, that very limitation is a pretty convincing demonstration in itself.

In what follows, I will rely much on a notion of dependency between words. The saying of one word ("jump") is dependent upon the saying of one or more other words in the same utterance ("child" or "fish", for example). I will clarify this notion presently, after some preliminary definitions. (The first indented paragraph within square braces, below.)

Analysis of an utterance yields a structure of dependencies among the particular words in it.

Words are classified according to these dependencies across many (by predictive claim, all) utterances of a language.

Def: Operator words (operators) are words the saying of which is dependent upon the saying of other words. Symbol: O.

Def: The words upon which an operator depends in a given utterance are (in) its argument, or are its argument words (arguments).

Def: Words with no such dependency are primitive arguments -- mostly what we think of as concrete nouns. Symbol: N.

Operators are further classified according to the above dependency classes of words in their argument. All that is needed is this simple dependency on dependency. Thus, On depends on one N (jump), Ono depends on one N and one O in that order (believe), Oon on the same argument word classes in the reverse order (surprise), and so on.

In addition to the operator and argument words, there are other morphemes that correlate with no category perception, but rather help to identify the dependency relations, morphemes like `-ing', `to' of the infinitive, `that' in e.g. `I think that you understand me'.

Words may be present in zero form if their presence is so strongly expectable in a given position relative to other words in the utterance that they need not actually be said in that position: John plays piano and Mary [plays] violin. Or expectable words may be given reduced form (e.g. pronouns).

[What is meant by dependency can be seen intuitively in perceptual terms. Perceptions that I imagine in association with the On operator "jump" necessarily involve something or someone (indicated by some word of class N) jumping. Perceptions that I imagine in association with "child" do not entail anything in particular about the child that I imagine--no word of some O subclass is required. However, the notion of dependency used here has its basis in the occurrence or occurrability of words in utterances. It may be that the latter reflects merely the former, but the privileges of occurrence of words are enormously easier to manage and analyze (in a way that enables agreement among investigators) than those of perceptions. Furthermore, it is not clear the extent to which words and discourses, and agreements reached and communicated by means of them, may determine which perceptual inputs humans pass along for higher-level perceptual control, which they supply by imagination, and which they ignore, and it seems wise not to beg the question.]

Given these preliminaries, I will summarize proposals I made in 1969-1970 (in my MA thesis at Penn).

At the first stage of analysis, the words in utterances are in linear sequence (some of them perhaps having their phonemic content reduced, even to zero, as noted).

```
a child
        \
           jumped
                  \
                   > over
                  /
            a book
```

A child jumped over a book. This is reduced from A child jumped; said jumping is over a book. The past tense is reduced from A child jumps; said jumping is before my saying this. I'll ignore these reductions for now. It's easier to draw the dependency graph at the next stage of analysis, undoing the linearization step of producing sentences.

At a further stage of analysis, the words/category perceptions are not in a particular linear order. (Alternative linearizations as sentences are possible.)

[Where (the saying of) an operator word is dependent upon (the saying of) two or more argument words, it is convenient for us to speak of the argument words as the first argument, second argument, etc., based upon the linearization felt to be most basic or most common, but the means for identifying and differentiating the words required in the argument of a given operator must be in other than sequential terms. I don't have a proposal at hand. Case grammar is an obvious candidate, but gets notoriously messy and fraught with disagreements between investigators. The simplest solution is to use the most basic or prevalent linear order as means for distinguishing multiple arguments prior to the linearization step. This runs into problems as dependencies in discourse pile up, and even in a single sentence such as the present example.]

```
child          _____________________
      \       /                     > over
         jump ---\              book
             \     \
              \     > before
               \   /
                say
              /
             I
```

Where at the earlier stage we often had repetitions of a word, where the word-occurrences were identified as "the same" (and often one occurrence was reduced to zero on that account), at this stage we may be able to have a single word/category perception signal satisfy the argument requirement under both operators (for both occurrences).

Q: do we postulate any mechanism in PCT with the plasticity required for this?

We no longer have a dependency tree for each sentence (a rooted, directed, oriented graph with a single root, one or a few branches at each node, and primitive argument words at the leaves). For sentences with modifiers or conjunctions we have something like a semilattice structure.

Continuing further with the analysis, we reach a mesh-like structure for a discourse. Central, topical, or thematic words have more dependencies than other words/categories occurring in the discourse.

Discourses in the same subject-matter domain use the same vocabulary, and they use it in the same way. Put in other terms, the "same" word may satisfy the input of different category perceptions, depending upon the subject-matter domain of the discourse and other perceptual context in which it occurs. "Wear" means one thing to a tailor and another to a sailor (in context of their respective specializations). One aspect of this is seen in the classifier vocabulary for a given domain. To take an example from work done over the past 20 years, in a physiology domain, "heart" is classified as a "body part". In a pharmacology domain, "heart" is not a primitive term but occurs only in phrases, and it is these phrases which are then treated as primitives for the domain and classified, as for example, "the beating of the heart" is classified as a "symptom". (In this way, a study of these matters can disclose, in the structures of linguistic information found in subfield discourses, a basis for logical precedence and other relations among different scientific subfields.)

One subject-matter domain may be distinguished from another by shared vocabulary across utterances in the subject matter, where the words satisfy classifier vocabulary for the domain, in a relatively small number of dependency relations specifyable among the classifier words. The shared vocabulary, recurrent dependencies, and the classifier vocabulary appear in the mesh of word dependencies in memory on the basis of discourses previously encountered. (Again: for "word dependencies" read perceptions of dependencies between words and between associated category perceptions.)

An example of classifier terminology for a domain:

> "Mail should not be thought of as an application, it's an
> enabling technology." --Eugene Lee, director of product
> planning for Email software vendor Beyond, quoted in Byte, 3/93:90.)

Some software packages are called applications. Some things (including but not limited to software packages) are called enabling technologies. Perceptions associated with the software packages called applications, especially those generalizable from the set of them with which a person is familiar, come to be treated as expectations about some new thing called an application. Similarly with the different perceptions associated with things called enabling technologies and generalizable from the known set of them. Mr. Lee tells us we ought to have the latter expectations of electronic mail, rather than the former. (Um--give us some examples of enabling technologies, please?)

Word dependencies are clearly also dependencies between category perceptions. I talk about them in terms of word relations because the words are for us a principal way of getting at the categories; and ex hypotheosi the two are equivalent for our purposes (see ref to Bill's proposal, above).

Now, if I have encountered a given operator-argument dependency o-a in a conversation or text, I have a basis in memory for finding that dependency acceptable in subsequent discourse. (I might have a basis in memory of nonverbal perceptual input for expecting the dependency between category perceptions associated with o and a respectively. Or I might not--I might only have imagined perceptual signals generalized from various experiences categorized as o and as a in the past, perhaps only vaguely, just enough to "make sense" of the utterance. I might have only a fuzzy notion of what an "enabling technology" is, or of what email is.) If o satisfies classifier word O and a satisfies classifier word A in the given subject-matter domain (in the mesh of word dependencies [perceptions of dependencies between words and between associated category perceptions] in memory on the basis of prior discourses), and if I have encountered other satisfiers of O and A (or even O and A themselves) in the given dependency relation, then I have a strong expectation that some (to me) novel o-a dependency is also acceptable.

A lot of learning a subject matter is learning of its classifier vocabulary. As mentioned, relatively few dependencies among classifier words recur with great regularity in discourses of a given domain. (Harris et al. _The Form of Information in Science_, Naomi Sager, _Natural Language Processing_). For nautical discourse, we might find the following sorts of dependencies (among others):

VESSEL	MANEUVERS	IN DIRECTION
ketch	wears away	around starboard tack
yawl	bends	to port
schooner	comes about	into the wind

On this basis, we would find acceptable something like "the schooner wore away to port", but "the starboard wore away around the yawl" is absolutely unacceptable -- nonsensical -- nautical usage. (Though a person ignorant of domain norms might well find this perfectly sayable and perfectly understandable -- say, the starbord side of the yawl rubbed against the pier during winter storage, and was eroded away, around the curve of the vessel on that side.) Similarly, "the [rubbing of the] rope wore away the paint" is acceptable not as specifically nautical usage but on the basis of dependencies in a more basic subject matter, perhaps mechanics, even though it may occur embedded in a nautical discussion.

As we learn the word dependencies for a domain, we learn to match our nonverbal perceptions to the mesh of dependencies remembered from prior discourse. We must do so for the words to "make sense". It is also true, though perhaps not so obvious, that we must do so for the nonverbal perceptions to "make sense", in particular for them to make the same sense to us as they do to those with whom we seek to cooperate and from whom we seek to learn. We must replace imagined perception with memory of actual perception, and so we become experienced and expert in the domain. But while learning (and perhaps while supposedly expert) we may prefer imagined perceptions that match domain norms shared with others over actual but discrepant perceptions, if control for the perception of agreement and cooperation has higher gain than control for discrepancies between imagined perception and actual perception, however that is managed.

And now once more beyond the pale: in learning to be people with those around us we learn to discount and ignore some perceptions, or to cover or reinterpret them with the aid of imagined perceptions. This is what has happened with socalled "psychic" perceptions, which have for many centuries been the occasion of painful and shameful death, and which apparently not all people access equally (variability analogous to color blindness).

Bruce Nevin bn@bbn.com

Date: Thu Mar 11, 1993 11:49 am PST
Subject: the central insight of PCT

[Hans Blom 931103] Rick Marken (930309.1500)

We seem to be getting to the point. How central is "the central insight of PCT" really?

>>If the sensor signal may be absent for shorter or longer periods, a model
>>(in the sense of a built-in or acquired approximation of the object to be
>>controlled) is required that temporarily provides an alternative means of
>>pseudo-feedback. Such models are usually not highly accurate and there-
>>fore drift will occur away from the optimal operating point.

>I don't get you point here, Hans. When a system is controlling (holding
>some aspect of the environment in a fixed or variable reference state
>while resisting the effects of disturbances) then it is the sensory
>(perceptual) signal (whatever its cause) is really controlled. ...
>The central insight of PCT is that WHEN THERE IS CONTROL THEN IT IS THE
>FUNCTIONAL EQUIVALENT OF THE SENSORY SIGNAL (PERCEPTION) THAT IS
>CONTROLLED.

If you consider the 'signals' that can be recovered from memory to be sensory signals as well, you are _logically_ right. SOMETHING is controlled by a control system, to be sure. Thus far, I understood 'sensory signals' to mean the information that reaches me through my sensors AT THIS MOMENT OF TIME. I'll keep to that notion. For me, memory is a different category: a concatenation and compaction of earlier sensory information combined with evolutionarily given information. The latter specify which of the sensory signals are IMPORTANT (related to e.g. bodily needs). The compaction process is not infallible given the limited resources, and therefore memory may contain misrepresentations, illusions, and what have you. If you use those as your "sensory signals", something occurs that I called "drift" above.

We come to the problem of what is REALLY controlled. It is not really the sensory functions. For example, in many engineering applications, a compro- mise is sought between optimal stabilization of some variable and minimum energy expenditure. Control requires fuel to be burnt, and fuel is expens- ive. Organisms, I suppose, have the same problem. Practice shows that you can conserve large amounts of fuel if the control of the variable to be controlled is just a tiny little bit sloppier (say 40% versus 2%). You could say that in that case there is an additional "energy consumption sensor", but engineering practice takes a different point of view. What is REALLY controlled is not sensory input but some grand overall one-dimen- sional control purpose that I call Q in my reply to Bill Powers (see below).

(Bill Powers (930309.1800))

Too much to answer, too little time. Just a few points, therefore. In both of your following remarks

>I would have to agree with your implication that organisms
>control as well as they can. That's a matter of definition. But
>in looking at the state of our world, I am not greatly impressed
>with the way people control for social harmony, economic
>viability, or maintenance of an environment fit to live in.

> Most of the people in
>the world live in poverty, hunger, and illness. I don't see how
>you can claim that they are optimal control systems.

you confuse "optimal" (an engineering word with an exact meaning) with "good" (a moral categorization). The "optimal" of engineering means only that some system reaches its grand overall goal as closely as possible, by definition. Engineering is not concerned with the question whether that goal is "good". Engineers are, though. In my own personal, idiosyncratic world model I tend to equate "optimal" with "good" (subjectively, for that person, given his/her opportunities, limitations and life plan). Maybe that provoked your remarks.

> I don't think that people are
>particularly adept at constructing systems of goals that hang
>together, are consistent with each other.

In optimal control theory, there is only one "supergoal" that can be controlled. There can be subgoals, however. It would be possible to declare the two (seemingly conflicting) goals "drive in the middle of the road" and also "drive one yard to the right of the middle". But then you would have to combine them into one goal. This can be done, for instance, by stating that goal #1 is twice as important as goal #2, or that goal #1 is 100% important during the first leg of the journey and 0% thereafter. No conflicts here. Again, I think, "conflict" is a uniquely human word with a moral implication.

About my "cliff avoiding" control strategy you say:

>This is a prime candidate for a working model. I claim that your
>design for a model of steering will end up with the driver
>oscillating between the edges of the road, whereas a PCT-type
>design will stabilize the car on some path well away from the
>edges.

Your intuition is right. Such simulations have been performed. The PCT-type control will indeed keep the car closer to the middle of the road on average. The avoiding-type control will keep the car on the road for a longer time, however.

>>My point is that the human perceptual + conceptual systems are
>>so beautifully designed that they even extract information from
>>very 'noisy' perceptions.

>They do that only as well as the statistics and the accuracy-time
>tradeoff permits. I don't worry much about extracting signal from
>noise; most of the behaviors we observe work at signal levels
>where noise can be neglected.

This is certainly true in the domain of muscle control. But is it also true in the other domains which concern you like "being a good person, making a decent living, and so forth"?

>>Control must always be limited; the world is just too complex
>>for our three pounds of brains to model it and our fifty pounds
>>or so of muscles to subdue it.

>Well, I won't be nasty and remind you of how wonderful our
>evolved control systems are supposed to me.

Again, "optimal" is not necessarily considered "wonderful" by most people if it is your goal to kill others. But if it IS, you'd better do it in an optimal manner.

>The PCT approach is to define the problem in terms of sensed
>variables: it is the sensed variable that will ultimately be
>controlled, so it should represent something specific in the
>environment to be controlled.

Modern control theory thinks differently. It is, of course, the sensed variables that are our only source of information about how our actions affect the objects that we want to control. But the control problem is not necessarily to bring some variables to some prescribed values and keep them there. That is, of course, a legitimate field for study, but control theory is far broader. By the way, I think that your use of the notion "reference level" confuses some psychologists and their ilk into having to think about "homeostasis". Recognition of this confusion might make the PCT-approach more accepatable to journal editors and referees...

>I think that a little more systematicity would also help control
>engineers, but that's their business.

The problem is difficult, but there ARE attempts to unify the whole field of optimal control systems. One of the earliest attempts can be found in Feldbaum's book "Optimal Control Systems" (Academic Press, 1965), trans- lated from Russian (funny how Russians dominated cybernetic theory for a while, not having computers...). The reasons he gives for his attempts are (I quote):

" (a) Any scientifically sound system is optimal, since in choosing some system by the same token we prefer it to another; hence, we regard it as better than other systems in some respect. The criteria with the aid of which the choice is made ("optimality criteria") can be different. But for any choice, in the final result a criterion for optimality does exist. Otherwise the selection of a system would be impossible."

This calls for the currently usually obeyed design principle of starting with an explicit optimality criterium. The second step is then to use some formalism to translate this criterium into an implementation. One well- known optimality criterium is the integral of the square of the deviation between an output and a setpoint over some period of time,

$$Q = \int_0^T (x(t) - s(t))^2 \, dt$$

where Q should be minimized. The minimum should be sought by whatever means are available for manipulation. If this is just one input u, it is the best u (t) we are after, for 0 < t < T. For linear systems it can be shown that 1) minimizing Q is mathematically tractable [moreover, there are well-known methods that prescribe how to do this] and that 2) there is no local extrema problem [the problem in searching in "rough" landscapes is how to find the GLOBAL extremum]; linear systems yield a smooth, single-peaked landscape where a steepest descent method os some such is guaranteed to find the minimum.

Another criterium, if energy expenditure is of concern, might be

$$Q = \int_0^T \left(a * (x (t) - s (t))^2 + b * u (t)^2 \right) dt$$

This also shows how to combine "subgoals" into one single "supergoal".

The other extreme is to kludge a system together and see how it performs. If the system has only a few tunable parameters, that is doable if tuning one parameter does not disturb the already performed tuning of the others too much. In control engineering terms, this is true only if the system is "linear in the parameters" or almost so (this reminds me of the terrible task of realigning one of those old-fashioned TV-sets with their twenty or so interacting adjustments!). In general, kludging can not be shown to yield a "best" system because "best (possible)" is not defined.

" (b) Any law of nature is a statement of limiting character, an assertion of what is possible and what is impossible in a certain field. Likewise the laws of the general type in cybernetics must give the possibility of making a judgment of what is attainable and what is impossible to achieve under definite real conditions. Therefore they can be formulated in the form of statements about the "ceiling of possibilities" for control. Indeed the finding of this "ceiling" among them is the problem of the theory of optimal systems."

This may explain the discrepancy between my statement that organisms show the best thinkable performance and yours that they sometimes do very badly indeed.

Feldbaum also discusses control systems that are optimal in speed of response. His example is a voltage-following servo-system where a servomotor has to turn a potentiometer such that the voltage at its rotor equals a given voltage. He uses simple arguments to show that the optimal controller must be non-linear. The argument proceeds as follows. "... assume at first that the system as a whole is described by a second-order linear differential equation with constant coefficients. For small damping coefficients ... the curve of U(t) [on a step change of the input E] has a sharply oscillatory character. In this case the control time T [the time required before the system's output is within, say, 5% of the input] is large. If the damping coefficient is made large, then the process U(t) acquires an aperiod [strongly damped] character. ... In this case the control time is also large. By means of an optimal setting of the damping coefficient (it is usually chosen barely less than critical) the control time can be reduced. ...

Let us make the damping coefficient depend on the difference between E and U. For large values of |E-U| let it be small, and the curve of U(t) ... will go along the oscillatory trajectory. But when the difference becomes small, let the damping coefficient be sharply increased. Then the "tail" of the curve of U(t) will show the behavior of the highly damped curve and approach the desired value smoothly. The control time T will turn out to be considerably smaller than with any linear system. These very simple experiments are justified by theory and experiment. It turns out that a system which is optimal in speed of response must be nonlinear, even in the very simple case under consideration".

In the remainder of the book Feldbaum goes on to show how optimal control can be implemented if less and less information is available about the object to be controlled, either because it is missing or because it is disturbed/noisy. Complete information consists of absolutely precise knowledge of:

- the function(s) F relating all different variables;
- the optimality criterium or control purpose Q;
- the values of the variables;
- the values of the disturbances;
- the values of the driving actions.

But systems can be designed also if one or more pieces of information are lacking or are random processes. The classification of systems is in terms of

- the characteristics of the object to be controlled (e.g. linear, quadratic, hysteretic);
- the demands on the performance ('behavior') of the object to be controlled;
- the information that is available to the controller about the object to be controlled.

By the way, the notion of a setpoint or reference level is missing in the text; neither word appears in the index of the book. The voltage-follower described above has an INPUT that should be followed by an OUTPUT. In the engineering literature the words 'setpoint' and 'reference level' appear to be used in the sense of '(more or less) CONSTANT inputs', i.e. in regulato- ry or compensatory control systems.

In his last chapter, Feldbaum notices a difficult problem. If the charac- teristics of the system to be controlled are badly known, control is bad as well. In most engineering situations, considerable time is therefore spent to get to know those characteristics through experiments, simulations and/or fine-tuning of the controller. An organism isn't in this luxurious position. The best it can do is, as Feldbaum shows, combine control with experimentation, discovery. In this part of the book, control theory almost merges with information theory. During experiments, control must necessari- ly be worse. But without experiments, control will be worse forever. Again, an optimal combination of experimenting and controlling should be sought (you surely remember the old-fashioned practice of adding a 'dither' to the input signal; in a primitive way, this serves the same function). Control systems that implement this are called "dual control" systems. Regrettably, so far they require too immense amounts of storage and computing capacities to be practical. I estimate that capabilities in the order of the human brain are required.

Dual control is especially useful if the characteristics of the system to be controlled change over time. For organisms, the object to be controlled is, of course, the world. And that world changes. Dual control is similar to adaptive control, although the latter almost always is NOT optimal but an approximation that allows implementation in existing machinery.

Maybe this explains something about apparently non-purposive behavior. Dual control explains that it is behavior that does not serve an explicit CUR- RENT goal, but has the meta-goal of producing better control in the future. Regrettably, the optimal strategy to follow in the experimentation can only be determined a posteriori, NOT a priori, since it necessarily contains a random component. The frequent human complaint "but I should have ..." is therefore quite understandable (and may even be correct) but at the point in time that the choice was made the information required to make the right choice was not yet there. The complaint is therefore (usually) illogical.

Maybe this gives a meta-perspective on science. Newton's laws are almost intuitive knowledge these days, but the process of discovering them was, to some extent, chance.

The best combination of experimentation and control exists when they coincide. Lucky we, who experiment for a living!

Best, Hans Blom

Date: Thu Mar 11, 1993 12:28 pm PST
From: William T. Powers
MBX: POWERS_W%FLC@vaxf.colorado.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Your video

Hi, Dag --

I've seen your video, and it's very well done. Your charisma is showing signs of coming to life -- there were, in fact, a couple of quite electrifying moments. I think there can be a lot more, but to get that freedom that is needed to carry you along spontaneously, you need to have the format more solidly under control. I'm going to offer suggestions for improvement, not because the presentation was bad but because I think you can make it far better. Your audience was learning well in some parts of the presentation, but not in all parts. I think you can get them to understand much more by reorganizing a bit.

The introductory parts of the talk were good and I wouldn't change them. You gave an interesting overview of what is to be learned and put it into perspective in a way that complimented the audience's intelligence.

When you got into control theory, however, I sensed that the presentation began to get scattered; you were trying to maintain two parallel lines of thought, one involving control theory but the other trying to reassure the audience that all this had something to do with their interests in Deming. It will be much more effective if you separate these goals and accomplish them one at a time. By jumping back and forth between these goals, you make both of them harder to achieve. You may not agree with my recommendations, but perhaps hearing them will suggest some new approaches to you.

I would start the actual introduction to control theory by laying out your strategy to the audience and getting their agreement with it. What I would say would go something like this:

===================

In the rest of this presentation, we're going to go through two stages of development. In the first part, I'm going to teach you the basic principles of perceptual control theory. To do this, it's best to focus on a simple example and make sure you understand every aspect of it, so this phenomenon becomes familiar to you and so you begin to know what to expect. Please don't worry about what this has to do with TQM [or whatever]. I promise that we'll get to that. What I say about the Deming approach will make a lot more sense to you later if you just focus for now on grasping certain relationships that are basic to perceptual control theory. I hope you will interrupt, ask questions, ask me to repeat anything, no matter how simple, that you're not perfectly clear about. The better you understand what you see in this segment, the easier it will be for you to see the parallels when we start talking about real life. So for about the next half hour, let's all concentrate on a single goal, together, which is to master some basic principles and make sure that everyone understands them. The payoff will come in the half hour that follows. I want to hear that you're willing to do this: to forget Deming for half an hour, and work only on understanding the basics of PCT. How about it?

============================

Then I would go directly to the rubber-band experiment. When you do it, you should have clearly in mind a sequence of basic principles that you want to demonstrate and explain. Don't worry about what you're going to say (output). Just be very clear at every stage exactly what you want the audience to understand. You are very good at this; you don't need to worry about your words.

The first part can go pretty much as you did in the video. Set up the task with a volunteer, and spend 20 seconds moving the rubber band around while the subject keeps the ping-pong ball on the target. Then ask the audience to explain it, as you did in the video. Be sure to emphasize that the question is, "What was the relationship between me, as the experimenter, and the volunteer, as the subject? What you would say is causing the subject to behave that way?"

With that finished, after no more than a minute, start explaining, and do so in much more detail than on the video. Say "Watch what my hand does, what his/her hand does, and what the ball in the middle does. Notice that as I pull gradually back on my end, the subject pulls gradually back on the other end. Notice that the ball stays pretty much in one place. Now as I raise (lower) my hand, notice that the subject lowers (raises) his/her hand. And notice that the result is ALWAYS that the ball remains in the same place."

Then explain, right there, that you asked the subject to keep the ball in a specific place. You did NOT ask the subject to move the hand in any particular way. You told the subject what to PERCEIVE, not how to ACT. And right at that point, PROVE that the subject is not reacting to your hand movements. Bring along a sheet of cardboard with a notch in one side. Hold the cardboard at right angles against the paper with the notch directly over the ball so the subject can see the ball, but not your end of the rubber band. And demonstrate that the subject can still keep the ball in one spot without being able to see what your hand is doing. When everyone in the audience agrees that the subject doesn't need to see the disturbance, hold the cardboard so the subject can see your hand but not the ball, and PROVE that control gets much worse when the subject can't see the ball but can see your hand. Check with the subject: "Can you see my hand? Can you see the ball?" Then check with the audience to be sure they get it: that the subject really has to see the ball in order to control it well. You're trying to establish some clear basic facts. This should have taken no more than 10 minutes.

Now you can start drawing your diagram (or have Christine start showing slides). The ball is there in the environment, so draw a ball. The subject has to perceive the ball, so draw a perceptual function and explain that it creates the perception inside the person of the ball outside the person. Now ask, where is the target? The audience will, of course, point to the target circle on the paper. But you say "wait a minute before you decide," and you whisper to the subject to keep the ball six inches to one side of the target, and spend 10 seconds showing the result. Then you ask the audience, "What do you think I told the volunteer?" Most of them will guess right; if they don't, tell them what you said.

Now ask again, what does the subject want and where is that want? What you want them to say, somehow, is that it's inside the subject. The subject is perceiving the distance between the ball and the target, and obviously wants to see a particular distance, not necessarily zero, as you have just shown. When you look at the piece of paper, you see the actual distance but you don't see the wanted distance. So where is it?

Now you go back to the diagram and you show arrows entering the perceptual function from both the ball and the target. You label the perceptual signal "perceived distance." And now you can add the reference signal, labelling it "wanted distance." Emphasize that this wanted distance is now inside the person's head. Then the question is, "So what?" You have here the perceived distance as it is at any moment. You also have here a specification for the wanted distance. Something has to happen right here if there's to be any basis for action. What operation has to be performed?

The answer you want to extort is "comparison." Somehow the person had to bring these two things together, the want and the perception, and judge how they are different. If the actual distance is greater than the wanted distance, the action has to make the distance smaller; if less, the action has to make it greater. So the action has to be based on the difference between the want and the perception. It doesn't depend just on the perception; it doesn't depend just on the want. It depends on the difference between them. Draw the comparator box, label the output "difference," and draw the arrow from the comparator to the output function.

Now ask what the rule has to be for converting the difference into an action, just in one dimension. This is not hard to figure out; if the perceived distance is less than the wanted distance, move your hand one way; if it's greater, move your hand the other way. Make sure everyone understands. Everyone should be nodding. If they aren't, ask what the problem is and fix it.

The last step is to close the loop. Notice that when the subject's hand moves, it moves in the right direction, and that the result is to return the ball to the target. You can illustrate this with the stimulus-response demonstration, suddenly pulling back on your end, suddenly relaxing again.

Now apply a very slow change in the disturbance and show that the ball remains near the target. "Notice that the perception and the action happen at the same time. You can't separate out the disturbance, the change in perception, the comparison, and the change in action. They're all happening at once. You understand how each part of this control system works; now when you see them all operating at the same time, you can see that the result is continuous control, in either one (pull back) or two (move up or down) dimensions. Or more."

Now show the relationship between the disturbance and the action. Explain why it now makes sense that keeping the ball over the target requires the subject's end of the rubber bands to move oppositely to yours. The subject is just correcting movements of the ball. This is the illusion of cause and effect. It seems that your hand mvoements are causing the subject's hand movements. But if you realize that the subject wants the ball to be in a specific place, in between the cause the the effect, of course you understand why the movements are as they are. When you understand what the subject perceives and wants, and what the subject has to do to make the perception match the want, you understand all the relationships between apparent causes and apparent effects.

At this point you can test their understanding. Either really, or as a thought experiment, ask them what will happen if Christine knots another rubber band near the ball and pulls upward on it. How will the subject's hand behave? And why? If they have any problem, ask what will happen to the ball when Christine pulls but the subject doesn't respond. Then ask what the subject has to do to get the ball over the target again (pull downward). It would be very nice to get the audience to make the prediction, and then actually do it and show that they are right. "Were you just guessing what would happen?" No. They KNEW what would happen. How did they know? Because they knew where the volunteer wanted -- intended -- the ball to be.

Now you can say, "When we started this demonstration, I asked you to explain what the subject did. A lot of suggestions were heard. Now, if I ask you again to explain what the subject did, what will you say? What caused the subject to behave that way?" And you should get nothing but right answers.

At this point you might take five seconds and ask "Have you ever heard of a theory of behavior that lets you explain what anyone is doing, however simple, and KNOW that you have the right explanation?"

The final point can now be made. Give the subject the pen, as in your video, and get a trace of the subject's actions as you create some random disturbances. Now you have to make this point very clearly, hammer it home, be sure that every person gets it. You say "If you had just walked into this room, and were told that this is an accurate trace of the subject's exact actions, what could you say about what the subject was doing?"

This is the conclusion of the demonstration. Make sure that all the people understand why observing actions doesn't tell you either what the person wants or what the person is perceiving. It doesn't tell you what the person is DOING -- what those actions are accomplishing that the person wants to perceive as being accomplished.

Close by telling them what comes next. "You now understand the basic concept called perceptual control theory. There is lots more to learn, but what you know now will always remain true. People act in order to make their perceptions of the world match what they want those perceptions to be. You can't understand their actions unless you know what they are controlling, and the specific target. When you do manage to figure out what they're controlling, you can explain an enormous number of cause-effect phenomena -- you can see what the cause is disturbing, and how the apparent effect, the actions of the person, are counteracting that disturbance."

"After the break, we're going to start applying what you know to the Deming Philosophy. We will look for parallels between what you saw in this very simple demonstration and what you see people doing in business management situations. We'll do a little role- play first, then apply perceptual control theory to the situation, and then do another role play later to show how a person who knows PCT will act differently. We'll talk about Profound Knowledge and what all this has to do with Deming's insights. We can't possibly cover all applications of PCT in the time left, but perhaps you are beginning to suspect even now that these applications will penetrate into every corner of life, in business and outside it. I have been developing my understanding of PCT for several years, and still have much to learn. I envy you, because the initial experience of seeing a real theory of behavior unfolding for the first time is one that can't be repeated."

Then the break; give them time to talk about it with each other and let it soak in for a short time.

===

I would really like to see one change in your presentation. It's all right at the beginning to say what PCT means to you and to claim that it's a scientific revolution. But after you have said that, drop it. No more side-comments about how wonderful PCT is or what marvels it will accomplish. To continue doing that is just to show that you aren't sure of the impact of this idea. By the time you get to the break, it will be crystal clear to the audience that they are learning about something very new and completely understandable. I predict that when they come back, you will find every eye on you and a dead silence waiting for you to fulfil your promises. And I think your own confidence and authority will carry you through the last segment, once you realize that you have well and truly hooked them on PCT, using only a few rubber bands.

Just make sure that every word you say after the break links the principles of PCT that your audience now understands to the subject in hand, with no generalizations that anyone off the street could have made (you tend to fall back on them when you're trying to think of what to say next) and no digressions except on purpose. And no matter what the temptation, don't hint at areas of PCT that the audience doesn't yet understand. Save that for the advanced seminar.

A little silence while you organize your thoughts is no crime. Don't say "and" to fill the time while your next sentence forms; just wait for it, and then say it. Finish one sentence, and when it's done, wait for the start of the next one, and say it. Don't let sentences trail off and be abandoned in the middle while you fiddle with the slides. Don't start a sentence before you know how it's going to end.

If you can show how just the few principles you have taught apply in a few real situations, you will have accomplished a great deal, and your audience will leave with a feeling of solid understanding. Don't try to cover life, the universe, and everything. Focus. If they see clearly how PCT works even in just a few simple situations, they will want to learn more, another time. All you need is for the audience to leave with the feeling that you are talking about something real -- they may still have a thousand questions, but the audience will be yours.

Best, Bill

Date: Thu Mar 11, 1993 1:35 pm PST
Subject: Stella Review

There has been discussion of Stella as a modelling tool. The Spring 1992 issue of Natural Resource Modelling (v. 6, n. 2, it just came out!) has a very informed review by Wayne Getz. If you are thinking of using Stella it may be worth looking this up.

-- Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: silvert@biome.bio.ns.ca
(the address bill@biome.bio.ns.ca is only for mailing lists)

Date: Thu Mar 11, 1993 4:05 pm PST
Subject: Behavior: The control of Q

[From Rick Marken (930311.1500)] Hans Blom (931103) --

>We seem to be getting to the point. How central is "the central insight
>of PCT" really?

I agree. Now we're getting to the nitty gritty. And your conclusion is quite straightforward:

>What is REALLY controlled is not sensory input but some grand overall one-dimen-
>sional control purpose that I call Q in my reply to Bill Powers (see below).

Well, that's clear; behavior is NOT the control of sensory input (perception); it is (at least sometimes) the control of Q. Going down a few paragraphs we find out about Q:

> One well-known optimality criterium is the integral of the square of the
>deviation between an output and a setpoint over some period of time,
>

$$Q = \int_0^T (x(t) - s(t))^2 \, dt$$

So this variable is what the system is controlling? And it is not a sensory variable? And if I apply disturbances to Q they will be counteracted by the system, keeping Q in a constant or variable reference state? So Q itself is what is controlled? Not Q as represented by a signal in the control system; the signal that is subtracted from the reference signal? Could you draw me a diagram of the Q control system? I really want to see how the Q control system works so I can simulate it on my computer.

Hm? "Behavior: The control of Q"? I really need that diagram.

Best Rick

Date: Thu Mar 11, 1993 4:50 pm PST
Subject: Optima control

[From Bill Powers (930311.1530)] Hans Blom (930311) --

(butting in on conversation with Rick Marken)

>If you consider the 'signals' that can be recovered from memory
>to be sensory signals as well, you are _logically_ right.
>SOMETHING is controlled by a control system, to be sure. Thus
>far, I understood 'sensory signals' to mean the information
>that reaches me through my sensors AT THIS MOMENT OF TIME.

That's what I mean, too. Control of imaginary information is not control in the same sense -- that is, a different sort of input is being controlled than when real-time perception is involved. My reason for proposing that perception is controlled is to contrast that idea with the externalized view that some external variable represented by the perceptual signal is controlled. The latter view holds up only as long as the perception maintains the same relationship to the external variable. In engineering that is a maintenance and calibration problem; in organisms it takes on a greater importance, because perceptual functions are largely learned, and are modifiable. It is only the signal that continues to match the reference specification, in that case; the external counterpart of the controlled signal changes as the perceptual function changes form.

>We come to the problem of what is REALLY controlled. It is not
>really the sensory functions. For example, in many engineering
>applications, a compromise is sought between optimal
>stabilization of some variable and minimum energy expenditure.
>Control requires fuel to be burnt, and fuel is expensive.
>Organisms, I suppose, have the same problem.

What do you think of our solution to problems like this, in the form of multiple control systems and hierarchies of control?

>What is REALLY controlled is not sensory input but some grand
>overall one-dimen- sional control purpose that I call Q in my
>reply to Bill Powers (see below).

If you're trying to wrap up an entire organism as a single hypercomplex control system, I suppose you would have to look for some grand overall system and a single overall purpose. That isn't the approach in HPCT. There may be many highest-level control systems acting in parallel, with relative independence. Of course there is an overall control system in my model, too, a reorganizing system, but it isn't concerned with learned behavior. Its reference levels and perceptual signals are built- in, and its mode of action is to reorganize the rest of the system. It isn't really a single entity, but a collection of control systems concerned with maintaining the life support systems, each one being concerned with a specific variable.

Writing to me:

>... you confuse "optimal" (an engineering word with an exact
>meaning) with "good" (a moral categorization). The "optimal" of

>engineering means only that some system reaches its grand
>overall goal as closely as possible, by definition.

I'm sort of between these meanings. If there are two control systems with incompatible goals inside the organism, clearly they are going to expend a lot of energy cancelling each other's efforts. This is suboptimal under certain assumptions: that energy expenditure is probably a cost to the whole organism, and that reduction of the control range resulting from conflict reduces the ability of both control systems to counteract disturbances. These losses of ability aren't "morally" bad, but the organism would be able to control over a wider range and for a longer time if they were not present. Of course given the conflict the control systems are in fact coming as close as possible to reaching their goals. But with a suitable adjustment of the system organization, they could come a lot closer. A great deal of psychotherapy is aimed at helping people resolve conflicts; I suppose you could say that helping them is a moral choice, but it does have engineering overtones.

>In optimal control theory, there is only one "supergoal" that can be controlled.

Can you explain why this has to be true? What if there is more than one control system operating at the highest level of organization? Of course you could make up some "supergoal" having to do with an optimal balance between these systems, but in that case the criterion of optimality would be in the eye of the beholder -- there would be nothing in the system itself trying to achieve that optimality.

>It would be possible to declare the two (seemingly conflicting)
>goals "drive in the middle of the road" and also "drive one
>yard to the right of the middle". But then you would have to
>combine them into one goal. This can be done, for instance, by
>stating that goal #1 is twice as important as goal #2, or that
>goal #1 is 100% important during the first leg of the journey
>and 0% thereafter.

This disembodied declarer has to live somewhere. What is it that sets these two disparate reference signals? What is it that decides to follow the rule you suggest? It can only be a superordinate system, and I see no necessity that there be just one such system.

>Your intuition is right. Such simulations have been performed.
>The PCT-type control will indeed keep the car closer to the
>middle of the road on average. The avoiding-type control will
>keep the car on the road for a longer time, however.

Longer than forever? I don't follow this at all. Why should the PCT-type control system EVER let the car go off the side of the road? I assume, of course, that it never malfunctions, but you have to assume the same thing for the edge-avoider. All things being equal, the PCT-type system would be able to resist larger disturbances than the edge-avoider, because the edge-avoider would keep shooting back and forth between the edges, and when near the edges would be more vulnerable to disturbances of a given size, assuming that it and the PCT controller had equal output capabilities.

>>... most of the behaviors we observe work at signal levels
>>where noise can be neglected.

> But is it also true in the other domains which concern you
>like "being a good person, making a decent living, and so forth"?

Yes, I think so. Most people have a pretty clear idea of what they mean by these things, and are quite sensitive to pressures that tend to cause errors, and act promptly and firmly to counteract any deviations from what they want to perceive.

I think that one of the legacies of traditional psychology is a general impression that human behavior is complex and chaotic, with regularities appearing only as statistical averages and with the future being a matter of rather shaky predictions. PCT, once you get used to seeing the things it calls to attention, shows a very different picture. Most behavior is highly regular and closely controlled; there is very little left to chance.

If this were not true, the world we experience would be very different. People would keep getting lost on the way to work; buildings and houses would constantly be falling down, or fail to have doors or windows, or be located in inaccessible places. Cars, if they ran at all, would always be crashing into each other or wandering off across fields. Nobody would know how to grow crops, or harvest them, or transport the food to some regular destination, or how to cook the food or keep it from spoiling. Most of the things that we use, encounter, or rely upon wouldn't even exist.

What astounds me is the way in which psychologists could have looked at the endless regularities of human existence, mostly maintained by and completely a product of human efforts, and failed to recognize them. It is terribly naive just to take the world the way you find it without asking how it could possibly be that way. Psychology has focussed on unusual side-effects, on tiny irregularities, and has failed to see the massive regularity that characterizes all living systems and the environments they have shaped to fit their wants. The signal-to-noise ratio in most aspects if life is very, very high. That has not prevented scientists from concentrating on the noise, and ignoring the signal.

>Modern control theory thinks differently. It is, of course, the
>sensed variables that are our only source of information about
>how our actions affect the objects that we want to control. But
>the control problem is not necessarily to bring some variables
>to some prescribed values and keep them there.

No, I have never said it was. PCT leads to HPCT, in which higher levels of control act by VARYING the reference signals for lower systems. They do so as their way of controlling derived perceptions, more generalized perceptions. Those systems in turn have their reference levels adjusted by still higher systems, concerned with still more abstract perceptual variables. The only dissonance between this view and your ideas of optimal control have to do with your assertion that at some level there is a single highest control system with a single highest goal.

As to your criteria of optimality, they are completely discretionary. I don't see any reason to suppose that organisms have adopted such criteria or seek to realize them. You're talking about engineers building control systems, not the processes by which living control systems evolve. The engineer can, by choice, combine all lower goals into supergoals, but there is nothing that compels us to suppose that organisms do the same thing -- except when they're trained as engineers.

>But systems can be designed also if one or more pieces of

>information are lacking or are random processes.

Yes, certainly. But the degree of control achieved under such conditions will not be impressive. People keep saying "But you have to use what's available!" My response is, "Sure, but don't expect much of the result if there's not enough information or regularity available."

Feldman seems to be saying something similar:

>If the characteristics of the system to be controlled are badly
>known, control is bad as well. In most engineering situations,
>considerable time is therefore spent to get to know those
>characteristics through experiments, simulations and/or fine-
>tuning of the controller. An organism isn't in this luxurious
>position. The best it can do is, as Feldbaum shows, combine
>control with experimentation, discovery.

That would seem to describe the situation for organisms perfectly. The characteristics of the system to be controlled -- the environment -- are not known at all, a priori, to any organism. For most organisms, they are never known in any systematic way -- most organisms don't "know" things in any cognitive human terms. Yet they do manage to acquire good control, so evidently cognition isn't the only way of knowing things.

Best,   Bill P.

Date: Thu Mar 11, 1993 8:25 pm PST
Subject: more more naive jacobianism, etc.

[Avery.Andrews 930312.1324]   Rick Marken (930311.0800)

I'm glad my little piece wasn't universally & completely incomprehensible. But..

>First, I would like to emphsize that the Jacobian approach you
>describe is CLOSED LOOP. So it can be thought of as one approach to
>the control of perception.

The closed-loopedness is also emphasized in the robotics textbook I got the terminology out of, so this is standard fare.

>The problems that Avery is running into with this approach are a
>bit esoteric for me -- but I think it has to do with the fact that
>certain fingertip targets (reference settings) result in
>outputs of the jacobian that specifiy physically incorrect or
>thematically impossible states of the arm components.

I hope they aren't really too esoteric, since I suspect that PCT-ers who don't understand the nature of this kind of problem in at least one area are sitting ducks for harsh criticism. The problem is that for certain arm configurations and error vectors, Cartesian xyz coordinates lead to Jacobians that don't give *useful* information for control: movement at both joints is orthogonal to the error vector, so this system thinks it's stuck. Nothing is wrong or mathematically impossible - the naive system is just too dumb to do anything appropriate with it.

One of the ways in which people are likely to try to rubbish PCT is knock together examples of systems with badly chosen perceptual dimensions, and then use the failure of these systems as evidence that PCT is naive. The best way to oppose this is to get in early with the point that choice of an appropriate perceptual system is a or the crucial step to the solution (`deciding what to do is easy when you know what's in front of you' is the way David Chapman put it).

>I think another way of saying this is that the success of the Little
>Man stems, in part, from incorporating appropriate perceptions to be controlled.

Yes. Finding appropriate perceptions is also a Gibson-Turvey et. al. theme for what that's worth.

>But an important implication of this is that the PCT
>model is HIERARCHICAL; the "error vector" in PCT specifies references
>for LOWER LEVEL PERCEPTIONS -- not for environmetal outputs, as
>in the jacobian approach.

I'm not sure about the `implication' bit - I don't see how anything I've said implies a hierarchy. Note that I've been pretending that perception in shoulder-based polar coordinates happens by magic (but it really can't be too hard, just a few trig and inverse trig functions stuck together, which some sort of neural circuit out to be able to compute). The key to Little Man's success is that it uses perceptual dimensions (`polaroid coordinates'), that have relatively smooth and monotonic relationships to changes in the relevant joint angles. One can certainly design manipulators that don't have such nice properties, but actual arms seem to be designed to be relatively `control-friendly'. My point is that this may well be an important part of the explanation of why we are pretty good at doing things with them.

Avery.Andrews@anu.edu.au

Date: Fri Mar 12, 1993 1:25 am PST
Subject: arm thoughts

[Avery.Andrews 930312.1912]

A comment or two on arm-ish issues before heading off on a little camping trip: the use of polar (head or shoulder centered) coordinates will prevent an arm from getting stuck in the way that the xy cartesian controller will, but it won't necessarily produce sensible behavior for large errors. So I suspect that if the hand is extended, and the target a bit above the elbow, you'd get some shoulder-elevation, since, initially, the hand is too low, and shoulder elevation is supposed to be the way to fix that. I haven't actually seen this behavior with a bill-style controller, though I have seen it with mine, when it is modified to control radial distance errors in addition to xyz coordinates.

Bizzi et. al. claim to have evidence that hand movements typically follow pre-determined trajectories, so that in ordinary manipulative activity, large errors would not occur, so things seem to hang together. Super-fast movements, like serious martial arts maneuvers, might be different, but they require heaps of practice, so I don't think much of interest follows from them, other than the fact that people can develop complicated Central Pattern Generation Circuits if they work hard enough at it.

Date: Fri Mar 12, 1993 7:16 am PST
Subject: Arms and stuff

[From Bill Powers (930312.0830)] Avery Andrews (930312.1912) --

Do you have a copy of Arm Version 2.0? It allows you to set up two target positions, and then have the target jump back and forth (instantly) between them at about 1-sec intervals, while the arm tracks. The initial error for each movement is as large as it can get, the entire magnitude of the target jump. The trajectory of the fingertip for such movements is a not-quite-straight line from the initial position to the final position (if the jump is 5 cm, the deviation from straight is roughly 2 or 3 mm). This is comparable to path curvatures in experiments with real human arms.

>Bizzi et. al. claim to have evidence that hand movements
>typically follow pre-determined trajectories ...

In what sort of situation? And what is doing the pre-determining? In Arm v.2, the paths for maximum-speed jumps are indeed predetermined, but they're not precalculated -- they simply emerge from the time-constants of the control systems and the dynamical properties of the arm. The trajectories for fast movements are often cited as evidence of a CPG, and their repeatability is taken as a sign that they're planned. But neither is necessarily true, or even possible.

For movements taking an appreciable length of time, the control systems simply follow the reference signals, and any path is possible. Put a dot on a piece of paper, and another dot about 3 inches to the right of it. Put a pencil on the first dot. Now write your name cursively so that the end of the last letter leaves the pencil on the second dot. How's that for a "predetermined" or "planned" trajectory?

>(`deciding what to do is easy when you
>know what's in front of you' is the way David Chapman put it).

Oh, cute. How easy is "knowing what's in front of you?"

Best, Bill P.

Date: Fri Mar 12, 1993 8:14 am PST
Subject: more more naive jacobianism, etc.

[From Rick Marken (930312.0730)] Avery.Andrews (930312.1324)--

>I hope they aren't really too esoteric, since I suspect that PCT-ers
>who don't understand the nature of this kind of problem in at least one
>area are sitting ducks for harsh criticism.

You ARE new to PCT. PCTers who DO understand the nature of the problem get harsh criticism ANYWAY -- and the better the PCTer understands it, the harsher the criticism (witness the Fowler-Turvey critique of PCT).

> The problem is that for
>certain arm configurations and error vectors, Cartesian xyz coordinates
>lead to Jacobians that don't give *useful* information for control:
>movement at both joints is orthogonal to the error vector, so this
>system thinks it's stuck.

This is similar to the problem Fowler and Turvey found with PCT; given their non-linear output function they showed that the error was often ambiguous about what to do to correct the error. This was true -- they just failed to notice that the same problem existed for their own "solution" (a version of a closed-loop control system that they did not recognize as such). What they found was simply that in some unusual cases (like focusing a camera or tuning a radio) a control system cannot know "what to do" until it does something and determines the perceptual effect of doing it. The fact that people can control in these situations is simply evidence for a hierarchy of control systems -- the very notion that Fowler and Turvey were trying to trash. Even the Jacobean problem (getting stuck) could be solved by adding another control system that perceives some measure of "stuckness" (possibly the derivatives of the finger tip) and "jogs" the components of the jacobean to keep that measure at some reference level.

I said:

> >model is HIERARCHICAL; the "error vector" in PCT specifies references
> >for LOWER LEVEL PERCEPTIONS -- not for environmetal outputs, as
> >in the jacobian approach.

You say:

>I'm not sure about the `implication' bit - I don't see how anything I've
>said implies a hierarchy.

Then I don't understand what you were saying.

My impression is that the jacobean has only one "level" of perception, the x,y,z or rho,theta1,theta2 coordinates of the finger tip. You can remap the coordinates of this point all you want -- Catersian, polar, city block, whatever -- but it probabaly won't make a big difference in the performance of your model. I thought you were suggesting that it is the coordinate space that is implicit in the components of the jacobean that would be mapped in useful ways in the PCT approach. This IS what the ARM demo does. Instead of a jacobean (or any other kind) of direct transform from finger error into arm position change outputs, the ARM transforms fingertip errors into references for several independent lower level perceptual variables -- this is the "coordinate space" in which the higher level systems live. These lower level perceptions (of arm angles and functions of several arm angles) are controlled by adjsting references to even lower level perceptions (the forces that change the angle and derived angle variables).

So you may not have implied that a hierarchy was part of your suggestion that finding appropriate perceptions is the PCT solution to complex control -- but if you didn't -- you should have. It's NOT just a matter of finding the "right" perceptions. The PCT solution to your jacobean problem is not to find the "right" coordinate space for the finger tip (the only controlled perceptual variable in your model). The PCT solution is to find lower level perceptual variables that can be used as part of the means to control finger tip position; PCT suggests that the jocobean matrix for computing output from error is the wrong approach; complex output computations like that are unlikely to be performed by the nervous system anyway, and they kind of bind you to a particular solution space (so that you get stuck when circumstances conspire to push you out of that space). The solution is to eschew output calculation and design the system so that errors are references for INPUT, not parameters for OUTPUT. That's why the ARM works; no complex output complex calculations; and no complex input calculations for that matter; it's the magic of the hierarchy. Do you have a copy of my spreadsheet model?

Best   Rick

Date: Fri Mar 12, 1993 8:51 am PST
From: Ken Mann / MCI ID: 507-3685
TO: * Dag Forssell / MCI ID: 474-2580
Subject: President Bill via MCI Mail

If you'd like to reach President Clinton, the White House now has
an MCI Mail mailbox. It will be downloaded and reviewed. Just
address your messages to.....

To: White House

Regards: Your MCI Mail Agent   Ken

Date: Fri Mar 12, 1993 9:10 am PST
Subject: Dag's video

[From Bill Powers (930312.0930)]

I sent the critique below to Dag Forssell after seeing the video tape of his "Purposeful Leadership" presentation to a group of Deming aficianados. Dag asked me to put it on the net, "warts and all," so I am doing as he asked. I suppose the principle is that any publicity is good publicity.

Hi, Dag --

See previous post......

Date: Fri Mar 12, 1993 9:28 am PST
Subject: Re: Information -- the challenge

[Allan Randall (930312.1200)]   Rick Marken (930310.1400) writes:

> First, you say that Information Theory (IT) is to PCT as
> calculus is to Newton's laws; IT is a tool like calculus.
> But calculus helps us make detailed predictions...
> ...at the end of your post you say:
>
> >you seemed to be asking for a *prediction* about which
> >(condition 2 or 3) will be better in a real-world situation.
> >This seems to be beyond the scope of information theory as it

> >stands now.
>
> Well, IT isn't much of a tool if it can't help us predict things;
> looks like PCT WITH IT is no better off than Isaac's brother Phil
> WITHOUT calculus.

But calculus does NOT allow you to make the kind of prediction that Bill Powers is asking for. What Bill is asking for is like asking calculus to predict the orbit of a planet, all by itself WITHOUT any of Newton's Laws. Bill is specifically asking, unless I am misunderstanding him, for a PCT-type prediction using information theory and ONLY information theory. I agree that this may not be possible. What I do not want to see him do is to reject what could turn out to be a valuable tool, on the basis that it cannot do the job all by itself. This is what I was getting at with the Newton analogy. Contrary to what you imply, calculus simply can't make physical predictions unless it is used in combination with a physical model of some kind. All by itself, it is just a mathematical technique, like information theory.

> >If no information about the
> >disturbance can be extracted from this data, then there is
> >no way the system can translate the error from this signal into
> >an action on the world that will counter the disturbance. Is this
> >or is this not true?
>
> This is NOT TRUE. Surprise!
>
> >If you agree, then you are agreeing with an
> >information theoretic analysis.
>
> So I guess I disagree with an IT analysis -- and if I'm right
> (which I am) then the IT analysis is wrong, right?

Right. Exactly. Here is something we can actually nail down. The statement I made, which you say is not true, is the whole crux of this controversy. It is a fundamentally information theoretic statement. The reason why I asked whether it was true or not was that I could not imagine a PCTer disagreeing with it (I was wrong - you did), but at the same time I considered it an information theoretic statement. Here you are claiming the statement is actually false - there is no information about the disturbance in the perceptual signal - a claim that is in direct contradiction with Ashby's Law of Requisite Information. If you are correct, Ashby's Law is completely unfounded. (However, I don't think you actually are correct.)

> It is this fact about control systems that nails everyone to the
> wall -- and proving it to myself is what turned me into a PCT "fanatic"
> ...In a high gain, negative feedback control loop, the
> output DOES NOT depend on the sensory input; rather, SENSORY INPUT
> IS CONTROLLED BY OUTPUT.

The second statement is true, but I think your first statement is false. The output DOES depend on the sensory input, and the sensory input DOES depend on the output. We do not have to choose between the two. It is better to say they are interdependent, than that one depends on the other. Isn't your response-stimulus description just as wrong as stimulus-response?

> What you do in a tracking task is NOT caused by what you see; ...

Magic, then?

>....there is a LOOP so that what you see is both a cause AND A RESULT
>of your output.

Now you are back to the closed loop and admitting that they are interdependent.

> The nearly perfect
> relationship between output and disturbance does NOT exist
> because the system has access to information (from the sensory input)
> about the disturbance.

This is silly. How can there be ANY relationship between the output and disturbance, let alone one that is "nearly perfect," if the system has no access to information about the disturbance? This is just physically impossible, Rick. I think either you are misinterpreting what I mean by "information," or you believe some kind of witchcraft is responsible for control.

> When control is good, there IS NO information
> about the disturbance in the stimulus -- NONE, ZILCH, NADA.

This is true NOT when control is *good* but when control is *perfect*, which is inherently impossible for an error-control system, as Ashby said. Let's be clear what we mean by "disturbance." Stop me when you disagree. The disturbance is the sum of all the various environmental influences impinging on the CEV. This CEV is not an absolute property of the environment, but is defined by the hierarchy. In information theory terms, the hierarchy is an encoding scheme for representing the environment. Each perceptual signal represents one CEV in the environment. This does *not* mean that the disturbance exists in the hierarchy and not in the world. In order to describe anything, says information theory, it must be described in some language. The disturbance is in the world, but the description that isolates it as an entity seperate from the rest of the environment is in the hierarchy, not the world.

Only that information which is relevant to control of the CEV is transmitted in the percept. But to say there is NO information at all, none, zilch, nada, is just incoherent. If I am driving my car down the road and there is a sudden huge gust of wind from the right and at the same time I start sliding on the ice, the "disturbance," in PCT terms, is the disturbance to the CEV that results from these forces. But an outside observer will tend to describe the disturbance in a different language (a different encoding scheme). They will probably describe the gust of wind, its force, and the sliding motion due to the ice as seperate, complex entities. But this description is no more an absolute depiction of reality than the perceptual one inside the driver! The "disturbance" relevant to PCT is the one described inside the organism that is controlling, not the one described dispassionately by an external observer. Both use an encoding scheme to describe the disturbance. One scheme requires many many bits, while the other requires very few.

This is what I mean when I say that the hierarchy brings the information content of the disturbance in line with the output capacity. An external observer describes the disturbance in a language that requires many bits (such as the detailed description of the molecular positions/momentums that the compensatory thermostat requires). The internal hierarchy describes the same real-world disturbance in a language requiring only a small number of bits - a number that can be handled by the capacity of the output channel (such as the much simpler description used by the error-control thermostat). This is Ashby's Law.

> The output
> mirrors the disturbance because this is what the output MUST DO in order
> to keep the input IN THE REFERENCE STATE; this is the magic of
> closed loop control.

Yes, this does sound like magic, and not a scientific explanation at all. There must be an explanation WHY the control system is able to do what it MUST DO. Just to say that it MUST DO it does not explain anything.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Fri Mar 12, 1993 9:53 am PST
Subject: Re: Optimal control

[Martin Taylor 930312 12:03] butting in recursively on
(Bill Powers 930311.1530) butting in to Hans Blom (930311)

>>What is REALLY controlled is not sensory input but some grand
>>overall one-dimen- sional control purpose that I call Q in my
>>reply to Bill Powers (see below).
>>...
>>In optimal control theory, there is only one "supergoal" that
>>can be controlled.
>
>If you're trying to wrap up an entire organism as a single
>hypercomplex control system, I suppose you would have to look for
>some grand overall system and a single overall purpose. That
>isn't the approach in HPCT. There may be many highest-level
>control systems acting in parallel, with relative independence.

In fact, exactly as many can be in active control at any one moment as there are output degrees of freedom. From moment to moment the set of grand overall "control purposes" may change, if the world happens to be kind enough not to disturb the CEV corresponding to perceptions not at a particular instant being controlled. Any (well, a very large) number of perceptions can be monitored, at any level, including the "supergoal" level, but only a reasonably small and well defined number can be simultaneously controlled.

>>It would be possible to declare the two (seemingly conflicting)
>>goals "drive in the middle of the road" and also "drive one
>>yard to the right of the middle". But then you would have to
>>combine them into one goal.

Only because they contain components that set disparate reference values that are intrinsically in the same perceptual dimension. That is the nature of "conflict" in PCT. In PCT "conflict" is a technical term without moral connotation, just as "optimal" is in engineering. "Conflict" in PCT refers to a situation in which there are two or more reference signals that cannot all simultaneously be matched by the corresponding perceptual signals. No more, no less.

Conflict is normal in a distributed perceptual control hierarchy, in which most, if not all elementary control systems (ECSs) have perceptual input functions that are non-orthogonal to those of many other ECSs. Conflict is abnormal in a discrete, designed hierarchy in which the various ECSs control perceptions that are substantially orthogonal. Most demonstration hierarchies are of this second kind.

>>Your intuition is right. Such simulations have been performed.
>>The PCT-type control will indeed keep the car closer to the
>>middle of the road on average. The avoiding-type control will
>>keep the car on the road for a longer time, however.
>
>Longer than forever?

I suspect Hans is talking about a linear system that will show a Gaussian distribution of error if the distribution of disturbances is Gaussian. The tail of a Gaussian distribution is infinitely long, so the PCT car will eventually go off a road of any width, no matter how high the gain. At some time, in this model, the car will be subject to a cross-wind of Mach 3, and even if it didn't slip or get blown off its wheels, the engine wouldn't be strong enough to compensate. I don't think the edge avoider would do any better, but even the edge-avoider is a perceptual controller, so there is no theoretical issue of PCT versus non-PCT control to separate the centre-seeker from the edge-avoider, In fact, the driver might control both perceptions simultaneously.

>The characteristics of the system to be controlled --
>the environment -- are not known at all, a priori, to any organism.

Oh, but I think they are, to a very large extent. Accepting:

>For most organisms, they are never known in any systematic way -
>- most organisms don't "know" things in any cognitive human terms.

what has been the effect of 4 billion years of evolutionary reorganization, if not to get each organism to "know" a great deal about the world in which it has to control?

Martin

Date: Fri Mar 12, 1993 12:54 pm PST
Subject: Dag's video

[From Dag Forssell (930312 12.15)]   Bill Powers (930312.0930)

>I suppose the principle is that any publicity is good publicity.

Some feedback is better than no feedback.

I asked Bill to put his long comment on the net before I had even read it, because I believe this kind of communication belongs on the net. I also wanted netters to become more aware of the existence of our video.

I was able to digest Bill's post for all of three hours before we put on a repeat performance last night. Bill's suggestions enhanced the rubber band exercises.

Bill, your post was much more than just a private comment. I am sure Gary will take note for his file on portable demonstrations. Greg appreciates it for the archive. I think several others share my interests in ways of selling PCT to markets OTHER THAN the Psychological Journals. Your suggestions are valuable and belong on the net.

I look forward to comments from others as you see the video. I am striving for "continuous improvement," which is a control process in and of itself and benefits from feedback.

Let me repeat my announcement of (930308-1):

On Feb 4, Christine and I presented an introduction to PCT to a Deming Users Group. This included a presentation of:

Results
Profound knowledge
Theory, Revolutions
Role play, before
Self-directing system / Control -- [see Bill's comments]
Role play, after
Systems thinking
Seven tools - Perception
The essence of TQM
Dr. Deming's 14 points - Integration

I will mail a video with supplemental booklet to anyone who asks. I'll ask for $10 for the tape and booklet with additional for postage as follows:

Postage to	Surface, 4th class	Air
USA	$2.-	$3.-
Canada	$3.-	$4.-
Europe	$4.-	$9.-
Pacific	$4.-	$11.-

I will honor E-mail, fax or letter request and trust that you add up and send U.S. funds by snail mail. State surface or air.

Note: The video is 119 minutes, 1/2 inch VHS, NTSC (U.S. video signal). Overhead slides are shown in closeups.

Dag & Christine Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808

Phone (805) 254-1195 Fax (805) 254-7956

Date: Fri Mar 12, 1993 1:16 pm PST
Subject: arms, coordinates, etc.

[Avery.Andrews 9303130630] Bill Powers (930312.0830)

>Do you have a copy of Arm Version 2.0? It allows you to set up
>two target positions, and then have the target jump back and ...

Yes, & I'll try it out, but Arm2 is quite complicated, & I was talking about the behavior of a very simple system. I'm not yet sure exactly what properties of the Arm2 model allow it to produce straight lines in this situation.

>In what sort of situation? And what is doing the pre-determining?
>In Arm v.2, the paths for maximum-speed jumps are indeed
>predetermined, but they're not precalculated -- they simply
>emerge from the time-constants of the control systems and the
>dynamical properties of the arm. The trajectories for fast

In pointing gestures. When arms are bopped towards the target, the return back to an appropriate point on the trajectory, & then proceed on it. No evidence for `planning' as opposed to `generation'. Looks to me as if slowers like those between the visual kinesthetic levels in Little Man shoudl be able to do this kind of generation.

>>(`deciding what to do is easy when you
>>know what's in front of you' is the way David Chapman put it).
>
>Oh, cute. How easy is "knowing what's in front of you?"

Hard. That that's the hard part was supposed to be one of the main points of his book.

(Rick Marken (930312.0730))

Just a few points - no time to say more.

>What they found was simply that in some unusual cases (like
>focusing a camera or tuning a radio) a control system cannot know "what
>to do" until it does something and determines the perceptual effect of
>doing it. The fact that people can control in these situations is simply
>evidence for a hierarchy of control systems -- the very notion that Fowler

This general point is important and well taken, but it's not a good solution to the problem of distance-from-shoulder control when the elbow is fully extended, since grownups in any case clearly know what to do in this situation without experimentation.

>You can remap the coordinates of this point all you want -- Catersian, polar,
>city block, whatever -- but it probabaly won't make a big difference in
>the performance of your model.

I have already seen this to be false - adding the polar distance to the cartesian Jacobian model improves performance considerably. I've also seen the resulting model get stuck in situations where polar coordinates would get it unstuck. There might be clever things to do with the cartesian jacobians that would help in these situation (whence the `naive' in my title), but on the whole, though I can't think of an good reason to continue developing this approach, when well chosen perceptual dimensions eliminate the problems without further cleverness.

> PCT suggests that the
>jocobean matrix for computing output from error is the wrong approach;
>complex output computations like that are unlikely to be performed by
>the nervous system anyway, and they kind of bind you to a particular solution
>space (so that you get stuck when circumstances conspire to push you
>out of that space).

Actually, they're pretty simple, just a few sines and cosines which neural circuits shouldn't have much trouble with. Maybe it would even be useful to have them around for certain purposes, tho at the moment I don't know what.

Avery.Andrews@anu.edu.au

Date: Fri Mar 12, 1993 1:33 pm PST
Subject: social categories

[From: Bruce Nevin (Thu 930311 15:20:09)]

Continuing with (Bill Powers (930303.0830)), which was in response to my (930302.1430) --

> Some people think (almost always mistakenly) that the way to
> entertain people around them is to tell jokes in dialect. They
> may not achieve an entertained state in those who hear the joke,
> but that is certainly the goal, and the correct explanation for
> the way the joke-teller is pronouncing words at the moment. I
> have observed that people who tell such jokes also perceive that
> the others are entertained, which explains why they persist. They
> perceive themselves in a social role that matches what they want
> it to be; the others perceive something else, but their self-
> images prevent them from being frank.

Are you sure that all the others present share your perceptions of the situation. You are inveighing against the joke-teller's delusion that he has privileged access into the perceptions of those around him. Whence your privileged access here? Specifically, how do you know that none of them is in fact entertained and that all of them feel constrained by their self images from being frank?

I assume that these are true statements about your experiences. Why does your self image prevent you from being frank? Does it have something to do with a relationship between how you imagine you would be perceived by others if you were frank, on the one hand, and how you want to be perceived by others, on the other hand? What basis do you have for these expectations?

Are you sure that the joke teller is controlling for a perception of "I am entertaining these people"? What else might he be controlling for? It is almost always a he, isn't it? What might that have to do with it? (Hint: look at the Tannen book I mentioned.)

> >And social dialect is an excellent index when you consider
> >natives of a place like NYC, for example.

> You can't run that logic both ways at once. If dialect is used to
> identify the subpopulation, then you can't prove that there is
> any shared characteristic other than the dialect, in which case
> all you've shown is that some people speak with that dialect and
> some don't. Either you pick the dialect as the discriminator, or
> some other set of characteristics like living in the inner city.
> To prove an association you'd have to show that not using the
> dialect means you don't live in the inner city, or that living in
> the inner city means that you have to speak with that dialect,
> and the two converse cases. If ALL the cases don't prove out
> every time, then all you've got is one of those statistical
> generalizations that's true of a population but not of any
> particular person. You certainly don't get any explanation for
> why one person speaks in a particular dialect on a particular
> occasion, and another doesn't.

In one study, Labov asked a large number of people a simple informational question in a department store. He arranged it so that it was a typical anonymous request for information, one on one, attended to and quickly forgotten. He arranged it so that the answer would display the speaker's pronunciation of a sound with respect to which two NYC social-class dialects contrast. For example, he asked sales clerks where a certain item was that he knew was on the fourth floor, and they answered "fourth floor." Then he went around the corner and recorded his transcription of the diagnostic sounds (I think is was /o/ and postvocalic /r/) for that person. In the lower-class department store, he uniformly got one pronunciation, no exceptions; in the upper-middle class department store he uniformly got the other pronunciation, no exceptions.

Now, he might have encountered a lower-class person in the wrong store, but in independent interviews they said they never shopped there because they couldn't afford it and they felt out of place there, so while shopping in the store cannot be ruled out for an individual it is unlikely. Similarly for those who could afford not to have to put up with the crowding and squalor and poor quality merchandise of the lower-class store, and who "coincidentally" all pronounced the other variant of the diagnostic sound (no exceptions).

Here, I'll venture a prediction: no one reared through puberty in a lower-class [an upper-class] family in NYC speaks an upper-class [a lower-class] NYC dialect. The only assumption here is that the way of talking have been learned in the usual way, without formal study and tutilage in the manner of My Fair Lady. In other words, apparent exceptions against which you would forthwith no doubt love to rail will have made an unusual and strenuous effort to learn the other dialect (and probably have exceptional acting or linguistic talent or they'll botch the attempt at mimicry). I don't know what the exceptional people will have been controlling for, but it must have been very important to them, because it's a lot of work. And even then, let the telephone ring with an old buddy or family member on the other end and something engaging to talk about, and effortful deviation from norms learned in childhood are out the window.

Middle-class (upwardly mobile, downwardly vulnerable) seem to be in some sense bi-dialectal, switching from one way of speaking to the other depending upon circumstances and the attention they give to how they are speaking. (See comments on Labov's somewhat humorously named measure of linguistic insecurity.) But the speech norms of middle-class people for formal talk are not identical to the norms for upper-class speech, nor are their norms for casual talk identical to the norms for lower-class speech. They are norms to which middle-class speakers make their speech conform. (They are closer to the norms of an earlier generation of upper-class and lower-class speakers, respectively.)

People who moved around a lot in childhood (like myself) obviously present a more complicated situation, and anyway fall outside the above prediction.

> Excuses, excuses. Let's face it: this way of
> understanding people just doesn't work. So-called social strata
> and groupings and norms exist mostly in the imagination of the
> statistician. They are perceptual prejudices. If you pick any
> individual at random and try to apply your generalization, you
> will find that you need a dozen excuses to explain why this
> person doesn't fit the category or behave like others whom you
> put in the same category. When you finish with the excuses (none
> of which is in fact investigated to see if it actually fits), all
> you have done is to discard from your sample all the individuals
> who don't fit the generalization. If you go through the
> population person by person in this way, you may end up with nobody left.

Seems to me you are flailing lots of straw around here. I'm not over in that direction where you are attacking the statisticians and the social determinists. I'm right here. Norms are not products of statistical analysis. Nor is the behavior of humans deterministically fixed by norms. For some norms (which appear to be enforcedly subconscious for reasons discussed on other occasions) people seem to choose to conform very closely. Under exceptional conditions, such as hypnosis or intensive study and practice, individuals may be able to shift to different norms.

> Society, in my view, is a continuum in which individuals vary
> along a hundred dimensions in ways independent of each other. Yet
> the members of the society like to look at superficial secondary
> characteristics like where one lives, how one dresses, how one
> speaks, what one does for a living, what one says aloud about
> beliefs, and so forth, and forms perceptions of groupings that
> have almost nothing to do with what makes each person that person
> and no other. Worst of all, the individuals then begin to deal
> with others whom they see as sharing these apparent patterns of
> secondary characteristics as if they were all right in the
> middle, and all alike. The use of formal statistics to define
> such groupings simply makes the prejudice harder to recognize for
> what it is.

I agree with what I perceive as the basic thrust of all this. Individuals participate in social patterning to varying degrees, and by their participation constantly help to re-institute it.

Shared prejudice, and control of one's appearance with respect to prejudices known to be shared, is a fact of social life. This includes of course controlling so as to appear to defy or spurn norms. Nonconformity typically depends upon the individual's perception of normative patterning just as strongly as conformity does, as indeed it must if it is to have an intended significance to others. Utter carelessness of what one may signify to one's fellows is I think beyond the scope of what we are considering, and not so common.

But the prominent characteristics (or caricatures of characteristics) that are the stuff of prejudices and shibboleths are a small and misleading subset (or intersecting set) of social norms. Are you aware of how much of your teeth you expose when you smile? Upper or lower or both? There are consistent regional differences.

I don't recall saying anything about formal statistics. Refresh my memory if I did.

Bruce   bn@bbn.com

Date: Fri Mar 12, 1993 1:55 pm PST
Subject: Information and control

[From Bill Powers (930312.1400)]   Allan Randall & Rick Marken (930312) --

The confusion here comes from not specifying what you mean by "information about" the disturbance. Obviously, if the controlled variable were not disturbed, the control system would do nothing. So in the variations of the controlled variable, there is "information" to the effect that there is a disturbance of some sort acting.

However, the control system reacts directly to changes in the controlled variable, regardless of what is causing them (even if the system itself is causing them). Also, it acts directly, at the same time, ON the controlled variable. The result is that the variations in the controlled quantity do not represent the state of the disturbance alone, but only the combination of the disturbance plus the system's own output actions. As these effects are almost exactly opposed when control is good, the variations in the controlled variable actually represent only the DIFFERENCE between the system's output and the disturbing variable.

From this we can perhaps clarify the argument.

The actual state of the original disturbing variable is NOT represented in the state of the controlled quantity. If at a given moment the disturbing variable has an amplitude of 100 units, the controlled quantity might actually be deviating from its undisturbed state by only 1 unit, or even by -1 unit. The explanation is, of course, that the system's own output is at the same time producing -99 units, or -101 units, of effect on the controlled quantity.

The output of the system can carry small random variations due to noise sources inside the system. Hence, when a nominal -99 units of output is required to cancel a disturbance of 100 units (in a control system with a loop gain of 100), the output can actually vary spontaneously by an amount greater than the amount of uncancelled disturbance. As a result, the controlled variable will reflect these endogenous noise variations, and those variations are normally comparable to the difference between the system's mean output and the disturbing variable's amplitude.

The practical result is that residual variations in the controlled variable will show only a very low correlation with variations in the disturbing variable, a correlation approaching zero. They will also show a low correlation with the system's own output. These low correlations are easily demonstrated in tracking experiments. With disturbances of low bandwidth, there is no apparent relationship between the observed variations in the controlled variable and variations in the disturbance. Statistical analysis shows correlations that range from +0.1 to - 0.1 (or so). Lagged correlations are not much larger.

If control were only a matter of information transmission from the disturbing variable into the control system, there would have to be enough information surviving the trip from the disturbing variable, through the controlled variable, through the sensors, to the output variable to account for the close match of the output effects to the disturbing effects. Yet what we observe is that the output can match the disturbance within one or two percent, while the controlled variable through which all this information has to pass shows random variations as large as the remaining variations due to the disturbing variable -- a noise content of at least 50%, within exactly the same bandwidth.

I think that what the information-theoretic approach may be missing is the fact that the only information path that really matters is the one around the closed loop: from controlled variable to sensor to output to controlled variable. The control system is organized so as not to depend on knowledge of the causes of disturbances. In fact, it rejects information from outside the control loop by systematically opposing the effects of any external influence. This rejection is not based on knowing the causes of those effects, but only on knowing the states of the variables inside the control loop. No matter where in the control loop an outside agency tries to disturb a variable, the remainder of the loop will respond to prevent such disturbances from having any important effect at the point of injection. Normally, of course, only the controlled variable in the external part of the loop is accessible to outside influences.

So I believe that the mistake Ashby made, and that information theorists in general have made, is to suppose that the action of a control loop is somehow driven by external influences. If that were true, you would have to explain the action of the system in terms of information passed to it from outside it. But in fact, control loops are driven by information completely inside the loop; they act to reject the effects of information from outside. This is essentially the same mistake made by stimulus-response theorists, when they saw disturbances being rejected by the actions of control systems, and concluded that the disturbances were stimuli and the actions were responses caused by the stimuli.

Best, Bill P.

Date: Fri Mar 12, 1993 3:35 pm PST
Subject: Re: Information -- the challenge

[From Rick Marken (930312.1300)] Allan Randall (930312.1200) --

>But calculus does NOT allow you to make the kind of prediction
>that Bill Powers is asking for. What Bill is asking for is like
>asking calculus to predict the orbit of a planet, all by itself
>WITHOUT any of Newton's Laws.

I don't think Bill or I know what the heck you folks think IT is good for. Feel free to use IT along with anything you like to make a prediction. All I want to see is how IT (calculus-like) can improve what we do in PCT -- which is try to discover controlled variables and model now they are controlled.

In a previous post you said:

>If no information about the
> disturbance can be extracted from this data, then there is
> no way the system can translate the error from this signal into
> an action on the world that will counter the disturbance. Is this
> or is this not true?

I said this is not true and you said:

>Right. Exactly. Here is something we can actually nail down. The
>statement I made, which you say is not true, is the whole crux
>of this controversy. It is a fundamentally information theoretic
>statement. The reason why I asked whether it was true or not was
>that I could not imagine a PCTer disagreeing with it (I was wrong -
>you did), but at the same time I considered it an information
>theoretic statement. Here you are claiming the statement is actually
>false - there is no information about the disturbance in the
>perceptual signal - a claim that is in direct contradiction with
>Ashby's Law of Requisite Information. If you are correct, Ashby's Law
>is completely unfounded. (However, I don't think you actually are
>correct.)

I'm correct.

OK. Let's get quantitative here. You can do the information analysis if you like; I'll just give you the results that I know. In a compensatory tracking task, when control is good, the correlation between the disturbance and the output (handle movement) is typically .99+. The correlation between input (cursor-target variations) and output is on the order of .02. In other words, what the subject sees (stimulus) has NO relationship to output; but the output is a perfect mirror of the disturbance. Now you are claiming that there is information in the variations in the stimulus that communicate information about the disturbance -- this, accoring to you, is why the output mirrors the disturbance. If you are right, then when the SAME disturbance is used on two different occasions and the output mirrors the disturbance in both cases then the information in the stimulus on both trials must be the same -- right? This is what the experiment in Ch 3 of "Mind readings" tests; the prediction is that the correlation between traces of the stimulus on two different trials using the SAME disturbance should be nearly the same; but the typical correlation between the stimulus traces was ALWAYS less than .2 -- in one case it was .0032.

So there is NO information about the disturbance in the stimulus (input) to a control system. My experiments show what I would interpret as nearly PERFECT transmission of information from the disturbance to the output -- with virtually NO information about the disturbance in the only channel that could be carrying that information -- the stimulus.

All this experimental rigamarole shouldn't really be necessary (except that the result is SO startling and unbelievable); all that should be necessary is the equations of control showing that o = -d (notice that the perceptual variable does NOT show up in this equation) and the fact that the stimulus input variable in a compensatory tracking task is i = o + d; that is, what you see (the input, i)at any instant is the COMBINED result of your actions (o) and disturbance (d). So information about d is not even physically EXPECTED to be visible in the input in control tasks like the compensatory tracking task. The only way a person could possibly get information about d from the input data is to know, at any instant, what EFFECT they themselves are having on the input (that is, know what o is) -- and this is impossible in principle. If they could get information about o then they could continuously solve for d = i-o (they can see i and continuosly subtract the informaiotn about o). But how could they POSSIBLY know o? DO you know what that would require?

>This is silly. How can there be ANY relationship between the output and
>disturbance, let alone one that is "nearly perfect," if the system has
>no access to information about the disturbance?

I suggest that you read chapter 3 of Mind Readings and ALL of Behavior: The control of perception. I think your question above is a good one; it reveals exactly WHY PCT has NOT gained much of a following in the behavioral sciences. You can't get many people to hop on your bandwagon when they think one of your fundemental observations is "silly". I know it seems silly -- how about crazy and impossible too. Unfortunately (for those of us who understand it and, hence, are alienated from mainstream behavioral science) it it TRUE.

>This is just physically impossible, Rick.

It is not only physically possible, is the only physically correct interpretation of the situation -- PCT or no PCT (remember that i = o + d). The interpretation that is actually WRONG is that one that seems right (and OBVIOUS) to you and 100,000 other behavioral scientists out there -- that there is information about the disturbance in the stimulus. There IS NO INFORMATION ABOUT THE DISTURBANCE IN THE STIMULUS. Goodby behavioral science as usual, hello looney bin.

>I think either you are misinterpreting what I mean by "information,"
>or you believe some kind of witchcraft is responsible for control.

Neither. I think I am correctly interpreting the meaning of information (though I'm willing to defer to you and Martin on that) and I believe that the nearly perfect transmission of information from disturbance to output despite the absense of ANY information about the disturbance in the input is a well understood result of the behavior of a negative feedback control loop -- no witchcraft necessary.

>Only that information which is relevant to control of the CEV is
>transmitted in the percept. But to say there is NO information at all,
>none, zilch, nada, is just incoherent.

I'm just like that -- incoherent but honest (sort of like the fool in Lear). Fact is, NO information about the disturbance is transmitted in the percept in a tracking task. But the ball is in your court; SHOW ME that I am wrong. I can take it.

>There must be an explanation WHY the control system is able to
>do what it MUST DO. Just to say that it MUST DO it does not explain anything.

Yes -- you are corrrect. I was trying to dodge a lot of explanation; but the explanation is simply that the output of a control system is continuously driven by an error variable (r-p) whose effect on output is REDUCED by that output ; ie. the explanation is the closed loop negative feedback process. That process, with NO information about the distubance or disturbances that are influening the perceptual input, produces outputs that are a perfect mirror image of the net disturbance -- and, hence, CONTROL.

Best Rick

Date: Fri Mar 12, 1993 8:54 pm PST
Subject: Information flow through control systems

[From Bill Powers (930312.2100)]

RE: information theory and control systems.

In my last post I came across an idea (control systems reject information from disturbances) that has now led to a further development that may (or may not) help resolve all these 'tis-so- 'taint-so squabbles.

In a very rough way, we can estimate the information capacity of a channel simply by comparing the input waveform with the output waveform. If the output bears a close resemblance to the input, we might guess with some confidence that the information capacity of the channel is at least as great as the amount of input information, because apparently little information is being lost. This is not an exact calculation, but a reasonable first approximation.

Now apply this to the key relationships in the behavior of a control system. There are two major channels we can observe or infer: one from the disturbance that affects the controlled variable to the action that opposes the disturbance, and one from the reference signal to the state of the controlled variable.

If we apply a varying disturbance to the controlled variable, we do not observe that the waveform of the disturbance is reproduced in the controlled variable. Instead, it is reproduced (inverted, but faithfully) in the action, the output, of the control system. As we increase the bandwidth of the disturbing variations, we observe faithful reproduction of the disturbance waveform in the output actions, up to some limiting bandwidth. After that bandwidth is exceeded, we find that high-frequency information in the disturbance waveform begins to be lost, failing to be reproduced in the output or action waveform. So we have an estimate of the information-carrying capacity of that channel.

In the same way, we can apply a varying waveform to the reference-signal input to the control system. Because this waveform is transmitted through the comparator to the output of the system, we would expect, perhaps, that the output waveform would closely resemble the waveform applied to the reference input -- but it does not (in the presence of disturbances, at least). It is the controlled variable that faithfully follows the waveform applied to the reference input, whether or not disturbances are present. By gradually increasing the bandwidth of the waveform applied at the reference input, we could find the bandwidth at which the waveform of the controlled variable begins to differ from that applied to the reference input, and again estimate the channel capacity.

We can therefore say that information impinging on the control system from outside, in the form of a disturbance waveform, shows up almost unattenuated in the waveform of the system's output, while information impinging on it in the form of a reference waveform shows up almost unattentuated in the variations of the controlled quantity -- in both cases, provided that the bandwidth of the driving waveform is not too great.

Information, therefore, is transmitted through a control system in two ways: from disturbance to output, and from reference signal to controlled variable. Within the channel capacities, information is not destroyed but transformed: the control system does not "reject" information as I proposed earlier, but routes it to different places, depending on its origin.

The mechanics of this dual transmission of information are contained in the inner workings of the control system. The closed-loop relationships that actually exist are quite different from the apparent input-output relationships we see in the two "channels." Any information-theoretic analysis of such a system must NOT use any simple straight-through causal calculations, but must be done with the closed loop taken into account.

So far I have not seen any information-theoretic analysis that takes the closed loop into account. The normal assumptions would seem not to hold inside the control system, for there is no simple path for either flow of information to follow. The information content of a signal inside the closed loop is partly determined by that signal itself; I have never seen this case treated. Perhaps someone else has.

The important insight here is that information from both sources
is not lost, but shows up in another system variable.

Best,   Bill P.

Date: Fri Mar 12, 1993 9:22 pm PST
From: Jackson
          EMS: ATTMail / MCI ID: 414-0940
          MBX: GI=Ray
          MBX: IN=RL
          MBX: DDA=ID=rljackson

TO: * Dag Forssell / MCI ID: 474-2580

Application message id: MAC-1.4-306690-rljackson-26
Grade of Delivery: Normal

Subject: Your page today 3/12   3/12/93   DIRECT

Hello Dag!
I received your page today, but I've been in meetings with customers and potential customers (Boeing and ATT), as well as conducting 20 hrs of teambuilding seminars this week.

I also got the video and intro package...I watched a few minutes of the video and am looking forward to dedicating some time to it soon. As a side note, I thought it was interesting to see what you looked like after all the communications we've had to this point. No real shock, but my first impression was: Why is an attractive, intelligent young lady like Christine hanging around a guy like you?...(joke).

Anyway, I have some things to discuss with you about not only the materials you've sent (which immensely please me), but also on some of the communications I've had and will continue to have with ATT & Boeing -- they were more excited about what I'm doing at MCG than the people I'm doing it for (isn't that usually the case with "internal" consultants?); hopefully there's a potential client for you (us?) there. I'll post this feedback to the CSGnet.

Finally I don't want you to think that I take your requests for input lightly; in fact, I'm always flattered you ask, and appreciative that you include me in your mailings. However, you know my schedule (for anther 8 weeks from tomorrow), but I hope to get you some quality feedback soon.

I'll be in touch.

Take Care, Ray

Date: Sat Mar 13, 1993 11:24 am PST
Subject: 'tis-so-'taint-so

[From Rick Marken (930313.1000)] Bill Powers (930312.2100) --

>In my last post I came across an idea (control systems reject
>information from disturbances) that has now led to a further
>development that may (or may not) help resolve all these 'tis-so-
>'taint-so squabbles.

I think you are a bit more optimistic than I am about resolving this matter. The only resolution I see is realizing that it 'tis so (sensory inputs contain no information about the disturbance), that it 'taint so (sensory inputs contain complete information about the disturbance) or that it doesn't matter. I think Allen Randall has stumbled upon the characteristic of PCT that makes it truly revolutionary and totally unacceptable in conventional behavioral science. I am enthusiastically picking at this "'tis-so-'taint-so" dispute, not because I am interested in "dissing" (as my daughter would say) IT but because I consider acceptance of the point (that the sensory input is a dependent variable; not an independent variable in control systems) to be the main stumbling block to correct understanding and acceptance of PCT by conventional behavioral scientists.

There are those on this net who don't like referring to PCT as "revolutionary". But Allen Randall has put his finger right on the revolutionary button of PCT. If there is ONE thing on which ALL conventional behavioral scientists agree it is that sensory input is an INDEPENDENT VARIABLE -- it is the beginning of the causal chain that ends in behavior. Even conventional behavioral scientists who believe that the feedback effects of behavior are important would still agree with this basic point -- sensory input is an independent variable. This is the absolutely fundemental assumption of ALL research in in the behavioral sciences -- ALL OF IT. What PCT says (in no uncertain terms) is that this assumption is FALSE -- sensory input is NOT an independent variable in a living control system (ie. all living organisms) -- it is a DEPENDENT variable (the independent variable being the reference signal inside the organism). I think it is clear that, if PCT is right about this then the WHOLE edifice of behavioral science comes crashing down -- somwthing up with which most behavioral scientists will unquestionably not put.

There is obviously not much room for a compromise resolution to this argument. PCT is just inherently revolutionary; and efforts to make it seem "not revolutionary" require tacit agreement with the fundemental assumption of the behavioral sciences -- that sensory input is an independent variable. This kind of compromise leads to the kind of PCT done by Carver/Scheier and their ilk. So while I agree that crying "revolutionary" all the time can be annoying, the fact is that PCT IS revolutionary -- whether we like it or not -- and failure to accept PCT as revolutionary prevents one from giving up (dropping, forgetting about, saying goodby to, etc) all the things that MUST be given up if one is going to do PCT correctly. You can't do space science (easily) with a Ptolmeic model; it is equally difficult (actually impossible) to study control in the context of a sensory input -- response output model of organisms (the one used when you do a conventional behavioral science experiment).

I don't expect anyone (especially conventional behavioral scientists) to just sign up to the "sensory input is not an independent variable" canon of PCT; I want people to test it as harshly as possible; that's why I welcome this debate with Allen Randall. I hope Allen sticks with it and tries like mad to show that the PCT view can't possibly be true. I want him to think of experimental tests, mathematical proofs -- whatever-- to try to convince ME that my argument (sensory input is NOT an independent variable) is false. I want him to push on this little point as hard and persistently as possible because this is where (for the behavioral sciences anyway) the rubber really meets the road (how's that Ed?). If I lose this debate -- and Allen is able to convince me that sensory input does carry information about the disturbance and, thus, is an independent variable that causes output -- then everyone will be happy because I will then have to shut up about PCT being revolutionary. If however, Allen relents and agrees that sensory input is always and only a dependent variable (in a closed loop negative feedback system) then maybe (but still unlikely) we can get some of the behavioral scientists who are looking in on this net to abandon their misguided efforts and start doing some PCT based research.

I know that you (Bill P.) are looking for ways to take what people can already accept and move them gradually and gently to the PCT perspective. That's fine with me; but I don't think one can make a gradual shift into PCT (any more than one can make a gradual shift from an earth to a sun centered view of the solar system). I don't think your approach will be any more or less successful than mine -- the only people who will really "get" PCT are the one's who want to -- no matter who is teaching it or how. But I want to pursue this argument about the "information in the stimulus" because 1) it's fun and 2) it's public so ANYONE on the net who thinks PCT is wrong about this can explain why -- and I really am interested in what they have to say.

So, again I claim 1) that there is typically no information about disturbances in sensory input (and when such information is available it is either unnecessary or, if used, it reduces the amount of information transferred from disturbance to output) even though system outputs perfectly mirror the information in these disturbances) and 2) sensory input is always a dependent variable, never an independent variable, in living sytems, thus there is virtually nothing that can be learned about the nature of organsisms in the framework of the typical (independent variable- dependent variable) behavioral science experiment.

'Tis so.

Best Rick

Date: Sat Mar 13, 1993 2:21 pm PST
Subject: BRAIN EVOLUTION - RKC

SUBJECT: BRAIN EVOLUTION - RKC

[From Bob Clark (930313.1700 EST)] GARY CZIKO (930307.0450 GMT)

A few weeks ago I looked at my Britannica (about 30 years old) for items that might relate neural sytems to PCT. I looked up various animal systems as well as Nervous Systems and Comparative Anatomy therof. Several were very interesting -- in non-PCT terms, of course. The descriptions seem quite complete and detailed. Several of the simpler forms seem to do quite well without any kind of control system. It seems that some form of Central Nervous System is probably essential for any PCS beyond a single level system with a fixed reference level. And no control system at all may be adequate for the simplest animals to "respond" to environmental disturbances such as changes in pH and variations in concentration of chemicals.

I would think that the U of I would have suitable experts who could summarize this information -- and might thereby become interested in PCT!

Regards, Bob Clark

Date: Sat Mar 13, 1993 2:30 pm PST
Subject: DAG'S VIDEO

SUBJECT:DAG'S VIDEO - RKC

[From Bob Clark (931306.???? EST)] Re: Dag Forssell (930312 12.15)

{Thanks for the quick "response" to my "stimulus!" That is, I sent out a signal and a package came back. From the User's viewpoint, that is sufficient.}

I can't make any detailed suggestions without knowing more about your audience and your purposes is making the presentation.

I assume that the tape consists of samples from various places in your actual performance. For example, I don't think the tape is intended as a sales piece. Rather, it is offers suggestions that other PCTers might use. Indeed, I am interested in your methods because I may have such an opportunity.

I also assume that you have had a considerable amount of experience in speaking to groups such as the one shown on your tape. Thus I will skip detailed items of manner etc.

I would like to suggest that a verbal presentation is much more effective if you work from a structured outline, making clear your overall purpose and breaking that down into secondary, supporting topics. It has been said, for verbal presentation, that it helps to start with the over-all topic, followed by development of the topic, concluding with summarizing the topic. This helps because listeners need to know "what to listen for," and also may be distracted from time to time, missing some of details. I notice that you provide written material both as a supplement to the actual presentation and for later review. To me, this seems essential for an unfamiliar subject.

Since my purposes in making such a presentation would be different from yours (and Bill's), I'll refrain from re-writing your material. (I'm too lazy anyway!) Rather, I will comment that your opening remarks about history of science and relations between Deming, PCT and yourself impress me as suitable for an audience already having some knowledge of both topics. Perhaps an academic group. An alternative might be to use Deming's 14 rules (or are they "principles?") as examples of PCT application rather than using PCT to support Deming.

But any of this must, of course, be adapted to your situation and goals.

You have changed my view of the rubber band demos. I had realized, of course, that they are useful and supplement the "Portable Demonstrator" that we used some 30 years ago. But when you introduced the "Double Ball" version, it added a new dimension. With the single ball the lower orders can be demonstrated by suitable adjustment of the timing and the patterns used by the demonstrator. However, without losing any of them, the second ball makes the subject a full participant. She is asked to "select which ball to control!" This can be carried further by suggesting that she change, from time to time, which ball she is controlling. Another step: switch who is the demonstrator and who the subject, done on the command of a third party. These observations have some additional consequences that I plan to discuss separately at another time.

Good show Dag, I appreciate the opportunity to know about your activities.

Regards, Bob Clark

Date: Sat Mar 13, 1993 2:56 pm PST
Subject: Re: Information and control

[Martin Taylor 930313 17:40] Bill Powers 930312.1400

I'm not sure whether I made the following point in a private posting to Bill or in public to CSG-L. I think the former, but if it was the latter, I apologise for the duplication.

>I think that what the information-theoretic approach may be
>missing is the fact that the only information path that really
>matters is the one around the closed loop: from controlled
>variable to sensor to output to controlled variable. The control
>system is organized so as not to depend on knowledge of the
>causes of disturbances. In fact, it rejects information from
>outside the control loop by systematically opposing the effects
>of any external influence.

The essential point about any stable system is that the uncertainty observed at any pointin it will be stable over time. The closed loop ensures this by opposing uncertainties that might be introduced by the disturbance.

Overall, what this means is that numerically there will be no information supplied by the disturbance to the preceptual signal. This may seem odd, or even magical, but it has to be so. If it were not, the whole system would have an ever increasing uncertainty wherever it might be observed. Entropy increases in a closed system. Closed systems are not stable. Stability comes from the flow of energy through a system. If the system is a control system, it is the use of this energy that allows the actions to oppose the disturbance. If the actions suffice, then the system is stable both mechanically and informationally. The uncertainties may on average reflect the uncertainties in the disturbance, in that control is imperfect, but that is a question of the dynamics of the system as a whole. Even Rick agreed that the dynamics of the system allow the dynamics of the disturbance to be detected, and the same is implicit in Bill's comment.

Sorry, the spurious character rate is getting to be higher than I can cope with. I'll stop here and try to edit out the ones I can see. But I think I've made most of my point.

Martin

Date: Sat Mar 13, 1993 3:43 pm PST
From: goldstein
Subject: your video

Dear Dag,

I would very much appreciate receiving a copy of your video to view.

My address is: David M. Goldstein, Ph.D.
                  801 Edgemoor Road
                  Cherry Hill, NJ 08034
                  (609) 667-0166

Thanks very much. I will send you a check as soon as I receive it.

Best regards, David Goldstein@saturn.rowan.edu

Date: Sat Mar 13, 1993 8:48 pm PST
Subject: Slipping Rubber Bands

[from Gary Cziko 930314.0410 GMT]

Dag Forssell and Bill Powers (yesterday or the day before):

As Dag predicted, I was interested in hearing of these new uses of the trusty old rubber band demo. Here's another twist (or rather, slip). In this demo, the subject is told beforehand to keep the knot over some inconspicuous (to the audience) target and the audience is trying to guess what the subject is "doing" as the demonstrator disturbs by pulling on his or her end (of the rubber band, that is).

Instead of using two rubber bands tied together, use three tied together end to end. Have the subject hold the end loop of one as usual, but you (the experimenter/demonstrator) hold on to the second knot, NOT the end of the third rubber band (got that)? Don't loop any fingers through, but hold on to the second knot between your thumb and index finger (as you would normally hold on to a string) with the third rubber band concealed within your hand.

Now do the demo as usual. When someone from the audience invariably says the the subject is simply mirroring your movements, stop at a position where there is good tension on the bands and then gradually let the third rubber band slip through your fingers and then hold again as the end arrives between your index finger and thumb. While the rubber band is slipping through your fingers the audience will see the subject move his or her hand toward yours WHILE YOUR HAND REMAINS STILL.

So much for the "experimenter's hand as stimulus" explanation of the subject's behavior.

===

Rick Marken and Allan Randall (many times):

Isn't it amazing that we have quite simple working model of what perceptual control is all about and yet it still often seems like magic? What's so hard about seeing that it is the difference between perception and reference level that "causes" behavior (while, of course, behavior at the same time "causes" this difference--when the reference level is unchanging, that is). Isn't it amazing that the S-R and "cognitive" people have no working model and yet it is much easier to "understand" the S-R and cognitive "models"?

I think it was Richard Feyneman who said that if you kept thinking about the absurdities of quantum mechanics (instead of just using it as a tool for work in particle physics) you would just "go down the tubes" and be through as a physicist. Maybe we have here something just as absurd but just as useful and "true" as the laws of particle physics. If physics survived the apparent absurdity of quantum mechanics, maybe there is hope that psychology can also survive the apparent absurdity of perceptual control.

Hm, I suppose this casts Bill Powers or Rick Marken as Niels Bohr and Martin Taylor or Allan Randall as Albert Einstein. Will Albert find the disturbance information hidden somewhere in the perceptual signal? Or will Niels's absurd magic triumph. I'm staying tuned!--Gary

Date: Sun Mar 14, 1993 9:22 am PST
Subject: Info theory: putting some ideas together

[From Bill Powers (930314.0900)]

Martin Taylor and I have been having a go-round privately. In one of his posts (930312) he said

>If I (the sensor) know that those equations are being used, and
>there is no added noise or phase uncertainty, then I need not
>look at the sample values. There is no uncertainty, and the
>signal conveys no information.

I was very puzzled about why the sensor had to know anything, but evidently the last sentence above stuck, and worked away in my head to lead to my post on the net yesterday, first about the control system rejecting information from the disturbance and then the understanding that there is one virtual channel from disturbance to output, and another from reference signal to controlled variable. The better the control system (the smaller it can keep its error signal and the wider the bandwidth over which it can do so), the more faithfully do these channels reproduce their input information at their outputs.

Information in higher levels of control inside the organism can be expressed in the environment as the behavior of visible controlled variables. That is the output part of communication. It represents a flow of information from the organism into the environment. But the obvious source of information flowing from the enviroment into the organism -- a disturbance acting on the controlled variable -- is not an actual source, because that information is cancelled by the actions of the control system, and does not show up as changes in the perception of the controlled variable. Environmental information from the disturbance shows up in the output actions of the organism, again in the environment, cancelling the effects of the disturbance. So disturbance information enters the system, but comes out again immediately as opposing actions.

So how DOES information get into the organism from the environment? Clearly, only through uncontrolled perceptions. Uncontrolled perceptions may be uncontrolled at many successive levels, but eventually they join with controlled perceptions to become inputs to a high-level control system. There they act as disturbances of the perception to which they contribute, _and their information content is therefore rejected by that system_. The actions of that system that oppose the effects of the uncontrolled perception result (after passing through layers of reference signals on their way toward the periphery) in variations of visible controlled variables. Those visible variations carry the information about what the highest-level system involved did to cancel the disturbing effect of the incoming information. And perception of them joins with the uncontrolled perceptions to leave the highest-level perception undisturbed.

This leads to a rather odd concept of communication: the output information from a system produced by intentional variations in controlled variables reflects the attempt of a higher-level control system to prevent uncontrolled input information from disturbing a higher-level controlled variable.

This has a negative ring to it, but perhaps this is not as negative a picture as it seems. Suppose that the highest-level system involved (which could be at any level) receives information via uncontrolled input channels that disturbs a world-view that is maintained by manipulating controlled perceptions. The first reaction to the uncontrolled information input will be an adjustment of lower-level reference signals that cancels the disturbance of the world view which the uncontrolled information tends to produce. If this is verbal communication, a statement that upsets the world view is countered by adjusting controllable details of perception. When I introduce a black rocket scientist to a bigot, the bigot immediately adjusts the controllable perceptions to explain how a dumb black person could be a rocket scientist. Oh, he's their token black, he's probably being carried by his colleagues so he appears brilliant, but the real work is probably done by someone else. This leaves the bigoted world-view of black people undisturbed; the information is rejected.

This works as long as the adjustments of controlled perceptions suffice as a way of rejecting the information and maintaining the world-view perception. But if the right kind of information is received, control may become more and more difficult, or fail. Then the person must start to reorganize -- find another world view and other ways of adjusting lower reference signals in which the adjustments will again succeed in preventing a disturbance from the uncontrolled information input. This can lead to a change in world view so that the disturbance is no longer a disturbance. The person now thinks, "Ho, hum, another black rocket scientist, so what's new?."

Or apply this to a conventional psychologist observing a rubber-band demonstration pf PCT. At first, this uncontrolled new information leads to strenuous attempts to adjust other perceptions and prevent a disturbance of the conventional view of how behavior works. But the longer additional information about control behavior enters the system, the more problems there are in finding adjustments of already-controlled perceptions that will in fact cancel the disturbance. Eventually the perception will change, and the former disturbance will be seen as a proper aspect of behavior, behavior the way it is now understood to work.

We say that we understand and accept communications from others when doing so does not constitute a disturbance of our controlled perceptions -- in other words, when no information is received from the communications. This is what makes gatherings of like-minded people so boring. None of their communications with each other cause any disturbances, or carry any information. It is what makes the CSG, and CSGnet, so interesting and attention-grabbing: there are plenty of disturbances, and world-views (or whatever) are undergoing continuous reorganization.

How'm I doing, you information theorists out there?

We can now see where Ashby made his mistake about "requisite variety." He reasoned that in an error-driven control system information from a disturbance must affect the regulator via its (the disturbance's) effects on the essential variable, so the disturbance information was used by the regulator as the basis for creating opposing actions. This led to a paradox, because the better the regulation, the less information about the disturbance was available to direct the actions of the regulator. This led to the idea that error-driven regulation must be inherently imperfect -- not just in the technical sense that a tiny bit of error had to remain, but in a much more serious sense. If the regulator were to lose as much as half of the information about the disturbance, how could it possibly cancel out more than half of the disturbing effects? Ashby didn't actually reason that way out loud, but the way he used the idea of error-driven regulation subsequent to that argument showed that he didn't expect much of it by way of precision.

Now we can see that the error-driven regulator is simply a different beast from the disturbance-driven regulator. The information that drives the output of the regulator comes from the essential variable, not from the disturbance. In fact, the effects of the disturbance are excluded from the control loop; the information about the disturbance is rejected. An information-theoretic analysis of the error-driven regulator must start froms scratch -- it isn't simply an extension of the analysis of the disturbance-driven regulator.

Best, Bill P.

Date: Sun Mar 14, 1993 10:31 am PST

Subject: Niels and Albert

[From RIck Marken (930314.1000)]   Gary Cziko (930314.0410 GMT) --

The "slipping rubber band" demo is beautiful; elegant and simple; you're getting to be awfully good at this.

>Hm, I suppose this casts Bill Powers or Rick Marken as Niels Bohr and
>Martin Taylor or Allan Randall as Albert Einstein.

I don't like that casting. I feel more like Albert trying to explain to a Newtonian that relativity predicts that the measured speed of light will be the same regardless of the direction of motion of the observer relative to that of the light. I don't think the PCT model involves magic (in the way quantum physics seems to); you gave a nice simple explanation of the PCT model in your post. What we are having trouble with is convincing people that a "non-prosaic" phenomenon (sensory input as DEPENDENT rather than INDEPENDENT variable) actually occurs and that it is predicted by the "prosaic" PCT model. I think relativity had an easier time than PCT because Albert was dealing with an actual science (physics) which expected (and got) precise predictions from and confirmations of the theory. PCT offers precise predictions and confirmations but conventional behavioral science offers statistical relationships in return ("well, that may happen in your little tracking task, but look at all the studies that have found a statistically significant relationship between independent and dependent variable"); not much you can do to fight that -- except possibly pray for the salvation of their immortal souls.

Best

Rick (sleeping peacefully now that Bill's in and the jerks are out) Marken

Date: Sun Mar 14, 1993 6:00 pm PST
Subject: Re: the challenge

[Allan Randall (930314.1830)]

Bill Powers (930310.1845 MST) writes:
> To make the parallel work better, however, you should think of
> something other than the calculus, which in fact CAN be applied
> to specific problems in mechanics, so Ike Newton would never have
> rejected it.

To my knowledge, calculus actually CANNOT make the kind of predictions you are asking of information theory. Calculus can be used to make predictions only in combination with a physical model such as Newtonian mechanics.

> I'm sure that AFTER I have designed a control system to behave in
> a certain way, an information theorist could estimate the
> information flows in the system, the entropies, and all that lot.
> If I were trying to model behavior that takes place under
> difficult conditions, this analysis might offer more of interest
> by way of predicting limits of performance.

You seem to be admitting here that information theory might have something useful to say about control systems. It just isn't something that interests you. However, in the past, I believe you have indicated otherwise, going so far as to say that information theory is invalid (I think you basically called it a bad analogy).

So I guess I still don't understand exactly where you stand. Is information theory completely wrong-headed or is it correct, but of little use to PCT?

I am now going to lay out a proposal for the challenge. This may not meet what you had in mind. Please let me know what fits and what doesn't. Forgive me for being so neurotically picky, but I do not wish to accept a challenge if it is unclear to me exactly what is being asked for and what the position is of the person doing the challenging.

THE CHALLENGE:
Using the three conditions (prior to the experiment being performed) make some assumptions about the bandwidths of the system, and compute in information theory terms the amount of control required for the three conditions. This will require PCT *and* information theory together. No prediction will arise from information theory alone. This is for the very reasons you have given: from Ashby's diagrams + information theory, one cannot predict what exactly R, the regulator, is doing. You cannot predict that R is going to oppose the disturbance. Whether this will meet your requirements for the challenge is the main point I'd like clarified before accepting.

The result might be something like: "Condition x requires such-and-so entropy flow, while condition y requires so-and-such entropies in order for control to occur. Condition x, as you can see, requires an output bandwidth much higher than that allowed for in our assumptions about the real-world experimental setup. Thus, we predict that condition x will not work, while condition y might."

> My point was that when you characterize signals in terms of
> information flow rather then in terms of amplitude and phase, it
> no longer is possible to predict the result of the above
> convergence. If ordinary statistical measures like variance are
> used, we would NOT in general expect the variance of E to be less
> than that of D or of R.

That is correct. It could be either. Again, I assume you must mean information when you say "variance." But why do you talk about characterizing "signals in terms of information flow rather than in terms of amplitude and phase"? Using information content as a measure does not justify throwing out such vital and related concepts as amplitude and phase. I don't think even Martin would be that radical!

> >...as long as the process gives off heat,
> >the resulting information channel may or may not have higher
> >entropy than the sum of the original two.
>
> So you appear to agree that we can't predict whether regulation
> will actually occur in the above arrangement on the basis of
> information theory alone.

Yes! We agree! This is exactly the point I've been trying to make.

> Where I predict a problem for an information theorist in trying
> to meet my challenge is in explaining, on IT grounds alone, why
> the amplitude of E becomes less than the amplitude of D.

It is the "IT grounds alone" that bothers me. I can't imagine why you would insist on this condition. You do not require planetary orbits to be predictable on "calculus grounds alone."

> The phenomenon of control depends on the oppositeness of the
> signals, not on their respective information contents.

The "oppositeness" can be seen in the blockage of the channel. This is an information loss - a decrease in entropy that requires work and heat production to accomplish. This is a classic kind of situation for an information theoretic description.

> The two signals could have the same information content by any common
> measure, yet their amplitudes would not, just because of that,
> have to cancel each other quantitatively.

Right. For instance, my apartment sometimes gets into a goshawful mess. Entropy just takes over - clothes and half-eaten pizzas knee-deep. To clean this mess up requires work to decrease the entropy. For that, I need information about the state of the mess. I can collect this information, but that fact alone does not mean that I will "cancel out" the mess. In fact, I could amplify it (mess it up even more). Or I might simply not use the information to do work at all. I have the information required to clean the apartment. The bandwidth is there in the output channel to allow the work to be done. That does not by itself mean I will actually do the work, thereby decreasing the entropy of the apartment and increasing the entropy of the universe. Usually, I just clean it in imagination. Sometimes I go one better, and simply change my basis for the entropy calculation, and suddenly - presto! - my apartment is already sparkling clean without lifting a finger! For example, one half eaten pizza two-thirds of the way from my desk to the sofa isn't "messy" - that's exactly where the pizza goes! (Of course, to change the basis, I still must do work to accomplish the necessary reorganisation of my PCT hierarchy).

> I sort of object to using a physical term like entropy in dealing
> with information in living systems. If you think that
> informational entropy is connected to physical entropy, you can't
> separate the formal entropy in a neural signal from the real
> physical entropy involved.

The very concept of a "real physical entropy" distinct from the "formal entropy" is quite meaningless. There is no reason to separate the physical entropy from the informational entropy. In order to process information, I need to do work and produce heat. Information theory is not an analogy to thermodynamics - it is the same theory. There is no distinction. Information = Physical Entropy. Both are subjective properties of a system requiring a basis (model, encoding scheme, subjective probability distribution) with which to describe them. Macroscopic thermodynamics is a form of information theory using temperature for this basis. But temperature is not an absolute property of a system! There is no inherent reason to consider temperature as an absolute basis for physical entropy. Believe it or not, this is widely recognized in standard textbook thermodynamics, although many physicists no doubt do not fully appreciate it (it is recognized in the Zeroth Law of Thermodynamics).

Quantum Mechanics digression (ignore if you don't care about physics): Interestingly enough, in quantum theory you need this arbitrary assumption just to separate our universe out from other alternate universes in the wave-form! In other words, there is no absolute standard by which our universe can be said to exist as a separate entity from all other possible universes. If you deny the objective existence of the other universes, and you wish to claim such an existence for our own universe, then you must apply an arbitrary subjective standard as an objective property of the universe and you end up in a hopeless paradox. In PCT terms, the subjective standard is the hierarchy. Thus, the PCT hierarchy is what actually defines the universe as a cohesive entity! In a very real sense, our universe is defined by our perception of it. For this reason, PCT may have a great contribution to make to quantum theory.

> I would guess that the entropy increase involved in simply
> transmitting a neural spike from one neuron to the next would be
> hundreds of times greater than the hypothetical decrease involved
> in the actual synaptic event - the transfer of a bit of information.

Yes, so what? The transfer of one bit *requires* an increase in entropy to accomplish the local decrease involved in the transmission of information. There is no requirement that this be anywhere near the theoretical lower limit (work must be done to transmit information, but it need not be done with near-perfect efficiency).

> ...but if they [Shannon and others]
> intended [negentropy] to mean some mysterious connection with
> dQ/T, they had it backward...
> The actual direction of information flow in the
> nervous system is opposite to the direction of energy flow. So
> there can't be any connection between entropy as physicists use
> the term and entropy as it appears in information theory.

Just to be clear, dQ is the heat flow, not the net energy flow. Heat flow is arbitrarily defined (in terms of temperature for macroscopic thermo). In terms of the nervous system, yes, of course entropy is increased by the transmission of information, but it is decreased locally (where it counts, in biological terms). This local decrease in entropy must result in a total increase in entropy in the organism, so the system must input low entropy energy as food. The connection with dQ/T is not "mysterious," but quite well defined. To really understand this subject, you need to put macroscopic thermo. aside and learn the full-blown statistical version. Otherwise, trying to match information theory to these macroscopic, temperature-based measures is just going to cause endless confusions.

Some of the confusion here may be simply the terminology. Negentropy is information content. Something not completely random is considered to have information content. Note that *high* information content in this sense means *fewer* bits to describe it, not more. This is one of the most frustrating confusions in trying to first learn information theory or thermo. It sounds so strange: how can high information content mean fewer bits? The result is a terminological confusion, where some people continually talk about "lots of information" as meaning more bits and thus higher entropy, while others mean fewer bits and lower entropy. My preference, and I think a fairly standard one, is to talk about "information" strictly as the entropy, or number of bits, and "information content" as negentropy, which is nonetheless still measured in bits. Information content is high when the "amount of information," or number of bits, is low! Its like a golf score - lower means more "golf-score content." It sounds contradictory, but its not. It is, however, unbelievably confusing when you are first learning information theory, since the experts tend to forget the distinction (taking it for granted) and they don't spell it out. People often lapse and fail to even make the distinction (this probably includes myself!). The physics and information science communities would do well to work for better standardisation of the terminology. Here's an example of something in your post that I think is nothing more than confusion over these terms:

> >The signal to E has very low entropy if the system is controlling.
>
> Now why would you say that? When the system is not controlling, E
> is varying exactly as D is, and if D is a signal within a certain
> bandwidth, then E is also a signal with that bandwidth. All the
> information in D is being transmitted to E.
> When the system is controlling, the information formerly reaching
> E is now very nearly cancelled; E now contains very little
> information. Would not this loss of negative entropy amount to a
> great INCREASE in entropy?

No. In fact, this is not a "loss of negative entropy," but a gain! When the information is cancelled, this means, in my terminology, that E has experienced "information loss," and thereby has high information content! After all, work is expended and total entropy increased in order to produce this blocking of the channel. We have to distinguish between information loss and decrease in information content. So yes, if the system is controlling, the disturbance will be transmitted to the percept (E), as you say. So since the percept is controlled, D is minimised, and has low entropy, as does the percept E.

The reason this sounds so contradictory is that high and low entropy can both seem intuitively like "lots of information," or "very little information," depending how you look at it. A complete random string of digits and a string of all zeros *both* seem intuitively "lacking in information." The common, everyday use of the word information is somewhat contradictory. Living systems seem to exist somewhere between these two extremes, and there are numerous attempts to define exactly where. This is an exciting and open area of research in information theory - one in which I think PCT has much to contribute.

> >Now surely you will agree that if it were impossible to encode
> >the disturbance D into a number of bits that could be handled
> >by the output channel, then the system could not control.
>
> If that number of bits could not be handled by the output channel
> (by which I presume you mean the path from T to E),...

No, I mean the path from R to T. I really prefer to talk in terms of PCT, as I find it to be a more elegant formulation. T is roughly the CEV and E is the perceptual signal. R is the comparator, of course, and Ashby's C is the reference.

> ...then the only
> the bits that survived passage through the channel would be able
> to disturb E.

Yes, and since these are not enough bits to describe the disturbance to the CEV, and hence E, the system will fail to control. Note that these surviving bits are not the disturbance - they are correcting for the disturbance by negating the representation of the disturbance contained in the percept E.

> All that the regulator R would have to do would be
> to provide bits to T that subtract, bit for bit, from those that
> survive the passage through T, and regulation would be perfect,

Again, you are confusing R->T and T->E. T->E is the input channel, while R->T is the output. R cannot provide these opposing bits because the output channel R->T cannot handle the information load.

> >...you need the perceptual functions to compress
> >the inputs into a single scalar value (for comparison) that
> >STILL RETAINS THE ESSENTIAL CHARACTERISTICS OF THE DISTURBANCE.
>
> No! It is not the disturbance that has to be represented, but the
> state of the controlled variable itself.

If the system is controlling, the percept signal carries few bits. The channel carries more bits only if there is a disturbance. The only way in which the percept is used (ignoring higher levels) is to compute an error which does directly represent the disturbance. So I agree that the percept directly represents the CEV, not the disturbance. But it is the disturbance to the CEV that is relevant to control. So the percept does indeed contain information of "the essential characteristics of the disturbance" - hence the ability to oppose the disturbance in the output. The "essential characteristics" part is important. D is transmitted to E, the percept, only through a transducer that represents the disturbance with a particular encoding scheme (there's that information theory again).

> The output only has to
> act on that variable with sufficient speed to keep the perceptual
> representation of that variable matching the reference signal.
> The disturbance itself (D above) could contain megahertz
> variations; most of those would disappear because E can't respond
> significantly to them. E, however, will still contain frequency
> components that are not represented in the perceptual signal.

Again, we're talking on different wavelengths. I consider E to be the perceptual signal, so obviously its hard to respond to this. E is the thing under control, according to Ashby, so it HAS to be the percept doesn't it? You seem to be placing the thing controlled in the external environment. Ashby places it clearly internal to the control system.

> Those frequency components will be uncontrolled (from the
> viewpoint of an external observer). The perceptual signal itself,
> however, will be controlled.

Of course, the "viewpoint of an external observer" is irrelevant to the organism. It is the percept, E, that counts. Components in the world that E can't respond to obviously are not characteristics essential to control. This is the whole idea of filtering the information

> You have to make up your mind whether you're talking about the
> variable responsible for the disturbances (D), or the effect of
> that variable (variations in E).

You are right - this is an important distinction and gets at the heart of what I'm trying to say. Perhaps we are using the term "disturbance" differently? I would define it as the net effect of things in the world on the CEV. The disturbance is not an absolutely defined entity - it depends on the model you use for your description. When you talk about "disturbing variables" that the control system has no information about, you probably mean "a description of the disturbance in an 'objective' language external to the organism." This is how an external observer is apt to view the disturbance - and its entropy under this basis is very high, call it H(D | observer). Now, the organism views D through a very different lense. The same real-world entities the observer just described are defined within the organism in terms of the hierarchy (i.e. in terms of the CEV). This description is much shorter, and has low entropy, call it H(D | organism). Note that H(D) is meaningless. According to Ashby's Law, the hierarchy must bring the entropy of the disturbance in line with the output capacity. If H(D | observer) is high, then we will tend to call the environment complex. If, in addition, H(D | organism) << H(D | observer), then we say that the organism controls in a complex world.

Note that I never said that the control system gets "direct information" about D - that whole notion is meaningless. But it most definitely DOES get information about D, however you choose to arbitrarily define it. This information is definitely contained in the perceptual signal, or the organism would be unable to control against the disturbance. How can you control against something you have no information about? The very notion is nonsensical.

> >So I ask you: where above did my reasoning go astray?
>
> In assuming that the control system needs information about the
> state of the disturbing variable D, and in assuming that control
> is in terms of the external variable E rather than the perceptual
> representation of it.

Again, E *is* the perceptual representation, and the only information about the "disturbing variable" relevent to control is what survives in E.

> The first is the most important error.
> There can be any number of independent disturbing variables
> acting through environmental links on E

Right, but the "disturbance", from the viewpoint of the organism, can only be the net affect of these forces on the CEV. You can't define disturbance otherwise without getting into an arbitrary mess that is completely irrelevant to control.

> ...The control system needs no information about them, singly or collectively.

How can you say this, when the sole purpose of the control system is to oppose the disturbance? It can't oppose something it has NO information about. It simply cannot.

> >If I can actually have complete knowledge of the disturbance D,

> >it is theoretically possible for me to respond appropriately
> >before it has had an effect on the controlled variable.
>
> What you actually should have said, to be precise, was "If I can
> actually have complete knowledge of the disturbance D, and if I
> can use this knowledge to produce actions having a precisely
> calculable effect on E, and if the actions actually produced
> affect E exactly as calculated so as to cancel the effects due to
> variations in D, then the actions will affect E so as to cancel
> the effect of variations in D."

No. This is a tautology, as you well know. I wasn't trying to say that if such a thing were possible, then it would be possible. I was simply trying to state that it *is* possible. Anything can be made into a tautology in the manner that you just did. If I state A, you reword it into "If A then A," and you have your tautology. This is a common debating tactic, but it holds no water. This is getting silly. I stand by my original wording. Do you deny that compensatory systems can in principle work? Do you deny that I could create a toy world in my computer with a compensatory controller? (The very idea of complex living organisms controlling via such a mechanism is, of course, absurd - but that is not the issue).

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Sun Mar 14, 1993 6:30 pm PST
Subject: Too many Karoly's, Info theory

[From Rick Marken (930314.1800)]

I made the mistake of going to the library and making a copy of Karoly's Annual Review of Psychology (1993) paper on self-regulatory mechanisms. Why didn't I listen to Greg? I could have just gone out and bought an emetic. Here is a representative piece of wisdom:

"Although from a purely engineering (cybernetic) perspective, the standard of correctness is 'physically embodied as a perfectly real reference signal inside the control system' (Powers, 1986), in human self-regulators the rule generation and rule following routines are variable and subject to moderating influences."

What in the world does this mean??? That reference signals are variable in humans? (If so, that's part of PCT and nothing in the Powers quote imples anything different). That program level references (rule generation "routines") are variable" (Still no problem). That the outputs resulting from comparison of perception and reference are variable and opposed to disturbance (moderating influences)? Again no problem. Throughout this article Karoly implies that Powers' approach to "self regulation" only applies to machines. I think the people working in this area (and it now apprently includes a LOT more than Carver and Scheier) are so far gone that learning PCT is not even a possibility. How depressing.

Bill Powers (930314.0900) --

WOW! What a wonderful antidote to the Karoly paper.

>So how DOES information get into the organism from the
>environment? Clearly, only through uncontrolled perceptions.

Brilliant; could this be one sideways contribution of IT to PCT?

>Uncontrolled perceptions may be uncontrolled at many successive
>levels, but eventually they join with controlled perceptions to
>become inputs to a high-level control system. There they act as
>disturbances of the perception to which they contribute, _and
>their information content is therefore rejected by that system_.

This is a fascinating idea; the only "informational" perceptions are uncontrolled perceptions -- which will tend to exist AS INFORMATION at the levels of perception at which they are uncontrolled. I wonder what this means phenomenologically?

Your post makes me realize that I should say that only CONTROLLED sensory input variables are dependent variables; uncontrolled sensory variables (like the moving target in a pursuit tracking task) are still not independent variables; maybe a good name for them is INFORMATIONAl VARIABLES; these are perceptions that are "just there". A perception (like the moving target in the traking task) can be an informational variable at one level (like the transition level; you just see a moving line) and a disturbance (and, thus, non- informational) at another -- like the relationship level where you are trying to control the difference relationship (target-cursor) At this level the information about the disturbance is cancelled out by the control process.

Your discussion reminds me that I am going to have to put some uncontrolled perceptions in the spreadsheet model; right now the model gives the impression that ALL sensory inputs are part of control processes at all levels; that is obviously NOT very realistic. Off to Excel.

Best Rick

PS. The "Bill" in my last signiture was Bill Clinton (I guess I just like Bills, unless they're from Buffalo); if you don't know who the "jerks" are then you are an incurable Republican and it's probably not worth providing information at that level.

Date: Sun Mar 14, 1993 7:32 pm PST
Subject: Re: the challenge

[From Rick Marken (93031.1900)]

Butting into my part of the discussion with Allan Randall:

Allan Randall (930314.1830)--

>Note that I never said that the control system gets "direct
>information" about D - that whole notion is meaningless. But it most
>definitely DOES get information about D, however you choose to
>arbitrarily define it. This information is definitely contained
>in the perceptual signal, or the organism would be unable to
>control against the disturbance.

Yes, you keep asserting this. I was hoping for some EVIDENCE. Can you show me how to FIND the information about the disturbance in the perceptual signal? I described experiments that seem to show pretty conclusively that there IS NO INFORMATION ABOUT THE DISTURBANCE IN THE PERCEPTUAL SIGNAL. Your only evidence that there is information about the disturbance in the perceptual signal is assertions like the following:

>How can you control against something
>you have no information about? The very notion is nonsensical.

The answer, by the way, is "by producing outputs proportional to the difference between perception and reference signal". But PLEASE tell me how I can find the information about the disturbance in the perceptual signal; how can I measure it? Without this, your assertion above is about as convincing to me as the following proof that god exists:

"How can life have a purpose if there is no god? The very notion is nonsensical."

Convinced?

Bill Powers said:

> ...The control system needs no
> information about them [disturbances], singly or collectively.

And you replied:

>How can you say this, when the sole purpose of the control system
>is to oppose the disturbance? It can't oppose something it has NO
>information about. It simply cannot.

Again an assertion; "It simply cannot". But it DOES. You have to show me the information in the perceptual signal that corresponds to the disturbance if you want to convince me that you are not just making religious claims. Saying over and over again that this information MUST be there just won't wash.

I ask you to consider once again the experiments described in chapter 3 of Mind Readings. I am also willing to send you HyperCard versions of these experiments so you can see that the results are what you get ALL THE TIME. When the same disturbance occurs on two different occasions you get almost exactly the same responses (r = .99+ between the response traces) on the two occasions. So there must be SOMETHING in the stimulus traces that is the same on both occasions -- the information about the disturbance is what I presume you would claim it is. This information should be the same on both occasions (even though we don't know what it is or how it's coded). I claim that it should reveal itself as a relatively high correlation between the stimulus traces on the two occasions. In fact, the correlation is very low -- sometimes as low as .0033 -- even though the response correlation is ALWAYS over .99. Perhaps you will argue that there is still information in the stimulus traces about the disturbances, it's just that the correlation won't pick it up. I then want to know WHY and HOW I CAN FIND IT. If IT can't tell me this then what the hell good is it. PCT, by the way, DOES predict the LACK of correlation between stimulus traces.

I have the distinct suspicion that you think the information about the disturbance MUST be in the stimulus because you think the stimulus (or perception thereof) is the cause of outputs that cancel the disturbance. BUT THEY ARE NOT. THAT IS THE VIEW THAT PCT SHOWS TO BE WRONG. But you keep asserting that PCT is wrong about this. Please SHOW ME (don't tell me or assert to me) why.

Thanks   Rick

Date: Mon Mar 15, 1993 7:54 am PST
Subject: sudden path / gradual path

[From: Bruce Nevin (Mon 930315 09:45:10)] Rick Marken (930313.1000)) --

>I know that you (Bill P.) are looking for ways to take what people can
>already accept and move them gradually and gently to the PCT perspective.
>That's fine with me; but I don't think one can make a gradual shift into
>PCT (any more than one can make a gradual shift from an earth to a sun
>centered view of the solar system).

Sounds like you two have a good cop - bad cop routine going.

I think of the analogy to a gestalt shift. The two images, damsel and crone, are discontinuous and the shift is abrupt and radical. However, teaching someone who can only see one to recognize the other in the same assemblage of visual shapes is a gradual process. The jawline of the damsel's averted face constitutes the nostril of the crone, her feathered hat the eye, etc. Then the abrupt "aha!"

The teacher with the gradual approach identifies the perceptions that the person distinguishes as constituting (higher-level perceptions of) various features in the old construal, and then shows what those same perceptions constitute in the new construal, until the person recognizes (aha!) the coherence of those reinterpretations in a new construal.

Coherence in a whole must be a controlled perception, right? All these perceptions constitute (higher-level perceptions of) features in a coherent whole, the befurred damsel with the averted face. Reinterpretation of individual features is resisted as a disturbance to the perception of a coherent whole. That rounded curve of white can't be a nostril, there's no way a nostril in that scale could appear there. You're nuts!

In the PCT/sensory-input-as-independent-variable confrontation, we don't have two equally valid construals. One is wrong. The way conventional psychology can get away with being wrong is presumably by supplying imagined perceptions to satisfy the higher-level perception of coherence in a whole. (Also by satisfying one's perception of establishing and furthering one's career, but that is a matter of each participating individual buying into social norms, that is, instituting social conventions in the values at which they set their own internal reference perceptions.)

An example of an imagined perception is the notion that sensory input is an independent variable. A perception that is imagined for the sake of reducing error in a higher-level perception of coherence (and error in the perception of establishing and furthering one's career, a perception that is presumably controlled with high gain) will be protected from counterevidence. Stones fall from the sky? Impossible! These reports of meteorites are popular delusions of simple country folk being exploited by publicity seekers.

What are good strategies when a controlled perception is mistaken, supported by misconstruals (mistaken categorizations?) and imagined perceptions on lower levels, and when the perceiver regards disturbance of the higher-level controlled perception as life threatening or almost so? I think that we can't say that imagined perceptions are controlled, so they can't be disturbed. Getting the other person to attend to actual sensory input that differs from the imagined perceptual signals supplied from on high results in something different, seems to me, from either disturbance or conflict. What is it, and what do we do about it?

Bruce Nevin   bn@bbn.com

Date: Mon Mar 15, 1993 8:29 am PST
Subject: Re: Niels and Albert

[Martin Taylor 930315 11:00]   Rick Marken 939314 10:00

None so blind as those who will not see, eh, Rick?

Trouble is, none of us know who that refers to. I think all of us (I mean you, me, Allan, Bill P.) understand and believe:

> a "non-prosaic" phenomenon (sensory input as DEPENDENT rather than INDEPENDENT
>variable) actually occurs and that it is predicted by the "prosaic" PCT model.

The comparison with Niels Bohr and Einstein is different from that between Newton and Einstein. (And I would prefer Heisenberg or Schroedinger to Bohr in the analogy). Both sides of that argument were putting something radical and initially non-intuitive in place of a long established doctrine. Both were right in their field of main concern, but there was (and I think still is) conflict and uncertainty about how to fix it at the border. Einstein based his concern about quantum physics on a notion of beauty. He had been right to use that criterion in developing relativity theory. Had he lived to see modern quantum electrodynamics, he might well have thought it beautiful, too. We will never know.

One doctrine I learned in graduate school and have not forsworn is that when two schools of thought contend, each claiming the other is dead wrong, they probably are both right except in that claim. It just takes seeing the problem in a new way. In our case, I don't think you are wrong, except in saying that I am. I'll leave Allan to fend for himself, since I can't see into his mind.

By the way, I'm not ignoring Bill's challenge, as Bill knows.

Martin

Date: Mon Mar 15, 1993 10:01 am PST
Subject: PCT, joy of man's desiring

[From Rick Marken (930315.0900)]   Bruce Nevin (Mon 930315 09:45:10) --

>Sounds like you two have a good cop - bad cop routine going.

Actually, this was not intentional; Bill P. is just a truly wonderful person and I am a schmuk.

Your description of learning to see the other figure in the "wife-cron" drawing was wonderful.

>What are good strategies when a controlled perception is
>mistaken, supported by misconstruals (mistaken categorizations?)
>and imagined perceptions on lower levels, and when the perceiver
>regards disturbance of the higher-level controlled perception as
>life threatening or almost so?

I don't think there are any. You might be able to drive a person into reorganization but the result of that process is unpredictable (you're as likely to end up with a muslim fundamentalist as a PCTer). I like Bill P.'s approach using poise and patience. I also like my approach because I'm a passionate fellow (sometimes). But I don't think it makes a twig of difference in terms of convincing people. Matter of fact, I'm SURE it doesn't make a twig of difference. I have watched at meetings as Bill has patiently taken a bright, curious and apparently very interested person through the demos and models -- the person always eagerly questioning, challenging, being astounded and, finally, being completely convinced-- and then we never hear from the person again (and when we do, it's in the pages of some journal explaining their coordinative structure model of behavior or something like that). There is no question that people will only learn PCT if they want to and (most important) if they are (for whatever reason) willing to accept that they don't have most of the answers already . This kind of person is a very rare bird -- especially in the halls of science.

Martin Taylor (930315 11:00)

>One doctrine I learned in graduate school and have not forsworn is
>that when two schools of thought contend, each claiming the other is
>dead wrong, they probably are both right except in that claim. It
>just takes seeing the problem in a new way.

When the schools of thought are the different schools of psychological thought that are taught in graduate schools to this day then that doctrine is probably correct.

>In our case, I don't think you are wrong, except in saying that I am.

If you said that there is information about the disturbance in the stimulus (or perception) then you are wrong. But I don't remember you making that claim. What did I say you were wrong about?

When it comes to the "information about the disturbance in the stimulus" claim, then I'm afraid we can't both be right about this. If you are right then it means that the stimulus in a control task has information that guides responses (so that they are opposed to the disturbance) -- making PCT consistent with conventional models of behavior. If I am right, then the stimulus in a control task cannot POSSIBLY guide responses -- and the whole process must be viewed (SERIOUSLY) in a new way -- as response guidance of the stimulus (actually, of the perception thereof). This fact about control is TOTALLY inconsistent with conventional approaches to behavior analysis and modelling. I think that there can be no compromise on this. If you try to compromise you end up producing the kind of mish-mush that you find in all the articles described in the Karoly paper. I don't want this do happen to you; save yourself while there is still time; understand that there IS NO INFORMATION ABOUT THE DISTURBANCE IN THE STIMULUS in control tasks and you will be able to make beutiful PCT music -- not an ugly cacophony like that produced by the Karolyans.

Best

Rick (canon on the loose again) Marken

Date: Mon Mar 15, 1993 10:57 am PST
Subject: USERS & ENGINEERS - RKC

[From Bob Clark (930315.13:15 EST)]

USER'S VIEWPOINT Rick Marken (930307.1500)
You refer to my remarks about error, illustrated with a two level system. My concern there was to point out the necessity for more than one level if "error" is to be perceived.

My view of the "User" is stated more completely in my previous post, Bob Clark (930301.1730). This is the view of the entity "riding around within the organism." This is like the view of the driver of a car, or pilot of an airplane. That is, he has little interest in the internal operations of the machine -- as long as it does what he wants. He "controls" it by using the levers, wheels, etc available. To him, it is pretty much a "S-R" system. "Push" here -- "something moves" over there. If the "movement" is not as desired, a different "push" is tried. When the User is familiar with the machine's operation, he selects "Pushes" from his remembered repertoire. Otherwise he uses a "trial and error," "experimental" method. Of course if he has more complete information about the workings of the machine (perhaps a remembered theoretical analysis), he can work out solutions to unusual situations with less experimentation.

To me, this is an "Internal Viewpoint." This is what is available from the "Inside." Note that this includes operation of the User's DME as it reviews available decisions from the User's Memory.

Your discussion of "spilled wine" illustrates the difference between the Viewpoints of the User and Observer -- as I define them. The User might have intended the "spill" -- you'd have to ask him. The Observer would be applying his own view of the "spilled wine event." I think this is consistent with the view indicated in your remarks:

*But we know that very often a perception that is an error to one person
*is not an error at all to another. The spilled wine, for example, might
*be just what someone wanted to see -- that's why they knocked over the glass.

The User is using his systems to bring his perceptions into agreement with the conditions that he selected.

OBSERVER'S VIEWPOINT
As you pointed out: all the Observer perceives is the spilled wine. (And the actions of the Bystanders and the Spiller himself.) The Observer is using his

systems to observe and remember (at least briefly) what is happening in the region to which he is paying attention.

*If the perception of spilled wine seems like an error to an observer, it
*must be because that perception deviates from some specification IN THE
*OBSERVER of what should be perceived. This seemingly obvious fact about
*perception is completely missed ....

Of course!! but people frequently get their "viewpoints" confused.

ENGINEER'S VIEWPOINT
The Engineer is using his knowledge and skills to make the system he is "building" (whether using hardware, software, or whatever) resemble the "human" systems he is interest in. The Engineer doesn't care whether or not the spilled wine is an "error." He observes that people make mistakes and he is interested in building a multi-level control system that can "perceive error" and correct it. An "Error" perceived by the system could be a "Disturbance" originating "Outside" the System. An "Error" could also originate inside the System in the form of a "Conflict." This form of "Error" perceived by the system could result from incompatible requirements derived from independent sets of "Goals."

Perhaps the following will clarify my earlier post, Bob Clark (930301.1730).

The Engineer's goal seems to be the construction, at least in principle, of an assembly of hardware (or equivalent computer-cum-software) that performs the same way that a human (or, perhaps, a simpler organism) does.

Some Engineers approach this in terms of levers, gears, pulleys, etc arranged so that inputs ("disturbances"?) at certain locations result in movements at other locations. By adding suitable "leading" terms (time derivatives) and "lagging" terms (time integrals) these systems can be quite effective for specified applications.

The PCT Engineer, if that is a suitable term, bases his design on the properties of negative feedback control systems. These are combined into a hierarchical structure, HPCT, assembled and modified to operate according to his desires. The Engineer proceeds by selecting from his inventory of memories, including physical and other principles, in order to bring his proposed structure into correspondence with his view of human behavior.

The design might include "recording and playback" capability as well as ability to "Reorganize" itself. In principle, these are both included in HPCT.

To the Engineer an "Error" is a "Mistake" in his design. This "Error" is revealed by an inconsistency between the performance of his "Engineered System" and the performance of the system he is trying to imitate.

Regards, Bob Clark

Date: Mon Mar 15, 1993 12:17 pm PST
Subject: Challenge: info theory & PCT

[Bill Powers (930315.0700)] Allan Randall (930314.1830) --

This post goes on for nearly a little more than 6 pages, and I have already cut out large parts of it. I don't seem to have had enough to do today. Just a warning to those who would like to use the delete key and make some space.

>To my knowledge, calculus actually CANNOT make the kind of
>predictions you are asking of information theory. Calculus can
>be used to make predictions only in combination with a physical
>model such as Newtonian mechanics.

What are physical models but mathematical forms, manipulated according to mathematical rules, that model or idealize observations? Given that each element of mass attracts each other element with a force proportional to (exactly) the product of the masses and (exactly) the inverse square of the distance between them, and given expressions for the conservation of potential + kinetic energy, one can apply the calculus and derive the fact that orbits are conic sections. That is the kind of prediction I am asking of information theory: a prediction of how the system will actually behave through time.

>> If I were trying to model behavior that takes place under
>> difficult conditions, this analysis might offer more of
>>interest by way of predicting limits of performance.

>You seem to be admitting here that information theory might
>have something useful to say about control systems.

Insofar as information theory could predict the limits of performance given signals and signal-handling devices with certain characteristics and in a known organization, sure.

>So I guess I still don't understand exactly where you stand. Is
>information theory completely wrong-headed or is it correct,
>but of little use to PCT?

Information theory rests on definitions and mathematical manipulations. Unless someone has made an undetected mathematical blunder, the calculations of information theory follow correctly from the premises. It's unlikely to be "incorrect" in those terms. The problems I see come not in the internal consistency of IT, but in its applications to observations. Premises can be wrong; when they are wrong, no amount of mathematical correctness will make the conclusions right. I don't yet see how IT is actually linked in any rigorous way to specific physical situations.

RE: statement of the challenge.

>Using the three conditions (prior to the experiment being
>performed) make some assumptions about the bandwidths of the
>system, and compute in information theory terms the amount of
>control required for the three conditions.

The challenge was a response to the assertion that PCT could be derived from information theory. The prediction I'm asking for is not how much control is required, but how much control there will be in the two situations. To use a theory to derive the fact that control will result from either arrangement means to make predictions by manipulations that follow the rules of the theory.

I'm going to skip a lot that I wrote here, because prediction is the real point.

> ... from Ashby's diagrams + information theory, one cannot
>predict what exactly R, the regulator, is doing. You cannot
>predict that R is going to oppose the disturbance. Whether this
>will meet your requirements for the challenge is the main point
>I'd like clarified before accepting.

If you stick with these conclusions, the challenge is unnecessary because you have agreed to my original claim. You are agreeing that information theory can't provide the predictions of behavior that control theory provides, but can only be applied once those predictions are known and verified.

>The very concept of a "real physical entropy" distinct from the
>"formal entropy" is quite meaningless. There is no reason to
>separate the physical entropy from the informational entropy.

You're just reasserting the claim that they are the same. This claim is based on nothing more (I claim) than a similarity in mathematical forms. In all the information-theoretic stuff I have seen, the assumption is that information and energy flow the same way. I was showing that energy does not go in the direction that is commonly assumed in the nervous system. But in addition to that, it's possible to show that messages can be sent with energy flowing EITHER way (for instance, sending a message to the guy operating a winch by intermittently putting a frictional drag on the cable). If information can flow one way with energy going either way, then the "entropy" involved in information flow is not physical entropy. Physical entropy is always transferred opposite to energy flow.

>When the information is cancelled, this means, in my
>terminology, that E has experienced "information loss," and
>thereby has high information content!

Are you saying that the system begins with a certain information content, and after losing information it has a greater information content? I don't follow your arithmetic.

>The transfer of one bit *requires* an increase in entropy to
>accomplish the local decrease involved in the transmission of information.

This seems to imply that at the terminus of the transmission, there is a local decrease in entropy. I would like to know where that occurs. I can see no point in the transmission of a neural impulse where there is anything but a loss of energy (and an increase in physical entropy) in the next place that the impulse reaches. Where is this final place where we suddenly get the opposite effect (taking energy FROM the arriving signal?).

This usage of "requires" is odd: I think you have the logical implication backward. Transmission of a bit PREDICTS a change of entropy, but it does not follow that a change of entropy PREDICTS transmission of a bit. Entropy is a not a cause, but an effect. The logical implication is "It is not the case that a bit is transferred and entropy does not change."

>There is no requirement that this be anywhere
>near the theoretical lower limit (work must be done to transmit
>information, but it need not be done with near-perfect efficiency).

But aren't you assuming that this work is done by the transmitter on the receiver? I can think of numerous methods to transmit a message in a way that requires the receiver to do work on the transmitter. For example, I can detect the presence of a control system by pushing and seeing if there is resistance to the push. I have to do work on the system to receive that message, and the message can be extracted from the amount of work I do. So by resisting and not resisting, you can send me 1's and 0's. 1 means I'm doing some work on you and you are controlling; 0 means I'm not and you're not.

>If the system is controlling, the percept signal carries
>few bits. The channel carries more bits only if there is a disturbance.

The latter is true only for homeostasis. Changing the reference signal to a new value will also cause more bits to be carried, for a moment, by the percept signal. Even without a disturbance, the system will then reduce those bits as closely as possible to zero. The action is based on the state of the percept, not the state of the disturbance. And the state of the percept does not depend on the state of the disturbance alone. This means that it does not depends on the disturbance in a known way AT ALL.

As Rick Marken has been trying to say, if you know only the sum of two numbers, you know nothing about either of the numbers by itself. The percept represents the sum of effects from the system's own output and from the disturbance. A percept of 3 units could represent the sum of 6 and - 3, 6000 and -5997, -200 and 203, and so forth. There is NO information in the percept about the disturbance, in a control system.

>E is the thing under control, according to Ashby, so it HAS to
>be the percept doesn't it? You seem to be placing the thing
>controlled in the external environment. Ashby places it clearly
>internal to the control system.

It's clearly inside the control system according to PCT. I doubt that Ashby saw it that way, as he made the "regulator" into a different unit. I was taking the external view in treating E as the environmental controlled quantity, with the sensor and percept being internal details of the regulator. But your interpretation is just as good.

>Perhaps we are using the term "disturbance" differently? I
>would define it as the net effect of things in the world on the CEV.

I have spoken of this ambiguity before. I use the term "disturbing variable" to mean the independent CAUSE of the disturbance. If control is tight, a 100-unit change in the disturbing variable might produce only a 1-unit disturbance of the controlled variable, where without the control it would produce a 100-unit disturbance. Just speaking of "the disturbance" is ambiguous -- do you mean the applied force, or the movement that results from it?

You can unambiguously specify the disturbing variable without considering what the system it affects is doing. I can specify that I am pushing on your arm with a force of 20 pounds, or a slowly-changing sine-wave force with an amplitude of 20 pounds and a certain frequency. By itself, this doesn't say anything about how your arm will move; that depends on what muscle forces are also acting at the same time and how they are changing. You can pull back on your end of the rubber bands, but you can't say how the knot will move as a result.

When multiple disturbing variables exist, you can reduce them to a single equivalent disturbing variable acting though some equivalent path on a one-dimensional controlled variable. But you can't say what changes in the controlled variable will actually occur without knowing how the output of the system is acting. That output also affects the controlled variable; it can cancel most of the effect that the disturbing variable would have had in the absence of control. You can't arbitrarily specify the actual disturbance of the controlled variable; that depends on the properties of the system being disturbed. All you can specify arbitrarily is the state of the disturbing variable.

When you model a control system, you MUST apply modelled disturbances via a disturbing variable. If you simply assume a given change in the controlled variable, you're breaking the loop: you're saying that the output makes no contribution to the state of the controlled variable. The amount of change in the controlled variable is one of the effects to be calculated, not an independent variable.

>Note that I never said that the control system gets "direct
>information" about D - that whole notion is meaningless. But it
>most definitely DOES get information about D, however you
>choose to arbitrarily define it.

This is not true. If I send you a Morse Code message and John simultaneously sends you a Morse Code message, and you receive only the OR of these two messages, how much information do you have about either John's message or mine? None at all. If the receiver gets only the OR of the messages, it has no way to sort out which dot or dash should be attributed to me or to John. The only information it can get is about the resulting combined message -- which in fact will be semantic gibberish.

A control system's input function receives information only about the controlled variable. It can't tell how much of the amplitude of that variable is due to an independent disturbance and how much is due to its own output function. It experiences only the sum.

>This information is definitely contained in the perceptual
>signal, or the organism would be unable to control against the disturbance.

You're toying with the same paradox that got Ashby. Suppose the disturbance is transmitting 100 bits per second to the controlled variable. According to the Law of Requisite Information, the output must also transmit 100 bits per second to the controlled variable if perfect control is to be achieved. This is clearly impossible, because then there would be zero bits per second coming from the controlled variable into the system's output function, while the output function is producing 100 bits per second of information. So what level of control would be possible? Suppose that the output function transmitted only 50 bits per second, the amount required by Law to "block" 50 bits per second from the disturbance. That leaves 50 bits per second unblocked, reaching the controlled variable, which is just sufficient to cause a perfect output function to produce the 50 bits per second assumed. On this basis you would predict that the error-driven control system could reduce the information flow from the disturbing variable to the controlled variable by at most one half.

At the same time, a compensating regulator could be perfect: 100 bits per second could pass from the disturbancing variable to the controlled variable and also to the regulator. A perfect regulator would pass the whole 100 bits per second to its output, which according to the Law is just sufficient to block the 100 bits from the disturbance entirely. So no bits would reach the controlled variable.

By this reasoning, the disturbance-based system has a wide margin of performance over the error-based system even with imperfect signal transmissions.

I claim that the experiment will show that this is not true. If it does, something is wrong with the Law of Requisite Information.

>> ...The control system needs no information about them
>>[multiple disturbances], singly or collectively.

>How can you say this, when the sole purpose of the control
>system is to oppose the disturbance? It can't oppose something
>it has NO information about. It simply cannot.

When you get this figured out, you can finally claim to understand PCT.

RE: tautology in defining compensating control

>No. This is a tautology, as you well know. I wasn't trying to
>say that if such a thing were possible, then it would be
>possible. I was simply trying to state that it *is* possible.

That's what I said. You asserted that it is possible, but you made your assertion sound like a deduction because you didn't fill in all the premises that the deduction would require.

You did present the statements as a (partial) deductive argument (930310):

>If I can actually have complete knowledge of the disturbance D,
>it is theoretically possible for me to respond appropriately
>before it has had an effect on the controlled variable.

I was pointing out that the single "if" you supplied was insufficient; you might have complete knowledge of D but be unable to calculate or produce an output of the exactly-required amount, or you might respond a little early or a little late, and so forth. To make your conclusion true (I could respond appropriately) you must supply all the required premises, among which are those that define what "appropriately" means.

If you had filled in all the premises, you would have found that your assertion had to be one of them. So you were simply making a groundless assertion, and the rest is window-dressing.

>Anything can be made into a tautology in the manner that you
>just did. If I state A, you reword it into "If A then A," and
>you have your tautology. This is a common debating tactic, but
>it holds no water.

It's a common debating tactic, but it's by no means possible to make ANY statement into a tautology:

If it rains tomorrow, my car will get wet. It will rain tomorrow; I conclude that my car will get wet.

It is not tautological to say that my car will get wet, because that statement does not depend on assuming that it will get wet. It depends on assuming a fact: that it will rain tomorrow. So the conclusion becomes testable; the car might not get wet tomorrow.

When you say that a perfect compensator compensates perfectly, you are just asserting the same statement twice. The premises on which a perfect compensator depends contain exactly, and only, the assumption of perfect compensation. The only way to make the argument non-tautological is to introduce factual premises: it is possible for a system to (fill in the requirements on each part). If those premises should hold true, then perfect compensation is possible. But perfect compensation may not be possible, because the assumed premises may not be factual. That's the only way to say something meaningful about perfect compensation.

A tautology is simply an argument that looks like a deductive argument but contains no possibility of being false. I agree that people often present such arguments, trying to make them seem to lead to a necessary conclusion and to disguise the hidden assertion of the conclusion. That man must have been guilty of something; the police arrested him, didn't they?.

Best, Bill P.

Date: Mon Mar 15, 1993 1:35 pm PST
Subject: Linguistics facts

[From Bill Powers (930315.1300)] Bruce Nevin (930311.1520) --

What's underneath our dispute here is a basic difference of interests.

You're an expert on how people usually or normally use language; I shouldn't even try to argue on that level because it's too easy for you to show that I'm ignorant. I'm interested in building a model of how the brain works, filling it in with as many facts from observed behavior as possible. You would think that I would welcome getting information about how people actually behave, but from my standpoint it's not easy to get the kind of information I need. I can't build a brain model that leads to 80% of the people doing one thing and 20% doing something quite different. A model does only one thing: what its design says it must do. This means that the model can only make one prediction at a time. If it predicts wrong an unpredictable 20% of the time, then it's just wrong.

The kind of data I need concern what people ALWAYS do. I don't think this kind of data is common in linquistics or any other conventional science of life.

When I say "always," I make reasonable allowances. But I do mean that exceptions should be very rare. I also mean that I expect the author of the data to have tried to discount it before publishing it -- why should I have to try to think up ways it could have been contaminated? Good data has already survived that process in the hands of the person who obtained it, the person most likely to think of things that are wrong with it.

Your report on Labov's data is, on the surface, the sort I like to get.

>In the lower-class department store, he uniformly got one
>pronunciation, no exceptions; in the upper-middle class
>department store he uniformly got the other pronunciation, no exceptions.

But the more I thought about that, the more questions I had. Perhaps Labov answered them and you just didn't get into that much detail.

For example, it seems to be an astonishing run of luck that NOBODY encountered in the "lower-class" department store was from out of town. Were there no commuters from New Jersey or the suburbs? From Brooklyn, or Queens, or White Plains, or Long Guyland? Were these people interviewed afterward to see what they actually had in common other than their pronunciation and an assumed (by association) membership in the "lower class?"

Another problem is in the upper-middle class store. Did Labov make sure that the salespeople he interviewed had not been systematically selected for speaking in a certain way, and specifically not in another? An interview with the personnel manager might have revealed that manner of speech was an important hiring criterion, so that Vassar graduates who spoke with the wrong accent would be kindly encouraged to seek employment elsewhere, or in a back room where they didn't meet the public -- or Labov. And followup interviews with the interviewees might have shown that some of them also worked in lower-class stores where they spoke differently, or that they lived in the same place where employees of the lower-class stores lived, or that they had upper-class accents for totally unexpected reasons.

An even more disturbing aspect of the data is the fact that Labov knew, when he applied his fine-tuned linguist's ear to his memory of the interview, where he was. He knew he was in what he categorized as a lower-class store, or an upper-middle class store. He could see the face and the clothing and the body language of the respondent. And he did his own evaluations.

In a case like this, where a highly subjective judgment is to be made, a common procedure is to have one person record the interviews and to have the expert do the evaluations in a different setting, without knowing where the recordings were obtained and without being influenced by the appearance of the interviewee. This eliminates at least any tendency to hear what one expects to hear on the basis of the surroundings and with knowledge of the theses one is trying to prove. It would be even better to rely on an expert who does not know or care whether the expectation is for the "class" of the store to have an effect on pronunciation or to have no effect.

I don't know what precautions Labov took or how deeply he looked into the origins and other characteristics of his interviewees. I don't know how much skeptical investigation he did, looking for selection factors that would change the meaning of the data. I don't know what he expected to find when he did this experiment. I don't know what he or other linguists would have made of the pronunciations if they had encountered a clerk from one store temporarily transplanted to the other store as a test for the effects of expectations.

If Labov discusses all these things and shows how he circumvented the implied problems, then I would definitely consider his data to be of interest. Otherwise it's just another guy trying to demonstrate that he's right -- or it might be. I wouldn't trust the data until I knew more about how it was obtained.

When I get into arguments with psychologists, I'm often bombarded with facts. "How do you explain ...?" they ask and cite some finding about human behavior. Psychologists are loaded with facts, which they have put together in their heads into some sort of coherent picture of human behavior (different pictures, of course, depending on the school). I get the same thing from linguists and practically everyone else associated with human behavior.

If I had the authority, I would ask all these people to gather up all the facts they have accumulated, trace them back to their sources, and put the papers describing the studies into three piles:

Pile 1: Studies in which, of the total tested population, at least 20% of the subjects failed to show the main effect or showed an opposite effect.

Pile 2: Studies in which between 5% and 20% of the subjects failed to behave according to hypothesis.

Pile 3: Studies in which 95% or more of the subjects behaved according to hypothesis.

I would pick my candidates for useful facts from Pile 3. The studies from Pile 2 I would send back to their authors with a request to try to find out why the deviant persons didn't behave as the rest did or to change their hypothesis, and to do another study to try to make it into Pile 3. Pile 1 would go in the wastebasket.

This would leave me with a small number of candidates for real facts. Those, if interesting, would be close enough to facts to justify putting in some labor trying to get rid of that last 5% of the deviants, probably by redesigning the experiment and the hypothesis and repeating the study a number of times. When the number of deviants was down to, say, 1% (or whatever the level where I got tired of being a good scientist), I would put the result in Pile 4: facts usable in a scientific model of behavior.

Only then would I put any serious effort into trying to find a model that could reproduce the reported behavior. Why go to the great labor of developing a model that will accurately predict the average behavior of a population in which you KNOW that 5% or 20% or 50% of the predictions made by this model will be wrong? I could build a model right now that would fit the prediction that "Mothers hold their babies on the left." All I would have to do would be to program it so it would never use its right arm to hold a baby. Then the model would perfectly fit the observed "fact" as remembered by the psychologist. Unfortunately, it would not fit the behavior of 20% or so of the mothers.

When you tell me linguistic facts, it would be helpful to my attitude if I knew which pile they were in.

Best, Bill P.

Date: Mon Mar 15, 1993 1:56 pm PST
Subject: Info about the disturbance (rated R)

[From Rick Marken (930315.1300)] Allan Randall (930314.1830) --

>How can you say this, when the sole purpose of the control
>system is to oppose the disturbance? It can't oppose something
>it has NO information about. It simply cannot.

Bill Powers (930315.0700) --

>When you get this figured out, you can finally claim to understand PCT.

And that is why I'm pushing and pushing on this point; whether or not Allen gets it figured out I think it is still valuable to keep this apparently picky little matter (whether there is information about the disturbance in the perceptual signal of a control system) right up there in from of CSG-L because this is the razor's edge -- where you either cross the threshold and "get it" (see what is going on here in PCTland -- output guides input, NOT vice versa)) or get spooked and return to the comfortably false assumptions of conventional behavioral science -- input guides output (perhaps retaining some of the language of PCT, like the Karolyans -- I like it, a new name for the false prophets of PCT).

For the benefit of those of you with children, the "information about the disturbance" discussion could begin a naked and violent process of REORGANIZATION in some of the participants-- it is therefore rated R.

Best Rick

Date: Mon Mar 15, 1993 2:20 pm PST
Subject: Re: Feynman

[Allan Randall (930315.1655 EST)] Gary Cziko (930314.0410 GMT) writes:

> I think it was Richard Feyneman who said that if you kept thinking about
> the absurdities of quantum mechanics (instead of just using it as a tool
> for work in particle physics) you would just "go down the tubes" and be
> through as a physicist. Maybe we have here something just as absurd but
> just as useful and "true" as the laws of particle physics. If physics
> survived the apparent absurdity of quantum mechanics, maybe there is hope
> that psychology can also survive the apparent absurdity of perceptual control.

This is kinda off the subject, but it hits a personal nerve. I am one of those people who disagree strongly with Feynman's "head in the sand" approach to quantum physics. I would rather risk going "down the tubes" in trying to understand something than just saying, "it's ununderstandable, so let's all just give up." If quantum mechanics seems absurd, you probably don't understand it yet. If trying again and again and again still sheds no light, then you may die not understanding it, but at least you died trying. Feynman actually had a very deep grasp of the theory (not surprizingly), but there was a point past which he was simply not willing to go. He seemed happy with a theory that was ultimately absurd.

> Hm, I suppose this casts Bill Powers or Rick Marken as Niels Bohr and
> Martin Taylor or Allan Randall as Albert Einstein. Will Albert find the
> disturbance information hidden somewhere in the perceptual signal? Or will
> Niels's absurd magic triumph. I'm staying tuned!--Gary

I have no trouble with this characterization. My personal opinion is that, within their understanding of the subject, Albert Einstein won the debate with Niels Bohr hands-down. I think we have insights today that they did not have, and I would not actually agree with Einstein - but his philosophical objections to Bohr were 100% on-target in my view. So yes, Bohr was doing what Marken seems to be doing - he was invoking magic to give his theory explanatory power. Bohr's view was inherently contradictory.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Mon Mar 15, 1993 3:28 pm PST
Subject: Re: PCT, joy of man's desiring

[Martin Taylor 930315 18:00] Rick Marken 930315.0900

>understand that there IS NO INFORMATION ABOUT THE DISTURBANCE IN THE STIMULUS
>in control tasks and you will be able to make beutiful PCT music

I understand that when the gain is infinite and there is zero transport lag around the loop, this is true. I understand that the gain cannot usefully be infinite if there is any noise in the perception, which means that the perception can correspond exactly and immediately to the real-world construct we call the CEV. I understand that such conditions do not exist in the real world.

I still hope to make beautiful, but less ethereal, PCT music.

Martin

Date: Mon Mar 15, 1993 3:36 pm PST
Subject: Do you believe in magic?

[From Rick Marken (930315.1500)] Allan Randall (930315.1655 EST) --

> So yes, Bohr was doing what Marken seems to be doing - he was invoking magic
> to give his theory explanatory power. Bohr's view was inherently
> contradictory.

I am not invoking magic AT ALL. I am invoking chapter (and verse) from "Mind Readings" (chapter 3, all verses)-- which describe what I consider to be experimental proof that THERE IS NO INFORMATION ABOUT THE DISTURBANCE IN THE PERCEPTUAL SIGNAL OF A CONTROL SYSTEM. This means that THE PERCEPTUAL INPUT TO A CONTROL SYSTEM CANNOT BE WHAT CAUSES THE OUTPUT OF THE SYSTEMN TO MIRROR THE DISTURBANCE. If you think that what I am describing is magic, then you must reveal the trick; that means, you must show me WHERE (IN THE INPUT TO A CONTROL SYSTEM) IS THE INFORMATION ABOUT THE DISTURBANCE? I believe this will be VERY difficult to do because the input to a control system is the net result of the simultaneous effects of disturbance(s) AND the output of the system itself.

Best Rick (Merlin) Marken

Date: Mon Mar 15, 1993 3:40 pm PST

Subject: Re: 'tis-so-'taint-so

[Avery Andrews 930316.0939] Rick Marken (930313.1000)

>There are those on this net who don't like referring to PCT as
>"revolutionary". But Allen Randall has put his finger right on the
>revolutionary button of PCT. If there is ONE thing on which ALL
>conventional behavioral scientists agree it is that sensory input
>is an INDEPENDENT VARIABLE

Well, Chomskyan linguists are different - they are quite uninterested in dependent vs. independent variables, and don't learn to analyse things this way in graduate school- they just ask `what's the pattern, and why'. I suspect that many non-psychologist `cognitive scientists' are similarly unschooled.

Avery.Andrews@anu.edu.au

Date: Mon Mar 15, 1993 3:48 pm PST
Subject: Still Going

[Rick Marken (930315.1515)] Martin Taylor (930315 18:00) --

>>understand that there IS NO INFORMATION
>>ABOUT THE DISTURBANCE IN THE STIMULUS in control tasks and you will
>>be able to make beutiful PCT music

>I understand that when the gain is infinite and there is zero transport
>lag around the loop, this is true.

Irrelevant. It is always true -- unless the control system is not acting at all (zero gain) in which case there is complete information about the disturbance at the input to the control system. Of course, now none of this information shows up in the output since the output is zero all the time.

>I still hope to make beautiful, but less ethereal, PCT music.

Your hope will be achieved when you can say -- confidently and with no reservations:

THERE IS NO INFORMATION ABOUT THE DISTURBANCE IN THE INPUT TO A CONTROL SYSTEM.

Amen. Rick

Date: Mon Mar 15, 1993 4:09 pm PST
Subject: Behavioral Scientists

[From Rick Marken (930315.1545)] Avery Andrews (930316.0939) --

>>There are those on this net who don't like referring to PCT as
>>"revolutionary". But Allen Randall has put his finger right on the
>>revolutionary button of PCT. If there is ONE thing on which ALL
>>conventional behavioral scientists agree it is that sensory input
>>is an INDEPENDENT VARIABLE

>Well, Chomskyan linguists are different - they are quite uninterested
>in dependent vs. independent variables, and don't learn to analyse things
>this way in graduate school- they just ask `what's the pattern,
>and why'. I suspect that many non-psychologist `cognitive scientists'
>are similarly unschooled.

I think many people would say that they don't believe in the dependent variable-independent variable "model" of behavior. This is one of the most irritating problems we have to deal with in PCT. If you ask people whether the IV-DV model would characterize their underlying assumptions about how behavior works they almost always say NO. But when they go out and do some research, what do they use -- the IV-DV framework. Maybe "real" Chomskyan linguists eschew interest in independent and dependent variables -- but when you put them in the lab (if you can get them there) what do they do? They vary an independent variable (type of sentence, for example) and measure a dependent variable (report of where they hear a click that was inserted in the sentence).

Chomskyans can probably rise above the claim that they assume that sensory input causes response output simply because they don't do any research -- at least of the experimental variaty. The same is true of "cognitive scientists" who just build AI systems. So how about this: my claim only applies to "behavioral scientists" who actually go into a lab and study behavior -- and do MORE than just observe it; I mean the people who TEST theories of behavior. It is these "behavioral scientists" to whom I refer when I say that they all operate under the assumption (conscious or not) that sensory input is the cause of response output. I'm willing to accept the notion that Chomskyans are above all that "lab" stuff.

Best Rick

Date: Mon Mar 15, 1993 4:23 pm PST
Subject: Re: Challenge: info theory & PCT

[Martin Taylor 930315 18:30] Bill Powers 930315.0700

Bill, In answering Allan, you say:

>The challenge was a response to the assertion that PCT could be
>derived from information theory.

Two points: (1) It's unfair to ask Allan to justify a claim of mine. He may or may not be able to do so, but it was my claim, and I told you I was working on a paper on it. I have shown the beginning of the paper, describing how "probability" should be interpreted, and he has problems with that. So if he could derive PCT from IT, he would probably do it differently from how I would do it.

(2) I see no conceivable way your challenge relates to a derivation of PCT from IT. All a solution to your challenge can do is compare the informational consequences of two circuits. My intention is to show that your circuits are the only reasonable way in which a system can maintain stability. I want to derive the fact of the hierarchy. I may not succeed in doing so in a non-circular way (as Tom Bourbon pointed out, many uses of IT have come to/from circular

conclusions/premises). The challenge is fun, and will be addressed. But it's irrelevant.

I said I hoped to get at it over the weekend, but I didn't get the review drafted completely until Sunday evening, and now Ina has had a go at editing it (we are joint reviewers of two books, not only one. Now we have all four drafts completed, and must coalesce them into two.) So it still isn't done. But the Control System Editor (Chris Love) is almost ready to send to Rick for beta testing.

Allan:
>> ... from Ashby's diagrams + information theory, one cannot
>>predict what exactly R, the regulator, is doing. You cannot
>>predict that R is going to oppose the disturbance.

Bill:
>If you stick with these conclusions, the challenge is unnecessary
>because you have agreed to my original claim. You are agreeing
>that information theory can't provide the predictions of behavior
>that control theory provides, but can only be applied once those
>predictions are known and verified.

That's a real non-sequitur. If I can see a coffee cup on the table I must drink it, or the theory falls apart? Why does R have to do anything particular just because it has some information available to it? All Allan is saying is that information theory doesn't tell whether I will drink the coffee, just that I will be able to (or otherwise) if I want.

>>When the information is cancelled, this means, in my
>>terminology, that E has experienced "information loss," and
>>thereby has high information content!
>
>Are you saying that the system begins with a certain information
>content, and after losing information it has a greater
>information content? I don't follow your arithmetic.

Terminology is confusion. I hate the wording "information content" as if it were a conserved quantity. One can talk about uncertainty, which is a property of a probability distribution. One can talk about information, which is a difference between two uncertainties. But "information content" misleads.

>>The transfer of one bit *requires* an increase in entropy to
>>accomplish the local decrease involved in the transmission of
>>information.
>
>This seems to imply that at the terminus of the transmission,
>there is a local decrease in entropy.

There may or may not be, depending on the uncertainties involved, but one thing is for sure: when a non-invertible event occurs in the universe, overall entropy increases. It may or may not increase in an open system.

>I can see no point in the transmission of a neural
>impulse where there is anything but a loss of energy (and an

>increase in physical entropy) in the next place that the impulse
>reaches. Where is this final place where we suddenly get the
>opposite effect (taking energy FROM the arriving signal?).

The energy comes from the neuron's "food" supply. The energy flow between the input and its excreta allows the local entropy to increase or decrease.

>...the state of the percept does not
>depend on the state of the disturbance alone. This means that it
>does not depends on the disturbance in a known way AT ALL.

Known where? If the ECS maintains a record of the change in the reference, it has the same potential sources of information as does one for which the reference level is always zero.

I sometimes use ALL CAPS to set off a word or comment, but I have the feeling that this capitalization of "AT ALL" is, like much of Rick's capitalization, more defensive than useful emphasis. If we get into a divide-by-zero situation (infinite gain, infinite speed), AT ALL may be justified. But I detect in a lot of both your comments a lack of middle ground between "totally perfect" and "non-existent and useless." Rick has said, without your correction, that if there is noise in the perception there is NO CONTROL. Or words with similar effect. Poor control is between the unachievable perfect control of a noiseless infinite gain system and the useless attempt to "control" without perception. Imperfect control is not no control. It is the real world.

>If I send you a Morse Code message and John
>simultaneously sends you a Morse Code message, and you receive
>only the OR of these two messages, how much information do you
>have about either John's message or mine? None at all. If the
>receiver gets only the OR of the messages, it has no way to sort
>out which dot or dash should be attributed to me or to John.

Not so. We are confronted with a very analogous situation almost all the time, with everything we hear. We do identify and interpret the music, the voice of our friend, the footsteps in the hall ... all at the same time. As a closer analogy, there is substantial research in the technology of recognizing one voice while a louder one is talking nearby. For the Morse code, there are relations among the keying "hand" characteristics, such as relative dot and dash length, that can start the separation of the two streams. There are redundancies both in the code itself and in the word stream (assuming it isn't encrypted to remove the redundancy). These can help sort out the two streams. It may not be possible to segregate and properly interpret both streams exactly, but to say that the amount of information you can get about either message is "None at all" is both theoretically and experientially wrong.

I'm not sure where this argument started from, but it doesn't feel to me as if it is going in a useful direction. Statements like the following don't help:

>>How can you say this, when the sole purpose of the control
>>system is to oppose the disturbance? It can't oppose something
>>it has NO information about. It simply cannot.
>

>When you get this figured out, you can finally claim to understand PCT.

Bill, your comment is exactly saying that PCT applies only to entities that operate in a universe thermodynamically isolated from the disturbing variable. It would make for a very uninteresting area of application for PCT. Let's deal in possible worlds, shall we? I think PCT applies to a very interesting part of the real world. If what you say is true, then I am much less interested in PCT, which has suddenly become a branch of abstract mathematics.

Martin

Date: Mon Mar 15, 1993 4:41 pm PST
Subject: Re: Linguistics facts

[Martin Taylor 930315 19:15]   Bill Powers 930315.1300

Here we go again ;-)

>I can't build a brain model that leads to 80% of the people
>doing one thing and 20% doing something quite different. A model
>does only one thing: what its design says it must do. This means
>that the model can only make one prediction at a time. If it
>predicts wrong an unpredictable 20% of the time, then it's just wrong.

I don't understand this. You've said similar things many times. But you also have talked about reorganization of the hierarchy as being necessarily random, and continuing until the hierarchy satisfactorily controls in the conditions to which it is exposed. Why should you assume that all people have reorganized the same, or that your model is wrong if it has one set of connections and not all of your experimental subjects have the same pattern? Why should you even assume that any of your subjects have completed all reorganization, and will control any particular variable. Some will, some won't, I would guess.

>The kind of data I need concern what people ALWAYS do. I don't
>think this kind of data is common in linquistics or any other
>conventional science of life.

There isn't much that anyone ALWAYS does, let alone that everyone does, if for no other reason that what people choose to control varies from moment to moment (I AM aware of the sloppy use of the word "choose"; it will suffice). Added to the probable variation among people among how they control when they do control what an observer might think is the same thing, I would think it no problem to even a very good model if it failed to predict some things 20% of the time.

>When I say "always," I make reasonable allowances. But I do mean
>that exceptions should be very rare.

"Reasonable," like so much else, is in the eye of the beholder. You perceive yourself as reasonable. I propose that your excessive stringency is due to the spectacular success you have had in your low-level modelling. You say you have never tried to fit a four-level hierarchy to data, yet you theorize (wonderfully) about an 11-level one. I know you apologize sometimes for doing so, and say it isn't science. But I venture to guess that in the more complex hierarchies, whose form is not dictated by the natural circumstances that face all humans most of their lives, you will never achieve the precision you seek, even though the models may be exactly correct.

There are two issues: Is the model correct? and do I have the right parameters for this particular event I am modelling? I think it will, in principle, not be possible to get the parameters correct for any reasonably complex hierarchic model.

Just to head off a possible misdirected counter, I am not saying that you cannot expect precise prediction of control from a high-level ECS that forms an element of a complex hierarchy (if they really do exist). I am arguing that you will not be able ever to get precise prediction from a complex hierarchy as a whole.

Martin

Date: Mon Mar 15, 1993 4:57 pm PST
Subject: Re: Behavioral Scientists

[Avery Andrews 930316.1051] Rick Marken (930315.1545)

Well, Chomskyans indeed don't spend much time in the lab, and I agree that when they do, they behave pretty much as you describe. But they do go out into to the field, listen to what people say, ask about how you say things, etc., and this, I would argue, is real research that doesn't fit the IV-DV mode. They also study texts (all we've got for extinct lgs. such as Old English or Ancient Greek), and draw occasionally interesting conclusions of what they find there, and I don't think this stuff fits the IV/DV mold either. In this respect, Chomskyans are no different from any other kind of linguist, except that novices in the field might do too much asking and too little listening.

Avery Andrews@anu.edu.au

Date: Mon Mar 15, 1993 4:59 pm PST
Subject: Re: Info theory: putting some ideas together

[Martin Taylor 930315 19:30] Bill Powers 930314.0900

Bill,

I think we need to stand back a bit. I see you as getting tangled in your own rhetoric. Probably you see me the same way. But your posting on how information about the environment gets into the organism seems to solve an unnecessary problem. Yes, we know that information gets in through uncontrolled perceptions, but the same is true of controlled perceptions. The information is about the state of the world, whether it be controlled or not. Whether this contains information about disturbances is irrelevant, but that is what has been giving you heartburn.

If a perception is controlled, there is information that it is so, or that there is error that will lead to action. But the perceptual signal flows up through the layers of ECSs whether or not they are actively controlling.

Did you read my "information for action" paper yet? In it, I describe four modes of perception, one of which would be controlled in PCT terms, the other three being through uncontrolled perceptions.   However, you are now expressing something of what I have been trying to get across from time to time when I return to the degrees-of-freedom argument:

>Uncontrolled perceptions may be uncontrolled at many successive
>levels, but eventually they join with controlled perceptions to
>become inputs to a high-level control system. There they act as
>disturbances of the perception to which they contribute, _and
>their information content is therefore rejected by that system_.

Except for that last (emphasized) clause, that's an aspect of what I have been saying. I'd even agree with the last clause, if "countered" were substituted for "rejected."

>The actions of that system that oppose the effects of the
>uncontrolled perception result (after passing through layers of
>reference signals on their way toward the periphery) in
>variations of visible controlled variables. Those visible
>variations carry the information about what the highest-level
>system involved did to cancel the disturbing effect of the
>incoming information. And perception of them joins with the
>uncontrolled perceptions to leave the highest-level perception undisturbed.

Wouldn't it be nice if we could resolve all our misunderstandings so easily?

>This works as long as the adjustments of controlled perceptions
>suffice as a way of rejecting the information and maintaining the
>world-view perception. But if the right kind of information is
>received, control may become more and more difficult, or fail.
>Then the person must start to reorganize -- find another world
>view and other ways of adjusting lower reference signals in which
>the adjustments will again succeed in preventing a disturbance
>from the uncontrolled information input.

Reorganization if necessary, but not necessarily reorganization ;-) (Sorry for the Canadiam in-joke).

Before reorganization, the normal effect would be to control some of those otherwise uncontrolled perceptions. That's part of the function of the alerting system. What we control changes all the time, out of our vast repertoire of what we can control. Reorganization presumably happens all the time when there is ongoing uncontrolled error, but if a shift in what is being controlled fixes the problem, reorganization will stop as if it had succeeded. My guess is that a reorganization solution to the problem is normally quite rare. But I don't know how I would test that.

Come to think of it, how does one test for reorganization in a human, anyway? And how does one fit the model to it so that less than 5% of the time it fails to predict the timing and the result of the reorganization event?

Martin

Date: Mon Mar 15, 1993 5:15 pm PST
Subject: information

[Avery Andrews 930316.1103]

I haven't had time to follow the information argument with the attention it deserves, but here's a thought anyway: maybe the problem is one of scales. A good control system presumably reduces error (reference - perception) to zero on an ecologically relevant scale of precision. E.g. the difference doesn't matter. But it can't do this on the physiologically relevant scale, since the error times the gain must produce enough output to oppose the disturbance.

The ecologically relevant scale is also the one relevant for control-systems higher-up the hierarchy - as long as the lower level systems are moving the steering wheel close enough their reference levels, the driver doesn't have to attend to wheel manipulation.

Avery.Andrews@anu.edu.au

Date: Mon Mar 15, 1993 6:47 pm PST
Subject: arms & coordinates

[Avery Andrews 930316.1200]

Perhaps one reason that the Arm demo hasn't made the impression that one might have hoped for is that the actual principles underlying its operation are somewhat more subtle than is immediately apparent, and there needs to be rather more in the way of purple prose than there is to explain how it works. And, I suspect, more on the way of simpler `sub-demos' to illustrate specific principles.

For example, after observing the rather stupid-looking trajectories created by my shoulder-centered polar coordinate arm (when the target is above the elbow the hand goes up to far, then circles back and over approaching the target in a sort of spiral), I thought, well what would happen if the polar coordinates had their origin in a more eye-like position, such as above the shoulder (at half the length of the forearm). Voila, the trajectories looked much more sensible, without any compensations. Whether they are realistic I don't know - haven't gotten around to checking out that part of the literature yet.

But the reason why they are more sensible is easy to see:

```
      x            T


      S --------- E -------- H
```

In a situation like this, the apparent angular error at the start of the movement will seem to be much less if the origin is at x (eye-position, sort of) rather than at S (shoulder position), so the inappropriate shoulder elevation will be reduced or even eliminated, depending on the details of the circumstances.

I think the shoulder-centered coordinates may have their virtues, however. For example, changes in shoulder yaw and pitch will have a very pleasant relationship to changes in the azimith and elevation of *any point* on the arm, namely, the identity relation. This isn't true for eye-centered polar coordinates. Coordinate transformations aren't really that hard to compute, & there's no reason why people couldn't make use of many coordinate systems simultaneously (indeed I think I recall an article in Synthese (or maybe BBS) arguing that they used six different ones at once).

More generally, I agree entirely with everything that Bruce Nevin said recently about PCT presentation - I don't think that people are all necessarily as block-headed as Rick Marken thinks they are, but they will act that way if the presentation isn't tailored to their current state of mind. In this vein, the theme that the eye-hand-arm system is set up so as to make effective perceptual control easy is one that might appeal to people who are attracted to Gibson, Kugler, Turvey, etc.

If anyone is interested, I can provide TC code for the `jacobian' and `polar' schemes, though it's pretty grotty. I would think that the basic points could be made with a spreadsheet, if trig functions are available (such as atan2, to go from xyz to polar).

Avery.Andrews@anu.edu.au

Date: Mon Mar 15, 1993 7:56 pm PST
Subject: Information about disturbance: basic analysis

[From Bill Powers (930315.1900 MST)]

RE: perceptual information about disturbances.

I agree with Martin Taylor that we need to stand back and try to find some new approach to this problem. All that's happening now is that positions are being hardened; pretty soon they will be fortified and that will be the end of that.

Let me try to analyze the problem without bringing in information theory at all.

If a disturbancing variable affects a controlled variable (to speak qualitatively for a moment), it is perfectly true that the perceptual signal will change a little. The perceptual signal being the route by which all "information" enters the control system, this means that the control system "has information about the disturbing variable." I am using the term "disturbing variable" to be perfectly sure it is understood that I am talking about the independent physical variable that contributes to the state of the controlled variable, not about its effects on the controlled variable, which must be deduced. When I use the term information here, I mean it only in the semantic sense.

The crucial question is what kind of knowledge that information provides. If the control system is reasonably good but not perfect, the change in the perceptual signal will give rise to a corresponding change in the output, which will reduce the unopposed effect of the disturbing variable by a large amount -- say 90% -- but will not remove it altogether. The net change in the perceptual signal, even so, is only 10% as large as the change that the disturbing variable would induce without the control.

Furthermore, as Martin pointed out, in real control systems there are lags. The output that is now being subtracted from the effect of the current amount of disturbance is the output that was generated by the perceptual signal of one lag-time ago. In addition, the output can change only at a certain speed, so it will rise and fall somewhat more slowly than the value of the disturbing variable rises and falls. Changes in the perceptual signal represent the current rate of change of the disturbance plus the slowed rate of change of the output generated by the perceptual signal of one lag-time ago.

So the perceptual signal at any moment represents the current magnitude of disturbing effects plus the (opposing) lagged and slowed influence of the just-previous perceptual signal's own effect on the output apparatus that is opposing the disturbance.

At a given moment, the perceptual signal, being a scalar, has just one magnitude, or it can be specified as having a mean magnitude with some uncertainty in both time and magnitude.

Maybe we can now make clearer what it means to say that the perceptual signal contains information about the disturbance. Given this perceptual signal, which contains all the information that the control system can get about the external world, how far could the control system go (even equipping its perceptual function with a supercomputer and a vast intelligent program) toward reconstructing the state of the independent disturbing variable itself?

For an external observer able to see all signals and variables, this is not a problem: simply subtract the current magnitude of the output from the current magnitude of the controlled variable, and you have the magnitude of the external disturbance, plus or minus random noise. The external observer is not limited to the information contained in the system's perceptual signal.

But what can the ECS itself determine about the disturbance? An ECS does not sense the state of its own output; it senses only the state of a controlled variable affected both by the output and by the disturbing variable. What can the information in the perceptual signal _alone_ represent about the disturbance?

The magnitude of the perceptual variable does not represent the magnitude of the disturbing variable averaged over the same short time. The disturbing variable's magnitude is many times that of the perceptual signal, by an amount that is larger as the control system's loop gain is larger. This would be seen immediately if somehow the output were forced to zero, so the full effect of the disturbing variable were passed through to the controlled variable. Then the perceptual signal would be a true (reasonably true) representation of the disturbing variable. With the feedback, its magnitude represents only a small part of the magnitude of the disturbing variable.

The waveform (time-variations) of the perceptual signal do not resemble the waveform of the disturbing variable. They represent only the difference between time-variations in the disturbing variable and time-variations in the slowed lagged output. So given only the time variations in the perceptual signal, there is no conceivable perceptual function that could calculate backward to isolate the variations in the disturbing variable from those in the output.

Neither would there be a channel to carry such information: in an ECS, the only output of the perceptual function represents the state of the controlled variable and nothing else.

So speaking strictly in terms of the effects of variables on signals and signals on variables, it seems clear that the perceptual signal in a single ECS is in fact somewhat affected by the disturbing variable, but in such a way that the state of the disturbing variable could not be uniquely reconstructed from the information in the perceptual signal alone.

In order to oppose most of the disturbance, the output must be adjusted nearly to the same magnitude as the disturbing variable, and its variations must not be so slowed or lagged that the controlled variable is allowed to change by large amounts. If control is to be good, the output must be a reasonably close mirror image of the amplitude and waveform of the disturbing variable.

Thus if we thought of the disturbing variable as responsible for the output actions, it would be necessary for the full magnitude of the disturbing variable and its actual variations through time to be represented in the signal that drives the output function. But we have just seen that the perceptual signal, the source of the driving signal, does not represent either the full magnitude of the disturbing variable or the actual waveform of the disturbing variable. The only conclusion possible is that the explanation in terms of the disturbing variable quantitatively causing the output is incorrect.

The correct explanation does not have to take the disturbing variable into account at all -- that is, the operation of the control system can be fully analyzed and explained without any disturbance present. Given this explanation, we can then predict the action of the system when any disturbance is introduced in any part of the system. Nothing has to be added to the explanation; the adjustment to the presence of the disturbance is a completely automatic outcome of the operation of the system by the same rules that apply when there is no disturbance. The procedure is identical to that of solving the homogeneous case of a differential equation first, and then completing the solution with a special case, an arbitrary driving function of known form.

I believe that I can justify and demonstrate each point made above. Rick Marken, in _Mind Readings_, has already demonstrated most of them. This discussion has said nothing about information theory or what it can or can't do: we've been working strictly with basic control theory.

Is this acceptable so far?

Best, Bill P.

Date: Mon Mar 15, 1993 8:34 pm PST
Subject: disturbances vs CEVs

[From Bill Powers (930315.2100)] Martin Taylor (930315.1800) --

>I understand that the gain cannot usefully be infinite if there
>is any noise in the perception, which means that the perception
>can correspond exactly and immediately to the real-world
>construct we call the CEV. I understand that such conditions
>do not exist in the real world.

I almost let this pass without comment -- I'm spending a whole lot of time at the keyboard and beginning to feel swamped.

Nobody - neither Rick nor I -- has ever said that the state of the CEV is not represented in the perceptual signal. Nor have we said that a change in the CEV is not (more or less) faithfully rendered as a change in the perceptual signal, or that the perceptual signal contains no information about the CEV or changes in the CEV. What we have said is that the perceptual signal does not contain (usable) information about the _disturbance_.

Just today, Allan Randall defined the disturbance as a change in the CEV, and I wrote a correction of that idea. Now I infer, perhaps incorrectly, that you, too, are defining the disturbance as a change in the CEV. Once before, I seem to remember, you expressed puzzlement as to why I insisted on introducing a separate disturbing variable -- you seemed to consider it superfluous. Is the same thing recurring?

Consider the rubber-band experiment. The disturbing variable is the position of the experimenter's end of the rubber-bands. The controlled variable -- the CEV -- is the position of the knot relative to the dot. The experimenter applies a disturbance by changing the state of the disturbing variable, his end of the rubber bands, not by grasping the knot, the CEV, and arbitrarily moving it.

This is critical. If the experimenter arbitrarily moves the knot, the subject will lose control completely and nothing will be learned about the control system. The subject will probably stop trying to control and complain, asking how it is possible to control the knot while the experimenter has hold of it.

When the experimenter varies the state of the disturbing variable by moving his end of the rubber bands in a reasonably smooth way, there is no way to predict how the CEV will vary; that is determined as much by the actions of the person on the other end of the experiment as by the experimenter. For any static state of the disturbing variable, the CEV might end up in any state, including no deflection at all.

The disturbing variable is an independent variable with respect to the control system. It can be set independently. The CEV is not an independent variable, at least while the control system is functioning normally.

Could this whole argument have arisen over a difference in understanding of what is meant by "disturbance?"

Best, Bill P.

Date: Mon Mar 15, 1993 9:51 pm PST
Subject: information in perception

[From Rick Marken (930315.2130)] Avery Andrews (930316.1200)

>I don't think that people are all necessarily as block-headed as
>Rick Marken thinks they are, but they will act that way if the
>presentation isn't tailored to their current state of mind.

If "block-headed" means stupid then I don't think that people who reject PCT are block-headed. In fact, I think most of them are very smart -- a lot smarter than me. If, however, "block-headed" means resistent to ideas that are disturbances to controlled perceptions then, yes, I think that people are "block-headed" simply because they are control systems. The peculiar disturbance resistant properties of a control system makes it impossible to get people to accept an idea that is a disturbance to a currently controlled variable. The reason that PCT has an acceptance problem is because most behavioral scientists (smart or not) are controlling for a variable that is disturbed by the concepts of PCT. I think the variable typically being controlled is something like "the input output model of behavior". But it might be something else.

Discovering this controlled variable is really why I am pushing this "information about disturbances" discussion with Martin and Allan. I believe that the idea that there is no information about disturbances in the perceptual input to a control system is very, very easy to understand -- even trivial. It requires no complex reasoning or the rejection of well understood physical principles. All it requires is a little, simple algebra (see below). But Martin and Allen are rejecting the idea that there is no information about disturbances in the perceptual input to a control system (by simply asserting that this just can't be so -- there must be such information) and I think they are doing this because this idea is a disturbance to their way of thinking about how living systems work. They believe that there must be something in the perceptual signal that can tell (inform, cause, guide, whatever) the output what to do. I think they consider this the only rational possibility because it is the only rational possibility if you are dealing with an input-output device (like a computer). In other words, I think they are not getting this very simple idea (no information about disturbances in the percetual input) -- not because they are stupid or bull headed -- but because this idea is a disturbance to their input-output view of behavior.

The fact that there is no information about the disturbance in the perceptual input to a control system can be easily seen from examining the formula for this input:

(1) $p(t) = d(t) + o(t)$

What the control system perceives (and controls) is a time varying signal, $p(t)$ that is at any instant the result of the combined effects of an indepedent disturbance variable, $d(t)$ and an output being generated by the system itself, $o(t)$. All the system perceives is $p(t)$; it has no way of knowing "how much" of $p(t)$ is at any instant the result of $d(t)$ or $o(t)$. In other words, it has no information about $d(t)$ -- all it has is $p(t)$. Nevertheless, it generates outputs, $o(t)$ that are a precise mirror or $d(t)$ -- something that can only be determined by an observer of the system:

(2) $o(t) = -k\, d(t)$

The system does this, not by "seeing" information about $d(t)$ in $p(t)$ and generating the appropriate output $o(t)$ in response to this information; that is, $o(t)$ is not generated the way it is described algebraically in equation (2). Rather, $o(t)$ is generated as the result of a comparison between $p(t)$ and a (possibly zero) reference value;

(3) $o(t) = g(r - p(t))$

This output is generated in a loop and each change in o(t) -- in a dynamically stable loop -- 'nudges' p(t) closer to r. So the loop is continuously generating o(t) values proportional to the existing error which tends to move the existing error to zero. A "side effect" of this process is that o(t) happens to be a mirror of d(t) (eq. 2). But the loop doesn't "know" anything about the form of the function d(t) -- this is the sense in which there is no information about d(t) in the input. There is no information about d(t) because the closed loop doesn't "care" what causes p(t) to vary ; all that the loop does is generate o(t) (which it also has no information about) which tends to cancel out whatever is causing p(t) to move away from r.

The appropriate way to describe this process is "control of perception". Perception, p(t), has no control over anything; it just IS. Perception does not "inform" the system about what to do; nor do variations in perception (when the perception is affected by the outputs of the system) contain any information about the causes of that variation -- that is, to what extent the variations are a result of disturbance or output.

I suspect that Martin and Allen will still want to defend the idea that there is information in p(t) which guides o(t). I, of course, will resist this defense because I am controlling for my understanding of the operation of a control system. All I can say is that their defense will be most effective if they can tell me how I can find the information in p(t) that will allow me to reconstruct the d(t) that was (partially) the cause of p(t) (see eq. 1). As I have said before, I have already tried one approach to finding the information in p(t); and I found no such information. What I found was that

o1(t) = o2(2) when d1(t) = d2(t) even though p1(t) <> p2(t)

where the 1's and 2's index the tracking trial during which these functions were obtained. In fact, the equality o1(t) = o2(2) was not perfect; the correlation between o1(t) and o2(2) was typically .997865 or so. Similarly, the inequality, p1(t) <> p2(t) is not perfect -- the correlation between p1(t) <>and p2(t) was not zero -- it was .0032 (or more, but never much greater than .2).

I think that Martin and Allan would have to claim that, in the tiny little correlation between p1(t) and p2(t) is the common "information" that is used to guide o1(t) and o2(t) -- making them approximately equal; I say that Martin's and Allen's position (that there is this information about d(t) in p(t)) invokes "magic" and, thus, I propose A NEW CHALLENGE:

I challenge Martin and Allan (and anyone else) to show me how to get the information about d(t) out of p(t) and reconstruct d(t) from it. I would bet a LOT of money that such a reconstruction algorithm will NOT be forthcoming --- but I think betting on the net is probably frowned upon -- so how about a big, embarassed concession from me when I get the algorithm and construct my first d(t) from the p(t) of your choice. The only caveat is that the reconstructed d(t) [call it d'(t)] represent the actual d(t) just as well as o(t) did in the actual run of the experiment. That is, the correlation between d'(t) and d(t) should be about the same as the correlation between d(t) and o(t).

Best Rick

Date: Mon Mar 15, 1993 10:50 pm PST

Subject: Fowler et. al.

[Avery Andrews 930316.1646]

Do people know about:

  Fowler, Rubin, Remez and Turvey (1980) Implications for Speech Production
  of a General Theory of Action, in Butterworth (ed) _Language Production_,
  pp. 371-420.

It contains, among other things, a discussion of BCP, including (pp 339-400) an argument that 2nd order control systems can't work.

At least, when Fowler & Turvey attack control theory, they go after Bill's version. I guess this is recognition of a sort.

 Avery.Andrews@anu.edu.au

Date: Tue Mar 16, 1993 3:06 am PST
Subject: Re: language norms

Bruce's posting on dialect, dialect-switching, etc. is reminiscent of William Labov's work in sociolinguistics. One other thing, people often deny that they say what they say. For example, mothers will insist that they pronounce 'can' with the same vowel as in 'hat,' but in normal speech it is more like a nasalized version of the vowel in 'let.' (Not all mothers -- this just comes from one of Labov's early studies.) Another example is 'gonna.' As an ESL professional I sometimes enjoy hearing people (teachers) say that they never say 'gonna' and that only 'going to' is correct. I rarely point it out to them when they later say 'gonna' quite naturally.

Eileen Prince   eprince@lynx.dac.northeastern.edu

Date: Tue Mar 16, 1993 4:25 am PST
Subject: back to basics

[Hans Blom, 930316]   Rick Marken (930311.1500)

>   Could you draw me a
>diagram of the Q control system? I really want to see how the Q
>control system works so I can simulate it on my computer.

Regrettably, diagrams are not a very convenient language to describe optimal control systems; mathematical formula's are the standard tools that are employed. Those, however, can certainly be translated into computer simulations. A classic is still Feldbaum's "Optimal Control Systems", but any more modern textbook on optimal control systems will provide you with the formulae. Translation into a computer simulation is straightforward.

Adaptive control theory, a simplification of optimal control theory, is easier to put into a diagram:

```
                  --------------------
         ------->|    system to      |-------
```

```
         |       |    be controlled  |       |
         |       ---------------------       |   outside world
         |                                   |  -----------------
         |       ---------------------       |  inside controller
 actions |------>| model of system   |<-----| perceptions
  u(k)   |       | to be controlled  |       |     x(k)
         |       ---------------------       |
         |                | |                |
         |       ---------------------       |
         |-------|  control system   |<-----|
                 |                   |
                 ---------------------
```

Much simplified (no-noise case): The control system issues actions u(k) and observes responses x(k). The model is any well-chosen formula such as

```
   x(k) = a + b*x(k-1)+c*x(k-2)+... + p*u(k-1)+q*u(k-2)+...
```

whose parameters a, b, ... p, q, ... must be tuned such that the model's response will approximate the system's response as closely as possible. This tuning is achieved through a correlation of actions and perceptions, both of which are provided to the model. This part is called "system iden- tification" or "parameter estimation"; an on-line technique for parameter estimation is, for example, the Extended Kalman Filter. Then, given the resulting 'knowledge' of the system to be controlled in the form of the estimated parameters, the model formula can be used either to predict future x'es given u's or future u's given x's.

```
                 --------------------- system response
         ------->|    system to      |----------->
         |       |    be controlled  |    |
         |       ---------------------   --- +
         |                ---------------| | comparator
         |                | tuning       --- -
         |       ---------------------    |
 actions |------>| model of system   |----------->
  u(k)   |       | to be controlled  | model response
         |       ---------------------
```

The above diagram may also help. The system's response is compared to the model's response and the difference is used to tune the parameters of the model. Thus, 'reorganisation' is considered to be a 'purposeful' feedback process guided by the error between expectations and what really goes on, rather than the 'random' process that Bill Powers proposes. System identi- fication in an adaptive control system is therefore a (learning) feedback loop within a (controlling) feedback loop.

The above scheme may explain what perceptual 'information' is used for: to build/tune the model, i.e. to learn. Initially, a baby's model is immature; the baby generates more or less random movements. But through observation of their perceptual effects, slowly the model is built up, and it is fine- tuned forever later in life. An important consequence is that parts of the model may become so perfect (because some aspects of life are so predict- able) that little or no information (perception) is required while yet (almost) perfect actions are generated.

Back to basics. I do not understand the following section of yours (930313.1000):

>What PCT says (in no uncertain terms) is that this assumption [that
>sensory input is an INDEPENDENT VARIABLE] is FALSE -- sensory input
>is NOT an independent variable in a living control system (ie. all
>living organisms) -- it is a DEPENDENT variable (the independent
>variable being the reference signal inside the organism). I think
>it is clear that, if PCT is right about this then the WHOLE edifice
>of behavioral science comes crashing down -- somwthing up with which
>most behavioral scientists will unquestionably not put.
>
>There is obviously not much room for a compromise resolution to this
>argument. PCT is just inherently revolutionary ...

This must have to do what you earlier called "the central insight of PCT". This position is not revolutionary, however. Bishop Berkeley already exposed the view that sensory input is a dependent variable in his "A treatise concerning the principles of human understanding", published in 1710. His philosophical position is called a strict empiricism: "esse est percipi", to be is to be observed. His position leads to propositions like "the tree that I see over there exists because I see it". Other philosophers have never found philosophical fault in this position: it is logically consistent and therefore both respectable and tenable. However, strict empiricists are not liked very much. People seem to resent being considered a figment of the other's perception.

That goes for me, too. I respect your point of view, but if you really believe that the universe up there at night has no existence independent from your perception, then there is very little that I can meaningfully discuss with you: there would be just too little common language/understanding.

However, in a later post (930314.1800) you react to Bill Powers (930314.0900):

>>So how DOES information get into the organism from the
>>environment? Clearly, only through uncontrolled perceptions.

>Brilliant; could this be one sideways contribution of IT to PCT?

>Your post makes me realize that I should say that only CONTROLLED
>sensory input variables are dependent variables; uncontrolled
>sensory variables (like the moving target in a pursuit tracking
>task) are still not independent variables; maybe a good name for
>them is INFORMATIONAl VARIABLES; these are perceptions that are "just
>there".

Now suddenly the view that there are also uncontrolled perceptions seems acceptable to you, even brilliant. How should I read this? To me it seems as if you are logically inconsistent here. Is that intentionally so or do you retract your earlier opinion at this point? It is important for me to know this, i.e. to understand you, i.e. to establish a correct model of you in my mind. Am I trying to have a scientific discussion with someone who has mutually inconsistent views, or are you in a learning process, where you change your views when error (what could that be? disagreement with Bill Powers?) develops? What IS the central insight of PCT for you?

Best,   Hans

Date: Tue Mar 16, 1993 7:04 am PST
Subject: Re: Linguistic facts

[From: Bruce Nevin (Tue 930316 09:04:52)] Bill Powers (930315.1300)

The kind of prediction you challenge me to make is a prediction about how reference perceptions are acquired, and which specific remembered perceptions will be made to constitute references for control. I think you must agree that that sort of prediction is subject to rather different standards than the more familiar kind of PCT modelling prediction, namely, given reference perception r (and gain g), a control system will control its perceptual input i to approximate r within a tolerance corresponding to g.

I did make a prediction, which it is perhaps just as well you ignored, because it was intemperate; there will of course always be exceptions, if only because of organic deficiency of perception or neurological function in some individuals.

Perhaps you will reconsider what you have written, given this clarification of our universe of discourse.

I know that you are interested in extending the modelling enterprise, from the modelling of control w.r.t. a given reference perception, to modelling the establishment of reference perceptions. You have been working on this in terms of reorganization.

I am saying that many reference perceptions, including surely those for language, are established not by random reorganization but by imitation of other individuals with whom one (primarily, but not exclusively, as a child) has much interaction that matters, and by a process thereby of instituting social norms in one's private reference perceptions.

I think you are vitally interested in these processes--backing off a bit, I think these processes are of vital interest for HPCT --because perceptions from the hypothesized category level on up depend strongly on language. Without an account of language in HPCT, the higher levels postulated for the hierarchy are hypothetical constructs with little more empirical basis than Freud's superego and id, or Eric Berne's TA constructs. The only evidence for them is the pragmatic observation that they are useful for organizing activities like teaching, therapy, and management in organizations, and the esthetic observation that they are a logical extension of constructs successfully modelled at the lowest levels of the hierarchy.

This is not a criticism of PCT. It is a science that is still in its infancy, as you and others have often said. On the contrary, I am urging how we must proceed if we are to progress.

I will omit detailed responses to your comments on Labov's work, what constitute useful data, etc. I believe, as I said, that you would say things quite differently given the distinction between the two kinds of prediction noted at the outset.

There are epistemological and ontological difficulties, to be sure, that arise in any attempt to reify `a language' or `a dialect' on the basis of generalizations over the speech of individuals. The appropriate approach would be I think as I have indicated, in terms of how individuals establish reference perceptions that they perceive as being socially standardized. The perception that another person's outputs deviate from one's own reference perceptions of the "same" outputs then constitutes a perception that the person is not a member of the familiar group of people perceived as sharing the same references as social norms. The perception of one's own group then becomes possible on the basis of the alienness of these others. This is the social function of prejudice, establishing a perception of whom (what kind of people) one may rely on most for cooperation, or most easily, with least preliminary negotiation of assumptions and expectations.

The usefulness of Labov's results for us is not as a basis for prediction that a social group in some way controls the behavior of individuals in it (the prediction that you seem to be demanding of me), but rather to show that there is a social reality that we must account for. I happen to believe that entities like "a language" and "a dialect" and "a social group" are byproducts of individuals controlling for cooperation with those with whom they have interactions that matter to them. There is feedback here, in the sense that people tend to restrict interactions that matter to them to people whom they perceive they can trust and can work with easily. I believe I have just sketched the principal basis for that perception. And people do come to have a perception of "our language" or "our dialect" or "our way of talking" as opposed to foreign ways, in a way exactly parallel to that sketched above in connection with the social function of prejudice. That hypostasis of social realities in individual's perceptions is not a statistical fiction, it is a perception as real as any other.

        Bruce   bn@bbn.com

Date: Tue Mar 16, 1993 7:23 am PST
Subject: Re: information

[Martin Taylor 930316 10:10]   Avery Andrews 930316.1103

On time scales, I think you are clearly getting the idea, just as you are with the "mechanical" side of PCT. You cut through to where I have for a long time been trying to guide the argument step by agreed step, but with no success because every smallest attempt meets with strong efforts to maintain the state of a controlled perception "that there is no information about the disturbance in the perceptual signal." Once you accept that there is, you ask about rates of information, and you find very quickly that the posting you made on time scales, precision and ecological relevance is right on target.

You see through problems to their heart, don't you.

Martin

Date: Tue Mar 16, 1993 7:25 am PST
Subject: Re: Information about disturbance: basic analysis

[Martin Taylor 930316 10:00]   Bill Powers 930315.1900

>RE: perceptual information about disturbances.
>

>I agree with Martin Taylor that we need to stand back and try to
>find some new approach to this problem. All that's happening now
>is that positions are being hardened; pretty soon they will be
>fortified and that will be the end of that.

I agree with everything in Bill's posting, at least at first reading. It describes exactly what I think happens, and is the basis of all my postings about information. Unless I spot some minor point in it later, and that point turns out to be important, we can take what Bill says to be the basis for further discussion on this topic.

If Rick also agrees with Bill's posting, we can see that our apparent disagreements are based more on misunderstandings of language than of principle. I do not have to change anything I have written (that I can remember) to accommodate what Bill said.

Martin

Date: Tue Mar 16, 1993 7:52 am PST
Subject: Re: information in perception

[Martin Taylor 930316 10:20] Rick Marken 930315.2130

>The peculiar disturbance resistant properties of a control system makes
>it impossible to get people to accept an idea that is a disturbance to a
>currently controlled variable.

We (Martin and Allan) are not in that group, but we recognize, all too well, that people do exhibit resistance to disturbances of cherished perceptions, and that we, like you, are people.

>But Martin and Allen
>are rejecting the idea that there is no information about disturbances
>in the perceptual input to a control system (by simply asserting that this
>just can't be so -- there must be such information) and I think they
>are doing this because this idea is a disturbance to their way of
>thinking about how living systems work.

I could very easily paraphrase this with only changes from "Martin and Allan" into "Rick" and deleting the word "no" in the second line. What good would it be to do so, except to point out the symmetry of the situation?

However, there are substantive comments that can be made.

>(1) p(t) = d(t) + o(t)
>
>What the control system perceives (and controls) is a time varying signal,
>p(t) that is at any instant the result of the combined effects of an
>indepedent disturbance variable, d(t) and an output being generated
>by the system itself, o(t). All the system perceives is p(t); it has no
>way of knowing "how much" of p(t) is at any instant the result of d(t)
>or o(t). In other words, it has no information about d(t) -- all it has is p(t)

Fine, except that the last sentence is a non-sequitur. It indicates that the core of the argument is a misunderstanding about the nature of information. p(t) = d(t) + o(t), and therefore p(t) contains information derived from d(t). This in no way implies that d(t) can be recovered from p(t) without knowledge of (as distinct from information from) o(t).

>Nevertheless, it generates outputs, o(t) that are a precise mirror or d(t)

That word "precise" is an argument killer. o(t) is NOT a precise mirror of d(t). It differs by an proportion that is roughly 1/Gain. In an integrating control, 1/Gain approaches infinity at infinite time, and then you can talk about o(t) being a precise mirror of d(t), but only if no disturbing "events" have happened in the interim. You can't use infinite gain without going to zero bandwidth, because otherwise you would be infinitely amplifying the least little noise in the comparator or perceptual input function. That again is a question of information.

>But the loop
>doesn't "know" anything about the form of the function d(t) -- this is
>the sense in which there is no information about d(t) in the input.

I'll certainly go along with this sense. But it is irrelevant to the points Allan and I have been trying to get across.

>All I can say is that their defense will
>be most effective if they can tell me how I can find the information
>in p(t) that will allow me to reconstruct the d(t) that was (partially) the
>cause of p(t) (see eq. 1)

I have no intention of defending the straw man Rick posts for me to defend, but as I mentioned to him privately, this question is quite interesting in its own right, and some day I'd like to look into it, if no-one else has done so and if it doesn't fall out of the information-theory analysis of the control loop.

If I could hazard a guess, the imagination loop is going to figure quite strongly in the solution.

Martin

Date: Tue Mar 16, 1993 9:20 am PST

[From Rick Marken (930316.0800)] Allan Randall said:

>>How can you say this, when the sole purpose of the control
>>system is to oppose the disturbance? It can't oppose something
>>it has NO information about. It simply cannot.

(Bill Powers 930315.0700) replied:

>When you get this figured out, you can finally claim to understand PCT.

and Martin Taylor (930315 18:30) says:

>Bill, your comment is exactly saying that PCT applies only to entities
>that operate in a universe thermodynamically isolated from the disturbing

>variable. It would make for a very uninteresting area of application
>for PCT. Let's deal in possible worlds, shall we? I think PCT applies
>to a very interesting part of the real world. If what you say is true,
>then I am much less interested in PCT, which has suddenly become a branch
>of abstract mathematics.

How could I let this statement go by last night; clearly, Martin has confronted my razor's edge, and balked.

Martin Taylor (930316 10:00) re:(Bill Powers 930315.1900)

>I agree with everything in Bill's posting, at least at first reading.

>If Rick also agrees with Bill's posting, we can see that our apparent
>disagreements are based more on misunderstandings of language than of
>principle.

I agree with Bill's posting, of course. I agree that there must be some confusion about language. That's why I tried to get quantitative in my last post. But that didn't seem to help. For example, I said:

>(1) p(t) = d(t) + o(t)
>
>What the control system perceives (and controls) is a time varying signal,
>p(t) that is at any instant the result of the combined effects of an
>indepedent disturbance variable, d(t) and an output being generated
>by the system itself, o(t). All the system perceives is p(t); it has no
>way of knowing "how much" of p(t) is at any instant the result of d(t)
>or o(t). In other words, it has no information about d(t) -- all it has
>is p(t)

And Martin Taylor (930316 10:20) replies:

>Fine, except that the last sentence is a non-sequitur. It indicates that
>the core of the argument is a misunderstanding about the nature of
>information.

OK. I'm waiting to have my understanding of "information" non-sequitized. But instead, here's what I get:

>p(t) = d(t) + o(t), and therefore p(t) contains information derived from d(t).

Another assertion; where the hell is this information? How can it be USED. p(t) is a series of numbers -- here's part of it 23,25,30,21, -10, 3, 8 ... Now, could you please explain to me where, in those numbers, is the information about d(t)? If you just say ITS THERE then I might break my computer screen. Isn't information a measure of what a message (which I take p(t) to be) tells you about the source, which I take d(t) to be. Or is this view of information a non-sequiter and thermodynamically impossible. I think your response to disturbance is showing.

>This in no way implies that d(t) can be recovered from p(t) without
>knowledge of (as distinct from information from) o(t).

So there is information about d(t) in p(t) but it cannot be recovered from p(t) without knowledge of o(t) -- is that it. So you are saying that there is information about d(t) in p(t) but it is useless without knowledge of o(t).

I'll buy this. So you are saying that the subject get's information about d(t) from p(t) because they know o(t)? Is that your position? If so, then I suppose I've got to develop a demonstration to show that the subject DOES NOT need to know about o(t) in order to produce o(t) = -d(t). I actually presented one non-experimental proof of this fact some time ago in the context of the "control of error" discussion. In that discussion, I showed that perception, p(t) was controlled (kept equal to r(t)) even when the disturbance was added before the trasformation of output into an effect on the controlled variable. That is, I added the disturbance to the error signal so that o(t) = k(r-p(t))+d(t). I showed both mathematically and with simulation that the disturbance is compensated for; in this case, by adjustments to the error signal. An interesting feature of this demonstration is that d(t) doesn't enter the loop directly through p(t). In fact, d(t) enters the loop as part of the effect of o(t), the very variable that we had to know in order to extract the information about d(t) from p(t).

So I'll tell you what -- how about this concession on MY part. There IS information about d(t) in p(t) [when p(t) = d(t) + o(t)] but it is useless and uninformative. In other words, there is uninformative information about d(t) in p(t). How's that?

I said:

>Nevertheless, it generates outputs, o(t) that are a precise mirror or d(t)

Martin says:

>That word "precise" is an argument killer. o(t) is NOT a precise mirror
>of d(t). It differs by an proportion that is roughly 1/Gain. In an
>integrating control, 1/Gain approaches infinity at infinite time, and
>then you can talk about o(t) being a precise mirror of d(t), but only
>if no disturbing "events" have happened in the interim. You can't use
>infinite gain without going to zero bandwidth, because otherwise you would
>be infinitely amplifying the least little noise in the comparator or
>perceptual input function. That again is a question of information.

Well my statement may have been an argument killer for you but your paragraph above is gobbldygook to me. What in the world is the point of all this? If it's a question of information then WHAT IS THE ANSWER. What does the lack of precision that comes from having a loop gain of 1,000,000 instead of infinity have to do with INFORMATION.

The best I can do with this is assume that you believe that the lack of precision of control leaves a remnent of variance in the input, p(t), that is the INFORMATION used to guide o(t). Thus, my challenge for you to find that remnent and reconstruct d(t) from it. But you say this is a straw man challenge. So please tell me what in the world is your point.

Best Rick

Date: Tue Mar 16, 1993 10:10 am PST

Subject: Language facts, PCT research

[From Bill Powers (930316.0900)] Bruce Nevin (930316.0904) --

>The kind of prediction you challenge me to make is a prediction
>about how reference perceptions are acquired, and which
>specific remembered perceptions will be made to constitute
>references for control.

I hadn't intended to ask for any predictions. Maybe this appeared to be my request because I said I was looking for observations that always hold true. I should have said "observations which have so far always, or very nearly always, held true." I'm just looking for behaviors that are very reproducible, like the way people move steering wheels when a gust of wind hits the car, or the way their muscles work when they're standing erect and someone pushes on them.

I'm a bit frustrated by our dealings, because I know that you understand PCT very well -- yet your approach to language seems still to be quite conventional. Maybe the problem is that PCT just doesn't have what it takes to organize a new approach to language. I'm probably asking too much, anyway. What I'm looking for is a linguist who has the kind of intimate and extensive acquaintance with human language that you have, but who is willing to chuck everything and start over, knowing only PCT. A pipe dream.

>I know that you are interested in extending the modelling enterprise,
>from the modelling of control w.r.t. a given reference perception, to
>modelling the establishment of reference perceptions. You have been
>working on this in terms of reorganization.

Reorganization and the acquisition of reference perceptions are only a small part of it. I should think that would pretty much explain itself once a basic PCT model of language existed. The really basic question is, what are people controlling for such that they do it by uttering words (and other means)? This question has many answers at many levels. In the course of trying to answer it, we should come across many control processes, and perhaps we will begin to see how they fit together. What we are looking for is not just patterns of linguistic behaviors -- outputs -- but relationships between those behaviors and controlled perceptions. Output patterns are bound to be variable; they change with disturbances. The controlled perceptions will be much less variable than the means by which they are controlled. The PCT approach is to try to explain variations in outputs by finding disturbances with which they correlate, and then using those hints to postulate and test potential controlled variables.

I don't see any sign of this sort of approach in the linguistic research of which I hear. Most of the effort seems to go into looking for regularities in language behavior. From the PCT viewpoint, this is something like trying to understand how people participate in sports by looking for regularities in the way they move their arms and legs. Such regularities are certainly findable, but they lead to statistical results, not clear and reproducible universal relationships. And they don't lead to an explanatory model.

>The appropriate approach would be I think as I have indicated,
>in terms of how individuals establish reference perceptions
>that they perceive as being socially standardized. The
>perception that another person's outputs deviate from one's own

>reference perceptions of the "same" outputs then constitutes a
>perception that the person is not a member of the familiar
>group of people perceived as sharing the same references as
>social norms. The perception of one's own group then becomes
>possible on the basis of the alienness of these others.

You could equally well have said

The appropriate approach would be to see how people set goals that they consider socially acceptable. When other people act differently, they are seen as belonging to a different group. So people define their own groups on the basis of differences from other groups.

You don't need PCT to make such a proposal. Plugging in PCT terms for their ordinary-language equivalents doesn't make this a PCT-type proposal. It's still just ordinary sociology, strung together using unexamined terms like "socially-acceptable" and "group" and "alikenness" and "familiar." If you had a PCT model of what is going on, such terms wouldn't even appear in the analysis or explanation. Those terms are simply a general fuzzy way of talking about the EFFECTS of controlling for specific perceptions by specific means. They aren't the terms in which any kind of explanation could be offered.

The problem with doing real PCT research is that we aren't ready to tackle the big global problems. Even stating such a problem reveals how superficial our understanding is. We would have to recognize that beneath a term like "socially acceptable" is a chasm of ignorance. What does a person actually perceive when that person refers to "society?" What are the actual controlled perceptions involved, the disturbances, the means of creating the perceptions and counteracting disturbances? There's only one way to find out: stop taking such terms for granted and start looking for the control processes that underlie them. And that means admitting that we can't make much progress at the interesting levels of human existence until we build a solid platform from which to work.

You say

>I think these processes are of vital interest for HPCT --because perceptions
>from the hypothesized category level on up depend strongly on language.

Now there's a statement that bears testing. Do perceptions from the category level on up depend strongly on language? What kind of control-system experiment could you set up that would show that this is a true statement -- and moreover, reveal exactly the form of the dependence? I know that others have done experiments purporting to test for such effects, but we have a different way of testing for perceptions, which others have never used. Maybe higher perceptions ARE linguistic. Maybe previous experimenters didn't really find out whether such an effect exists. The best approach for us would be to devise a PCT experiment in which we could measure the way a person is controlling a perception, and a way in which language could be used to try to modify the controlled perception. Maybe, using PCT, we could get results that are much clearer and more reliable, so the probability attached to this statement moves into a new region (upward or downward). In fact, before we went any further with using this idea, I would greatly prefer that we establish its truth.

Best, Bill P.

Date: Tue Mar 16, 1993 11:23 am PST
Subject: back to basics

[From Rick Marken (930316.1100)]   Hans Blom (930316) --

I'll let the good cop handle most of this; I do have a job, you know. But this stuff is just to juicy to lay off of completely.

I said:
>   Could you draw me a
>diagram of the Q control system? I really want to see how the Q
>control system works so I can simulate it on my computer

I made this request after Hans said that perception is not always what is controlled by a control system. So Hans has complied with some diagrams. Unfortunately, Q never appears in these diagrams -- and Q was the non-perceptual variable that he said was controlled. All that Hans sent were diagrams of an optimal control system which controls a perception that represents the difference between the actual and modelled effects of a system on a controlled variable. Nu? Looks like control of perception to me.

>The above scheme may explain what perceptual 'information' is used for: to
>build/tune the model, i.e. to learn.

Oy vay. Another control theorist who thinks that there is perceptual "information" that can be used to guide ("build/tune") behavior. I won't even ask ("where is it?", that is).

Hans says:

>Back to basics. I do not understand the following section of yours (930313.1000)

>>What PCT says (in no uncertain terms) is that this assumption [that
>>sensory input is an INDEPENDENT VARIABLE] is FALSE -- sensory input
>>is NOT an independent variable in a living control system (ie. all
>>living organisms) -- it is a DEPENDENT variable (the independent
>>variable being the reference signal inside the organism). I think
>>it is clear that, if PCT is right about this then the WHOLE edifice
>>of behavioral science comes crashing down -- somwthing up with which
>>most behavioral scientists will unquestionably not put.
>>
>>There is obviously not much room for a compromise resolution to this
>>argument. PCT is just inherently revolutionary ...

>This must have to do what you earlier called "the central insight of PCT".

You got it.

>This position is not revolutionary, however. Bishop Berkeley already
>exposed the view that sensory input is a dependent variable in his "A
>treatise concerning the principles of human understanding", published in 1710.

Really? I thought Berkeley was saying that perception was all there is; that we have no direct access to "boss reality" (which I happen to agree with, by the way). But this has nothing to do with the idea that sensory input is a dependent variable.

In conventional behavioral science (and in your statement about information in perception) perception is treated as an independent variable (independent of the actions of the system) and behavior is the dependent variable so:

b = f(p)

behavior, b, is a function of input perception, p.

PCT says NO; p is a dependent variable -- viz.

p = f(r)

where r is a reference input, usually inside the organism itself. This formulation is sometimes verbally accepted but, as is obvious from the ongoing conversation about information in perception, not thoroughly understood.

>That goes for me, too. I respect your point of view, but if you really
>believe that the universe up there at night has no existence independent
>from your perception, then there is very little that I can meaningfully
>discuss with you: there would be just too little common language/under-
>standing.

This is all completely irrelevant; nothing I ever said could possibly be construed as advocacy of solipcism (though you have apparently managed to so construe it -- so it is possible); and solipcism has nothing to do with PCT. Just for the record, I believe very firmly that there is a "boss reality" determining my perceptions and constraining the effects I can have on them. Our current best notion of the idea of this boss reality is embodied in our models of physics and chemistry.

No stone kicking necessary, see.

>Now suddenly the view that there are also uncontrolled perceptions seems
>acceptable to you, even brilliant. How should I read this?

I really don't know Hans; when did I ever give the impression that that I didn't believe that there were uncontrolled perceptions?

Best Rick

Date: Tue Mar 16, 1993 11:54 am PST
Subject: Re: error, information theory

I posted something on information a few days ago, which Gary INFORMED ;-> me didn't make it out to the list. Since then the argument's developed well beyond my post, and I'm enjoying it immensely. This issue is of key importance to me, as both an information scientist and control philosopher (if not theorist).

First a comment on semantic information:

> [Martin Taylor 930308 12:00] > I have to
> redevelop information theory using subjective probability, in such a way

> that I do not have to put up with claims such as "Shannon information has
> been shown not to relate to the everyday notion of information or meaning."

What do you think of Fred Dretske's Semantic Information Theory? His view takes similar stances: that information IS about individual messages, not ensembles with frequencies, and that it relates directly to common language senses of information.

In the end, I suspect that what's at stake is a philosophical dispute about the nature of CAUSAL vs. INFORMATIONAL processes. Clearly D is CAUSALLY linked to O, but the INFORMATIONAL link is hardly (!) clear. This is a distinction that Dretske also clearly makes, finding only a contingent relation between the two kinds of links (as Bill does). This is contrary to received wisdom, which makes causal and informational links essentially equivalent.

The whole issue is fascinating, espeically Bill's observations about the "virtual communication channels" between (RL and CEV) and (D and O). But I also agree w/Randall that IT and PCT are at qualitatively different levels of description. wrt/the original challenge, there seems to be agreement that no matter how you wire up the diagrams, that IT can be used to describe the informational properties of ANY communications channel. And that's all that the links in a PCT diagram are: links in a communications network among environment, controlled variable, perceived variable, reference level, error signal, output fnuction, etc. Shannon's theorem applies at each link, and Ashby's Law (which is isomorphic, I believe) as well: that each link must have sufficient variety to encode the signal without error in the presence of expected noise.

The question is of the SIGNIFICANCE of this point: just because IT is LOGICALLY PRIOR to PCT, doesn't necessarily mean that it's FUNDAMENTAL to it, but rather it may be IRRELEVANT to it. This is what I understand Bill's position to be now. IT makes no prediction about one particular network topology over another, only what the bandwidths of each link must be to maintain good control. IT strikes me as a very modest and local theory, nothing like the grand sweep of PCT or Ashby-en (or anyone else's) Cybernetics in general.

Please don't stop, people. Even if it's "just" an argument about semantics, those tend to be far more important than they're given credit for.

O-->
| Cliff Joslyn, Cybernetician at Large, 327 Spring St #2 Portland ME 04102 USA
| Systems Science, SUNY Binghamton   NASA Goddard Space Flight Center
| cjoslyn@bingsuns.cc.binghamton.edu   joslyn@kong.gsfc.nasa.gov
V All the world is biscuit shaped. . .

Date: Tue Mar 16, 1993 12:53 pm PST

[Martin Taylor 930316 15:30] Rick Marken 930316 08:00

I give up. There's an impenetrable wall, or the resistance to disturbance of a converted missionary. Every attempt at explanation is gobbledygook or plain wrong. It's obvious that there is NO explanation that will suffice, other than simple capitulation. So, for the purposes of clearing the CSG airwaves, I capitulate, and will continue this information-theory discussion with Bill Powers separately. At least Bill and I can see common ground, even if we don't yet have a common understanding. With Rick, as soon as there appears to be the possibility of common ground, the rug gets whipped away to reveal a chasm. I don't like that style of argument. T'is so, t'aint so. Stupid.

For an example, one final quote from Rick:

>What does the lack of precision that comes from having a loop gain
>of 1,000,000 instead of infinity have to do with INFORMATION.

Everything.

Bye (on this topic).   Martin

Date: Tue Mar 16, 1993 2:13 pm PST

[From Rick Marken (939316.1200)]

Let me try to articulate what I think is really at stake in this "information about the disturbance" discussion.

The PCT view of perception is that it is just there; it doesn't "communicate information", it has no "cues" or "affordances", it is not a stimulus (in the sense of making anything happen), it is not a "reinforcement". Perception consists of variable aspects of experience (from the point of view of the behaving system) or simply variable signals in the nervous system (from the point of view of a brain scientist).

Perceptions are just "layed out" there for the perceptual control systems to deal with. These perceptual control systems "care" about some of these perceptual variables -- these are the controlled perceptions; perceptual control systems act to get these pereptual variables to match reference specifications. Other perceptual variables are not under control -- they are uncontrolled perceptions. These perceptions can apparently become the objects of consciousness -- but there are no perceptual control systems trying to get these variables to any particular reference level.

So perception (from a PCT perspective) is just raw material; control systems act to make some of this material match internal references; these references are selected (in theory) because by making perceptions match these references, higher level perceptions are made to match their references specifications.

Clearly, this view of perception is quite unusual. In conventional psychology we often hear of perception providing "cues" for behavior (or cues to the "real" identity of the object represented by the perception); Gibson (and D. Norman) talk about objects having "affordances" -- perceptual characteristics that "tell" the perceiver how the perception can be used (the perception of a "chair" for example "affords" sitting -- yes, these people are serious). We often hear about perception providing "information" about the appropriate response to make; information that can be "used" to guide behavior (this does imply an "active" role for the behaving system, but it suggests that perceptions have some- thing in them -- information -- besides the perception itself; something that can "tell" something to the behaving system about what to do).

The PCT view is that perceptions have nothing to tell -- unless you call the result of the comparison of a perception to a reference "telling"; but in this comparison the perception is just what it is -- ultimately a number, say 10. The perception "tells" nothing in and of itself. If the reference for that perception is 5 then the comparison "tells" the system (as an error signal) that the perception is "too big". If the reference for the perception is 15 then the comparison tells the system that the perception is "too small". So the same perception can "tell" the system different things; it depends on what the system wants the perception to be. So it is not the perception that is "telling" anything; it is the result of the comparisom between the perception and the reference that "tells" something.

So from a PCT perspective, perception, in itself, does not "tell" a control system anything; it just IS. This is what I mean when I say that there is no information in perception. Information is not in the perception; information is in the difference between perception and reference; information is in how the world as it is (perception) differs from the world as it should be (the reference).

Best Rick

Date: Tue Mar 16, 1993 3:19 pm PST
Subject: Language xpt, PCT research

[Martin Taylor 930316 16:45]
Triggered by, but not a response to (Bill Powers 930316.0900)

As you may know, Allan Randall is working on a contract for me, as part of which he produced the review/bibliography of PCT that some people seemed to appreciate. A long-term goal of this work (not achievable within the life of this contract) is to develop "intelligent" computer interfaces, based on Layered Protocol Theory, which we see as a two-person application of PCT. A stage in this (also not achievable within this contract) is the development of a control network for speech recognition and synthesis. What I want to talk about is a stage leading to that-- an experimental study of PCT and grammar.

When someone is exposed to linguistic material, whether it be speech or text, the classical view is that the input is analyzed and matched with some kind of stored pattern at a rather low level, to produce, say, letter identities or phonemes or demisyllables or even words. This is an entirely passive process, and the recognition process is divorced from the synthesis of language (speech or writing). We feel that this is perhaps not the best way to approach the problem, either when thinking about how people do it or in constructing a silicon-based way to do it.

The basic idea we have is one that has been batted around on the net from time to time, mainly by Bruce Nevin and Bill Powers initially (at least initially in terms of my acquaintance with CSG-L). It is that we are always controlling many of our perceptions, including the language ones, even when we are just passively listening (or reading). There is a school of thought in the speech recognition community called "analysis-by-synthesis" (ABS) which is not what we are proposing. But we are proposing that the recognition elements of the control net learn though ABS, using parts of the control hierarchy that are also involved on language production.

Consider a simple elementary control system (ECS), classical kind:

```
 ------|-------
 |    ref      |
 |  ->cmp --   |
 |  |      |   |
 | perc    out |
 |  |      V   |
  --|-------|--
    ^       |
    |       |
     --cev--
       |
      dist
```

A generalized diagram. CEV is the complex of states or events in the world that the perceptual function turns into a perceptual signal. It has no independent existence, but is affected by the actions of some disturber. The output signal counters those effects to maintain the percept near the reference, so (as we all agree), the disturbance is hardly if at all evident in the perceptual signal.

(I put the side-bar paragraph in partly to indicate that I am not carrying forward the information-theory discussion here).

If the link from output to cev is broken, the perceptual signal is a transformation of the disturbance, which is a straightforward view of what is happening in "passive" listening or reading. But is this view reasonable? There's all sorts of evidence (that Bill probably would call unscientific, but nevertheless suggestive) that people tend to track and to predict when they are listening or reading. They may not know what word, phrase, or concept (different levels of abstraction) is coming up, but they often are close. This seems very like the imagination loop in action (Hans Blom calls it the system model, or something of the kind).

What we propose is a study of reorganization (remember the 12 varieties). This is a straw man I'd like to have knocked down or strengthened.

The disturbance we propose will be a super-simplified version of text: a sequence of characters generated by a Markov chain or a generative grammar. Each character will be allocated a 3-vector (for example, A might be 1,2,1, and B at 3, -1, 5). This 3-vector constitutes the disturbance to the CEV, which can be visualized as a point in 3-space. The sensory inputs are the current location of the CEV point in 3-space, in any suitable coordinates, not necessarily those defining the letter vectors. The disturbance generator might produce ABEZCEBF... but the PIFs would generate waveforms. There would have to be at least 3 ECSs to allow the control system as a whole to "cover" the space. Let's say there are 3 exactly, in which case their PIFs define the perceived coordinate system of the space.

The outputs of the 3 ECSs define a 3-vector that is used to influence the CEV (actually added to the current value of the disturbance). There is a transport lag in the system, as well as various slowing factors and the like, which may participate in the learning process. This means that the character now disturbing the CEV has not had any influence on the output vector.

The objective of the study is to see whether, and to what degree, the control system can learn to keep its perceptual signals near zero. Obviously, if the grammar is just a random sequence of characters, it will be totally unable to control, and any output it makes will be added in quadrature to the disturbance; its best performance will be attained at a gain of zero. But if the grammar has sequential constraints, a sufficient hierarchy on top of the 3 peripheral ECSs should be able to do somewhat better than that.

The forms of both the PIFs and the output transformation are the questions of interest. We can subsume the sense and relative magnitude of the outputs in each dimension into the form of the output function, mimicking the link-alteration aspect of standard reorganization.

That's the skeleton of the idea, and I'd appreciate any comments anyone might have.

Martin

Date: Tue Mar 16, 1993 3:36 pm PST
Subject: Control of perception: Berkeley be damned

[From Bill Powers (930316.1500)] Hans Blom (930316) --

The second of your adaptive control diagrams is essentially the same scheme I was proposing a while ago for model-based control. As you say, model-based control would have to use some systematic algorithm for making the model's response agree with the environment's -- not a matter for random reorganization.

I don't think that model-based control is very practical for real behavior unless there is also a provision for making corrections in the action based on the deviation of the model from the real environment's response. All that the model-based scheme can do for you is carry the system through periods when real-time feedback is unavailable. This, of course, can be critical, but no organism could rely on it for long. There has to be some way to reset the model at frequent intervals, because it will always drift away from reality rather rapidly. Just think about walking around in a completely dark room after a power failure at night. You can go a short distance without running into anything, but you can't go far without at least touching something identifiable to reorient your imaginary model of you in the room. It's interesting how little tactile information is needed to accomplish the resetting, but without any such information at all, no mental model stays in calibration for more than a few minutes (if that long). And of course even the best mental model fails immediately if the environment is disturbed.

Your response to Rick Marken:

>This must have to do what you earlier called "the central
>insight of PCT". This position is not revolutionary, however.
>Bishop Berkeley already exposed the view that sensory input is
>a dependent variable in his "A treatise concerning the
>principles of human understanding", published in 1710. His
>philosophical position is called a strict empiricism: "esse est
>percipi", to be is to be observed. His position leads to
>propositions like "the tree that I see over there exists
>because I see it".

If you think this central insight is philosophical or mystical, you've missed the point (your quoting the good Bishop makes me sure you have missed it). PCT says nothing about the existence of the tree over there. Let the philosophers worry about that. All that PCT says is that you can't avoid bumping into the tree unless you can see your relationship to it. That means seeing it and its distance from you. If you're wearing trick glasses, you may bump into the tree anyway, because you can control only the relationship you perceive, which may be different from the one that actually exists. If the tree is really a holographic projection, you'll still steer to walk around it because you can control only what you perceive; as far as your perceptions can tell your brain it's a real tree and walking into it will hurt.

All that the brain knows about the external world comes to it in the form of perceptual signals in the afferent neural pathways. There is no other way for that information to get into the brain. If the brain wants to control the position of a real glass on a real table, it's out of luck: it doesn't have any way to know about the real glass and the real table. It can, however, adjust its output signals so that a neural signal representing the glass can be manipulated to achieve a certain relationship with a neural signal representing the table. THAT the brain CAN do.

I should think that all this would be self-evident to any engineer who has ever actually built a working control system. A real hardware control system can't interact directly with the physical plant it is controlling. All it can do is alter its electronic output signals and see what happens to the signals being generated by its sensors. That's all it knows about what is happening outside it. If the sensors jump out of calibration, the control system will happily continue controlling the miscalibrated perception, while the technician in charge rushes to hit the STOP button. What is controlled is ONLY what is perceived. One hopes that what is perceived has some relationship to what IS, but that is something that has to be determined indirectly.

This simple concept which should cause no problems for any control engineer causes immense problems for conventional sciences of life. The reason is that these conventional sciences ignore the difference between what is perceived and what is -- at least when they're trying to explain behavior. And not having any experience with real system design, it seems perfectly reasonable to such conventional scientists that a stimulus input from real objects in the environment should be able to cause motor outputs that steer the organism through a variable environment along a path to the cheese or the mate or whatever. What's the problem? You can see them doing it, so it must be easy.

If you're an engineer watching an organism behave, you will have a hard time making your mental model behave in this simple cause- effect way. You will notice that the eyes keep moving around, that the head moves and bobs up and down, that the steps are a little imprecise and slightly wobbly, that things in the environment are shifting around. Being a person who is charged with making systems actually work, you will wonder how the organism gets away with such imprecision of action -- where are all the stimuli coming from that cause the corrections of the little mistakes and overshoots and hesitations? How does the environment know that it should stimulate the organism just in the right way to correct for a previous stumble? How does that little unevenness in the path send just the right stimulus up the spine to make just the right muscles change their tension to keep the leg from jamming into the ground or flailing in empty air on the next step? Any engineer who pays attention in a professional way to the claims of s-r theorists would soon walk away shaking his head. No way!

Unfortunately, engineers seem to abandon their normal professional attitudes when they start trying to explain behavior. They start listening to the psychologists and physiologists and neurologists who think that behavior can just be "generated," open-loop. Perhaps they're just being polite because they're on another scientist's turf. They say, "Oh, is that how it works? OK, you must know what you're talking about; I'll see if I can make that work."

And of course they can make it work. A good engineer can make any damn fool idea work. They can build an arm that's as solid as the front end of a Mack truck, equip it with precision bearings and gears and stepper motors, compute the driving signals using 80- bit floating point arithmetic, and make the arm move exactly as wanted. The smart ones must surely realize that this is NOTHING like the way a human arm works. But the psychologists see what they've done, and nod wisely. It works just the way they expected.

PCT is all about the realization that human systems simply can't work that way. Their outputs are rubbery and imprecise; their neural computers are good to maybe 1% at best; they don't sense everything in the environment that might interfere with the action. Yet they work precisely and well, for four score years and six. A person with his little 1% analog computers can get out of bed in the morning, and perform one action after another all day long, each action starting where the last one left off, and 16 hours later end up exactly at the side of the same bed, with no cumulative errors at all. Only one kind of system can accomplish that sort of behavior: a negative feedback control system.

>I respect your point of view, but if you really believe that
>the universe up there at night has no existence independent
>from your perception, then there is very little that I can
>meaningfully discuss with you: there would be just too little
>common language/under- standing.

That kind of respect I can do without. It's condescending. Actually, I do believe that the universe up there at night has an existence independent of my perception. I am also convinced that that universe is nothing like my perception of it. The stars are not tiny pinpoints against a black background of nothingness. The Milky Way is not a vague sheet of haze. When Mars slows in its journey among the stars, reverses its path, and then starts forward again, I know that my perception is not what is actually happening. Nevertheless, I can find my way somewhere by keeping my perception of the North Star in a certain relationship to my perceived direction of travel, and I know which parts of the sky to center in my visual field so that Mars will appear in the middle. Even though my entire education has taught me that the physical universe is not at all the way I perceive it, I can still act to make my perceptions behave in ways that have predictable outcomes, and know that behind the scenes I am having unseen effects in the unperceived universe that are causing my perceptions to behave as they do. I can turn a knob-thing on a radio with full confidence that I will shortly hear voices or music -- even without knowing what actually happened behind the knob. Yes, there is a real reality. I am convinced of that. But it is not the reality I experience. If you gave this ten minutes' thought, you would realize the same thing.

Best,   Bill P.

Date: Tue Mar 16, 1993 4:07 pm PST
Subject: information in perception

[From Rick Marken (930316.1400)] Martin Taylor (930316 15:30)--

>I give up. Aw. Common. We're workin' on it.

> Every attempt at explanation is gobbledygook or plain wrong.

Well, every time I say something you don't understand it's a non-sequiter -- but I keep tryin'.

>It's obvious that there is NO explanation that will suffice,
>other than simple capitulation.

That is not true (for me). I just want an explanation of where the information is in p(t). If there is no way to get at that information or if it is unusable then I want to know why you think it is so important. I just want to know what you (and Allan) are trying to say. Believe me, your claim that there is information about the disturbance in p(t) sounds just as crazy to me as my claim (and Bill's) that there is NO such information in p(t) appears to sound to you.

>So, for the purposes of clearing the CSG airwaves, I capitulate,

That avoids it; it doesn't clear it.

>and will continue this information-theory
>discussion with Bill Powers separately. At least Bill and I can see
>common ground, even if we don't yet have a common understanding.

I look forward to that discussion; but I warn you; Bill P. may seem seem a lot nicer than me, but he understands PCT just as well as I do (he he). So you will still be stuck with the problem of explaining why you think there is information about anything (let alone the disturbance) in p(t).

>With Rick, as soon as there appears to be the possibility of common
>ground, the rug gets whipped away to reveal a chasm. I don't like
>that style of argument. T'is so, t'aint so. Stupid.

It's not a style of argument; it's an approach to knowledge. I'm all for common ground when that ground is, indeed, common. But I can't agree with something that just flat out contradicts a basic tenet of the PCT model. If what I hear from you just seems wrong because of our disparate ways of using language then let's try to work that out -- and the only way I know of to get past language is to agree on what experiences we expect our models to produce in agreed on circumstances. I've described my attempt at a demo of lack of information in p(t). You have never said what is wrong with that demo (described in chapter 3 of Mind Readings) or what the appropriate demo would be to show that there is information in p(t).

>For an example, one final quote from Rick:

>>What does the lack of precision that comes from having a loop gain
>>of 1,000,000 instead of infinity have to do with INFORMATION.

>Everything.

I'm willing to believe this, really. It's just not obvious at all to me. I am really trying to understand you. Could you please get quantitative; what is the information in p(t) that exists when the loop gain is less than infinity but is gone when it is infinity. Since most real control systems have loop gains that are considerably less than infinity it should be easy to explain how I could set up a control demo that would convince me that there is information in p(t); show me how to do it.

Best Rick

Date: Tue Mar 16, 1993 8:38 pm PST
Subject: information, psychics, linguistics.

[Avery Andrews 930317.1404]

Spending a bit more time on the information controversy, it seems pretty clear to me that Rick Marken & Bill Powers have run afoul of the bizarre fact that `information theory' has nothing to do with content. This seems strange, but is true, and is the motivation for a lot of what has been going on in places like Palo Alto for the last ten years ago (so Dretske and lots of other people want a `semantic information theory' to address issues of what content is). The fact that information theory is not about content is probably one of the reasons why linguists such as myself know virtually nothing about it, but one thing I do think I know is that if the info coming over one wire is P1, and that coming over another is P2, the info coming over both is P1+P2. Now knowing only p(t) or o(t), we can't recover d(t), but miraculously, knowing both, we can (if we pretend that there is no noise), so in some sense or other there must be info pertaining to d(t) in both p(t) and o(t), since 0 + 0 = 0. But, equally mysteriously, in the realm of content (i.e. informative information), 0 + 0 != 0. What we have to do is figure out how the information can in some sense be there, and yet be uninformative. Alternatively, we need to understand the sense in which it is there, and the sense in which it is not.

This requires that people actively *think* and try to make sense out of what other people are saying, rather than relentlessly defending their positions. You can't expect the other guys to give you their essential insights on a silver platter. To paraphrase the Godfather, real knowledge can't be given, it must be taken.

I'd also like to point out that the situation with d(t), o(t) and p(t) is basically the same as encryption with a one-time pad: the pad (o(t)) certainly has no content at all, the message (p(t)) has no accessible content for people who don't have the pad, but if you have the message and the pad then you get the content. Maybe the philosophy of cryptography, if there is such a thing, has some application here.

[Bruce Nevin (Tue 93039 12:17:00, Thu 930311 12:08:14)]

>And now once more beyond the pale: in learning to be people with
>those around us we learn to discount and ignore some perceptions,
>or to cover or reinterpret them with the aid of imagined
>This is what has happened with socalled "psychic"
>perceptions, which have for many centuries been the occasion of

>painful and shameful death, and which apparently not all people
>access equally (variability analogous to color blindness).

This is also how people learn GB syntax, covering up the perceptions to the effect that the analyses don't actually work. But I'm not so sure about the psychic stuff. One reason I don't believe in it, in spite of the fact that my own mother claims to have telepathic experiences all the time, is from reading a book about a Crow Indian warrior called Two-Leggings (can't remember the exact title, unfortunately). He and his friends seemed to use clairvoyance in an extremely matter-of-fact way to locate game, enemy war-parties, etc., yet these skills were neither salable to the US Army, nor usable against them, to prevent massacres & surprise attacks such as Wounded Knee. Somehow, it vanishes into smoke when confronted with our stuff. So I think that nonexistence is the best explanation. But who knows. Maybe it's just that our songs are stronger.

(Bill Powers (930316.0900))

>Most of the effort seems to go into
>looking for regularities in language behavior. From the PCT
>viewpoint, this is something like trying to understand how people
>participate in sports by looking for regularities in the way they
>move their arms and legs.

This is because you don't have to go looking for these regularities - there just out there, in wierd and wonderful profusion, like the fossils that S.J. Gould likes to go on about. Real linguists are just the people who are compelled to think about them, which is why you won't ever find the linguist you're looking for. This is also why linguistics is a lousy model for cognitive science in general, and why it runs into the sand when it gets to discourse - linguistics is about stuff that happens on a small time scale, where interaction doesn't matter much, whereas discourse and most other cognitive things are highly interactive (Martin Taylor has some stuff about this in his papers). To get a good PCT linguistics, you have to start with both more or less as they are, and resolve the contradictions, rather than just blank out one of the subjects. You might wind up totally redoing it of course (like when Phlogiston went into the scrap heap, but you start with it as it is). I think that the things that Bruce says about norms are on the right track, tho there's clearly a long way to go. Martin's recent post also looks interesting, tho I haven't figured out anything useful to say about it yet.

Avery.Andrews@anu.edu.au

Date: Tue Mar 16, 1993 9:10 pm PST
Subject: information, research proposals

[From Rick Marken (930316.2000)] Cliff Joslyn --

>In the end, I suspect that waht's at stake is a philosophical
>dispute about the nature of CAUSAL vs. INFORMATIONAL processes.
>Clearly D is CAUSALLY linked to O, but the INFORMATIONAL link is
>hardly (!) clear.

Nicely put, Cliff. There is no question that d(t) is causally linked to p(t); I'm arguing that there is no informational link. Martin's and Allan's (where is Allan?) chagrin is understandable if they thought I was denying the causal link from d(t) to p(t). Maybe that's why Martin thought that denying the informational link from d(t) to p(t) made PCT a mathematical abstraction; maybe he thought that my use of the term "information" was synonymous with "causal". Is that right, Martin?

>Please don't stop, people. Even if it's "just" an argument about
>semantics, those tend to be far more important than they're given credit for.

I'm willing to go on until I'm satisfied that I understand what Martin and Allan are talking about; but I don't know if I can get anyone to come out and play.

Martin Taylor (930316 16:45) --

> What I want to talk about is a stage leading to that--
>an experimental study of PCT and grammar.

Before this "information in perception" debate broke out I was planning to submit a research proposal to the net also; I don't have time to do much research (and post to the net too -- he he; well, actually, I'm still trying to write -- I'm even slower at it than Gary Cziko). But I think this is a good idea; there may be some graduate students looking in (or graduate advisors) who might be able to run with some of these research suggestions.

>The objective of the study is to see whether, and to what degree, the
>control system can learn to keep its perceptual signals near zero.

I really didn't understand your description of the study. But after reading this far I am getting the distinct impression that the control systems you plan to study are simulations. Why not do a study with real control systems -- like people? If you are intending to study people then I need some clarification of the vector representation of the letters and all that. I sure hope you are planning to do this with living control systems. What we really need in PCT is good data; who cares how a particular implementation of an algorithm works? Unless your working in AI or Artificial Life, of course, in which case you only care about the behavior of the model itself.

Research Proposal

My proposal is to study the human ability to control sequences (I use letters, just like you do, Martin). The studies would be an extension of the perceptual studies described in my "Hierarchical behavior of perception" paper. I think there are at least two very good reasons for doing these studies; 1) they would demonstrate the applicability of PCT to a complex behavior-- "sequence produciton" and 2) they would show that limitations in what appear to be the ability to generate "outputs" may actually be limitations in the ability to perceive the intended perceptual consequences of those outputs.

I have already set up and done the basic experiment and it works like a charm -- but there are many variations and some modelling work to do, and that's where the PhD thesis student would come in.

The basic task is for the subject to control a sequence of letters. The letters appear at alternating positions on the screen like so (with time moving from left to right:

```
A  C  E  B  D
 B  D  A  C  ....
```

The subject is to keep the ABCDE sequece occuring on the screen. He or she does by pressing the mouse button when necessary -- a button press becomes necessary when the sequence is "disturbed" which means that a new transition rule between letters comes into effect such as

```
ABCDEADBECADBEC...
      ^
```

The new transition rule occurred at the ^; note that the new sequence has the same letters but none of the transitions are the same as in the "reference" sequence. Pressing the button returns the "reference" transitions. If the button is pressed when the reference transitions are in effect the "disturbance" transition sequence is implemented. So, in order to control the sequence, you have to push the button at the right time.

Some quick preliminary (but very highly reliable) results are as follows: When you speed up the rate at which the letters in the sequence occur, control (measured as number of correct transitions in a run of 100 or so letters) decreases. I would argue that this is because you can't perceive the sequence when it is going too fast; but it could also be argued that this is a reaction time limit. So I did the same experiment but changed the disturbance; instead of a new transition rule using the same letters I just put in a new set of letters as the disturbance so:

ABCDEABUYTSWUYTSW...

Now the subject doesn't have to perceive the sequence in order to control the ABCDE pattern; the disturbance can be detected at the configuration level. And sure enough, the ABCDE pattern can be controlled at about twice the rate relative to the maximum rate for sequence control. SO the limitation on sequence control was not due to a limitation in how quickly the subject can push the mouse button when the disturbance is detected; at the fast rates the disturbance to teh sequence is not even detected -- because you can't detect the sequence (this is subjectively obvious, by the way).

There are many variations on this kind of experiment; but the goal is to show
1) that complex variables (like sequences) can be controlled;
2) that control of these variables requires that they be perceived and
3) that what seems like an output limit (the speed with which you can move your fingers to type letters, for example) may not be what is limiting the ability to control higher order perceptual variables like producing typed sequences.

I have not even tried to model this situation yet; it would sure be nice to have a grad student or three.

Best   Rick

Date: Tue Mar 16, 1993 9:40 pm PST
Subject: Fowler, Turvey, etc.

[Avery Andrews 930317.1514]

I can't find my copy of Allan's WTP citation list, so don't know if the Fowler et. al. paper I mentioned last night is on it, but it does seem pretty clear that the specific criticism of PCT is just a repeat of that in Fowlder & Turvey 1978, and that the authors are suffering from a case of double-standardism, in that the same kind of problem that sinks PCT-without-reorganization sinks their concept of `coordinative structure' in exactly the same way, since `coordinative structures' don't differ in any way I can perceive from higher-level control systems.

In fact, F,T et al. may have done us a service by producing some specific examples of useful systems of this kind, since BCP is a bit sketch on the details. One of their examples is that of an airplane, where a joystick is connected to ailerons, wing-flaps, & a rudder, so that moving the stick to the left lifts the left flap, lowers the right and waggles the rudder to the left, etc., thereby folding three degrees of freedom into one.

Now if the physical systems aren't mechanically linked, the question arises of how this `folding' is to be achieved. F&T don't have much of anything to say about this, but one signal setting multiple reference levels via an output fan seems the obvious way to do it. Now if these reference levels are for forces, problems will arise if there is turbulence, etc., since different opposing forces on the various airfoils will result in different positions being obtained, with presumably suboptimal results. So the answer is to have the left-right `lean' of the joystick set position reference level for the three movable airfoils, & then we have a nice hierarchical control system doing something useful.

To finish off the whole person-airplane system, we can say that an error in heading (angular deviation of nose from target object) determines a reference level for change-of-heading, which determines a reference level for joystick position, which determines an assortment of reference levels for airfoil positions, which determines reference levels for forces applied airfoils.

Btw, while I certainly agree that Fowler & Turvey get PCT seriously wrong, I don't think they are particularly heinous by academic standards - people just aren't expected to work that hard to understand the other guy's positions. You just gotta put together a reply and politely straighten them out. It also seems pretty clear to me that F, T & Co have some kind of serious ideological obstacle to accepting control theory (Dionysius of Halicarnassus (c. 0 AD) tells us that ignorance will be right by chance some of the time, but only perversity will be wrong *all* the time), but that doesn't mean that there's no point in answering their arguments. People do occasionally change their minds, and then there are who knows how many uncommitteds wandering around out there.

People just don't care about the input-output model of behavior as much as Rick thinks they do. Indeed, I suspect that one of the reasons linguists don't spend much time in the psych lab is that the input-ouput, IV-DV stuff just seems stupid and boring to them, w.r.t. the things they are interested in, but that's the only thing that people seem to know how to do in labs.

Avery.Andrews@anu.edu.au

Date: Tue Mar 16, 1993 9:51 pm PST

Subject: Re: information

[From RIck Marken (930316.2100)] Avery Andrews (930317.1404) --

>Spending a bit more time on the information controversy, it seems
>pretty clear to me that Rick Marken & Bill Powers have run afoul
>of the bizarre fact that `information theory' has nothing to do with content.

So that's why it seemed to have nothing to say.

> Now knowing only p(t) or o(t), we
>can't recover d(t), but miraculously, knowing both, we can

Vat miracle? I already admitted that, when p(t) = d(t) + o(t) you can recover d(t) from p(t) if you know o(t) also -- it's called algebra (not information theory) and algebra is also pretty content free.

>What we have to do is figure out how the information can in some sense
>be there, and yet be uninformative.

Not a problem. If d(t) is in p(t) (which it is) and you don't have o(t) then the information about d(t) in p(t) is uninformative; about as informative as having hoeroglyphics and no rosetta stone. I already conceeded to Martin that if all he means is that d(t) is part of p(t) then, no problem, I admit it; I assume it. But it's not much good to know that (from the control system's point of view) because the control system has no access to o(t) -- which, by the way, is the variable that is presumably being generated from the information in p(t) about d(t) -- information that is only "informative" when o(t) is available (don't forget, this is a CIRCLE).

> Alternatively, we need to understand the sense in
>which it is there, and the sense in which it is not.

Hallelujah. That's what I've been trying to find out. My intuition is that the information is not there (even if d(t) is part of p(t)) since there is no way that the system can get to it.

>This requires that people actively *think* and try to make sense out of what
>other people are saying, rather than relentlessly defending their positions.

OK. Let's try this. If the fact that p(t) = d(t) + o(t) means that there is information in p(t) about d(t) then YES, THERE IS INFORMATION ABOUT d(t) IN p(t). So it's there but it is of absolutely no use to the system. So the fact that o(t) = -d(t) has nothing to do with the fact that there IS this precious information about d(t) in p(t) -- because the system can't use it. The relationship of the information in p(t) to the output o(t) is purely coincidental -- its like me reading the gettysburg address in hebrew. I can't read any hebrew but the information ("forescore and seven ...") is there, encoded in those funny square squiggles. I can "read" it because I have the gettyburg address memorized. My output o(t) would be seen as matching the input p(t) by a bilingual (hebrew-english) observer. But my output is not based on the information in p(t) -- even though it's there, for me, it's not.

So if admitting that there is uninformative information in p(t) will help settle this then I'll sign up to it; somehow I don't feel like I've given up much ground, though.

Best Rick

Date: Wed Mar 17, 1993 12:36 am PST
Subject: information

[Avery Andrews 930317.1830] RIck Marken (930316.2100)

>So if admitting that there is uninformative information in p(t) will
>help settle this then I'll sign up to it; somehow I don't feel like
>I've given up much ground, though.

I'm not sure whether you have to give up any ground at all - the issue as I understand it is whether so-called `information theory' (which maybe should have been called `channel capacity theory' or something like that' has some relevance to PCT. This debate can't even start until the content issue is set aside as the terminological kafuffle that it is. But the idea that channel capacity limitations have some relevance to the design and function of nervous systems seems eminently plausible to me.

Here's an argument something vaguely deserving of the name `information about d(t)' is present in p(t) and o(t). Suppose we add to our system two random noise generators, one into p(t), one into o(t), both downstream from where we are taking our measurements. Switching either of these generators on will clearly degrade our information about d(t), and switching both on will degrade it more than switching one on (to the same level of noisiness). The idea being if we can damage the info by injecting noise into a channel, it must in some sense be there. But in what sense is kinda mysterious, isn't it.

Avery.Andrews@anu.edu.au

Date: Wed Mar 17, 1993 7:18 am PST
Subject: Information, uncertainty,sensing disturbances

[From Bill Powers (930317.0700)]

Reorganization always entails disorganization. One of my statistical nonreliable rules of thumb is that whenever two (or more) opponents in an argument finally lose control and start shrieking and stomping away, agreement is just around the corner (if they stay in contact and don't kill each other). Reorganization always involves affect.

So now we have Rick Marken, Martin Taylor, and me all coming up with ideas about information theory and control, reaching a peak of frustration, and along come Cliff Joslyn and Avery Andrews -- our two coolest heads -- to drop in the crystals that will bring a new order into the whole thing. Now it's time for all of us to start saying, "OOOOH! I thought you meant ..."

I will summarize what today's posts have allowed me to understand.

First problem: the word information. Martin keeps saying "the system receives information -- that is, its uncertainty is reduced ---". He has never said, "The system receives information -- that is, it knows more about the source of the signal." So clearly, when Martin uses the term information, he is not talking about finding out something about the source of the information (semantics). He is talking about the ability to construct a perception with the least noise in it (statistics).

This whole aspect of the argument with Rick, therefore, hinges entirely on the use of a single word, information, as if it had a single meaning, which it does not. If Martin would simply drop the use of this ambiguous term, and substitute the proper usage of "uncertainty" or "noise" in its place, Rick would have no further basis for arguing with him. Or Rick, of course, could use a different word.

Having said that:

The combination of Cliff's and Avery's posts suddenly showed me that my own view has been too narrow. It is perfectly possible for a _hierarchical_ control system to deduce the actual state of a disturbance of one lower-level control system. I described the method myself, yesterday, without realizing that I was exemplifying the very kind of higher-level system that is needed. Now I realize that Wayne Hershberger came within a hair's breadth of making exactly the same proposal.

The method, once you admit the possibility, is obvious. The lower-level system acts by sending reference signals to still-lower systems. Those lower systems respond by matching a perceptual signal representing the action to the reference signal. If the output represented by the perceptual signal is a reasonable representation of the action, then knowledge of the system's own output is available in that perceptual signal.

All that remains is to form a new perceptual function with two inputs: one is the perception of the lower system's output, and the other is a copy of the lower system's perceptual signal. Subtracting the output representation from the perceptual signal yields a new perceptual signal that is proportional to the total disturbance. This signal does not do the system being directly disturbed any good, but it promises the possibility that the disturbing variable itself might be brought under control. All that would remain to construct would be a comparator and an output function that can affect the source of the disturbance. Note the fascinating possibility here. The system as a whole has no sensor capable of detecting the disturbing variable directly. Yet by the method just outlined, that disturbing variable can be represented as a perceptual signal, and if an output can be found that affects that perceptual signal, the DEDUCED disturbing variable can in fact be controlled. To the extent that the deduction is objectively correct, the external disturbing variable will also be controlled. Rick, you should have no problem modeling this in your spreadsheet.

The signal representing the output of the system may not measure the actual effect of that output at the point where the effect cancels the disturbance. For example, the environment might insert one or two integrations (as when masses are involved). In that case, the perceptual function where the disturbance is perceived will have to apply that same function to the raw perception of the output, before subtracting from the lower perceptual signal. The nature of the required function is not known _a priori_ to the organism. Thus one may have to LEARN to perceive the disturbance -- that is, to find a perceptual function that will give the most consistent results. This connects with Martin's concept of information as something conditioned by the form of the receiver. The receiver can learn to minimize its uncertainty -- that is, the noise level of the resulting perceptual signal.

It is probably not possible for a single pair of systems, one of lower level and one of higher level, to find a unique representation of the state of the disturbing variable. It is certainly not possible to find a unique representation if more than one disturbing variable is acting at the same time. However, multiple control systems at two levels should be able to sort out multiple disturbing variables if each disturbing variable has a slightly different effect on each lower-level controlled variable. At least it will be possible to represent the environment as containing a set of independent disturbing variables, even through the elements of this set may be further subdivisible.

I think that out of our mutual tumult there has come an idea of extraordinary importance. This idea could easily brush aside my feeble attempts to characterize levels of control. The process of deducing disturbances on the basis of sensed output and lower-level controlled perceptions may be the key to formation of the hierarchy of control. It is _necessary_ that lower-level control systems be present and working for a remote disturbance to be deduced and then put under control -- but that is not _sufficient_. The higher-level control system must also discover a transformation of the perception of the output that, when added to the perception already controlled at the lower level, will yield a controllable perception of higher level. That transformation is a model of the part of the environment that lies between the output and the lower-level controlled variable.

It is still probably true that the higher-level perceptions found and put under control in this way have an unprovable connection to the external world. There may be alternative interpretations that are equally feasable and lead equally to valid control processes of higher level. But my impression is that the connection between the perceived world at higher levels and the external reality now has a much more systematic basis. There is now a rationale by which a system can reach deeper and deeper into the reality beyond its receptors and discover -- or at least construct -- new aspects of that reality that can be brought under control.

Note: Bringing disturbances under direct control naturally makes the operation of lower-level systems more reliable. One might then say, if the disturbances are under control, why do we still need any lower-level control systems? Why can't they be S-R systems in the old style?

The answer, of course, is that the very deduction of the state of the disturbance depends on the presence of a working lower-level control system. That control system is all that creates an output that mirrors the state of the disturbing variable, give or take a systematic transformation. There is no other way to perceive the state of the disturbing variable (if we have sensors for the disturbing variable, this approach is unnecessary). This is a very feedbacky answer, and it helps to reassure me that we are on the right track.

Best, Bill P.

Date: Wed Mar 17, 1993 9:11 am PST
Subject: Re: information

[Allan Randall (930317.1030 EST)]

Before continuing too far with this discussion, I'd like to lay out some points that I propose we all agree to first, so as to minimise the terminological confusion that has been rampant in this discussion so far.

I think we are all realizing that we must be sure that we agree on a definition of "disturbance" before we continue this discussion - otherwise we'll just be talking past each other. So, is everybody agreed on the following defintion?:

disturbance: the total sum environmental influence on the CEV

This is my understanding of the word and what it means, and I think it is what is indicated by the usual PCT diagrams. The other possible definition, is to talk about "perceived disturbance," if I may call it that. This is the sum total disturbance to the CEV, and thus includes the output of the control system. It is easy to confuse the two, since the "disturbance" is seen from an outsider's point of view and "perceived disturbance" is seen from the control system's point of view. I think we can agree that both are reasonable uses of the everyday word "disturbance." But "disturbance" in this discussion should refer to the *external* environmental influences, completely separated out from the output of the control system itself. Agreed? Are we also agreed that this disturbance, while defined in this external point of view, is nonetheless defined in terms of the CEV, which is defined according to the internal point of view? This seems to be the meaning of disturbance as it appears in most of the PCT diagrams: it inputs into the CEV (defined internally) but excludes the output of the control system (defined externally).

Shall we also agree that the hypothetical entities out there in the universe that actually cause the disturbance are to be called the "disturbing variables"?

The problem with all this, and something that must be addressed, is this: At what point in time, if any, do we include the effects of *past* output as part of the *current* environmental disturbance? Once the control system outputs to the environment, it can become quite intractable to isolate the environmental influences from the past output of the control system. Even our hypothetical external observer would not be likely to make such an absolute separation between control system output and environmental influences. The more time that goes by, the less tractable it is to separate the two. If the environment is chaotic, as our universe is, then the trajectory of the disturbing variables in their phase space will exponentially diverge with even the tiniest deviation in the output of the control system. Because of quantum effects, it will at some point become impossible, even in principle, for our external observer to separate the disturbing variables from the past output of the control system. At this point, at the latest, the information to do the separation is truly lost and I think we should agree that the effects of the past output be included in the disturbing variables. The other extreme would be to include *all* past output (up until a single iteration ago, or some tiny dt) in the current disturbing variables. I would find it preferable, however, to recognize that this separation can be to some degree arbitrary. So long as we realize that there are external "disturbing variables" that can, for some arbitrary time window, be considered external to the organism.

Now we need to agree on a working definition of "information." Can everyone agree that if, by making use of B, it is possible to describe A with fewer bits, then B contains information about A? In this context, the percept P contains information about disturbance D if using P would allow a more compact description of D (with fewer bits) than not using P. This is as opposed to the complete reconstruction of D from P, which should not be required to say that P "has information about" D. In other words, "having information about" does not mean having *complete* information.

I think we should also decide to stop using the term "information content" and "negentropy." These terms tend to be endlessly confusing, as seen in my discussion with Bill Powers, and they are not necessary. Instead, we will talk in terms of "amount of information," "number of bits," "entropy," "information loss and gain," and similar terms.

Are we also agreed that the reference signal can be considered, for the purposes of this discussion, to be constant?

Are we also agreed that the *output*, if not the percept, contains information about D (however we end up defining the time window of the disturbing variables)?

I think these are things we need to agree on. If anyone disagrees with any of these points, then THAT argument will have to be settled before the current debate will go anywhere.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Wed Mar 17, 1993 10:06 am PST
Subject: Uncommon ground

[From Rick Marken (930317.0800)] Avery Andrews (930317.1514) --

>People just don't care about the input-output model of behavior as
>much as Rick thinks they do.

Multi-millions of dollars are spent in the US (and probably Australia too) in support of behavioral science research (psych, sociology, econ, poli sci, etc) where the data is collected and analyzed in the context of the general linear model; multiple regression, ANOVA, etc. I bet few of these people would consciously say "I assume that the basic model underlying behavior is a cause-effect model" but they sure ACT like this is what they assume-- and big bucks are being spent in tacit support of this assumption. I believe it is important to know that this is the model that behaviooral scientists are "controlling for" -- consciously or not -- because I am sure that it is the reason why PCT -- after, what, 30 or so years on the scene -- has made virtually NO headway in the behavioral sciences. Its either that PCT is just a stupid model and all the behavioral scientists have been smart enough to notice that (but there are some other obviously stupid "models" running around in the behavioral sciences and, nevertheless, they get a lot of attention) or it is because of active resistence. I think PCT has made no headway because there IS active resistence and I think the underlying reason for this resistence is that PCT is a disturbance to the assumption in the behavioral sciences of a cause-effect model.

>Indeed, I suspect that one of the reasons
>linguists don't spend much time in the psych lab is that the
>input-ouput, IV-DV stuff just seems stupid and boring to them, w.r.t.

>the things they are interested in, but that's the only thing that people
>seem to know how to do in labs.

Well, I don't know if I would brag about not having any model at all. Just observing is very genteel and all -- but unless you try to predict and explain what you see with a model, what have you learned? I think that linguists do have implicit models -- cause effect models. If they don't want to test them then that's their problem. I don't care if they test them in labs or in the real world; but I don't think you have much of a science unless you test models. If linguists don't know how to do anything other than IV-DV research when they do it it's because their basic (unconscious model) is cause-effect. If they don't go into the lab to do IV-DV research it must be because they don't like the implications of their own models. If their models were not cause-effect -- if, for example, they were control of input models -- I'm sure these bright folks would have noticed very quickly that there is an alternative to the IV-DV approach to research and they would have very quickly understood research based on testing for controlled variables.

Avery Andrews (930317.1830) --

>Suppose we add to our system two random noise generators, one into p(t), one
>into o(t), both downstream from where we are taking our measurements.
>Switching either of these generators on will clearly degrade our
>information about d(t),

Wait a minute. What do you mean "clearly degrade"? I thought we finally agreed that d(t) is "information" only in the sense that it is part of the perceptual signal. When you add noise to p(t) you now have p(t) = d(t)+o(t)+ e(t) where e(t) is the noise. Notice that d(t) is still part of p(t) in all it's glory and o(t) will be proportional to d(t) + e(t). Now all you can find out from p(t) by knowing o(t) is the sum, d(t) + e(t) which means that information about d(t) is not degraded -- it is eliminated (in the sense of usable information). But the information about d(t) -- in the sense in which we agreed that the information about d(t) exists -- is still there.

>The idea being if we can
>damage the info by injecting noise into a channel, it must in some
>sense be there. But in what sense is kinda mysterious, isn't it.

Well, now you've confused things completely. How do you measure the amount by which the information (about d(t) I presume) is "damaged" by injecting noise? We seem to be back to square one (despite Bill P.'s diligent attempts at reconciliation this morning -- I'll get to that in a second).

You are saying that the information ABOUT d(t) is DAMAGED by addition of noise, e(t), to the perceptual signal. This is completely inconsistent with my understanding of what we had agreed to be the information available about d(t) in p(t). I thought we agreed 1) that information about d(t) is available in p(t) simply because p(t) = o(t)+d(t) 2) you can recover this information if you know o(t) and 3) since the control system itself does not have access to information about its own o(t) [Bill's proposal this morning requires multiple control systems, some of which have, as input, the o(t) of other control systems] the information about d(t) that is in p(t) in not informative -- IT MIGHT AS WELL NOT EVEN BE THERE BECAUSE THE CONTROL SYSTEM CAN'T (AND DOESN'T) USE IT.

Now you are saying that information about d(t) is degraded by addition of e(t) to p(t). If you go by meaning 2) above of inform- ation then the best you can get out of p(t) is d(t) + e(t) -- in other words a signal that is the sum of two signals. So, unless you can now get a hold of e(t), you are in the same position as you are given definition 3) of information -- the information about d(t) in d(t) + e(t) is uninformative.

Bill Powers (930317.0700) --

>Now it's time for all of us to start saying, "OOOOH! I thought you meant ..."

Well, after Avery's last couple of posts its clear to me that what I thought the "opposition" meant by information is what they mean by information. Avery clearly believes (see above) that there is USABLE information about d(t) in p(t)-- otherwise, why would noise "degrade" that information.

>The combination of Cliff's and Avery's posts suddenly showed me
>that my own view has been too narrow. It is perfectly possible
>for a _hierarchical_ control system to deduce the actual state of
>a disturbance of one lower-level control system.

This was a very interesting discussion -- and I think it is definitely worth more development. I never doubted that a hierarchical system could deduce d(t); since it could get information about p(t) and o(t) (well, actually, as you said, the error signal -- the function relating output to input -- the physical world tranformation -- would have to be a guess). But I don't think that is what Avery or Martin have in mind when they talk about the information in p(t); although I think it would be a nice plateau on which to settle their reorganization.

Based on Avery's comments above, it seems to me that there is still some cause-effect thinking about control system lurking around in the background. I think that some people still think that, in a control system, there is USABLE information about d(t) in p(t). I have conceded that information about d(t) EXISTS in p(t). But I am arguing that that is a matter of supreme irrelevance to the system that is controlling p(t) because there is no way (or need) for that system (in and of itself) to get the information about d(t) that is in p(t) -- because it has no way (or need) to know o(t). So there IS information about d(t) in p(t) but, as I said about a week ago to Allan Randall -- SO WHAT? It is completely invisible to the control system; its just like it wasn't there. It can be MADE visible by the clever hierarchical combination of several control systems (that's why clever hierarchical control systems like us can do science).

I'm pushing on this because I want Avery to appreciate the magnitude of what Uncle Bill hath wrought.

Best Rick

Date: Wed Mar 17, 1993 11:42 am PST
Subject: defining information

[From Rick Marken (930317.1000)] Allan Randall (930317.1030 EST) --

>disturbance: the total sum environmental influence on the CEV

Check.

>"perceived disturbance," if I may call
>it that. This is the sum total disturbance to the CEV, and thus
>includes the output of the control system.

I don't like that term -- it gives the impression that the disturbance itself CAN be perceived. How about calling it what it is -- the CEV (or the perceptual variable) -- after all, that's exactly what you described above:

CEV(t) = d(t) + o(t)

Calling this the "perceived disturbance" may be the simple linguistic basis for all the misunderstandings. There is no such thing as a "pereived disturbance" in a control system; just that good ol' perceptual input variable or, if you like, the CEV.

>It is easy to confuse the two, since the "disturbance" is seen from an
>outsider's point of view and "perceived disturbance" is seen from the
>control system's point of view.

It was easy for you to confuse them because you used the word "disturbance" in both cases. Very unfortunate.

>I think we can agree that both are reasonable uses of the
>everyday word "disturbance."

Absolutely not! Disturbance refers to something that "interferes with" something else; in the equation CEV(t) = o(t) + d(t), d(t) can be thought of as adding "interference" to the effects of o(t) on the CEV. But variations in CEV(t) are NOT interfering with anything; they just EXIST. They are not interfering, for example, with the reference signal, r(t) -- the variations in CEV(t) may not match the variations in r(t) but they don't affect those variations (that's why r(t) is considered an INDEPENDENT VARIABLE by the way). In fact, in a closed loop system, r(t) might be though of as a "disturbance" to the natural variations in the CEV. After all, the closed loop makes variations in CEV(t) match variations in r(t).

>But "disturbance" in this discussion
>should refer to the *external* environmental influences, completely
>separated out from the output of the control system itself. Agreed?

You betcha! I agree.

>Are we also agreed that this disturbance, while defined in this
>external point of view, is nonetheless defined in terms of the
>CEV, which is defined according to the internal point of view?

Say what? Why not just say CEV(t) = d(t) + o(t). If that's what the above sentence means then I agree with it.

>This seems to be the meaning of disturbance as it appears in most of the
>PCT diagrams: it inputs into the CEV (defined internally) but excludes
>the output of the control system (defined externally).

The CEV refers to an external variable; p(t) refers to the perceptual representation of CEV(t). For simplicity we usually assume that p(t)= k(CEV(t)); this works fine in our tracking tasks. CEV is the actual position of the line on the screen; p is the perception of that position. The disturbance and output variables are numbers in the computer (thus, they are external to the behaving system); the d(t) numbers are generated independently by the computer; the o(t) numbers are a result of the subject's joystick movememts.

>Shall we also agree that the hypothetical entities out there in the
>universe that actually cause the disturbance are to be called the
>"disturbing variables"?

Yah. Shure.

>The problem with all this, and something that must be addressed, is
>this: At what point in time, if any, do we include the effects of
>*past* output as part of the *current* environmental disturbance?

This is only a problem for those who think of control loops sequentially. In the formula CEV(t) = d(t) + o(t) it is always the current value of the output (occuring at time t) that is combined with the current value of the disturbance. The fact that o(t) might be the result of processes occuring earlier in time is of absolutely no consequence. The physical fact of the matter is that the current state of the CEV is determined, simultaneously, by the current value of the disturbance and the current value of the output.

>Because of quantum effects, it will at some point
>become impossible, even in principle, for our external observer
>to separate the disturbing variables from the past output of the control system.

It's really not necessary to fly before you can walk; we're not even close to the point were we need to worry about such esoteric phenomena. Maybe you guys are having a problem with this because it is too simple. Don't worry, once you get the basics you can start pushing the theoretical envelope.

>Now we need to agree on a working definition of "information." Can
>everyone agree that if, by making use of B, it is possible to describe
>A with fewer bits, then B contains information about A? In this context,
>the percept P contains information about disturbance D if using P
>would allow a more compact description of D (with fewer bits) than
>not using P.

This is a bit fuzzy for me. How do I know how many bits I need to describe D without P? If D is digitized and there are 100 8 bit samples then do I need 800 bits to characterize D? Without knowing P I could probably think of some compression schemes all on my own. Would P improve my ability to compress D? What are the coding rules that I can use? I guess the answer to the above question is "no" from me; I need it to be a bit simpler. Why not just measure information in the old fashioned way --

H = log2 (variance of D).

Then we can measure the information transmitted by P about D as the proportion of variance of D accounted for by variance in P (with the appropriate log2 transformations). How about that?

>This is as opposed to the complete reconstruction of
>D from P, which should not be required to say that P "has information
>about" D. In other words, "having information about" does not mean
>having *complete* information.

I never assumed that "having information about" means having *complete* information. Information (whatever it is) is a VARIABLE; all I've been saying is that the information (however it is measured) about d(t) that exists in p(t) is very close to zero as long as you have no information about o(t).

>Are we also agreed that the reference signal can be considered, for
>the purposes of this discussion, to be constant?

No. It can be a constant or a variable. This should not influence the our conclusions about the infornation in p(t) -- or CEV(t).

>Are we also agreed that the *output*, if not the percept, contains
>information about D (however we end up defining the time window of
>the disturbing variables)?

Given my definition of information transmittion (variance in one signal accounted for by variance in another signal) the *output* (o(t)) contains nearly 100% of the information about D. So, yes, we are definitelely agreed on this one.

>I think these are things we need to agree on. If anyone disagrees
>with any of these points, then THAT argument will have to be settled
>before the current debate will go anywhere.

You can see where I agree and disagree above. I think the really big problems are 1) calling the CEV the "perceived disturbance" -- big mistake and 2) how to measure the information about d(t) in p(t) and how to measure the information transmitted about d(t) by p(t).

I'm ready to negotiate.

Best Rick

Date: Wed Mar 17, 1993 12:20 pm PST
Subject: Disturbances that don't disturb

[From Bill Powers (930317.1200)] Allan Randall (930317.1030 EST) --

Looks like we are getting to the nub of the definitional problem here.

>So, is everybody agreed on the following definition?:

>disturbance: the total sum environmental influence on the CEV.

No, and the reason we don't agree, once understood, should clear everything up. The word "influence" is ambiguous for the very same reason that "disturbance" is ambiguous.

Suppose we have two objects connected by a spring. Initially, the spring is just slack.

```
          A --- /////////// --- B
```

By moving A to the left, we can stretch the spring, which in turn will apply a force to B.

Let's define the CEV of a control system as the _position_ of B along the x axis. Let's also define the _disturbing variable_ as the _position_ of A along the x-axis. We can now add an object C, which is the output of the control system, also connected to B (the CEV) by a spring:

```
          A --- /////////// --- B --- \\\\\\\\\\ --- C
```

The output action consists of moving the object C along the x axis, so we say that the output is measured by the _position_ of C.

Now let's get the language straight. What are "the two influences on B?" They are A and C, the positions of the two objects. So here we are defining an "influence" as "something capable of affecting something else." If we know the spring constant, we can say that a movement of A of 1 cm will initially result in k units of force applied to B, and similarly for C.

Influence, however, has another sense. We can ask, "How much influence do the positions of A and C have on B?" Now the problem becomes clearer. If A moves 1 cm to the left while C simultaneously moves one cm to the right, and the springs are identical, the answer is NONE. Neither A nor C has, in fact, any influence on B _in terms of an effect on B_. The total force acting on B will remain zero, and B will not move.

A and C are classed as being among the objects whose position is capable of influencing or disturbing B. In normal parlance, we would class them as influences or disturbances. But that does not tell us whether, in fact, changing the position of either A or C will produce any actual change in B. So we can't say _a priori_ whether an influence will have an influence, or whether a disturbance will cause a disturbance!

When we draw a disturbing variable in a control-system diagram, we are drawing something analogous to A. The output is analogous to C.

```
     Disturbing
      Variable ----////////--- CEV ---- \\\\\\\\ --- Output
```

Without knowing anything about the state of the CEV, we can specify the state of the disturbing variable. If the units involved are distances, we can specify the location of the disturbing variable. If we observe the system for a while with the disturbing variable in the position that has no effect on the CEV, then suddenly move the disturbing variable to a new position and keep it there, we may well observe something like this:

****

```
Disturbing variable

*********************

-------------------------------------------------------------
                       *
                        *
CEV                      *
                           *
***********************       *          ************
                                  *  *

-------------------------------------------------------------

                                  *  *
                              *          *************
                           *
(Opposing) Output        *
                        *
                       *
**********************
--------------------------------------------------------------
```

It's clear that the behavior of the CEV is not like the behavior of the disturbing variable. Neither is it like the behavior of the control system's output. We would certainly say that the disturbing variable influences or disturbs the CEV, but we can't say that the disturbance or influence that actually occurs reflects the behavior of the disturbing variable.

This is where we have been sliding past each other. It makes no sense to say that the control system's perceptual signal contains no information about the CEV. That is why it has seemed so self-evidently true to you that a disturbance (meaning an actual change in the CEV) conveys information to the control system -- and so stupid of us to claim that it does not.

But now interpret "disturbance" to mean "state of the disturbing variable." The diagrams above show (roughly) the behavior of the system with an integrating control system. Before the change in the disturbing variable, the CEV had a value of zero. Some time after the change, the CEV is again approaching a value of zero. But the states of the disturbing variable and the output are quite different between these two times. This is why we have claimed that the control system gets no information about the disturbing variable: the CEV itself doesn't reflect the state of the disturbing variable. Knowing that the CEV has a value of zero tells you nothing about the value of the disturbing variable.

Remember that in an ECS, the perceptual function does not receive information about the output of the system. It senses ONLY the CEV. Thus a given state of the CEV is always caused to an unknowable degree by the state of the output, the rest being due to the state of the disturbing variable. Not knowing the state of the output, the perceptual function has no way to know the state of the disturbing variable, either. For any value of the CEV within its possible range, the disturbing variable might have any value within its possible range. There is total uncertainty about the state of the disturbing variable.

This may have been overkill, but if it leads to mutual understanding it was worth the effort.

Best, Bill P.

Date: Wed Mar 17, 1993 1:10 pm PST
Subject: Re: information

[Martin Taylor 930317 14:20] Avery Andrews 930317.1830

I said I wasn't going to post any more publicly about information theory for a while. But it's a bit like a New Year's resolution. Hard to renounce the fun things in life for too long.

>I'm not sure whether you [Rick] have to give up any ground at all.

The concept of "giving up ground" suggests a contest where there is a winner and a loser. The metaphor seems very wrong. When someone has an insight, that person wins, and if the insight can be communicated, other people win as well. If they change their previous views (as I suppose they must if they have an insight) they still win. They haven't "given up ground" by changing their opinion unless maintenance of that opinion was more important than the truth of that opinion.

>the issue as I understand it is whether so-called `information theory' (which
>maybe should have been called `channel capacity theory' or something
>like that' has some relevance to PCT.

I think the major reason why Information Theory has got a bad name is that it is confused with channel capacities. Shannon, for sure, developed it in that context, but the metaphor just makes no coherent sense other than from the Engineer's viewpoint (Bob Clark). When you take the position on which both relativity theory and PCT are built (that you can only do something about what you can observe) there are at least three different information numbers relating to a "channel," and one of them really does deal with content.

One thing to keep in mind is that information is a differential quantity. It makes no more sense to talk about the information content of something (an event, a message) than it does to talk about the "velocity content" of an object. The conserved quantity is sometimes called "uncertainty" and is a function of the distribution of subjective probabilities of some state or event. The absolute uncertainty usually depends on the resolution of measurement, so it is not a unique number, but if the resolution is constant, the uncertainty after an event happens can be legitimately subtracted from the incertainty before the event. The difference is the information provided by the event about the uncertain situation. It can be positive or negative.

The three uncertainties that could be of interest in a communication channel are (1) the transmitter's uncertainty derived from the probability of sending any of the possible symbols or signal values; (2) the receiver's uncertainty about what symbols or signal values will be sent; (3) the receivers uncertainty about some aspect of the world that might be affected by the symbols or signals sent. (It is easy to add more, such as the receiver's uncertainty about what will be received, and there are indefinitely many aspects of the world that could be considered under type 3). None of these represent the classic "channel capacity" except in limiting cases in which the prior uncertainty is that all symbols (signal values) of which the channel is capable are equiprobable at every sampling moment. (A samplng moment is defined by the electronic or physical bandwidth of the channel, and limits the rate at which independent information can be obtained through the channel).

It is a mistake to see the circuits of PCT as "channels" with "capacities" and to try to derive the behaviour of the circuit from those capacities. The capacities impose limits, for sure, but they don't necessarily indicate rates of information acquisition at any point in the system.

When you talk of information "about" something, you are talking of a type 3 number. When you talk of "channel capacity theory" you are talking about the relation between type 1 and type 2 numbers as seen by a third party. Channel capacity numbers do not touch the idea of semantic content. "Information about" does.

By the way, these numbers are invented for the purposes of this posting. You won't find them in any published stuff on information theory that I know of.

I'll shut up again, until I next break my resolution.

Martin

Date: Wed Mar 17, 1993 1:20 pm PST
Subject: information, modelling

[Avery Andrews 930318.0700] Avery Andrews (930317.1514)

>Well, I don't know if I would brag about not having any model
>at all. Just observing is very genteel and all -- but unless you

There are times when a good description is a lot better than a lousy model, Newton & Maxwell providing famous examples, pre-molecular genetics another. This may or may not be true for linguistics one (I suspect that the time for modelling is a lot closer than it used to be, but it does happen occasionally.

>>Suppose we add to our system
>>two random noise generators, one into p(t), one into o(t), both
>>downstream from where we are taking our measurements. Switching
>>either of these generators on will clearly degrade our information
>>about d(t),

>Wait a minute. What do you mean "clearly degrade"? I thought we
>finally agreed that d(t) is "information" only in the sense that
>it is part of the perceptual signal. When you add noise to p(t)

Due to noise in the channels, our reconstruction of d(t) from p(t) and o(t) will always be imperfect - that is, it wil differ from the actual d(t) in a way that one might quantify by integrating the square of the difference between the actual and reconstructed d(t) or something like that. Injecting more noise into the channel increases the divergence, e.g. makes our reconstructed approximation to d(t) less like the original one. I don't recall agreeing to anything to the effect that d(t) was or wasn't information.

> Now all you can find out

>from p(t) by knowing o(t) is the sum, d(t) + e(t) which means
>that information about d(t) is not degraded

If you can get perfect information about e(t), which is unrealistic.

In general, you still seem to be ignoring the point that there are two senses of the term `information' floating around, the `technical' sense of information in `information theory' that neither of us knows anything much about, and the ordinary sense, which nobody has any real theories about (but some people, like Fred Dretske, do have interesting ideas). In the ordinary sense p(t) alone does not provide info about d(t), but there seems to be a more mysterious and technical sense in which the information is `there'. But sorting this out would require us to actually learn this theory, or at least allow Martin and Allan to explain the high points to us, without firing off broadsides of objections to everything they say based on terminological misunderstanding.

Avery.Andrews@anu.edu.au

Date: Wed Mar 17, 1993 1:51 pm PST
Subject: research proposals

[From Richard Thurman (930317.1400)] Rick Marken (930316.2000)

Rick your research proposal looks interesting!

>I have already set up and done the basic experiment and it works like
>a charm -- but there are many variations and some modeling work to
>do, and that's where the PhD thesis student would come in.

I may have a (partial) solution for you concerning people to help you setup and run the experiment(s). This summer Dr. Tom Hancock will be a visiting professor at this lab. He is under contract to research "adaptive feedback based on Perceptual Control Theory" from mid April to the end of August. I think that the kinds of experiments and and modeling you are describing would fit in with what he has in mind.

In addition I would be willing to help on an as needed basis. As you know (from private posts) I am interested in pursuing PCT research further. While I understand that you are actively seeking grad student help, please don't discount people who are past that whole mess. I for one am willing to assume the lowly position of research assistant if it will help me get a handle on this PCT stuff. If it will help me reorganize faster to get up close and personal with the data then I'm willing to do it. (I can't stand the flip-flop perceptions of going from a cause-effect orientation to a negative feedback closed-loop orientation. Just when I think I'm getting it right I realize that I was thinking about the situation from the incorrect perspective.)

The only stipulation I think I would need to put on the Lab doing this type of research is that it needs to be couched in terms of training and education. That is, any technical reports or published papers would need to have a training spin to them.

Interested? Rich

Richard Thurman
Air Force Armstrong Lab   Aircrew Training Research Division
BLDG. 558   Williams AFB AZ. 85240-6457   (602) 988-6561
Internet: Thurman%HRLOT1.Decnet@EIS.Brooks.AF.Mil
or   Thurman@192.207.189.65

Date: Wed Mar 17, 1993 3:19 pm PST
Subject: infomystery, research proposal

[From (the indefatiguable) Rick Marken (930317.1400)]

Avery Andrews (930318.0700)--

>In general, you still seem to be ignoring the point that there are two
>senses of the term `information' floating around, the `technical'
>sense of information in `information theory' ...
>and the ordinary sense, which nobody has any real theories about

> In the ordinary sense p(t) alone does not provide info
>about d(t), but there seems to be a more mysterious and technical
>sense in which the information is `there'.

I think this is a bluff. I have asked for an explanation of how, in the "techical" sense, information about d(t) is "there" in p(t). I have not heard anything technically convincing. In a private post, Martin Taylor gave the following techical definition of information:

>the magnitude of change in uncertainty, and uncertainty is a
>property of a subjective probability distribution at a location.

Based on this definition I proposed the following approach to measuring the information about d(t) communicated in p(t):

>I would say that the location of
>uncertainty about the disturbance is "inside the control system".
>I'll assume that inside the control system is a subjective probability
>distribution characterizing the probability that a particular value
>of the disturbance will occur at any time, t. So at each time, t,
>the subjective probability distribution at t defines the control
>system's uncertainly about which value of d(t) will occur at that
>time. This uncertainty is changed (hopefully reduced) by the
>the perception at time t (p(t)) which presumably contains information
>about d(t).

>It should be a pretty easy matter to make some
>assumptions about the subjective probability distribution of d(t) at
>each instant (my guess is that it would always be normal about the
>mean expected disturbance value) and then compute the gain in
>information that results from being given p(t) at each instant.

I think, from here, I could write a program to compute the information about d(t) communicated to the destination by p(t). But there are some details still needed -- I could make some reasonable assumptions, but one man's reasonable assumptions are another's "see, you don't understand information theory". So I am asking the information theory experts to bless, improve or reject (with clearly explained reasons why) the approach to measuring the information about d(t) in p(t) that I described above. OK.

Richard Thurman (930317.1400) --

>I may have a (partial) solution for you concerning people to help you
>setup and run the experiment(s). This summer Dr. Tom Hancock will be
>a visiting professor at this lab. He is under contract to research
>"adaptive feedback based on Perceptual Control Theory" from mid April
>to the end of August. I think that the kinds of experiments and
>and modeling you are describing would fit in with what he has in mind.

>The only stipulation I think I would need to put on the Lab doing this
>type of research is that it needs to be couched in terms of training
>and education. That is, any technical reports or published papers would
>need to have a training spin to them.

>Interested?   Sold.

And I think the training emphasis would be great. The idea would be to show that training is largely a matter of learning which perceptions to control, not which "outputs" to generate.

Keep in touch on this. If we do it over CSG-L (instead of in private) maybe we can benefit from the advice of others.

Best   Rick

Date: Wed Mar 17, 1993 3:55 pm PST
Subject: mystery journal

[Avery Andrews 930318.0929]

Does anyone know what journal title might be abbreviated by:

  Autom. Rem. Cntrl.

(it's the `Rem' that I'm stuck on, of course).

This contains articles (late sixties and early seventies) by people such as Aizerman, Andreeva, Chernov and Litventsev, which according to Keslo, Holt, Kuger and Turvey (1980) (Tut. in Mot. Beh.) propose neural organizations which are supposed to achieve stabilization without perceptual control (in posture maintenance, tracking, etc.).

Hopefully somebody is confused, but we better check it out.

Avery.Andrews@anu.edu.au

Date: Wed Mar 17, 1993 4:12 pm PST
Subject: Back to Basics II

[From Rick Marken (930317.1530)]

Here are some replies to a private post from Martin that he said I could reply to on the net.
Rick,

>You always post the equation as p(t) = o(t) + d(t), but in discussion
>everyone, including you most of the time acknowledge that there is a
>not-well-known function relating the output to the effects on the CEV.
>If there weren't, wouldn't the cognitive outflow people be right?

This might be my fault; the letters p, o and d refer to variables, not functions, and the t in parenthesis is an index, not an operand. This equation just decribes physical reality in a compensatory tracking task (if we think of p as the number representing the CEV -- position of a line -- rather than the perception of that line; as I said in another post, there is very good reason to suspect that the perception of the CEV (in a tracking task) is linearly proportional to the physical measure of the CEV).

The output function is really the function that transforms the error signal into the output variable, o. So the output function, O, is o(t) = O(e(t)) and O is likely to be highly non-linear. This is indeed one reason why cognitive outflow models can't work -- but another reason is that the intended result of the cognitive outflow (the CEV) also depends on disturbances (actually, on their effects, d(t)).

I said:

>This is only a problem for those who think of control loops
>sequentially. In the formula CEV(t) = d(t) + o(t) it is
>always the current value of the output (occuring at time t) that
>is combined with the current value of the disturbance. The fact
>that o(t) might be the result of processes occuring earlier in time
>is of absolutely no consequence. The physical fact of the matter
>is that the current state of the CEV is determined, simultaneously,
>by the current value of the disturbance and the current value of the
>output.

Martin replies:

>This cannot be a correct interpretation of the equation. The CEV does
>not react instantaneously to output. The equation talks about signals
>which add to form a perceptual signal, or else it talks about physical
>effects that add to change the CEV and therefore the perceptual signal.
>Either way, the three terms in the equation must be of the same form,
>and that form is not the form of the output signal of the ECS.

>The equation should read p(t) = P sub t (F(o) + D(d)) where P sub t is
>function P evaluated at time t. F is some function of o, where o is the
>history of all output, D is some function of d, where d is the history
>of all disturbance. One can simplify this by saying simply that the
>"disturbance" is D, rather than D(d), since perceptual control doesn't
>care about the distinction. And if we reference p(t) to the CEV rather
>than to the perceptual signal, we can eliminate the function P, just

>evaluating F at time t. But we can't ignore F, as you and Bill have
>both pointed out today. Allan is saying that F is uncertain. So do you.
>He says past actions have side effects. That's not controversial in any
>version of PCT that I know.

This confusion may all result from my notation. Think of p(t) as the sequence of numbers (over time) that represent the position of the cursor on a computer screen, o(t) are the numbers coming from the joystick, d(t) are just smoothly varying random numbers. At any time t, p(t) = o(t) + d(t); that is just the way the physical situation is set up. I agree that o(t) at a particular time might be the response to a perceptual input from time t-tau. But what is currently on the screen, p(t) is always the simultaneous result of the current disturbance number, d(t) and output number, o(t).

So time delays certainly do exist in control loops -- but they are not pertinant to the question of whether or not there is information about d(t) in p(t). The fact of the matter is that p(t) is always a JOINT result of o(t) and d(t) (see Bill's earlier post today on disturbances). This is a VERY important point to understand. If we can't agree on this (that the perceptual input to a control system is at all times the simultaneous, joint result of both disturbance and output) then we are really thermodynamically isolated from each other.

Best Rick

Date: Wed Mar 17, 1993 5:39 pm PST
Subject: infomysteries

[Avery Andrews 930318.1121] Rick Marken (930317.1400)

I don't think I'm bluffing about anything (tho have done a bit of guessing, some of which was a bit wide of the mark, as Martin recently indicated). I'm with you in wanting to know how knowing p(t) lets you describe d(t) with fewer bits. I think my little story about noise generators is a sufficient basis for saying that information about d(t) is somehow present in p(t), but if you don't, fine. After all, I do say that it is present in `some mysterious sense' which means I don't claim to know what's going on, but only to suspect that something interesting is.

On a completely different note, I've just written a little implementation of a simple `coordinative structure' via perceptual control, vaguely inspired by the Abbs & Winstein work on lip movements. There are two little points, driven by thrusters through a medium. And there is a reference level for separation of the points, so that when they are too far apart the thrusters drive them together, & vice versa. The reference level is produced by a square wave generator. Finally either of the two points can be `frozen', whereupon the other goes further to compensate.

In the grab-bag of possible paper topics, perceptual control as an implementation theory of coordinate structures strikes me as having some promise.

Avery.Andrews@anu.edu.au

Date: Wed Mar 17, 1993 7:19 pm PST
Subject: Gone 'til Tuesday

[Allan Randall (930317.2150 EST)]

Hi. To all who are participating in the information theory debate, you might like to know I'll be away until Tuesday. Maybe by the time I get back Rick Marken will be trying to convince Bill Powers that there is tons and tons of information about the disturbance in the perceptual signal. Hopefully, pigs will also be flying by then. Hey, a lot can happen in five days. :-) :-)

-- Allan

Date: Wed Mar 17, 1993 8:54 pm PST
Subject: infomysteries

[From Rick Marken (930317.2030)]   Avery Andrews (930318.1121)--

>I don't think I'm bluffing about anything (tho have done a bit of
>guessing, some of which was a bit wide of the mark, as Martin recently
>indicated).

Sorry. I didn't mean to say that you, personally, were bluffing; I was crying out to the vast, faceless, nameless group of experts in "technical" information theory; I think a school of knowledge (like IT) -- even if it can provide no predictions (as Allan notes) -- should at least be able to tell me how to measure the main variable in its arsenal -- information. PCT tells us how to measure control, at least.

The program sounds fun, by the way.

>In the grab-bag of possible paper topics, perceptual control as an
>implementation theory of coordinate structures strikes me as having
>some promise.

Pardon another self-promotion but two such papers (describing simple control model implementations of coordinative strutures -- along with experimental tests with humans) have already been published (and reprinted in Mind Readings -- both in chapter 6 on Coordination). Both appeared in major journals; neither has recieved ANY attention from the coordinative structure crowd -- positive or negative. But, hey, they were published; sorry about having to sacrifice all those trees for nothing, though.

G'day   Rick

Date: Wed Mar 17, 1993 9:45 pm PST
Subject: Re: challenge

[Allan Randall (930317.1200 EST)]   Bill Powers (930315.0700)

This posting deals with the challenge, not the information in disturbance or definitional stuff - I will probably not have time to get back to that until I return from my trip early next week.

Okay, this challenge thing. I decided from your last response not to formally accept your challenge, since it seemed directed at Martin specifically, and not at myself or other Ashby-type information theorists. The reason is that you pretty clearly stated that information theory would have to supply a prediction that control would occur, from information theory and Ashby's diagrams alone. If this is your position, then I have no disagreement with you.

However, Martin has encouraged me to try one more time to arrive at a mutually agreeable form for the challenge, as he suggested I may be misinterpreting you. So here goes.

> >You seem to be admitting here that information theory might
> >have something useful to say about control systems.
>
> Insofar as information theory could predict the limits of
> performance given signals and signal-handling devices with
> certain characteristics and in a known organization, sure.

Hmmm, again maybe we have no argument. Anything that does what you just described sounds pretty darn relevant and applicable to me. Perhaps we just value different things. It would seem hard to believe that something that could tell you about limits of performance would not also be useful in designing a control system. Ashby's Law could be viewed as a statement about limits of performance. Statements about such limits can be quite fundamental. So if this is your position, then we differ only in the degree to which we think information theory is relevant. This is hardly a fundamental disagreement, and so the challenge is indeed not directed at me, but solely at Martin.

> >So I guess I still don't understand exactly where you stand. Is
> >information theory completely wrong-headed or is it correct,
> >but of little use to PCT?
>
> Information theory rests on definitions and mathematical
> manipulations...It's unlikely to be "incorrect" in those
> terms... I don't yet see how IT is
> actually linked in any rigorous way to specific physical
> situations.

Oops. Now you seem to question its validity again (at least as something that can be applied to physical situations). Is it valid to talk about information transmission in a control system as the mathematical measure called entropy? That is the question. If using information theory in this sense is not valid in the first place, then any "limits of performance" measures you get will be utterly useless. So I am *still* confused as to where you stand.

> The prediction I'm asking for is
> not how much control is required, but how much control there will
> be in the two situations. To use a theory to derive the fact that
> control will result from either arrangement means to make
> predictions by manipulations that follow the rules of the theory.

What information theory will actually tell you is that there is *more* control in the compensatory system. In fact, there is so much control going on in the compensatory system that it would be ludicrous to even suggest a real device or organism achieving it. Information theory could tell you that the compensatory system will do very poorly because it cannot be given the output capacity or the processing power it requires. This kind of prediction is what I think you have ruled out by saying:

> > ... from Ashby's diagrams + information theory, one cannot
> >predict what exactly R, the regulator, is doing. You cannot
> >predict that R is going to oppose the disturbance. Whether this
> >will meet your requirements for the challenge is the main point
> >I'd like clarified before accepting.
>
> If you stick with these conclusions, the challenge is unnecessary
> because you have agreed to my original claim. You are agreeing
> that information theory can't provide the predictions of behavior
> that control theory provides, but can only be applied once those
> predictions are known and verified.

Not quite. It can be applied before anything is verified. But it cannot be applied to predict that control *will* happen - only that it could (or could not).

So my version of the challenge would take Ashby's compensatory and error-driven control systems and, assuming they were both designed to control, make a prediction concerning which would control better. I would not be able to say that either system, from Ashby's diagrams alone, *would* control. Maybe they will both play "Mary Had A Little Lamb," and completely ignore their inputs, for all I know. But I *can* tell you which is more *capable* of control.

If this does not satisfy your requirements for the challenge, then I think we can all agree that you are specifically challenging Martin's derivation claim and *not* Ashby's quantification claim. However, I will do my version of the challenge all the same, as I think it could be useful. I'm just trying to determine here whether I can give you a formal acceptance or not.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc.   Toronto, ON

Date: Thu Mar 18, 1993 2:57 am PST
Subject: help!

[Hans Blom, 930317]

Back to square one. I'm afraid I have to admit that I don't understand what Rick Marken (930313.1000) called "the central insight of PCT". Bill Powers, please help me to understand. Rick says:

>What PCT says (in no uncertain terms) is that this assumption [that
>sensory input is an INDEPENDENT VARIABLE] is FALSE -- sensory input
>is NOT an independent variable in a living control system (ie. all
>living organisms) -- it is a DEPENDENT variable (the independent
>variable being the reference signal inside the organism). I think
>it is clear that, if PCT is right about this then the WHOLE edifice
>of behavioral science comes crashing down -- something up with which
>most behavioral scientists will unquestionably not put.

So sensory input is a dependent variable. I take this to mean that sensory input has no degrees of freedom of its own; its is dependent UPON something, and that something is the reference signal. The reference signal therefore determines the sensory input (perceptions). Is this the same as saying that the reference signal controls the perceptions? If not, what is the difference? Are ALL sensory inputs dependent? The following [Bill Powers (930314.0900)] says:

>>So how DOES information get into the organism from the
>>environment? Clearly, only through uncontrolled perceptions.

Originally, the title "Behavior: the control of perception" gave me the impression that all perceptions are controlled. Now I understand that there are also uncontrolled perceptions. Is it therefore "Behavior: the control of SOME perceptions"?

In control theory, we sometimes draw the following diagram of a (any) system:

```
       input    -----------
     ---------->|    1    |
         |      -----------  output
         ------>|    2    |---------->
                -----------     |
                |    3    |------
                -----------
                |    4    |
                -----------
```

A system can (this is pure logic) be thought to be split up into 4 compart- ments. In the diagram above, compartments 1 and 2 are called controllable, compartments 2 and 3 are called observable. Is that a fair summary of the above?

Now, if some perceptions are controlled and some are not, are there also intermediates like:

- some perceptions are sometimes controlled, but not at other times;
- some perceptions are partially or approximately controlled;
- some perceptions are controlled in some degrees of freedom (dimensions)
  but not in others?

And how is it the other way around? Is behavior FULLY in the service of the control of perceptions or could there also be behavior that is not?

Best,   Hans Blom

Date: Thu Mar 18, 1993 3:16 am PST
Subject: Re: Bohr's Philosophy of Physics

I'm sorry to have missed the previous discussion (Cziko & Randall) regarding Bohr, Einstein and Feynman (it takes so long to sift the voluminous material for this board that too much gets tossed prematurely).

I'm afraid I don't see how anyone could call Bohr's approach "invoking magic", when what he was calling for was the primacy of (verifiable) results of observations and calculations over tacit images of the nature of an underlying

"Reality". At root the problems of quantum mechanics are arguments over rival epistemologies and conflicting notions of what it is to make and interpret a scientific model. Murdoch's book, The Philosophy of Niels Bohr (Cambridge U Press, 1987) is the clearest treatment of the Bohr-Einstein debate that I have seen, casting Bohr as a pragmatist (in the style of Wm James, Mach, Schlick) and Einstein as a realist. If anyone was invoking unexplained supernatural entities or predilections, it was Einstein.

It is not the case that Bohr's philosophy of science condemns us to give up on looking for other sets of observables -- it merely forces us to keep the realm of scientific imagination and theory-construction (and religion and esthetics and politics and whim) separate from the means (measurement, computation) by which we evaluate particular models.

This division is necessary for intellectual hygene. It is just too easy to be seduced by ones' own vision of how "things might really be". The only strategy that keeps us honest is this verificationist, "show me" approach, where all elements in a model are either directly measured or computed from measured initial conditions. The alternative is a "possible worlds" approach, in which the imagination, liberated from specifics, is free to construct one ad hoc move after another, whenever an objection is lodged. If we're going to do that, we might as well become lit-crits and/or psychoanalytics (at least they get to talk about sex and their psyches while they go around in circles).

So much of contemporary mathematical physics (and the current wave of pop-physics pulps), having adopted a platonic-realist approach, no longer seriously attempts to connect theory with observation. One of the great intellectual tragedies of the late 20th century has been this infusion of platonic mysticism (following Godel, the later Carnap, and Tarski) into philosophy, the foundations of mathematics, physics, linguistics, and the cognitive sciences. We are still dealing with the wreckage.

Enough ranting for now, Let's get back to the subject (whatever it was).

Peter Cariani eplunix!peter@eddie.mit.edu

By now I should know better than to get involved with discussions like this. :)

Date: Thu Mar 18, 1993 4:10 am PST
Subject: Re: mystery journal

[Hans Blom, 930318] Avery Andrews 930318.0929

>Does anyone know what journal title might be abbreviated by:
>
> Autom. Rem. Cntrl.
>
>(it's the `Rem' that I'm stuck on, of course).
>
>This contains articles (late sixties and early seventies) by people such
>as Aizerman, Andreeva, Chernov and Litventsev, which according
>to Keslo, Holt, Kuger and Turvey (1980) (Tut. in Mot. Beh.) propose
>neural organizations whichm are supposed to achieve stabilization
>without perceptual control (in posture maintenance, tracking, etc.).

"Automation and Remote Control" is a Russian periodical that is also translated into English. It features Russian authors almost exclusively. Articles are often high-quality, though difficult to read because of the different jargon/terminology (it takes some time to get used to). It is not an obscure periodical; our Electrical Engineering library has it. I suppose you could find it at most universities that have an Electrical Engineering department.

Best, Hans

Date: Thu Mar 18, 1993 7:34 am PST
Subject: Feedback in the Gobi

[Martin Taylor 930318 10:00]

An interesting little item in today's paper about a social feedback system.

Apparently there is a battalion of Chinese soldiers in the Gobi desert whose only task is to keep the sand shovelled off a particular railway line. The only purpose of the railway line is to supply these same soldiers.

I leave it to others to draw obvious (or less obvious, perhaps) conclusions.

Martin

Date: Thu Mar 18, 1993 8:58 am PST
Subject: Feedback in the Gobi

Tom Baines [930318 10:25] RE: [Martin Taylor 930318 10:00]

Hey! So what? During peacetime ALL soldiers shovel something, and I'd much prefer sand to other things I've experienced.

Tom

Date: Thu Mar 18, 1993 9:41 am PST
Subject: Control of Music

[from Gary Cziko 930318.1648 GMT] Rick Marken and Richard Thurman:

I've been thinking about the proposed control of sequence experiment. But instead of flashing letters on a computer screen, I think it might be a better test to use a sequence of tones like "do-re-mi-fa-sol-fa-me-re-do."

I think this would be a better test for none other than the intuitive feeling that it would be easier to perceive disturbances to this auditory sequence tan to the visual one using letters. Other conditions could also be used: one in which a tone is played with a different timbre (which is a configuration of frequencies) and dynamics (loudness or intensity).

So if I understand Rick's thinking, he would predict that the subject could accurately push a key to restore the correct intensity ("dyamics") at a quite quick tempo, control timbre ("get that trumpet out of there, it's supposed to sound like a flute") at a slower tempo, and sequence at a still slower tempo.

Now that I've started to think about this, I can see music perception as a way of getting to quite high levels of perception--Key modulations as transitions, perhaps musical phrases as events, musical style (e.g., baroque, classical, romantic, modern) as category, etc.

But I nonetheless have the intuition that once the melody (sequence) is well known), a person could react to a wrong note as quickly as he or she could to a wrong loudness or wrong timbre. But I hope I'm wrong.--Gary

Date: Thu Mar 18, 1993 9:52 am PST
Subject: I-less in Gobi

[From: Bruce Nevin (Thu 930318 11:52:38)]

I assume that the only *current* purpose of the railroad in the Gobi is supplying the soldiers who maintain it, and that some other purpose is foreseen for it (or was at one time, allowing thus for bureaucratic inertia).

Much the same criticism could be made of recessive traits in genetics.

Something like an analog to recessive traits has been discussed in anthropology. A river-dwelling tribe shifts from a fishing and hunting economy to farming. Circumstances change, farming is a bust. Only a few cranky old people have bothered to retain knowledge of the "old ways" and they were always regarded as peculiar and marginal, but now young folks learn from them and take up fishing and hunting again.

Conversely, "primitive" tribes in South America that hitherto have been described as representing an earlier stage of human cultural evolution turn out (it is now recognized) to have gone from farming to hunting/gathering under the impact (sic) of European contact and conquest. Perhaps they did not reinvent the necessary primitive technologies entirely de novo if some aspects were learned and remembered by some individuals. And whether reinvented or recovered from hitherto oddballs with quirky hobbies, they modernized them e.g. with bright dyes on arrow shafts to make them easier to find. (This converse case in a recent issue of _New Scientist_.)

Aaargh! I wasn't going to put any time into CSG today, too much to do!

Bruce   bn@bbn.com

Date: Thu Mar 18, 1993 9:59 am PST
From: CHARLES W. TUCKER   MBX: N050024@univscvm.csd.scarolina.edu
Subject: COMMENTS ON VIDEO

ROM CHUCK TUCKER (930317)   Dear Dag and Christine, Moments ago I was at the end of a very long post about your video and I hit a key and disconnected my computer from the mainframe (the first time I have done it accidently) and I could not retrieve it so I am writing it over again on a file that I will not lose unless I hit the wrong key in WS.

First, I want to thank you for the video. I had read Bill's comments about it before I received it but after viewing it I re- read them and found them to be excellent. As you noted they should be in the archives with the demos since they are very good instructions (the best that I have seen Bill give for these demos) and I plan to use them and would suggest that you do the same. I have a few suggestions for the video that I hope you will take not as criticism but in the spirit of improving upon your presentation.

I don't see this video as one which can be used as a demonstration but it is, for me, a good record of what you do when you give a presentation of PCT. If you wanted to develop it into a "demonstration video" I would suggest that you pay very close attention to Ed's video that was done with PBS. It seems to me that that video was actually designed to be a demonstration and produced with that as its purpose (Ed also tried to make videos of Bill's presentation but he did not have the equipment and other production facilities to do what PBS did). I do use Bill's video in my classes and they are useful. One of the major changes that you would have to do to make you video into a demonstration is to change the camera angle. But much more has to done as you probably know.

I would suggest that you attempt to eliminate at least three behaviors from your presentation: "I believe in my heart," "It is the nature of the beast," and the word 'feedback' as used in the statement that you made "Thanks for the feedback." The first two statements I believe are habitual for you but for me they mainly say "I firmly believe what I am telling you." It is fine to be confident and you should be but the repetition of it may make others believe that you are not sure of your self. The use of the word 'feedback' in the above statement you should recognize as a common but an inappropriate use when discussing PCT - feedback should be reserved for the actions one takes toward self not between people - no one gives you feedback.

I would also suggest that you reduce your criticism of S-R approaches and linear causality. I say this because I have found with my students that they begin to wonder if I am protesting to much and I also find myself talking about ideas which I claim are worthless and less about ideas that I claim are worthwhile. I don't think we want to remind others of those ideas that they believe in while we are trying to get them to create a disturbance of those very ideas so they can begin to reorganize themselves to adopt a new and different and an odd (to them) approach. This is mainly tactics but I think that it makes sense in terms of the model.

One of the most difficult ideas to get others to grasp about PCT is the statement that I put on the net several weeks ago as sort of a snide remark to Bill and Martin - you can't tell what a person is doing by paying attention to what she is doing. Bill in his 930312.0930 post states it again " . . . why observing actions doesn't tell you either what the person wants or what the person is perceiving. It doesn't tell you what the person is DOING - what those actions are accomplishing that the person wants to perceive as being accomplished." It becomes very confusing when you (any of us state) "it is action that is important not behavior" or you separate 'action' from 'consequences ' in your diagram. I have read statements on the net which point out that people who pay attention to behavior will never find out anything important about what people do. ALL OF THESE STATEMENTS MAKE SENSE IF YOU ALREADY KNOW ABOUT PCT BUT (AS BILL NOTED) THEY ARE "TESTS" TO OTHERS. It becomes especially confusing when you (as we all do who do demos) spend a great deal of effort getting the person to move around and make traces of their hand movements on a piece of paper. You point out what the person is doing by noting to the audience what she does when you pull quickly on the rubber band. THEN you say "you can't tell what she is doing by paying attention to what she is doing." Well, that statement is accurate ONLY in a technical sense and ONLY within PCT.

Actions are those activities (behaviors, movements) which can be used to find out what the person is controlling for when you apply the test of performing acts to see if they are disturbances while there are many other movements performed by the person (the movement of Cathy's skirt) which are not relevant to controlling a reference signal BUT we can't ignore ALL behavior (which would mean not observe anything that a person does) and still have some way of testing what he/she is controlling for within the act. In fact, we are very precise about what we observe and carefully perform several "tests" (note the Coin Game) to see what the person does to correct for disturbances. This has to be made clear or people will come away from your presentations like they come away from my classes - very confused.

There is a problem with the word 'consequence' that I think that Bill has noted in some of his statements on Skinner. If you use the word 'consequence' to mean all behaviors then that is not proper within PCT but if it is use as "outcome" or "desired consequence" or as "wanted consequence" I believe it is appropriate in PCT. I am wondering whether it would not be simpler to note that this is probably what Skinner meant by the word 'consequence' most of the time that he used it so that much of this data could be appropriated rather than constantly being conflict with his view - just wondering.

Finally, I want to bring to your attention a book that Rick mentioned on the net sometime ago and dismissed as not relevant to PCT - I think that there is something useful about this book and I would like to see if you agree. The book is Donald H. Ford. 1987. <<Humans as Self-Construction Living Systems.>> New Jersey: Lawrence Erlbaum. BF 38 .F66 It is a review of many literatures with a model which is very similar to PCT (surely quite compatible and not contrary in major ways). I note part of what Ford states in his final chapter about organizations. "If employess are viewed are self-governing components of a large living system, rather than as mechanical components of a machine, then the challenge is to create work situations in which (a) the employee's personal goals and the organizations goal's are linked in mutually beneficial ways, and (b) employees evaluate what they do as imoportant to both their personal goals and those of the organization, and believe their efforts are valued by their supervisors. In this way, people's self-organization and self- construction can be accomplished in significant part through directing their efforts toward facilitating the success of the larger organiztion. ... Moreover, as living systems, each employee has some potential for constructing new ideas about how to improve organizational fuctioning, and finds interest and satisfaction in life through making progress towards personally constructed or chosen goals. .... (660)." Doesn't this sound like what you might be interested in promoting when you tell people how PCT will be useful for organizations? It seems to me that this is the case. I suggest that you look at this book since it might be helpful for your work.

Hope what I have said is useful to you and if you find that I do not understand either what you are doing or saying please correct me before I go one in life with such misimpressions.

Best Regards, Chuck
Date: Thu Mar 18, 1993 11:51 am PST
Subject: Misc subjects

[From Bill Powers (930318.1015)]

Wolfgang Zocher (930318) --
(response to private post)
Sorry about that very nasty illness. I hope you will give a higher priority to getting well than to writing programs! I'm sure I convey the hopes of everyone else on CSGnet that you will recover soon. Two weeks of pneumonia is no joke.

Allan Randall (930317) --

>So my version of the challenge would take Ashby's compensatory
>and error-driven control systems and, assuming they were both
>designed to control, make a prediction concerning which would
>control better.

OK. For ideal components and linear channels, the prediction is pretty clear. But I will be interested to see the analysis you use for the error-driven system, in which E is partly a function of itself (via R).

Once you have the analysis for the two cases done, it would be useful to derive the real-world requirements on both systems for achieving a certain degree of regulation in the presence of noise. One channel common to both arrangements is that from R to T. How much difference would it make to each system if the output of R contained some specifiable amount of random variation?

Outside the scope of this challenge, there is a factor that Ashby didn't take into account: the possibility of disturbances that act directly on E, and are not detected by R. Under those conditions, disturbance-driven regulation is impossible, while error-driven regulation continues as before.
Meeting the "challenge" is less important than producing an actual analysis that I might be able to use! Keep in mind that I am only a humble engineer, and need to have everything spelled out in babytalk.

Hans Blom (930317) --

>So sensory input is a dependent variable. I take this to mean
>that sensory input has no degrees of freedom of its own; its is
>dependent UPON something, and that something is the reference signal.

What we say about dependent and independent variables is always in the context of the particular model we are proposing. In a single elementary control system in the PCT model, the perceptual signal is ALWAYS a one-dimensional scalar. If there is a multidimensional external quantity being controlled, in our model the way it currently stands more than one control system would be needed to keep all its degrees of freedom under control. I think we all recognize that this conception has some failings -- handling sideways interactions among control systems of the same level would be very awkward, for example. But at our present stage of experimental sophistication, this simple model seems to handle everything we can understand with satisfactory precision.

So having reduced the problem to one dimension per control system, we can ask what determines the state of a single input variable for a single control system. By definition, the input variable has no way of altering itself. As per Newton's laws, it changes only when the sum of all effects on it is nonzero.

There are two determinants of an input variable: the sum of all independent environmental physical effects acting directly on the variable (of which the variable is a function), and the output of the control system. As we are dealing only with one-dimensional variables, this means that no matter how many independent disturbances there are and by what paths they affect the input variable, we can always express the result as a single equivalent disturbance acting through a single equivalent path. This leaves only two influences on the input variable: "the" disturbance, and the control system's output. The net disturbing influence is arbitrary and independent of the operation of the system. All the variables in the loop, including the output quantity and the input variable, can be solved for using the closed-loop equations -- they are all dependent variables. As you note, the reference signal is also an independent variable relative to the control loop, and hence relative to the input variable.

>Is this the same as saying that the reference signal controls the perceptions?

Yes, given that the control system is capable of maintaining its error signal very small. The action of the system will almost completely cancel the effects of independent environmental disturbances on the sensory input, and at the same time force the sensory input to track the varying reference value established by a varying reference signal.

>Originally, the title "Behavior: the control of perception"
>gave me the impression that all perceptions are controlled. Now
>I understand that there are also uncontrolled perceptions. Is
>it therefore "Behavior: the control of SOME perceptions"?

That wouldn't have made a very catchy title, but you're right. Behavior controls only some perceptions, those that can be systematically affected by output actions and for which the organism has reference signals and control systems. The remainder can be controlled in trivial ways (not looking at the moon keeps the perception of the moon at zero), but for the most part simply make up the world within which the things we care about happen.

>Now, if some perceptions are controlled and some are not, are
>there also intermediates like:

>- some perceptions are sometimes controlled, but not at other times;

>- some perceptions are partially or approximately controlled;
>- some perceptions are controlled in some degrees of freedom
>(dimensions) but not in others?

In this model, an elementary control system does not decide whether or not to control. It simply controls. If some variables are controlled only part of the time, the explanation has to be sought at a higher level in the hierarchical model. As part of a higher-level control process, a higher-level system may change which lower-level control systems it is using to control its own (derived) perceptions. It must have some way, therefore, of turning lower systems on and off. There are several ways, which have different implications. But the main thing is that when a control system is turned on, it controls ALL of the time. It can't turn itself on or off: something else must do that. That's just my basic design principle. If the external part of the loop is lost, the control system will frantically crank up its output trying to correct the error. It will continue to do this until a higher-level system notices something amiss and makes the required adjustments. Rick Marken has shown that when the sign of the external feedback is reversed, the control system that had been tracking runs away on an exponential curve -- for about half a second. The curve closely matches that of the model when feedback is reversed. Then (according to the model), a higher- level system reverses the sign of the control system's error or output connection and it regains control.

As to partial or approximate control, that is only a question of how well the control system works. There is a complete spectrum of control ranging from hardly any to very precise. If we make the reasonable assumption that control systems evolved because it was in the species' interest to determine for itself how certain parts of the local environment behave, we can assume that the less error is allowed by a control system, the greater the advantage to the organism.

On the other hand, there are specific circumstances in which very tight control could be a disadvantage -- a waste of energy, for example, considering the benefit to be gained. You have mentioned something like this. Once again, my basic design principle applies. A control system does not decide for itself how well to control (assuming there is any choice). If its loop gain is lowered under certain circumstances, a higher-level system is doing the adjustment of gain, as part of maintaining control of higher-order perceptions.

A specific example of this appeared in my model of operant conditioning three or four years ago. One level of control had a reference signal set by a control system for body weight. The reference signal specified the level of a perception of short- term nutritional state that was immediately affected by the rate at which food was ingested (body weight was a long-term function of average nutritional input). This short-term state decayed fairly rapidly with time. The action of the system was to vary the frequency at which a bar was pressed, producing food input through a schedule of reinforcement and thus maintaining the perception of nutritional input level matching its given reference signal from the weight-control system.

Another higher-level control system, acting at the same time, compared a cost of bar-pressing proportional to the rate of pressing with a benefit of nutritional input proportional to the rate of ingesting food. As the cost rose above the benefit, the output gain of the bar-pressing system was lowered to keep the benefit at least as high as the cost. I'm sure you'll recognize this as a primitive form of optimal control (a one-way control system in this case).

This model did very well in fitting the bar-pressing behavior of rats over a wide range of schedules of reinforcement and two conditions of body-weight (forced by withholding food between experiments in the real studies).

I was more or less forced into this model, because no matter how I tried to make the bar-pressing control system vary its own gain with nutritional input (still remaining an elementary control system), I could not reproduce the double-valued function relating the schedule of reinforcement (bar-presses per reward, which ranged from 1 to 160) to the rate of bar-pressing. Only when the cost-benefit control system was introduced was I able to make the curve reverse at the right place. Then the model came very close to all the data points from the real rats.

>Is behavior FULLY in the service of the control of perceptions
>or could there also be behavior that is not?

One has to wonder (a) why an organism would learn to produce behavior that never had any feedback effects on that organism, and (2) how any organized behavior could reliably be produced, in a variable environment, without feedback control. My hunch is that essentially all behaviors (that is, outputs) are learned in order to control some perception -- that in organisms there is no open-loop behavior of any significance.

It's possible that evolution might have created some spontaneous emission of actions without any feedback effects on the organism doing the acting, as a benefit to the species. But such open-loop acts would have to be very simple and noncritical, because to reproduce the effect of any act in a normal environment would be almost impossible without feedback from the actual effect created. This is not to say that a feedback control action couldn't be inherited because of a side-effect it has on other organisms, with evolutionary consequences. To reproduce that side-effect in a variable environment, however, the organism would have to control for the effect of motor acts, not the acts themselves. There's just too much chaos and interference out there to make any totally open-loop behavior feasible. When a peacock spreads his tail, the actual spreading must be a control process, and perhaps even the subsequent response of a mate is also controlled for -- but I'm sure that the side-effect of making more peacocks is NOT a controlled variable.

Peter Cariani (930318) --

Hi, Peter, long time no hear.

>I'm afraid I don't see how anyone could call Bohr's approach
>"invoking magic", when what he was calling for was the primacy
>of (verifiable) results of observations and calculations over
>tacit images of the nature of an underlying "Reality".

I don't come down on either Bohr's side or Heisenberg's. Bohr's view is extreme, and teeters on the edge of solipsism. Heisenberg's is naive, attributing uncertainty to the wrong entity. I think control theory gives us a third alternative, which I'm surprised has not shown up in physics (maybe it has).

We can easily say that our perceptions of reality (read: instrument readings and interpretations) are a formal system that we made up ourselves, based only on what we can perceive, not on any objective "reality." But we don't have to stop there as Bohr did. We also can act, produce outputs (read: experimental manipulations). The effects of our actions are related to the perceptions we get back only in the most indirect way. But such effects DO OCCUR, even though we can't perceive how our output is affecting our input. Furthermore, our perceptions often change when we have performed no act: there are agencies out there.

To me this is a proof of existence: there is a reality out there and it contains active agents. Unfortunately, we have to guess at its details -- propose models of what MIGHT be there that would account for the effects our actions have on our perceptions and predict new effects of new actions. This guessing game works extraordinarily well when the demands on models are exacting enough: namely, that prediction errors should be no worse than measurement errors. It works so well that one can reasonably suppose that the resulting models are not inconsistent with what is really going on. This doesn't mean they're isomorphic to reality; it means only that something true is captured in them.

An epistemology that is based on observation alone can't lead to such a conclusion. When you include action in the picture, and close the loop, something different emerges.

>So much of contemporary mathematical physics (and the current
>wave of pop-physics pulps), having adopted a platonic-realist
>approach, no longer seriously attempts to connect theory with
>observation. One of the great intellectual tragedies of the
>late 20th century has been this infusion of platonic mysticism
>(following Godel, the later Carnap, and Tarski) into
>philosophy, the foundations of mathematics, physics,
>linguistics, and the cognitive sciences. We are still dealing
>with the wreckage.

Platonic mysticism! Bravo. But the other side is anti-platonic scholasticism, the triumph of pure reason over experiment. The antidote to both sides is to include action in the picture as well as perception.

Best to all, Bill P.

Date: Thu Mar 18, 1993 12:15 pm PST
Subject: Re: help!

[Martin Taylor 930318 14:40] Hans Blom, 930317

>Now, if some perceptions are controlled and some are not, are there also
>intermediates like:
>
>- some perceptions are sometimes controlled, but not at other times;

Necessarily true. This is a consequence of the fact that we have far more sensory degrees of freedom than motor degrees of freedom.

>- some perceptions are partially or approximately controlled;

That's a fuzzy concept. Does it mean the same as "some perceptions are controlled at low gain" or "some perceptions are controlled with a zero-point dead zone?" Either way, I think the answer is "yes." But I have a feeling there is something else behind this question.

>- some perceptions are controlled in some degrees of freedom (dimensions)
> but not in others?

In PCT (classic form), a perception has only a scalar value (one degree of freedom). It is not a necessary element of PCT, I think (though I will probably get a disagreement from Bill Powers on that). But whatever I think, I have not yet seen any instance in which there is a requirement for a perceptual signal that is not a scalar.

>And how is it the other way around? Is behavior FULLY in the service of the
>control of perceptions or could there also be behavior that is not?

Moot point. All actions have side effects, and if you are talking about actions as equivalent to behaviour, then behaviour is not fully in the service of perception control. But in the jargon of PCT, the word "behaviour" is used for actions executed because some perception deviates from its reference. If the hierarchy is correctly linked, actions do serve to reduce perception-reference error.

On a wider scale, I think there is a whole class of behaviour that might be called "exploratory." Actions in this class serve to determine under relatively safe conditions how those actions are likely to affect perception, and thereby to assist in reorganization (or, if you like, in system modelling). But though these actions at one level are not currently controlling any perception, at a higher level they are, and they are certainly being performed under control at lower levels.

From an evolutionary point of view, any action that has no discernable effect on perception is wasted energy and would be selected against. Actions that do have discernable effects can be used in control of those perceptions they affect, and if there is need for those perceptions to be controlled, those actions will be available for the purpose. So, in the long run, you probably have to say that any actions that are not part of a functioning perceptual control loop are either the actions of a still-reorganizing hierarchy or are undetectable by the actor.

Martin

Date: Thu Mar 18, 1993 1:35 pm PST
Subject: Re: infomystery, research proposal

[Martin Taylor 930318 15:20]
(Rick Marken and Bill Powers, various messages 930317)

Let's look at the mechanics of what happens in a control loop, taking Bill's example as a starting point.

In the following, B represents a CEV, in the specific example an object connected by springs to A and C. The movements of A and C affect the movement of B in opposing directions.

```
>
>           A --- /////////// --- B --- \\\\\\\\\\ --- C
>
```

A control system has sensory input representing the position of B (and nothing else). The position of C (or, equivalently, a force applied to C) is the output of the control system. (The "equivalence" statement is there to point out that the output is normally of a kind quite different from the sensed variable, and there is a functional relation, perhaps stochastic, between the output and its effect on the CEV.)

```
>
>When we draw a disturbing variable in a control-system diagram,
>we are drawing something analogous to A. The output is analogous
>to C.
>
>     Disturbing
>      Variable ----////////--- CEV ---- \\\\\\\\ --- Output
>
>Without knowing anything about the state of the CEV, we can
>specify the state of the disturbing variable. If the units
>involved are distances, we can specify the location of the
>disturbing variable. If we observe the system for a while with
>the disturbing variable in the position that has no effect on the
>CEV, then suddenly move the disturbing variable to a new position
>and keep it there, we may well observe something like this:
>
>---------------------T----------------------------------------
>
>                       **************************
>
>Disturbing variable
>
>*********************
>
>---------------------T--------------------------------------
>                       *
>                        *
>CEV                      *
>                           *
>***********************       *          ************
>                                  *  *
>
>---------------------T---------------------------------------
>
>                                  *  *
>                              *          *************
>                           *
>(Opposing) Output        *
>                        *
>                       *
>**********************
>----------------------T---------------------------------------
```

(I have inserted time markers "T" in Bill's diagram axes).

At this point, my interpretation of the situation differs from Bill's, at least I think it does. The following is a little obscure.

> It makes no
>sense to say that the control system's perceptual signal contains
>no information about the CEV. That is why it has seemed so self-
>evidently true to you that a disturbance (meaning an actual
>change in the CEV) conveys information to the control system --

>and so stupid of us to claim that it does not.

In the diagrams, at the moment the disturbance step occurs, it is fully and completely represented in the CEV (and, by hypothesis, in the perceptual signal, though that is a point we will wish to dispute in later postings). As the effect of the disturbance is used by the ECS, through changes in its output, the representation of the disturbance is, so to speak, "bled off" the perceptual signal, to find representation in the output signal. To the extent that the effect of the disturbance remains in the perceptual signal, it is not in the output signal. To the extent that it has been used to affect the output, it is no longer in the perceptual signal.

Now consider the same situation, except that the ECS is blinded for a period, so that even though A is moving, and thus affecting B, the ECS does not move C to compensate. A moves left and right, slowly and smoothly, so that B also moves in a continuous fashion (as only a third party could see, the ECS being blinded). At time T in the diagram, the ECS is again sighted, to detect that during the blinded time B has moved to the position noted. This is not the result of a step disturbance introduced by a step movement of A, but is the integral of all the effects of A's movement during the blinded period. It is the integral of micro-steps that would have appeared momentarily in the perceptual signal before being "bled off" into the output, had the ECS not been blinded. To a first approximation, it might well have seemed that those smooth changes in A would have had NO effect on the perceptual signal, if (1) the gain were high, and (2) the transport lag were low.

The point of all this is to try to demystify the "magical" fact that there is no representation of the disturbance in the perceptual signal, although there is an almost perfect representation of it in the output signal. Given this way of looking at the situation, one can see there is no reason at all to expect that the momentary control actions will be the same on successive trials using the same disturbance pattern. Near identity of that pattern would, as Rick has pointed out, be evidence that the disturbance is in some way represented in the perceptual signal, but unless the disturbance has components faster than the possibility of control, that won't happen. Unless the disturbance has some macro-step events (which have high-frequency components) that are not quickly compensated by the output actions, the micro-step actions would be in quadrature on successive runs. (By micro-step, I mean the results one might get in sampling a continuous motion of a low enough bandwidth).

Is this an agreeable reformulation of what Bill wrote a couple of days ago explaining the events going on around a real loop with finite transport lag?

Martin

Date: Thu Mar 18, 1993 1:28 pm PST
Subject: Perception as Interaction

[from Mark Bickhard (NOT on CSGnet) via Gary Cziko (930318):

Bill Powers (930318.1015):

>Platonic mysticism! Bravo. But the other side is anti-platonic
>scholasticism, the triumph of pure reason over experiment. The

>antidote to both sides is to include action in the picture as
>well as perception.

I think I mostly agree - the caveats due to the clear historicity of the dialogue which makes it unclear just what commitments all of these comments might be presupposing.

One demurral I would make, however, is to point out that perception is itself already ALSO (inter)action.

Mark

Mark H. Bickhard   Department of Psychology
17 Memorial Drive East   Lehigh University
Bethlehem, PA 18015   215-758-3633   MHB0@LEHIGH.EDU

Date: Thu Mar 18, 1993 2:17 pm PST
Subject: Re: research proposal

[Martin Taylor 930318 16:00]   (Rick Marken 930316.2000)

>I really didn't understand your description of the study. But after
>reading this far I am getting the distinct impression that the
>control systems you plan to study are simulations.

Not really. The end (far distant) goal is an "intelligent" computer interface. If it works, a byproduct may be better understanding of human use of language.

>Why not do a
>study with real control systems -- like people? If you are intending
>to study people then I need some clarification of the vector
>representation of the letters and all that.

Yes, so would I. In my mind (and certainly not in any way as a simulation) the vectors represent the degree to which the symbols have attributes. If we were simulating a recognizer of printed words, the features might be letters contained, plus such things as length, porportional position of ascenders and descenders, location of desnity maxima of marks on the page, and so forth. If it were spoken words, we might have spectral vectors and their time derivatives (or any of a host of other features). If it were concepts, the attributes might be featheriness, weight, linearity, patchiness of colour, ... In the experiment we propose, our intention is to assign letters arbitrarily to locations in the vector space (we are going to cheat, though, and mentally group the letters into sets such as vowels, liquids, stops, voiced ... and colocate them according to our mental groupings. The control system won't know anything about that, though it may learn).

> I sure hope you are planning
>to do this with living control systems. What we really need in PCT
>is good data; who cares how a particular implementation of an algorithm
>works? Unless your working in AI or Artificial Life, of course, in which
>case you only care about the behavior of the model itself.

I agree with the need, but that's not what we are doing. As I said, the ultimate goal is a human factors one, but the immediate goal is to study the properties of different forms of reorganization, to see whether simple control system can learn to control under the artificial world conditions to which they will be exposed.

You may not have noticed, but the task world is that of the Little Man, and the letter symbols simply drive the cursor to be tracked by the finger. Can a Little Man that reorganizes (a.k.a. the Little Baby) learn where the finger is likely to have to go if the motions of the tracking cursor are directed by a synthetic grammar? Will it learn to predict the letter group before it learns to predict the actual letter (or over)? The latter question is tantamount to asking whether a control hierarchy will (as I think it will) learn to paraphrase.

In respect of your proposed study, I am reminded of a peculiar result obtained by somebody here when I was a summer student around 1958. So far as I know, it was never published because no-one knew what to do with it. The experiment was very simple. There were 25 slides, 23 of which showed in big letters ABDCE (or some such). Two showed ADBCE. These were placed 18th and 23rd (roughly) in the sequence of 25. The subjects had to report for each slide the sequence of letters. Easy? No subject reported the deviant sequence at slide 18. All subjects reported the deviant sequence at slide 23. When asked afterward whether they had noticed any other deviant slide than 23, no-one acknowledged having done so. But I'll bet they would have noticed if slide 18 had GVLYK. I'm not sure if this relates to your result of faster control when the letters change than when the sequence changes, but it has lain in the back of my mind for many years as something needing explanation. Martin

Date: Thu Mar 18, 1993 3:31 pm PST
Subject: Re: Language facts, PCT research

[Martin Taylor 930318 17:30] Bill Powers 930316.0900

>What we are looking for is not just patterns of linguistic behaviors --
>outputs -- but relationships between those behaviors and
>controlled perceptions. Output patterns are bound to be variable;
>they change with disturbances. The controlled perceptions will be
>much less variable than the means by which they are controlled.
>The PCT approach is to try to explain variations in outputs by
>finding disturbances with which they correlate, and then using
>those hints to postulate and test potential controlled variables.
>
>I don't see any sign of this sort of approach in the linguistic
>research of which I hear. Most of the effort seems to go into
>looking for regularities in language behavior. From the PCT
>viewpoint, this is something like trying to understand how people
>participate in sports by looking for regularities in the way they
>move their arms and legs.

I think a special approach is needed for analysis of social interactive behaviour.

There's nothing wrong with these paragraphs as such, but things are different when you are using the output of other control systems (all one can observe) and the main reference that you are using is for the others to act as you wish, using your output as their sensory input. Your CEVs are their outputs. Theirs are your outputs. So for social interactions such as language, it does make sense to look at regularities in output. It's not like that when you are analyzing sport behaviour.

Inanimate objects in the world do not "observe" your outputs. They do respond to them. You can treat another person as an inanimate object as one way of controlling a perception you have, but it is usually easier to recognize their control system properties. That means providing for them disturbances (your outputs) that they know how to counter (interpret). They can control their perceptions relating to recognition of what you want to communicate, if you make outputs that have particular shapes (sound patterns, word patterns, gesture patterns, argument styles). If your outputs have other shapes, they have a harder time establishing control of their perceptions relating to references for acting in accord with your wishes. You control your perceptions better if they do act in accord with your wishes by observing your outputs, too (given that you perceive them to be cooperative.

You control your perceptions, all right, but to do so in social interactions you must also control for your perceptions of your externally observable actions. In non-social circumstances you would not normally do that.

Martin

Date: Thu Mar 18, 1993 4:05 pm PST
Subject: PCT research

[From Richard Thurman (930318.1630)]   Rick Marken (930317.1400)

>And I think the training emphasis would be great. The idea
>would be to show that training is largely a matter of learning
>which perceptions to control, not which "outputs" to generate.

>Keep in touch on this. If we do it over CSG-L (instead of
>in private) maybe we can benefit from the advice of others.

I just got off the phone with Dr. Hancock (Grand Canyon Univ) and told him about the idea. He is amenable to it and may have grad students who may also wish to participate. In addition he said that if you want to start running subjects he has some available for the next few weeks (from some of his classes). He's got two Mac labs where he can run subjects.

So far we have been able to run a number of studies dealing with training and education. For example we set up a HyperCard based drill and practice program for teaching "radar signature recognition." (I had to make it look like a task that would have an Air Force application but the bottom line is that the program can be set up to drill just about anything.) Some of the interesting preliminary findings are that we could fairly easily separate out three distinct groups of subjects. Some students were controlling for understanding the information. Some students were controlling for getting a correct answer (that is they didn't care much about why they were correct, they simply wanted to see the word correct on the screen). And some students were controlling for getting out of the experiment. Its interesting data (in a group-means statistics sort of way).

We will be presenting our findings at AERA (American Educational Research Association) in April ("Instructional Feedback in a Servo- control Theory Framework") and at APA in August ("Student Modeling and Perceptual Control for

Intelligent Tutoring Systems"). If any body is interested I will send out drafts of the papers.

Gary Cziko 930318.1648

>I've been thinking about the proposed control of sequence experiment. But
>instead of flashing letters on a computer screen, I think it might be a
>better test to use a sequence of tones like "do-re-mi-fa-sol-fa-me-re-do."

Rick can probably set up the HyperCard stack so that it will be a simple adjustment to go from visual to auditory presentation. (I'm assuming he is going to use HyperCard.) It would be interesting to see if different patterns emerge with different perceptual modalities.

>I think this would be a better test for none other than the intuitive
>feeling that it would be easier to perceive disturbances to this auditory
>sequence tan to the visual one using letters.

I'm not sure about that. Some of us (I guess meaning me) are not very good at listening to music. I'm not sure that an auditory modality is easier to perceive sequence or not.

>Now that I've started to think about this, I can see music perception as a
>way of getting to quite high levels of perception--Key modulations as
>transitions, perhaps musical phrases as events, musical style (e.g.,
>baroque, classical, romantic, modern) as category, etc.

>But I nonetheless have the intuition that once the melody (sequence) is
>well known), a person could react to a wrong note as quickly as he or she
>could to a wrong loudness or wrong timbre. But I hope I'm wrong.

Doesn't PCT explain that intuition. Once a melody (or at least a certain phrase) is known it hovers around the Event level of perception. If a wrong note is sounded then its perceived "not the event." So it should be reacted to faster than a sequence. Is that right?

Rich Richard Thurman

Date: Thu Mar 18, 1993 9:26 pm PST
Subject: Video

[From Dag Forssell (930318)]

Thanks Chuck! Re: Direct post (930317)

You made a number of comments about your own experiences and perspective that others on the net would benefit from. I would be pleased to see it on the net.

In the next few days, Christine and I are assembling the next tape, from our March 11 performance, complete with a different camera angle. Since you, Bill and Bob have given us feedback, you qualify (along with the archive) for complimentary copies of the next version. Others may want to stand by for progress reports from you.

Thanks, Dag

Date: Thu Mar 18, 1993 10:29 pm PST
Subject: where's the disturbance?

[From Rick Marken (930318.2200)]   Martin Taylor (930318 15:20) --

>Let's look at the mechanics of what happens in a control loop, taking
>Bill's example as a starting point.

>In the following, B represents a CEV, in the specific example an object
>connected by springs to A and C. The movements of A and C affect the
>movement of B in opposing directions.

>A control system has sensory input representing the position of B (and
>nothing else).

Right!!

>In the diagrams, at the moment the disturbance step occurs, it is fully
>and completely represented in the CEV (and, by hypothesis, in the
>perceptual signal, though that is a point we will wish to dispute in
>later postings).

I am willing to agree that, at the moment the disturbance step occurs, the disturbance IS COMPLETELY REPRESENTED IN THE CEV. I'm even willing to assume that the control system is both sluggish and that it has a substantial trasport lag -- so the output doesn't even start to compensate for the disturbance for some time after the start of the step. So for some time (as long as you like) the behavior of the CEV is due solely and completely to the disturbance. Now the only problem is HOW DOES THE CONTROL SYSTEM KNOW THAT THIS IS THE CASE? That is, how does the control system know when the disturbance step started? How does the control system know when the variations in the CEV that it is looking at are due only to the disturbance and when they are the combined result of its own effects and those of the disturbance?

In order for the control system to know when the disturbance has started or stopped it would have to be able to perceive the disturbance. But you have assumed (above) that "the control system has sensory information about the position of B [the CEV] (and nothing else)". So the only person who knows that the CEV (during your example) is the result of the disturbance alone is the observer of the control system -- ie. you.

> As the effect of the disturbance is used by the
>ECS, through changes in its output, the representation of the disturbance
>is, so to speak, "bled off" the perceptual signal, to find representation
>in the output signal.

This is a nice try at saving the S-R point of view but it can't work; the control system cannot possibly (for the reasons I mentioned above) know when to do this "bleeding off". Exactly the same changes in the CEV could be the result of the combined effect of the disturbance and the control control system's own outputs or of those outputs only (with the disturbance constant) -- in fact the response to disturbance that is reflected in the CEV in your example is exactly the same as the transient response of the CEV to an output resulting from the addition of an electrial impulse "disturbace" to the error signal in the control system; if the control system "bled off" the same output that it made when the source of this variation in the CEV was caused by the disturbance, there would be positive feedback -- but, as I have shown, the addition of such a disturbance to the error signal is "compensated for" just as effectively as if the disturbance had come from the outside world.

>The point of all this is to try to demystify the "magical" fact that
>there is no representation of the disturbance in the perceptual signal,
>although there is an almost perfect representation of it in the output signal.

It seems magical only when the control loop is looked at in S-R terms (with the CEV being the S and the output being the R). There was never any magic from a PCT perspective -- there is no perceptible representation of the disturbance in the perceptual signal; output is non-magically made to mirror the disturbance by a system the generates outputs that are proportional to r-p, the difference between current perception and a reference for that perception.

The world can't tell you what to do -- and, frankly, it probably doesn't much care.

>Is this an agreeable reformulation of what Bill wrote a couple of days
>ago explaining the events going on around a real loop with finite transport lag?

Nope (see above). But I'll let Bill speak for himself.

Bottom line: the dynamics of the effects of disturbance and output on the controlled variable don't bail out what is basically the S-R view of control. There is NO information about the disturbance in the CEV (or the perception thereof). Thus, the basic assumption of scientific psychology for the last 100 plus years is out the window; this is why the idea that there is no representation of the disturbance in the perception is being so strongly resisted (and will continue to be so -- despite the evidence); no one wants to give up comfortable assumptions -- but facts is facts. I'm willing to give up my comfortable assumption (that there is no recoverable representation of the disturbance in the CEV) is someone would just show me how to recover that representation -- and, hence, the disturbance (without being given outside information about the output and/or the disturbance itself).

Best   Rick

Date: Fri Mar 19, 1993 1:01 am PST
Subject: Re: where's the disturbance?

[From Oded Maler (930319.0915 ET]

* [From Rick Marken (930318.2200)]   Martin Taylor (930318 15:20) -- Etc.

I have a lot to ketchup with the info stuff, but let me try to formulate it the way I see it. Suppose there is one CEV x in one dimension. It was at position x0 at t=0 (suppose for the moment that perception is precise), and for a period t the control system pushed it toward the right (positive). Then at time t, the variable is observed to be in the left (x<x0). In this case you can be sure that during

that interval of time the net effect of the disturbance (its integral or whatever) was to push it to the left. Of course you cannot exactly know which of the infinitely many curves having the same integral was the actual disturbance, but you can exclude many of the other possibilities (you can be sure that the disturbance was not in the same direction of you efforts). This, I think, is the sense in which the perceptual signal *does* provide information about the disturbance. All this can be done in a probabilistic framework, with noise in the perception, action etc., and with some prior probabilities on the form of the disturbance. Even if you cannot measure your own actions, just know that they are bounded, still the velue of the perceptual signal can tell you something about more likely or less likely form of the disturbance.

Btw, I like the idea of calling all the real external world a "disturbance", it's anthropocentrism in its best! ("Life is what happens to you when you try to do other things" -J. Lennon).

--Oded

p.s.   I talked to a friend who is a mathemtical control theorist (= the abstract version of a control engineer) and to his mind, every control is feed-back control.

Date: Fri Mar 19, 1993 6:00 am PST
Subject: Lng fcts (or was that aua a?)

[From: Bruce Nevin ()Fri 930319 08:43:06]

[Martin Taylor 930318 17:30] --   (Bill Powers 930316.0900)

Thanks, Martin. That is an important part of what I have been saying. Maybe having it in a different voice will help.

Swamped,   Bruce   bn@bbn.com

Date: Fri Mar 19, 1993 8:04 am PST
Subject: Perceptions, outputs and disturbances.

[From Bill Powers (930319.0730)]   Rick Marken (930318.2200)--

I almost fell for Martin Taylor's explanation, Rick. It's only your persistence that put me back on the track. Martin is a silver-tongued devil.

The point is that when a perceptual signal changes, the control system in which it exists hasn't the least idea WHY it changed, nor does it even have any idea that it's a perceptual signal. It's just something being experienced. The control system experiences only this:

```
>                       *
>                        *
>Perception               *
>                           *
>***********************       *          ************
```

> * *
>
No disturbance. No output. Just the perception.

Now, what could have caused the perception to behave like this? The answer is, an infinity of combinations of other variations in the environment. But that is a question that only a higher-level system could ask, or answer. Only a higher-level system could receive a copy of this perceptual signal, AND OF OTHER PERCEPTUAL SIGNALS OF THE SAME ORDER, and perceive a causal relationship between the signals.

The next time you're in an elevator (alone), just before the doors close, hold your hand outstretched and watch it. When the acceleration begins, you'll see your hand sightly dip (rise) for a moment and come back to the former position. You will also hear the machinery start, feel increased (decreased) pressure on the soles of your feet, and feel increased (decreased) effort in muscles all over your body, particularly those holding the arm out. Use the parenthesed expressions if the elevator starts downward. You will see the hand doing approximately what the above diagram shows, although without the sharp peak (which will be upside down if the elevator descends).

The "you" in this case is clearly involved with a perceptual system of higher level than those in the spinal and lower brain loops that are involved in maintaining the arm in position. This higher perceptual system can perceive the visual and kinesthetic position of the arm, and can also perceive copies of the stretch and tendon signals that the spinal control loops are controlling. It is only in the perceptions of the higher system that the covariance of these signals has any experiential significance. Only the higher systems could conclude that the sound of the elevator has something causal to do with the sudden feeling of a greater (lesser) weight, which is felt as pressure changes on the feet and greater (lesser) effort in the muscles. There is no indication that the feeling of increased (decreased) weight in the outstretched arm is actually a measure of the output effort in the lower systems as they try to maintain the arm in the specified position.

And there is no perception of the disturbance at all, by any system at any level. Only _effects_ of the disturbance are sensed. All our perceptions of causal relationships are in terms of these _effects_ -- there is no perception of inertial acceleration itself. The disturbing variable is and remains unsensed.

Even if this is an open-cage elevator, from which you can observe the scenery outside suddenly start moving upward or downward, you do not experience the disturbance itself. You experience only visual perceptions that change in a way correlated with kinesthetic and auditory perceptions. The disturbance itself is a law of physics: F = MA. The "F" in that equation is the disturbance. Almost all that you experience of it is your own muscle tensions and other sensations that result from counteracting it. The only exception is the pressure on the soles of your feet. And if you didn't resist the new upward force, even that pressure wouldn't change -- you would collapse.

So, Rick, you are completely right in firmly insisting that the perceptual signal contains no information about the disturbance in the environment. When one takes the proper position from which to understand this system, which is INSIDE it, everything falls properly into place.

Once we are clear that control systems experience only their own perceptions, we can then return to the Engineer's position, in which we imagine a real outside world full of physical entities and physical connections, and propose an environmental model to go with the experiential model. In that Engineer's world, the F in F = MA is a perfectly real variable. We can see, in that modeled world, all the influences that add up to the state of the perceptual signal in the control system -- each of the disturbances that contributes and the state of the output variable. We can examine the physical link connecting these variables linearly or nonlinearly or dynamically to the sensor, and we can see that the sensor responds according to its inner construction by generating a perceptual signal that depends continuously on all these external variables. None of that is known to the control system in question.

From the external point of view we can explain the perceptual signal in terms of the various effects on it. We can use the language of analog computing, or differential equations, or probabilities, or information, or sports psychology -- whatever we are used to. This is all part of the process of modeling. It is a process that makes us feel that we, as external observers, understand how the other system works. We can explain to the person in the elevator why it is that all these different perceptions seem to undergo similar waveforms of change at the same time, and we can sort out causal relationships that the person could not find in the perceptions themselves. We can also point out that what seem to be causal relationships, such as the feeling of pressure on the feet seeming to cause the body image to move in the world reported as a visual image, are really concurrent perceptual effects of something else, the acceleration of the elevator.

When we, ourselves, are in the elevator, however, that explanation lives in a different world -- unless we remind ourselves of it constantly until it begins to seem natural. If we aren't thinking of this model, it seems that the floor of the elevator pushes against our feet and carries us upward, and that our bodies and the outstretched arm become heavier and tend to sag downward. We have it all sorted out into causes and effects -- quite incorrectly in terms of a carefully constructed model.

Oded Maler (930319.0915 ET) --

>Suppose there is one CEV x in one dimension. It was at position
>x0 at t=0 (suppose for the moment that perception is precise),
>and for a period t the control system pushed it toward the
>right (positive). Then at time t, the variable is observed to
>be in the left (x<x0). In this case you can be sure that during
>that interval of time the net effect of the disturbance (its
>integral or whatever) was to push it to the left.

This doesn't work as a method of analysis. By dividing the action into discrete episodes, you are taking partial derivatives that aren't available to the system when it is working. In the real system, you don't see all other variables being held constant while one is changed. The sequence of events that you describe would be neccessary in order to separate the effects of a disturbance from the effects of an opposing output. And furthermore, as I just pointed out above, you would have to be able to perceive the disturbance itself, and the output itself -- not just the resulting perceptual signal.

To grasp how the control system really works, you have to operate under the same handicap under which it operates: you have input information ONLY about the sensor signal. The system has to be able to work without knowing how many disturbances are acting or at what angles or with what magnitudes, and without knowing the state of its own output. You are sitting in a little room at a control console which has on it one uncalibrated meter and one handle that you can pull. That is all. And that is all you need to keep the meter needle in any position you choose.

>Btw, I like the idea of calling all the real external world a
>"disturbance", it's anthropocentrism in its best! ("Life is
>what happens to you when you try to do other things" -J. Lennon).

Neat. There is nothing wrong with anthropocentrism when you're trying to explain the world as it is experienced by human beings.

>p.s.

>I talked to a friend who is a mathemtical control theorist (=
>the abstract version of a control engineer) and to his mind,
>every control is feed-back control.

Hooray for your friend. Sign him up. I'm not kidding -- we mathematical cripples need all the help we can get.

Best to all, Bill P.

Date: Fri Mar 19, 1993 9:47 am PST
Subject: The PCT Revolution

[From Rick Marken (930319.0830)] Bill Powers (930319.0730) --

>I almost fell for Martin Taylor's explanation, Rick. It's only
>your persistence that put me back on the track. Martin is a
>silver-tongued devil.

Thanks for your support; I'm sure that I must seem like an uncompromising maniac unwilling to give an inch over a seemingly trivial point. But to me, understanding that there is no information about (or a usable representation of) the disturbance in the perception (or CEV) is the central (revolutionary) fact of PCT (with respect to the behavioral sciences). It is the fact that PROVES that the input-output model of behavior -- the framework within which all work in the behavioral sciences is done-- is WRONG. PCT even shows why it has seemed like the input-output model works -- disturbance resistance; in the process of controlling (uninformative) perception, control systems generate outputs that counteract disturbances to the controlled perception; if you see the disturbance (like a light shined in the eye) then the output (pupil size variations) seem to be caused by it. But control systems don't really work that way -- when organisms are in a closed negative feedback relationship to their own sensory inputs, then (most of) these sensory inputs are NOT the cause of outputs; in fact, they are controlled by outputs, and this control is not accomplished because of anything informative about the sensory inputs themselves.

Note that those aligned against me on this point are people who "accept" or at least are sympathetic to the PCT model of behavior -- Martin, Allan, Oded, Avery. I'm sure there are many others listening who feel the same way "PCT is great but this Marken guy is pretty extreme -- imagine, saying that perceptual input doesn't contain any information about the disturbance. A guy like that just gives PCT a bad name -- what does he expect me to believe, that our perceptions don't provide any information about how we should deal with the world. I like the idea of control of perception and feedback and all that but I don't see why a PCTer would have to deny physical reality to try to sell the theory. Having loonies like Marken around just gives PCT a bad name".

I (Marken) understand how you feel. All I ask is that those who are interested go over all the discussion and try to understand the points and the justifications for them. I think you will find (especially if you do the relevant demos and the relevant math) that crazy Marken is right. There is no compromise on this issue -- not because I don't want to (I'm all for cooperation) but because I am correctly describing the operation of a control system; and if people are organized as control systems (and I think its fairly safe to say that they are, simply because they are locked into a negative feedback situation with respect to their own perception) then what I am saying is true; people compensate perfectly for disturbances to their own perceptual input, not because there is information or some usable representation about the disturbance in the input that allows people to compute, bleed off or whatever, the appropriate output but because their output is proportional to the difference between perception and a reference for that perception -- that's how it works. And the result is that the perecption is made to match the reference. Ahd since the reference is ultimately determine by the system (person) themselves, it means that people determine "what they do" meaning what they experimence so, in this sense, people are completely autonomous.

Ultimately, I don't think you can fully appreciate the significance and depth of PCT unless you can understand this apparently trivial little point:

In a control system, the perceptual signal contains no information about the environmental disturbances (or system outputs) that are acting to cause the variance in that perceptual signal. This means that the perecptual signal cannot be the cause, guide, source of information, or anything else regarding the appropriate outputs that should be made by the control system in order to produce the intended result (the reference value of the perecptual input).

Once you understand this, you can kiss all versions of the input- output approach to living systems goodby. It might be a painful divorce but, believe me, it opens the way to a MUCH better life of science.

Best Richard S. Marken

"Religion is too important to be left to religions"

Date: Fri Mar 19, 1993 1:12 pm PST
Subject: Re: Perceptions, outputs and disturbances.

[Martin Taylor 930319 14:30]
(Bill Powers 930319.0730 and Rick Marken 930318.2200)

This has to be very quick (and I've already been interrupted once).

Earthquake alert -- Earthquake alert!!

The ground shifts under our feet once more. Now I am being asked not only to show that the disturbance is at some point represented in the perceptual signal but also that a scalar variable is supposed to be capable of answering "why" questions about the disturbing variable!!!!

All I presented was a mechanistic picture (seen from outside, to be sure) of what happens in the time period around a step in the disturbance, and then to show that this is indistinguishable from the integral of what happens during a period in which the sensory input to an ECS is cut off. So now I am a "silver toungued devil" (Bill) who is making "a nice try at saving the S-R point of view" (Rick). I used only Bill's diagrams for what happens, paraphrasing Bill's writing of 930315.1900, with which we (all three) agreed. The result is to show that to the extent the disturbance is represented in the perceptual signal it isn't represented in the output signal, and vice versa. That's all.

Why is that controversial? More to the point, why is it alarming to you?

The following is irrelevant to that argument, but relevant to another. Is it really necessary to argue about technical matters by reference to imputed motives? You well know I hold no brief for S-R approaches. Neither, if I interpret "devil" aright, have I any interest in introducing confusion into people's understanding of PCT. Nothing I have written requires any change in PCT except to note that the statement that "there is no representation of the disturbance in the perceptual signal" should have appended to it "to the extent that the perceptual signal is under control."

With infinite precision perceptual signals and zero transport lag around the loop, the perceptual signal is always completely under control and the disturbance is never represented there. If either of these conditions fails, the passage of the information "about" the disturbance through the perceptual signal to the output signal can be detected, and was beautifully illustrated in Bill's diagrams that I quoted.

Rick brings up the question of viewpoint when he says:

>So for some time (as long as you like) the behavior of the
>CEV is due solely and completely to the disturbance. Now the only
>problem is HOW DOES THE CONTROL SYSTEM KNOW THAT THIS IS THE CASE?
>That is, how does the control system know when the disturbance step
>started? How does the control system know when the variations in the
>CEV that it is looking at are due only to the disturbance and when they
>are the combined result of its own effects and those of the disturbance?

As I see it, knowing is not the province of an ECS. Acting is. It is by taking the Engineer's viewpoint that one can see what is happening. Scalar signals are scalar signals, period. They take on significance by how they relate to the world in other eyes.

>if the control system "bled off" the same output that it made when the
>source of this variation in the CEV was caused by the disturbance, there
>would be positive feedback

I don't understand this. The control system doesn't DO anything, unless you mean keeping the perceptual signal under control. It doesn't DO anything to itself, because itself is not an object of its own perception. What is "bled off" is the deviation of the perceptual signal from its reference, which the Engineer can see has been due to the disturbance. As far as the control system is concerned, that deviation might well be due to its own output (as when you experiment with changing the sign of the effect of an action). Allan considered that this kind of inappropriate (i.e. not contributing to negative feedback) effect of output ought to be considered as part of the disturbance. I think that's a semantic matter.

Now, Bill. I hope my silver tongue can convince you that I don't really think that there is anything in the perceptual signal that would enable any observer of it to distinguish between a motorized cam-driven spring-puller and a strong wind that moves the CEV on the end of a spring I use to control it. I think you are being pretty slippery when you substitute "disturbing variable" for "disturbance" in order to disagree with me.

>And there is no perception of the disturbance at all, by any
>system at any level. Only _effects_ of the disturbance are sensed.

Sophistry. How could anyone disagree with that, except to say that the effect of the disturbance that is sensed is what gets into the sensor system, and the effect that is perceived is an effect that comes from some transformation of the outputs of the sensor system. To retain this, you have to say that nothing in the Boss world is perceived. Only the effects of things are perceived. Of course that's true. So what?

>So, Rick, you are completely right in firmly insisting that the
>perceptual signal contains no information about the disturbance
>in the environment.

Not so, unless the perceptual signal is precisely under control.

>When one takes the proper position from which to understand this system,
>which is INSIDE it, everything falls properly into place.

Exactly. I've actually been much more rigid on this, separating out what is available from different viewpoints, than you have been. I've been chivvying you for some time to be consistent about maintaining this view. It's the only one that is coherent. Not taking that view is the reason most appeals to Information Theory have failed. You ALWAYS wind up with an inconsistency if you try to take multiple viewpoints at the same time.

(Rick)
>>The point of all this is to try to demystify the "magical" fact that
>>there is no representation of the disturbance in the perceptual signal,
>>although there is an almost perfect representation of it in the output signal.
>
>It seems magical only when the control loop is looked at in
>S-R terms (with the CEV being the S and the output being the R).
>There was never any magic from a PCT perspective -- there is no
>perceptible representation of the disturbance in the perceptual
>signal; output is non-magically made to mirror the disturbance
>by a system the generates outputs that are proportional to r-p,
>the difference between current perception and a reference for
>that perception.

The stage magician knows thee is no magic. To everyone else, it looks as if there is, though they know there cannot be. I only show the trick, even though I realize that's outside the union rules. Sorry.

Martin

Date: Fri Mar 19, 1993 2:10 pm PST
Subject: Dag,s Video - comments posted as per request of Dag

************** FROM CHUCK TUCKER 930317 ******************

Date: Fri Mar 19, 1993 3:34 pm PST
Subject: Why it's important

[From Rick Marken (930319.1300)] Martin Taylor (930319 14:30)--

>Now I am being asked not
>only to show that the disturbance is at some point represented in the
>perceptual signal but also that a scalar variable is supposed to be
>capable of answering "why" questions about the disturbing variable!!!!

No. What is being pointed out is that, although the disturbance may, at some point, be represented in the perceptual signal (it is, for example, so represented while p(t) = d(t) + zero) this can be of no functional significance to the control system, which simply gets p(t) -- and cannot respond any differently to p(t) values resulting from the fact that p(t) = d(t) + zero, p(t) = zero + o(t) or p(t) = d(t) + o(t).

>The result is to show that to the extent the
>disturbance is represented in the perceptual signal it isn't represented
>in the output signal, and vice versa. That's all.

>Why is that controversial? More to the point, why is it alarming to you?

If that is all you wanted to say -- which is really just saying "when p(t) = d(t) it is because o(t) is zero" then there would have been no controversy or alarm. I am prepared to agree that, if a=b+c and c = zero then a=b. But you are saying a bit more, as revealed in the following:

>What is "bled off" is the deviation of the perceptual signal from its
>reference, which the Engineer can see has been due to the disturbance.

This is only true during the transport lag (and after some period of no disturbance since the output is an integral); so it is an unusual circumstance. The engineer would usually see that what is continually "bled off" is the deviation between reference signal and a perceptual signal that is the COMBINED result of the disturbance AND system outputs. What you imply in the above is that the "disturbance produced" deviation of perception from reference is what determines the output. But this is typically not the case; what produces the deviation of perception from reference is the combined effects of disturbance and output. You are implying that there is some important functional significance to the fact that the disturbance is occasionally represented almost perfectly (limited by noise) in the perceptual signal. But this is WRONG. Whether or not the disturbance is EVER represented in the perceptual signal is not important to the functioning of the control loop -- in fact, in all the control loops we build, the disturbance is virtually NEVER represented ALONE in the perceptual signal -- it is only so represented when o(t) hits exactly zero or is constant for a couple samples -- pretty rare events.

>As far as the control system is concerned, that
>deviation might well be due to its own output (as when you experiment
>with changing the sign of the effect of an action). Allan considered
>that this kind of inappropriate (i.e. not contributing to negative
>feedback) effect of output ought to be considered as part of the
>disturbance. I think that's a semantic matter.

Can we agree that perception is ALMOST ALWAYS a SIMULTANEOUS result of output and disturbance ? I am beginning to realize that our problems may result from your apparent belief that it is necessary that sometimes only the disturbance (or "inappropriate output") cause the perceptual signal to deviate from the reference signal. You apparently believe that the only way for disturbance countering output to be generated is for a disturbance produced error is generated. But this is NOT TRUE.

NOTA BENE. The size and sign of the difference between perception and reference is at (virtually) all times the SIMULTANEOUS result of disturbances AND the effect of the system's own output.

This is why I say that there is no information in the disturbance that can determine output. The determiner of output is the size of the error signal; can we agree on that? But the error signal is virtually NEVER determined by the net effect of only the disturbance to the perceptual variable; the error signal depends on BOTH disturbance and output.

So, the reason this topic is controversial is because it seems to me that what you are saying about the operation of a control system is incorrect.

Why is this alarming? Because it seems to me that you are trying to squeeze the beautiful concept of control of perception into the Procrustean bed of conventional linear cause effect thinking. PCT demands that we start psychology all over again -- psychology beautifully reborn. But, if we keep trying to see control in linear cause effect terms (instead of in terms of the circular causality that is actually involved) this renaissance will be continually be delayed.

Perhaps we can focus this discussion for a moment on the main point I made in this post -- about the cause of error in a control system. Let me ask you the following questions; I think your answers could really help us move towards a solution to this apparent disagreement:

1) Do you believe that the disturbance to a CEV (and the perecption thereof) typically causes the error (discrepency between reference and perception) that leads to the output that opposes that disturbance?

2) Do you believe that it is necessary that the disturbance at least occasionally be represented in the CEV?

3) If you answered "yes" to (2), why do you believe this?

By the way, just by using the term CEV I assume we are discussing a variable that is at least somewhat under control -- like the cursor in a compensatory tracking task. So let's make these questions tangible and assume that they relate to that tracking task.

Best Rick

PS. Gary and Rich -- I appreciate your posts on my research proposal and I will respond to them ASAP. But my priority at the moment (obviosuly) is to get to the bottom of this "information about the disturbance" controversy.

Date: Fri Mar 19, 1993 3:46 pm PST
Subject: One more thing

[From Rick Marken (930319.1500)] Martin Taylor (930319 14:30)--

>With infinite precision perceptual signals and zero transport lag around
>the loop, the perceptual signal is always completely under control and
>the disturbance is never represented there. If either of these conditions
>fails, the passage of the information "about" the disturbance through
>the perceptual signal to the output signal can be detected, and was
>beautifully illustrated in Bill's diagrams that I quoted.

This paragraph is controversial because it is FALSE. If it is not false, Martin should at least be able to show us how to detect the information about the disturbance that passes through the perceptual signal.

It is alarming because it suggests that the output is driven by the the information in the perceptual signal -- a lineal cause effect notion inconsistent with the behavior of a control system.

Best Rick

Date: Sat Mar 20, 1993 9:18 am PST
Subject: Re: One more thing

[Martin Taylor 930320 11:00] (Rick 930319.1300, 930319.0830, 930319.1500)

With the first of these postings (by receipt, as listed above) I thought that Rick and I were well on the way to convergence and agreement. With the second, I was in pretty complete agreement, but with the third, quoted here, I find we seem to be back to square one:

>Martin Taylor (930319 14:30)--
>
>>With infinite precision perceptual signals and zero transport lag around
>>the loop, the perceptual signal is always completely under control and
>>the disturbance is never represented there. If either of these conditions
>>fails, the passage of the information "about" the disturbance through
>>the perceptual signal to the output signal can be detected, and was
>>beautifully illustrated in Bill's diagrams that I quoted.
>

>This paragraph is controversial because it is FALSE. If it is not false,
>Martin should at least be able to show us how to detect the information
>about the disturbance that passes through the perceptual signal.

Exactly what I did in my posting using Bill's diagrams. I can do it no more directly and simplistically than that. At present. I may some day be able to find a demonstration in words of half a syllable, but so far I haven't managed to learn to talk that way.

>It is alarming because it suggests that the output is driven by the
>the information in the perceptual signal -- a lineal cause effect notion
>inconsistent with the behavior of a control system.

I have a feeling that Rick sees S-R psychologists under the bed at night. I find no suggestion of a lineal cause-effect relation here, or that the output is driven by the perceptual signal. It isn't. It's driven by the ERROR signal. The ONLY reason that the information from the disturbance is undetectable in the perceptual signal is that it has been used by the control LOOP and is therefore gone. The output is only the mechanism whereby the CEV is deprived of whatever influence the disturbing variable might have had. Once the information has been used, it's gone (though it can be partially recovered from the output signal if the ECS's imagination loop contains a good system model.) Whatever hasn't been used in control (because of low gain or transport lag) is still there.

In 930319.1300 Rick acknowledges that the information about the disturbance can be found in the perceptual signal if the output signal is zero. I'm not sure why this is relevant, but at least it's a move in the right direction. What I have been trying to show is that CHANGES in the disturbance are detected as CHANGES in the perceptual signal until CHANGES in the output serve to cancel them. And yes, there's no problem in my mind with the changes in the CEV having been due to changes in the reference signal that caused changes in the output signal. A disturbance, to me anyway, is how the CEV differs from what the ECS would like it to be.

If you think this concept of a disturbance is odd, consider again Bill's spring diagram:

```
        A--//////////--B--\\\\\\\\\\--C
```

where A is the disturbing variable, the position of B is the CEV, and C is the output of an ECS.

Now think of the condition before C was connected, i.e. before this ECS began to assume control. A was pulling on B, and provided there is some friction, there is tension on the spring. What is the disturbance? I claim it is undefined, inasmuch as B is "somewhere," a somewhere that is related to an environemental context not needed in the discussion.

That somewhere is the same place it would be if A were not pulling. Maybe friction is insufficient to hold B in place, and B is moving. It is still "somewhere," at any moment. What is the disturbance in this situation? The tension in the A-B spring? No-one is sensing that, either before or after the ECS begins to control the position of B. It seems to me that the concept of disturbance cannot be meaningful until there is control going on (successful or not), which is to say, that there is a perception of the position of B relative to some reference point, which could hardly be anything other than the point corresponding to the reference signal in the ECS whose output is C.

The output signal itself does not mirror the disturbance. Suppose A is pulling, but B is stationary because of static friction. Now the ECS begins to attempt to control, by making C pull if B is to the left of its reference position. It takes a finite pull by C to balance the pull by A, and at this balance point B still does not move. More pull by C does not move B until C pulls hard enough to overcome the static friction. The output of the ECS is a force, but the disturbance is an error of position. And the same applies throughout the period of control. The output is a force, which CANNOT be added to the disturbance (whether it be considered as a force or a position) to produce a perception of position. There isn't even a one-to-one functional relation between the output signal and the perceptual variable in this case, so $p(t) = d(t) + o(t)$ cannot be a meaningful statement. $p(t)$ is a combination of the EFFECTS of the disturbing variable and the output signal on the perceptual signal. Now let's answer the "focus" questions:

>1) Do you believe that the disturbance to a CEV (and the perecption
>thereof) typically causes the error (discrepency between reference and
>perception) that leads to the output that opposes that disturbance?

I think one has to define the disturbance in respect to the error exactly as one defines the CEV in respect to the percept. The perceptual input function defines the CEV (provided you include the operations of all the PIFs in lower ECSs that provide the sensory signals to this ECS). The perceptual signal is the current representation of the state of the CEV. From an outsider's viewpoint, in which the perceptual signal and the CEV can be independently measured, they may differ. The same applies to the error and the disturbance. From inside the ECS, the perceptual signal is all there is.

>2) Do you believe that it is necessary that the disturbance at least
>occasionally be represented in the CEV?

The disturbance is not "represented" in the CEV. The disturbing variable may affect the CEV, but representation occurs within the ECS.

I suppose one could say that the reference signal can be represented (by an outsider, not the ECS) in the CEV space, but this is dubious, since the PIF may well not be invertible.

(Interesting, come to think of it; this is exactly the inverse problem of the classical roboticists. They worry about the non-uniqueness of the inverse kinematic problem, and here I am saying that "many means can achieve the same goals" can be seen as the same inverse perceptual problem).

If in some way the reference signal could be represented in the CEV, then I guess one could say the disturbance could be represented in the CEV, but at the moment this seems to me to be making a VERY big stretch.

>3) If you answered "yes" to (2), why do you believe this?

I didn't say "yes."

>...it seems to me that you are trying to

>squeeze the beautiful concept of control of perception into the
>Procrustean bed of conventional linear cause effect thinking.

I find this ironic. Long before I heard of PCT, I was a missionary for eliminating ALL cause-effect claims, on the ground that in the whole universe everything is interrelated; the effects of actions react on the actor, environmental conditions must be right for "causes" to have "effects," everything is under a multitude of influences. Cause-effect thinking ruins our justice system and our economy just as surely as it is a baneful influence in psychology. (Cause-effect thinking leads to the notion that government deficits can be cut by reducing spending, for example). So, for my attempts to analyze the control LOOP to be characterized (repeatedly) as a support for S-R approaches or as an attempt to maintain an S-R view strikes me as a bit ludicrous. But I can accept that you might see my comments as lending aid and comfort to the enemy. There may be something in the wording that triggers your "Reds under the bed" detector. I don't see it, so I have a hard time controlling it.

I am beginning, at last, to have the feeling that we really do see in the same way what happens in a control loop. I don't know whether I am changing my opinions (it doesn't feel as if I am), whether you are (it doesn't seem so from here) or whether I am simply coming to a better understanding of what has been underlying your objections to my postings (I hope that's it). Whenever you say positively what occurs, I have no problems. It is always as I have understood the situation. I have problems only with the negative things. The differences are in what to take from that common understanding.

But I also see that most times somebody else (I mean other than us) says that they "agree" with another disputant, I get the feeling that they don't.

General note. This is likely to be my last posting before at least Thursday, (Tuesday is possible but unlikely) and maybe until the following week. So barbed arrows may have a long travel time, and earthquakes will remain undetected for a while.

Martin

Date: Sat Mar 20, 1993 11:21 am PST
Subject: information

[Avery Andrews 930321.0509]

I'd agree entirely that the ECS `knows' nothing about the disturbance, but not with the claim that the perceptual signal contains no information about it. After all, we might claim that a video cassete contains the information that Rosemary erased the tape, even though both it and the VCR it gets played on aren't able to extract this information. If the reference signal, output system, and connection between output and perception are constant, we can reconstruct the disturbance from the perceptual signal, so some info is there, I'd say (but not for the ECS).

Avery Andrews@anu.edu.au

Date: Sat Mar 20, 1993 11:31 am PST
Subject: coordinative structures

Looking at Rick's papers `The Perceptual Organization of Behavior' and `Degrees of Freedom in Behavior', the first one seems to me to present a real implementation of a `coordinative structure', while the latter seems to show a way in which a coordinate structure can superficially appear to exist but not actually be there (if the disturbance is coordinated in 2 dfs, so will be the output, even if control of the df's is independent.

Thinking about P.O.B. and my own airplane example, I also find that the coord. struct. contains a lot of ambiguity. Rick's system and my (hypothetical airplane), for example, control for several quantities that happen to be the same, so that if one of them meets an overwhelming disturbance, the others will follow their reference levels, & coordination will be destroyed. One can imagine systems where the fact of coordination is more important that the reference level (two arms, each holding one side of a tray: the main thing is that they stay level, the exact height of the tray off the ground doesn't matter so much. My aperture-control example (the one with the thrusters) is similar in that it controls for an (objective, analyst's perspective) relationship between two things (the distance between two points) without caring about actual position. I think something like this is more like what coordinative structures are supposed to be about, but the prose coming with the concept is murky enough so that I am not sure.

Avery.Andrews@anu.edu.au

Date: Sat Mar 20, 1993 3:41 pm PST
Subject: Re: Why it's important

From Ken Hacker [932003]

The discussions about information are fascinating since in communication studies we study information, how it is produced, used, etc. etc. I have just one question from the peanut gallery at this time: How is possible to have a signal, any signal, whether electrical impulse or human sign, without some information value or content to the signal? In other words, a signal, signum, is related to the notion of sign -- something indicating something else, by definition. Any signals on this will be appreciated. BTW, feedback IS given by other people!
Ken Hacker

Date: Sat Mar 20, 1993 8:33 pm PST
Subject: Reconstructing disturbances; information

[From Bill Powers (930320.2100)] Avery Andrews (930321.0509) --

>After all, we might claim that a video cassete contains the
>information that Rosemary erased the tape, even though both it
>and the VCR it gets played on aren't able to extract this information.

This is a fine example, but not of the point you want to make. If I hand you a tape and on playing it you find that there's nothing recorded on it, all you know is that you have either a blank or an erased tape. The state of the tape contains no information about the process by which it got that way, much less the agent involved.

>If the reference signal, output system, and connection between
>output and perception are constant, we can reconstruct the
>disturbance from the perceptual signal, so some info is there,
>I'd say (but not for the ECS).

Suppose the disturbing variable is a constant 10 units, and the output is a constant -9.9 units, both measured in terms of effect on the CEV when acting alone. The perception, referred to the environment, is 0.1 units. In general, a value of the perception of x units predicts that the output is x units less than the disturbance, both referred to effects on the perceptual signal. It does not predict either what the output is or what the disturbance is. You must know either the disturbance value or the output value to reconstruct the other from the value of the perceptual signal. If you know neither the output value nor the disturbance value, the value of the perceptual signal predicts nothing. I emphasize that by "disturbance" I mean the external variable acting on the CEV, not the CEV itself. The same holds true whether you talk about the CEV or the perceptual signal.

This is a such a simple point! I think that the high-powered brains out there are trying to make the argument far more complicated than it is.

If I tell you that the value of the perceptual signal (or CEV) is a constant 3 units, what is the value of the disturbing variable? This is much like the old joke: "Here is a late score from the big game: 3 to 2."

>My aperture-control example (the one with the thrusters) is
>similar in that it controls for an (objective, analyst's
>perspective) relationship between two things (the distance
>between two points) without caring about actual position.

This is really quite easy to accomplish using a control system that perceive the difference in position between the two objects. The same system could be used to hold a tray level. When Wolfgang Zocher recovers enough to put the last touches on SIMCON, I'll post a program that shows how it works.

From Ken Hacker [932003]

>I have just one question from the peanut gallery at this time: How is
>possible to have a signal, any signal, whether electrical impulse or
>human sign, without some information value or content to the signal?

The information "contained" in the signal depends on the receiver more than the transmitter. When you tune a radio away from a station, you lose the information from the station, but you gain information about the ambient electronic noise level -- if that's what you're perceiving. If what you want to perceive is the electronic noise level, then accidentally tuning in a station giving the news amounts to drowning the "signal" in "noise." If you want to hear the news around here, but tune in to one or two of the stations at the wrong time of the day, you will get no information at all -- unless you understand Navaho.

The PCT view of "transmitting information" is (if I may presume) that verbal or other communication does not _carry_ meaning; it _evokes_ meanings from the receiver's own memories, and then only in terms of the perceptual functions the person brings to the communication. Martin Taylor, on other grounds, has also concluded that the idea of transmitting information is a misunderstanding.

You only have to follow the arguments on this net for a while to realize that what one person writes is not what another person reads.

>BTW, feedback IS given by other people!

No, _input_ is given to you by other people. It isn't feedback unless your actions gave rise to it. You provide feedback only for yourself (although it may not always match your reference signals).

Best Bill P.

Date: Sat Mar 20, 1993 11:06 pm PST
Subject: feedback, coordinative structures, mass spring

[Avery Andrews 930321.1445]

Another aspect of the coordinative-structures `vs.' feedback issue that occurs to me is that neither Rick Marken nor the mass-spring weenies seem to point out that, viewed on a suitable time-scale, a feedback system simply *is* a mass-spring system - when a disturbance pushes the perception away from the reference level, the output of the system exerts a restoring force, opposite and at least varying monotonically w.r.t. the extent of the departure.

A point on the Kelso, Turvey, etc. side is that mass-spring behavior can appear without neural feedback circuits. When you tense your muscles, for example (co-contracting agonist and antagonist), you get a system with a rather high spring-constant and extremely fast response, since the restoring force is produced by the elasticity of various tissues. But this fast response is rather expensive to maintain, since co-contracting muscles uses up a substantial amount of energy. So for slower or more predictable disturbances (such lifting coffee-mugs or books of varying weights, or holding toddlers by the hand, dogs on the leash, etc.), mechanically derived spring-behavior is not very appropriate, and the feedback-mediated version is called for.

Avery.Andrews@anu.edu.au

Date: Sun Mar 21, 1993 6:15 am PST
Subject: Dag's video

To: Dag Forssell and interested others From: David Goldstein
Subject: Introduction video tape Date: 03/21/93

Thank you for letting me review your video. I was hoping that I could use it at the place where I work which is a residential treatment center for adolescents. I see that it was tailored to the Deming group you addressed and more generally, to a business/industry audience. Nevertheless, I think there are parts of the presentation I thought which would be appropriate for any audience.

I will leave it to Bill, Rick and others to comment about the technical side of your talk. The use of the word "instruction" for error signals bothers me a little bit. People might get the idea that an instruction is a command to do something.

Ed Ford's phrase perceptual difference avoids this danger. Perhaps calling it the perception/comparison difference would also work.

The role playing was a particularly effective part of the presentation. It concretly showed the results of one way versus the HPCT way of doing things. I think that people will remember this.

The story about your versus Christine's idea of marriage was very enjoyable and effective. It really illustrates the difference between two people's wants and how that can lead to interpersonal conflict. The role play applies this to the working situation. As part of the role play, a discussion takes place of the clerk's idea of being a good worker versus the manager's idea. The manager makes it clear that being on time is part of the company's definition. The manager is being supportive by helping the clerk find a plan to meet this part of the definition.

Towards the end of the tape, you present a summary of the major implications of HPCT for managers. I suggest that it might be a good idea to present this at the beginning even though it will not make complete sense at that time.

In closing, I wish you the best of success in your enterprise. I would be happy to review any other tapes connected with your program. From the written material provided with the tape, I would especially be interested in how you present some of the topics in days 2 and 3, namely, the topics of: feelings, want selection.

Warmest regards, David Goldstein goldstein@saturn.rowan.edu

Date: Sun Mar 21, 1993 7:26 am PST
Subject: Mass-spring model: impossible.

[From Bill Powers (930321.0730)] Avery Andrews (930321.1445) --

>A point on the Kelso, Turvey, etc. side is that mass-spring
>behavior can appear without neural feedback circuits.

What do you recommend -- neurosurgery to remove the feedback connections? They exist. You can't ignore them.

>When you tense your muscles, for example (co-contracting
>agonist and antagonist), you get a system with a rather high
>spring-constant and extremely fast response, since the
>restoring force is produced by the elasticity of various tissues.

The high spring constant is produced at high tension because of the quadratic relationship between force and the elongation in muscle. It is just as present in a control-system model, which uses the same muscle model that the mass-on-a-spring model uses. In fact, the control system model IS a mass-spring model embedded in a feedback loop.

I have shown that the quadratic muscle model leads to a LINEAR spring constant when the common-mode excitation to the opposing muscles is held constant, while the driving signal consists of one signal rising above the mean value while the other falls below it by the same amount. Thus the Little Man v. 2 model can be set to reproduce the behavior of the mass-spring model under any assumed constant average muscle tension. You could even set the parameters of the lower-level systems to eliminate the feedback (set the sensitivity to zero) and look just at the effect of relaying the reference signal directly to the muscles without feedback. Then you would have exactly the mass-spring model.

The principal difference between the PCT model and the MS model is in the way it explains where the driving signals come from. In the MS model, the signal waveforms entering the muscles must be computed and emitted blindly. To move the arm rapidly from one stationary position to another, the mass-spring model must receive a signal waveform (a balanced push-pull waveform) that starts at some initial level required to support the arm, then

1. Rises abruptly to a large value to accelerate the arm to a high angular velocity,

2. Declines to zero and goes to a smaller negative peak to decelerate the arm as it nears the final position, and

3. Makes one or two small adustments before returning to the initial state to hold the arm in the final position.

The waveform of the driving signal simply reflects what is required, according to the physiological properties of muscle and the inertial properties of the arm, to reproduce the trajectory of movement that is observed. No matter what model you use, it must generate exactly the same signals with the same waveforms to operate the muscles, assuming that the muscle model is reasonably correct.

The control-system model derives the driving signals from sensors and input functions with very simple dynamic filtering, plus comparison with a reference signal specifying the desired final position. The positions and movements of the arm themselves generate the required variations in the driving waveforms. When the reference signal is abruptly set to its final value (where it remains constant during the movement), the arm lags behind, leading to a large peak in the error signal and hence a peak in the driving signal which gets the arm moving rapidly. As the arm velocity builds up, the rate feedback quickly reduces the initial peak acceleration; as the position error declines, the rate feedback causes the driving signal to reverse, producing deceleration. As the final position is approached, the falling velocity reduces the rate feedback, decreasing the deceleration to zero just as the target position is reached and the velocity reaches zero. All this happens without any explicit computation of the waveforms and without any special provision for timing of the signals. Both the position and the velocity feedback signals return to their initial states automatically, just at the right time, as a natural consequence of the feedback process. From reading Kelso, Bizzi, et. al., you might get the impression that all that is required to move an arm to a new position is to set the driving signal to a new value. Various authors like to point out that the movement of a simple one-joint arm is described by a second-order differential equation, and that with proper choice of parameters the equation describes an optimally- damped movement to a new position. All that is quite true.

However, the behavior of a control system using the same muscle and arm model is ALSO described by a second-order differential equation, and with proper choice of parameters it, too, predicts an optimally damped movement to a new position when (now) the reference signal is set to a new value. So what's the difference?

The difference is that Kelso et. al. assume that when they find the parameters (spring constant and damping) needed to make the open-loop model behave correctly, those are the parameters of the muscle. This is extremely unlikely to be true. The natural muscle spring constant, the mass of the arm, and the natural damping due to the muscle would combine to produce a series of damped oscillations if the driving signal were simply switched to a new value.

In fact it's perfectly clear that when an arm moves from one position to another at maximum speed, the muscles do NOT simply come to a new state of tension with the arm then finding a new equilibrium position between them. The agonists tense at first and the antagonists relax, as one would expect from a new value of driving signal. But during the movement, the agonists then relax and the antagonists tense to produce decleration, and finally both sets of muscles return to a constant new state of excitation. Furthermore, the initial force generated in the agonists is many times the final force needed to suspend the arm in its final position -- judging from the accelerations observed, 10 to 20 times as high. And the time-integral of all these effects leaves the arm stationary exactly at the desired new position.

These variations in muscle contraction during the movement mean that the driving signal does NOT just come to a new constant value and stay there. The signals reaching the muscles are adjusted during the movement, by amounts many times their final values. From this it follows that the driving signals must be patterned through time, and in an open-loop model this leads inevitably to the idea of a pattern generator that uses computations of the inverse dynamics of the arm to deduce what pattern is necessary to reproduce a given trajectory and velocity of movement.

Having refused to take the feedback path of explanation, the theorist is forced to follow the open-loop path step by logical step until the absurdity of the central pattern generator and the impossibility of the inverse kinematic computations are revealed in all their horrifying detail. At this point, continuing down this blind alley becomes a matter of pride, stubbornness, and defense of reputation.

In the control-system model, changing arm position is simply a matter of altering the value of the reference signal to new amount and leaving it there. All the required variations in the signal waveforms entering the muscle are created by position and rate feedback during the movement. The arm is very stiff against rapid disturbances, far more than it would be without the rate feedback -- although, like real human arms, it exhibits a delay that a passive muscle model couldn't possibly explain.

I don't see that the passive mass-spring model is in the running at all, for fast movements, slow movements, or posture maintenance. It leads to absurdities; it ignores the fact that muscles change their states during a movement, or if that fact is taken into account, it forces the argument to deduce a pattern generator, which in turn least to the ridiculous notion of computing inverse dynamics in real time. That model is just wrong.

Best, Bill P.

Date: Sun Mar 21, 1993 9:29 am PST
Subject: Dag's video

[From Dag Forssell (930121 09.20)]   Re: David Goldstein (930321)

Appreciate your post. You now qualify for the update, which is twice as educational and, I hope, more enticing.

I did not label errors "instruction." I talked of them as dissatisfaction. It is output I have labeled "instruction." The revised video will be much more clear on these things.

Let me know what your charges think of the update.

Best, Dag

Date: Sun Mar 21, 1993 11:19 am PST
Subject: Reconstructing disturbances; information

From Ken Hacker [932103]   Bill argues the following (932103) --

>The PCT view of "transmitting information" is (if I may presume)
>that verbal or other communication does not _carry_ meaning; it
>_evokes_ meanings from the receiver's own memories, and then only
>in terms of the perceptual functions the person brings to the
>communication. Martin Taylor, on other grounds, has also
>concluded that the idea of transmitting information is a misunderstanding.

I agree totally (as do most communication theorists). Information is produced in relation to receivers and what questions and ideas they have. Still, any signal, sign, symbol, or message varies in its ability to evoke certain feelings, interpretations, thoughts, etc.   All signals are not created equally!

>No, _input_ is given to you by other people. It isn't feedback
>unless your actions gave rise to it. You provide feedback only
>for yourself (although it may not always match your reference signals).

This is splitting hairs on another word that can never be perfectly unambiguous. However, I agree that feedback is the result of a feedforward process initiated by the receiver signals from others. But there is no good reason to think that any human cannot stimulate feedback to the signals sent to other humans. This is how human interaction begins --usually with nonverbal signals followed by verbal message exchange, but always with loops of messages sent and messages fed back which the sender uses to regulate the sending process. Of course, the regulation at an internal level, i.e. control, is the produce of self-signalling processes. But those processes are not independent from the loop of messaging sending and conscious monitoring of messages sent (fed) back.

Thanks for the feedback, Bill.   Ken Hacker

Date: Sun Mar 21, 1993 11:37 am PST
Subject: Need help, Rick!

from Ed Ford (930321:0625)

Rick, I think I understand what you are saying about the disturbance and the input signal. It would be helpful if you would use some examples, especially using people, to illustrate your point. Can what you are saying be experienced when doing the ruuber band demo, especially the one I suggested using the pointed finger?

Joel, did you get my post. I need your address and phone number.

Dag, I'll look at your tape in the next three or four days, I'm presently playing catch up.

Best, Ed

Date: Sun Mar 21, 1993 1:01 pm PST
Subject: Basics, disturbance causes error?

[From Rick Marken (930321.1200)] Martin Taylor (930320 11:00) --

>so p(t) = d(t) + o(t) cannot be a meaningful statement.

If you really think that p(t) = d(t) + o(t) is somehow not a meaningful statement, then we really are poles apart; your statement seems to me to be a denial of the physical situation in a compensatory tracking task. If you take a look at the HyperCard conflict stack that I sent you, you will find in the tracking loop the following code

```
put trunc (A1*H1+B1*H2+distx) into xp
```

xp is a number representing the horizontal position of the point to be plotted; it corresponds to p(t) -- with t being the index of the iteration on which this is happening. A1 and B1 are constant multipliers. H1 and H2 are the numbers corresponding to the x and y positions of the mouse on the current iteration -- they are equivalent to o(t); finally, distx is the value of the disturbance on this iteration; of course, it corresponds to d(t). So, on each iteration of the program, the horizontal position of the cursor on the screen is proportional to the sum of the current disturbance value (distx) and the current value of the mouse output (H1 and H2). You keep saying that there might be a lag in the production of o(t) and I keep trying to point out that this is not important. I hope thinking in terms of the computer program will help you understand this. H1 and H2 are the CURRENT values of the mouse; they might have been caused (inside the subject) by processes that happened hours earlier -- but that doesn't matter in terms of what the subject is SEEING NOW; what the subject is seeing (p(t)) is the combined effects of the PRESENT distrubance (distx) and the PRESENT output (H1 and H2). So can we please agree that, at least in a compensatory tracking task (like the one in the conflict stack) p(t) (or, at least, the CEV -- the position of the cursor on the screen) is ALWAYS proportional to the SUM of dirturbance and output -- ie. p(t) = o(t) + d(t)? Run the HyperCard stack and see if you can convince yourself this is really what is happening -- what you are seeing when you do the tracking task (unless you don't move the mouse at all) is o(t) + d(t). Of course, when you are not moving the mouse at all you are not controlling -- so there is information about d(t) in uncontrolled variables, as Bill P. has so politely pointed out.

This remarkably simple (but apparently unpleasant) little fact is the reason why I keep saying that in a control system "there is no information about the disturbance in p(t)". All you are ever perceiving (in tracking task or "real life") is the combined result of what is "happening" (d(t)) and what you are doing while it's happening (o(t)).

Martin says:

>Now let's answer the "focus" questions:

>>1) Do you believe that the disturbance to a CEV (and the perecption
>>thereof) typically causes the error (discrepency between reference and
>>perception) that leads to the output that opposes that disturbance?

>I think one has to define the disturbance in respect to the error exactly
>as one defines the CEV in respect to the percept. The perceptual input
>function defines the CEV (provided you include the operations of all
>the PIFs in lower ECSs that provide the sensory signals to this ECS).
>The perceptual signal is the current representation of the state of the
>CEV. From an outsider's viewpoint, in which the perceptual signal and
>the CEV can be independently measured, they may differ. The same applies
>to the error and the disturbance. From inside the ECS, the perceptual
>signal is all there is.

Did anyone see an answer to my question in there? I still don't know what you believe, but the correct answer is very simple: No. Since e = r-p and p = o + d, e = r - o + d so the error is determined by the combination of output and disturbance (assuming constant r). The disturbance DOES NOT cause the error in a control system.

>>2) Do you believe that it is necessary that the disturbance at least
>>occasionally be represented in the CEV?

>The disturbance is not "represented" in the CEV. The disturbing variable
>may affect the CEV, but representation occurs within the ECS.

So the answer is no? Good. But what's this "the representation occurs within the ECS"? You don't mean that it is represented in p(t)? Good. So we have two remaining signals in an ECS -- reference and error. Where is the disturbance represented? The correct answer is? In neither of them. It's obviously not in r(t) because d(t) does not show up in it's calculation. It's not in e(t) either because e(t) = r - o(t) + d(t) so again d(t) itself is not represented in the signal -- just the sum -- d(t) + o(t). So there is no representation of d(t) ANYWHERE in a control system -- not in p(t), r(t) or e(t); it only exists in o(t) as the integrated, amplified error -- even though d(t) is NOT represented in the error itself.

>>3) If you answered "yes" to (2), why do you believe this?

>I didn't say "yes."

Correct. So I still don't know what you believe.

> So, for my attempts to analyze the control LOOP to be
>characterized (repeatedly) as a support for S-R approaches or as an

>attempt to maintain an S-R view strikes me as a bit ludicrous.

I'm not trying to make a political point; I think it's wonderful that you don't support S-R approaches to behavior. It's just that you keep describing the operation of a control system incorrectly, and the mistakes you make (such as saying that there is information about the disturbance in p(t) or, now that you have apparently backed off of that claim (see above), that such information resides somewhere inside the control system) are precisely those that are necessary to maintain an input-output view of the operation of such a system. It makes it possible to view the ultimate cause of output (o(t)) as something outside the system (the information about d(t)).

> But I can accept that you might see my comments as lending aid and comfort
>to the enemy.

I think we have different "enemies". You have problems with SR theories of behavior. I have problems with those ideas too -- but, more importantly, I have problems with the whole conventional approach to studying behavior -- both its goals (finding relationships between IVs and DVs) and it's assumptions (that such relationship reveal something about how behavior works). I am arguing so much with you because I think your view leaves open the door to conventioal research methodology. After all, if there is information about d(t) at some point along the path from p(t) to o(t) then it seems reasonable to assume that one could learn something about how that information is "processed" by looking at relationships between d(t) and o(t). But this, according to PCT, is what the conventional IV-DV approach is already looking at -- and finding at least statistical relationships between these variables. But we (PCTers) already know that if the system you are studying is a control system then the function relating d(t) to o(t) is the inverse of the feedback function relating o(t) to p(t) -- it has NOTHING TO DO with the processes inside the organism that convert d(t) into o(t) (in fact, they are converting p(t) into o(t)).

>I am beginning, at last, to have the feeling that we really do see in
>the same way what happens in a control loop.

I hope so. There are some easy ways to test this. You could try again to answer my question 1) above so I could see how you think the nitty gritty of a control loop works -- is error caused by the disturbance in a control loop? Or you could tell me (as an opponent of SR theory) whether the typical IV-DV experiment reveals anything about the workings of the control systems under study (assuming that they are control systems); for example, what do we learn about the subject in an experiment in which we manipulate the amount of information in the stimulus -- the IV -- and measure the time to respond to each stimulus -- DV)? My answer would be "nearly nothing".

I am adamant about this because I know (from experience with the Karolyans) that unless people start doing their research in a new way (trying to test for controlled variables, for example) PCT will just become a new language for talking about the SOC (same old crap). PCT is not a new language; it is a completely new model of behavior in which behavior (as dependent variable) is controlled perception, not caused output. In order to test this model, a whole new approach to research is needed -- one that is not only non-statistical (one subject at a time) and highly precise (accept only findings with less than 1% error) but based a proper understanding of the control model.

Best   Rick

Date: Sun Mar 21, 1993 5:49 pm PST
Subject: Sequence control, info demo

[From Rick Marken (930321.1700)]

At long last, here are some brief responses to the comments on my reserach proposal on the control of sequences.

Gary Cziko (930318.1648 GMT) --

>So if I understand Rick's thinking, he would predict that the subject could
>accurately push a key to restore the correct intensity ("dyamics") at a
>quite quick tempo, control timbre ("get that trumpet out of there, it's
>supposed to sound like a flute") at a slower tempo, and sequence at a still
>slower tempo.

Yes. I also think it's a great idea to do this test in different sensory "modalities" -- vision, hearing, touch. I presented perceptual evidence in the "Hierarchical behavior of perception" paper that the "speed limit" (bandwidth for all you real engineers out there) for sequence perception was about the same whether it was a visual, auditory or touch sequence. HPCT (as currently constructed) would predict that a sequence is a sequence; the types of lower order sensations (modalities) or configurations that make up the sequence should not matter. So I predict that a sequence of notes will have the same speed limit for control as a sequence of letters. The problem with notes (as you mention) is that a fast sequence of notes can be perceived and controlled as an event. In the perceptual studies, what was used were different types of sounds in sequence (buzz, hiss, tone). But there is still the possibility of this sequence becoming an event; even a sequence of letters could conceivably become an event; that's why I projected the letters at alternating positions on the screen.

Experiments is experiments -- even when testing for controlled variables you must try to eliminate confounding variables; psychological science will still be a fun game, even when it's done right -- using PCT.

Richard Thurman (930318.1630) --

>Rick can probably set up the HyperCard stack so that it will be a simple
>adjustment to go from visual to auditory presentation. (I'm assuming
>he is going to use HyperCard.)

Actually, I wrote the latest version in QuickBasic. But the whole point if for you guys to do the work!! I just want to stand on the side and kibbitz (make irritating suggestions about what to do based on nothing better than the fact that I'm not actually doing anything).

>It would be interesting to see if different
>patterns emerge with different perceptual modalities.

Indeed. I'm hoping for similar patterns but I'll take what I get.

Ed Ford (930321:0625) --

>Rick, I think I understand what you are saying about the
>disutrbance and the input signal. It would be helpful if you
>would use some examples, especially using people, to illustrate
>your point. Can what you are saying be experienced when doing the
>ruuber band demo, especially the one I suggested using the pointed finger?

Yes and no.

Yes, using the rubber bands you can show that the position of the knot (p(t)) is always a result of what the subject is doing to his or her end of the rubber (o(t)) and what you are doing to the "distrubing" end of the rubber band (d(t)). You might be able to show (when you move the distrubance real slowly) that the position of the knot does not change as you might expect it to if just the disturbance were acting. For example, when you pull gently to the left, the knot might be expected to move correspondingly to the left. But the subject might be able to notice that sometimes the knot is actually moving to the right(due to the added effects of their own actions) while you are pulling to the left. This means that the position of the knot is not a stimulus that "tells" the subject how to pull on their rubber band to correct the disturbance. So the stimulus response view of control cannot be preserved even when the actual variable (the knot) that the subject is controlling. is discovered.

No, because there is a somewhat deeper point in the "no information about the disturbance in perception" discussions. One point (made above) is that the controlled variable [the knot] (or the perception thereof) is not a stimulus for output. But the deeper point is that perception is just there -- it is neither right no wrong, good or bad, in error or out of error -- it is not INFORMATIVE; it just IS. The position of the knot is just the position of a knot -- BUT once you have a reference regarding where it should be then it seems like some knot positions are definitely "wrong" and one particular one is "right". This is a tough point to demonstrate because people don't care much about the position of knots and when you tell them this it seems pretty trivial. But try to explain that this applies to ALL perceptions that are controlled and you will get some strong reactions. People who are controlling for the neatness of their house have a difficult time believing that the neatness of the house is just a perception -- when the house looks "messy", that perception seems just plain wrong. It is difficult to demonstrate that perceptions are just perceptions -- and that they only become "good" or "bad" or "right" or "wrong" (ie. they only become informative) with respect to ones owns references for them. That fact is easy to demonstrate with "knot" positions but a hell of a lot more difficult to demonstrate with political, religious and economic "positions".

Ken Hacker [932103]--

>I agree totally (as do most communication theorists). Information is
>produced in relation to receivers and what questions and ideas they have.
>Still, any signal, sign, symbol, or message varies in its ability to
>evoke certain feelings, interpretations, thoughts, etc. All signals are not
>created equally!

Yes, but this is because all receivers are not created equally. The differential ability of any signal, sign, symbol of message to evoke certain feelings, interpretations,thoughts etc is a result of differences in the characteristics (perceptual functions, reference signals) of the receivers -- not of the signals, etc, themselves. At least, this would be the PCT perspective.

Best Rick

Date: Mon Mar 22, 1993 10:17 am PST
Subject: Misc. from Mary

[from Mary Powers]

Gary: re Closed Loop as a journal, and the idea that progress reports, etc. can be read on the net - you may not realize that HALF of our paid-up membership is NOT on the net - CL is their only way of knowing what's going on.

Lars: I sent you the PCT bibliography. Didn't you get it?

Dick R.: You can get lodging info and tourist brochures from Durango by calling 1-800-525-8855.

Dag: Can you send me Toto's current address?

Ken Hacker: I think you've expressed the Carver/Scheier/Karoly/etc/etc version of control theory pretty clearly. This is how to talk about control theory without having to change your mind about traditional psychology.

>any signal, sign, symbol or message varies in its ability to
>evoke...
>there is no good reason to think that any human cannot stimulate
>feedback...

Ascribing "ability" to a message and "stimulating" feedback are precisely the brick walls PCT is up against. All the evoking is in the control system, not the message, which is why the message's ability seems to vary. Feedback is a property of the system, it does not get stimulated or turned on (or off) by an outside agency.

>However, I agree that feedback is the result of a feedforward
>process initiated by the receiver signals from others.

Agree with who? Feedforward, as far as I can tell, is a buzzword that vaguely resembles the reference signal. It is by comparing received signals to the reference signal that action is initiated. This gives the appearance that the action is a response to the "initiating" signals. The reason the input signals seem to vary in their ability to evoke a response is because they may or may not produce a perceptual signal that matters (differs from the reference signal).

It is not splitting hairs to define feedback as Bill Powers does. It is fundamental to PCT. The difference between feedback and input is fudged over in the self-regulation literature. It is a consequence of using control theory metaphorically rather than rigorously. I'm afraid most psychologists do not know the difference, having never before encountered a real model.

Mary Powers

Date: Mon Mar 22, 1993 11:27 am PST
From: CZIKO Gary
          MBX: g-cziko@uiuc.edu
Subject: Video

Dag:

I would like to have your revised video sent to me via airmail.

I will put $13 in the mail to you today or tomorrow (if I can find your snail address).--Gary

Gary Cziko   Telephone: 217-333-8527
Educational Psychology   FAX: 217-244-7620
University of Illinois   E-mail: g-cziko@uiuc.edu
1310 S. Sixth Street   Radio: N9MJZ
210 Education Building
Champaign, Illinois 61820-6990

Date: Mon Mar 22, 1993 12:20 pm PST
Subject: lobsters too

[From: Bruce Nevin (Mon 930322 14:47:19)]

Belated note about an article on nervous systems of lobsters, plus some stuff on leeches. Looks like a good place to look for useful neurobiology literature because of the simplicity of these systems, e.g. about 1000 neurons in the lobster's stomach, in four ganglia.

The article is in _New Scientist_ for 27 February 1993, 24-28. Although it is marred by a PopSci tone and lack of references, and even more by commitment to presuppositions about "neural commands" it does name major players to follow up in the literature, and describes some intriguing findings. They were looking for pattern generators for stomach rhythms (lobster) and heart rhythms (leech). In each case, they talk about how neurotransmitters affect how motor neurons "reinterpret" the "commands" issued to them from the CNS. It would take someone better versed by far than I in the issues to fathom what might be going on in PCT terms. There does seem to be a CPG in the lobster for its pylorus, which functions as a kind of filter between the teeth in the first part of the stomach (the "gastric mill") and the midgut, a single neuron called the "anterior burster":

> The main reason the neuron fires rhythmically is that its membrane contains a complement of different ion channels that open and close in a sequence that alterntely drives the neuron to fire impulse bursts and then suppresses this activity. (26)

In the leech heart, on the other hand, there are pairs of neurons, each of which in turn inhibits the other. "In effect, the neurons act as a biological version of a simple electric oscillator." (27)

What they go on most about is variation of these rhythms determined by neurotransmitters, some with long-lasting effect in the environment of many cells in a network, others with short-term effects at a synapse (I may be overstating the short-term/nearby long-term/widespread correlation). And some hints about multifunctionality of supposedly specialized neurons.

Authors: Michael Nusbaum, Physiology & Biophysics at U. Alabama Birmingham, Ronald Calabrese, Biophysics, Emory, Atlanta, GA.

        Bruce   bn@bbn.com

Date: Mon Mar 22, 1993 12:28 pm PST
Subject: Saying it like it is

[From Rick Marken (930322.1100)]   Mary Powers --

>Ken Hacker: I think you've expressed the
>Carver/Scheier/Karoly/etc/etc version of control theory pretty
>clearly. This is how to talk about control theory without having
>to change your mind about traditional psychology.

>Ascribing "ability" to a message and "stimulating" [to] feedback are
>precisely the brick walls PCT is up against. All the evoking is
>in the control system, not the message, which is why the
>message's ability seems to vary. Feedback is a property of the
>system, it does not get stimulated or turned on (or off) by an outside agency.

Mary, Oh Mary, Why don't you post more often? In a few graceful words you have explained the whole point of this "information about the disturbance" series clearly and precisely. I only disagree with you about one thing -- what we run into with the Karolyans is not a brick wall, it's an anti-matter force field.

>The difference between feedback and
>input is fudged over in the self-regulation literature. It is a
>consequence of using control theory metaphorically rather than
>rigorously. I'm afraid most psychologists do not know the
>difference, having never before encountered a real model.

Again, right on target. But I'm sure glad you said it instead of me.

Love

Rick (ready to let anyone else take the heat for saying it like it is) Marken

Date: Mon Mar 22, 1993 1:22 pm PST
Subject: Misc. from Mary

From Ken Hacker [932203]   In response to Mary Powers:

Your response to my message (note how my messages GENERATED your GENERATION of messages; they did not arise from thin air) is an interesting illustration of why PCT is up against so many barriers. There is more semantic tussling going on than real explanation in terms of processes we can all appreciate from various non-One Way perspectives.

>Ken Hacker: I think you've expressed the
>Carver/Scheier/Karoly/etc/etc version of control theory pretty
>clearly. This is how to talk about control theory without having
>to change your mind about traditional psychology.

I am not interested in defending traditional psychology and any suggestion that I am, is patently false.

>Ascribing "ability" to a message and "stimulating" feedback are
>precisely the brick walls PCT is up against. All the evoking is
>in the control system, not the message, which is why the
>message's ability seems to vary. Feedback is a property of the
>system, it does not get stimulated or turned on (or off) by an outside agency.

Of course the evoking is in the control system, but that is not the issue. In communication, one control system is providing disturbances to another. Those disturbances, in the form of messages, vary greatly in responses or reguation by receivers. The messages mean different things to different control systems. Certainly. But for any single control system, the variability of messages and message responses is important. For example, I would respond much more kindly to PCT if the messages here were less of what I PERCEIVE as arrogance. I would think and respond differently to something as simply as tone. Is tone information; I think it is. Is tone made up by me? No, it is perceived by me and I will have different perceptions for different messages. The problem with PCT as you describe it is that it denies the social aspects of what is social: human INTERaction.

>The reason the input
>signals seem to vary in their ability to evoke a response is
>because they may or may not produce a perceptual signal that
>matters (differs from the reference signal).

I agree. Remember, now, from your argument, if you are taken back by any of my comments, it has nothing to do with what I am saying; it is all inside of you!! From my argument, we both have reference signals, comparators, etc. etc., and will compare and adjust what we are reading and writing and negotiate messages and interpretations to help us understand more about ourselves and each other.
Ken Hacker

Date: Mon Mar 22, 1993 9:43 pm PST
Subject: Information is in the world?

[From Rick Marken (930322.2100)]

Chuck Tucker -- I've been wrestling with you private posts all day. Let me just respond quickly to one of your comments; it took me some time to finally get it right (I think).

>It seems to me that I can state with great confidence and you would agree
>with it that statements which appear to you to transform PCT into an
>S-R like formulation are disturbing to you (I can think like you so I
>have an advantage).

This is a colloquial meaning of disturbance. The statements which "appear to me to transform PCT into an S-R like formulation" are just perceptions (or the cause of perceptions, if you like); They are not necessarily disturbances in the sense that they tend to push a controlled perception away from a reference. The perceptions resulting from these statements are probably uncontrolled, in the sense that they depend ONLY on the statements, and not on any of my own output effects. That is, the resulting perceptions might be described as p(t) = d(t) -- where d(t) is the statements -- but not p(t) = d(t) + o(t). If there is no effect of the organism on a perception VIA THE EXTERNAL ENVIRONMENT -- that is, if there is no contribution of o(t) to a perception, then that perception CANNOT be controlled.

I think the perceptions resulting from the statements you describe are uncontrolled -- my outputs have no effect on them. This just means that there is no output I can produce that will change the way I perceive what Martin, Allan or anyone else is saying; I just perceive what they say (or have perceptions based on what they say). But these perceptions do cause a "disturbance" in the colloquial sense; that is, they do make me feel uncomfortable -- like something is amiss. I think this is because we have REFERENCES FOR UNCONTROLLED PERCEPTIONS (Bill P. mentioned this possibility long ago). All this means is that certain states of our perceptions (uncontrolled) produce error that we can do nothing about; this must lead to reorganization (trying to find variables to control that might eliminate this error). I think this is what we are seeing a lot of in this "information about the disturbance" discussion -- on all sides.

But it is important to understand that, with uncontrolled perceptions, the effect of the distrubance (technical sense) is completly "visible" as the perception (p(t) = d(t)) -- and any discrepency between this uncontrolled perception and a reference for its preferred value (r) is due completely to variations in d(t); these discerpencies (r-d(t)) are disturbances in the colloquial sense because they are unpleasant (if we perceive these discrepencies, they are what feels "amiss") and there is nothing we can do to reduce this error. So with uncontrolled perceptions, disturbances in the technical sense are highly correlated with disturbances in the colloquial sense; it's the correlation between d(t) and r-d(t).

Of course, this also means that all the information about uncontrolled variables is represented in p(t) and maintained in the error "signal" (there may be no explicit signal), r-p(t).

Before Chuck sent his note, my sermon today was going to be from p. 56, paragraph 2 of "The psychology of everyday things" by D. Norman (Basic Books, 1988). Norman is a leader in the field of cognitive psychology and this book is about how to design products that cognitive systems (people) can use. The section I wanted to discuss is called " Information is in the world". Now that I think of it, maybe I need say no more -- for now.

Finally, a quick note to Ken Hacker (932203) --

> The problem with PCT as you describe it is that it denies the social
>aspects of what is social: human INTERaction.

Could you explain how it does that? Seems like from the relationship level on up there are perceptions to be controlled that could reasonably be described as involving human interaction (relationships like "husband", "father"; programs like "helping a co-worker"; principles like "do unto others"; system concepts like "society").

>Remember, now, from your argument, if you are taken back by
>any of my comments, it has nothing to do with what I am saying; it is
>all inside of you!!

This is partially correct. Let's say that your comments are a disturbance, d(t), that contributes to some perception I am controlling. The perception (if it is controlled -- see above) MUST be a simultaneous result of both d(t) -- you comments -- and o(t) -- some effect I have on p(t) via the external environment; let's say that being "taken back" is one value of the output, o(t). Since

o(t) = r - d(t)

you can generate an o(t) = "taken aback" if you generate the appropriate level of d(t). That is, you can get me to do "taken back" if you say the appropriate comment. But the appropriate comment (value of d(t)) depends on r, which is, indeed, inside of me (and its value is determined by me). So the same comment, d(t), might lead to a "taken back" value of o(t) when my reference is set at one level and to no response (or a more positive one) when it is set at another. But your comment, d(t), has its influence on o(t) via p(t) -- which has no information in it about your comment, d(t).

Best   Rick

Date: Tue Mar 23, 1993 4:59 am PST
Subject: Simcon version 4.4 on biome

[From Bill Silvert (930323.0900)]   Followup to Bill Power's posting:

>[From Bill Powers (930322.1600)]
>
>Here is the documentation that goes with Wolfgang Zocher's
>beautiful analog-computer-on-a-PC program. I've uploaded some
>ascii-converted versions of the program to Bill Silvert's file
>server at biome.bio.ns.ca; Bill will no doubt check them and put
>them in the pubs/csg/simcon directory. Better wait for Bill's
>notification that the files have been put in the correct
>directory. I've also asked Gary Cziko to download them and put a
>binary version in the same directory (I still can't upload binary
>files). I would like a very brief note from everyone who
>downloads this program and is prepared to use it. Anyone who
>can't download it can send me a self-addressed STAMPED disk
>mailer with a formatted disc in it (360K or more, 5-1/4 or 3-1/2
>inch), and I will send all the stuff by return mail.

These files are now available in pub/csg/simcon (that is "pub", not "pubs", unfortunately). I combined the three UUE files into one and extracted the binary, so SIMCONZ.EXE is the file you want. --

Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2

InterNet Address: silvert@biome.bio.ns.ca
(the address bill@biome.bio.ns.ca is only for mailing lists)

Date: Tue Mar 23, 1993 8:43 am PST
Subject: Re: uncontrolled perception of statements

[Martin Taylor playing hooky 930323 11:30] Rick Marken 930322.2100

I think there is a "talking past" going on with Chuck Tucker.

>Chuck Tucker -- I've been wrestling with you private posts all
>day. Let me just respond quickly to one of your comments; it took
>me some time to finally get it right (I think).
>
>>It seems to me that I can state with great confidence and you would agree
>>with it that statements which appear to you to transform PCT into an
>>S-R like formulation are disturbing to you (I can think like you so I
>>have an advantage).
>
>This is a colloquial meaning of disturbance. The statements
>which "appear to me to transform PCT into an S-R like formulation"
>are just perceptions (or the cause of perceptions, if you like);
>They are not necessarily disturbances in the sense that they tend to
>push a controlled perception away from a reference.

I would assume (and have assumed) that the disturbed perception is your perception of other people's beliefs, and that your own postings are actions in a control loop whose reference level is something like "X should have a correct understanding of PCT." (Obviously that can't be literally a reference level, but it'll do for a colloquial rendition of one or some). The statements themselves, as you correctly observe, are uncontrolled perceptions or disturbances, if you do not have an ECS at a higher level whose CEV relates to the disturbing variable whose actions are those statements (see Powers' posting that you acclaimed last week).

I'd like to comment on other parts of this posting, but no time to do so. Maybe at the end of the week, if I'm not swamped again (still).

See you Thursday, perhaps.

Martin

Date: Tue Mar 23, 1993 9:00 am PST
Subject: Arrogance

[From Rick Marken (930323.0830)]   Ken Hacker (932203) --

>I would respond much more kindly to PCT if the messages here
>were less of what I PERCEIVE as arrogance.

I take it you mean the messages coming from those describing the PCT model. I would appreciate it if you could point to some of the comments that might contribute to your perception of arrogance. I have never seen anything that I would describe with the term "arrogance" coming from the PCT side; "exasperation", sure; "impatience", maybe. But "arrogance"? I actually haven't seen much of what I would call "arrogance" coming from either side of the dialog -- even though it might be an understandably tempting attitude for those who are already on the "winning" (non-PCT, of course) side.

I think that, to some extent, your perception might depend on the medium itself (e-mail). I think that if you met Bill and Mary Powers, for example, "arrogant" would be just about the very last adjective that you might want to use to describe them. On the other hand, I might be a different story.

Best Rick

Date: Tue Mar 23, 1993 11:16 am PST
Subject: Bobbing Heads

[from Gary Cziko 930323.1700 GMT] Bill Powers (930319.0730) said:

>The next time you're in an elevator (alone), just before the
>doors close, hold your hand outstretched and watch it. When the
>acceleration begins, you'll see your hand sightly dip (rise) for
>a moment and come back to the former position. You will also hear
>the machinery start, feel increased (decreased) pressure on the
>soles of your feet, and feel increased (decreased) effort in
>muscles all over your body, particularly those holding the arm
>out. Use the parenthesed expressions if the elevator starts
>downward. You will see the hand doing approximately what the
>above diagram shows, although without the sharp peak (which will
>be upside down if the elevator descends).

It's quite a coincidence that Bill should bring this up just as I was observing the same phenomenon--but travelling sideways, not up and down.

You don't need an elevator and don't need to travel alone. Just watch people's head bob back and forth in a car (or bus or train or whatever) in "reponse" to changes in acceleration. As driver when I slow down to make a stop I see my passengers' heads move forward then back again to where it should be. As I accelerate from the stop, the heads move backwards before coming to rest again a la verticale. No problem with explaining this using PCT.

But what is particularly intriguing is that the driver's head does not seem to move nearly as much. When I'm driving, I control the acceleration of the car and so I seem to use this advance knowledge of impending accelerations to minimize my head bobbing. This phenomenon is what I'm having some trouble understanding as in PCT terms since it appears to be a good example of what a "normal" psychologist would probably refer to as FEEDFORWARD.

I can nicely demonstrate the difference in head bobbing between driver and passenger by becoming a passenger myself as I drive. I can do this by using my cruise control in third gear (manual tranmission) set at about 30 miles per hour. At this low speed and gear ration my cruise control is not very stable and keeps swinging above and below the target speeds--and my head bobs back and forth like everyone else's in the car.

(Bill, I suspect that my cruise control does this because in third gear there is not enough slowing for the loop gain of the system?).--Gary

Date: Tue Mar 23, 1993 11:31 am PST
Subject: p(t) = d(t) + o(t)

[From Rick Marken (930323.1000)] Martin Taylor (930323 11:30) --

>I would assume (and have assumed) that the disturbed perception is your
>perception of other people's beliefs, and that your own postings are actions
>in a control loop whose reference level is something like "X should have
>a correct understanding of PCT."

I agree. This is what I was assuming when I started to reply to Chuck. But then I realized that your postings (qua postings) can be considered disturbances in Chuck's sense (not effects on a controlled variable but deviations from a reference level). Your postings can be considered perceptions that I cannot control -- and, in this sense, to the extent that I have references for these uncontrolled perceptions -- they are disturbances.

Whatever perception I am controlling in this "information about the disturbance in percpetion" discussion, it must be a function of both your posts (d(t)) and my posts (o(t)) because, if a perception is controlled, it MUST be a function of BOTH disturbances and outputs SIMULTANEOUSLY (simultaneous within the integration window of the controlled perception) so that p(t) = d(t) + o(t). This means that the perception I am controlling contains no information about your posts, d(t) -- so, with respect to the perception I am controlling, your posts are NOT a disturbance in the way Chuck was referring to a disturbance.

Let's say I am controlling for a variable that might be described as "correct description of the PCT model". This is p(t) and, I presume, it can have various levels, from "not very correct" to "very correct" (as I perceive it, of course). I have a reference level, r, for this perception, p(t), and I am trying to keep p(t) = r.

The important point is that p(t) depends on BOTH d(t) -- your posts; and o(t), my posts. But at the level at which I am controlling for p(t) I have NO information about your posts themselves; all I know is the current value of p(t) -- "degree of correctness of the description of the PCT model". At this level I have no information about why I am experiencing this particular level of p(t); that is, I have no idea how much the level of "degree of correctness of the description of the PCT model" is determined by the contents of your posts or mine. All I know is my perception of "degree of correctness of the description of the PCT model" and I generate outputs (posts) to keep this perception at my reference level.

Now this may seem screwy, I know. What I am saying is that, at the level of the perception I have described as the controlled variable I have no idea what was "said" in your posts OR mine. But I would argue that this seems screwy only because we can both move our consciousness to the level of our own perceptual systems that perceive what was "said" in the individual posts; at that level I can see that it was obviously your post that made my perception of the "degree of correctness of the description of the PCT model" become lower (let's say) -- or that my post obviously made that perception become higher. That is, at that level I can perceive the information in d(t) AND o(t). But to do that, I had to "get out" of the perspective of the control system that is controlling the the variable that is based on d(t) and o(t) -- the variable called "degree of correctness of the description of the PCT model" -- and look at the perceptual outputs of these other perceptual systems; systems which are not perceiving the controlled variable -- "degree of correctness of the description of the PCT model".

This is what happens in a pursuit tracking task, for example. At the level of the system (in the subject) that is perceiving and controlling the difference between the cursor and the (moving) target, all that is known is p(t) -- the perception of the difference between target, d(t), and cursor, o(t). So in putsuit tracking

p(t) = d(t) - o(t).

But when you are the subject in this experiment you CAN perceive the movements of d(t); they are clear as day (not "hidden" as they are in compensatory tracking); The subject can see d(t) moving back and forth, clearly disrupting ("distrubing") his/her ability to keep the cursor (o(t)) on the target. But this observation is made from the perspective of a different system in the subject's control hierarchy -- one that is NOT controlling p(t) = d(t) - o(t). The system that is controlling p(t) has NO information about d(t) -- though other systems can detect d(t) and help the system controlling p(t) take advantage of any regularities in d(t).

I think the problem with believing the "no information about the disturbance in perception" argument comes from the fact that we typically DO have information about disturbances in our perceptions because we (people) consist of multitude of layers of perceptions -- some of of these perceptions are controlled variables and some are not. The d(t) that is a non- informative aspect of a controlled perception (p(t)) is likely to also be an informative aspect of an UNcontrolled perception at another level of the hierarchy; this is what happens in pursuit tracking; I am sure it is what is happening in this discussion of information about the disturbance in perception.

Remember, the fact that there is NO information about disturbances in perception applies ONLY to CONTROLLED perceptions -- that is, to perceptions that are the JOINT result of disturbance and output -- perceptions which can be described as p(t) = d(t) + o(t). The components of p(t) -- o(t) and d(t) -- may be (and very often are) UNCONTROLLED perceptions at lower levels of the control hierarchy. So information about o(t) and d(t) may be available as UNCONTROLLED perceptions in the person controlling p(t) -- but p(t) itself contains NO information about d(t) or o(t).

Best Rick

Date: Tue Mar 23, 1993 12:43 pm PST
Subject: Hacker hassels (Mary)

[Mary Powers 9303.23] Ken Hacker 9303.22

No, my message to you did not come out of thin air, but what do you mean by

>your response to my message (note how my messages GENERATED your
>generation of messages; they did not arise from thin air)...

This implies that you caused me to send the message. I say that you contributed to my message-sending by providing an input. I compared that input to my reference signals and the resultant error is why I produced my message. If I had no reference signals for PCT your input would not have mattered enough for me to act.

The point is, the PCT model explains the "variability" of the effect of messages. Saying that a message from you generates a message from me does not.

>There is more semantic tussling going on than real explanation...

Last message was splitting hairs. Now semantic tussling. I think you are resisting a disturbance to some precious ideas by dismissing a different viewpoint as trivial.

>I am not interested in defending traditional psychology and any
>suggestion that I am is patently false.

Let's clarify our definitions of traditional psychology. I include as traditional: trendy half-assed self-regulatory versions of control theory that talk about feedforward, about people giving (or withholding) feedback from other people, about loops of messages sent and messages sent back, about people generating messages in other people, about ideas like stimulating feedback, and sign-offs like "thanks for the feedback, Bill".

I was taken back by your comments, because they constituted a disturbance, both intellectually (this guy doesn't understand PCT at all) and at a gut level (ohmygod, he's mad at me). It IS all inside of me, Ken. Having reference signals, comparators, etc. is not some cool, dispassionate arrangement; it's where we live. Emotions are how we know what's going on. Better to locate their origins inside than impute them elsewhere - by saying I am arrogant you can blame me for saying you don't know as much as you think you do, instead of entertaining that disturbing, errorproducing possibility and maybe having to change some cherished concepts you don't even realize you are defending.

I am as prone to blame and name-calling as anyone else in order to protect really important reference perceptions like my good opinion of myself. PCT is trying to explain not just some idealized system but the way people actually work, warts and all, including externalizing by blaming and other defenses. I've had a very interesting number of hours locating my disturbance at being called arrogant in me rather than as being caused by you being nasty. I'd much rather attribute my discomfort to you - it's how we are trained in this culture to handle attacks and it protects against uncomfortable self-examination. However, without getting into it too far, I seem to find the idea that I am arrogant absolutely terrifying, judging from the jolt of adrenaline I got when I read your message. I don't think the intent of your message was to scare me to death - that was the spin I put on it. In an alternate universe, equipped with different reference levels, I would shrug it off as irrelevant, or be mildly ticked off, or whatever was appropriate. The point is, your message was indeed, simply input, and how I interpreted it was ENTIRELY up to me, as a result of a lifetime of social interactions, including, apparently, times when being thought arrogant was pretty dangerous.

The point of this self-analysis is a) to show that what I made of your message is probably not what you expected and intended by your "tone", and b) to suggest that you examine all those things you call split hairs and semantics and arrogant statements to see why you don't want them to interfere with your perfectly satisfactory and correct understanding of control theory.

Mary Powers

Date: Tue Mar 23, 1993 1:43 pm PST
Subject: Simcon version 4.4 on biome

[from Gary Cziko 930323.2130 GMT} Bill Silvert (930323.0900) said of Simcon 4:

>These files are now available in pub/csg/simcon (that is "pub", not
>"pubs", unfortunately). I combined the three UUE files into one and
>extracted the binary, so SIMCONZ.EXE is the file you want.

Great--that's a big help.

I downloaded the simconz.exe file, unzipped on my computer by typing simconz, ran the masspring demo and got the plots on my screen. So it seems to work fine.--Gary

Date: Tue Mar 23, 1993 2:22 pm PST
Subject: CLOSED LOOP; Mill's "method of difference"

From Greg Williams (930323)

The next CLOSED LOOP (due out 4-15 or thereabouts) will be devoted to Portable PCT Demos. Any comments for possible inclusion are welcome on the net through 3-31.

Something for the cause-effect "hassel" (how about "hassle"?): John Stuart Mill's SYSTEM OF LOGIC, 1919, says, "If an instance in which the phenomenon under investigation occurs, and an instance in which it does not occur, have every circumstance in common save one, that one occurring only in the former; the circumstance in which alone the two instances differ is the effect, or the cause, or an indispensable part of the cause, of the phenomenon." Note that Mill doesn't say "or both the cause and the effect." So maybe the cybernetic revolution reaches down past the biology devils and even the physics devils to the logic devils?

As Jerry Lee Lewis says, "Think about it."

As ever (now that my modem is working again!), Greg

Date: Tue Mar 23, 1993 7:30 pm PST
Subject: charming cooperation

[from i.n.kurtzer(930323.2110)]
hello again. today i saw a chamng example of cooperative behavior. (i realize this is only distally related to the ongoing discussion but the situation deserves comment). pairs of young children were playing a game that required them, each pair, to carry a inflatable ball across a plaza with their noses. to do this the children would stand face to face and press the ball between themselves with their

noses. being nosey myself (i couldn't resist) i watched them walk sideways to move the ball while pinching proboscises. of course, if one pressed too hard, or too lightly, or walked too quickly or failed to keep pace the ball couldn't be carried across. also, i once again highly suggest that anyone interested in cooperative behavior should read Kropotkin's MUTUAL AID.
i.n.kurtzer

Date: Tue Mar 23, 1993 8:27 pm PST
Subject: Arrogance

From Ken Hacker [932303]

Rick Marken, I perceive arrogance in any form of intellectual discussion which begins to sound like ideology instead of exploration. I believe that human explanation of existence is light years from any unified theories of anything and that we need to ask better questions before we congratulate ourselves for elegant answers.
As for Bill Powers PCT, I believe that Bill is an outstanding scholar and scientist. I have nothing but respect for him and I appreciate all of his comments (in my perception, his feedback) regarding my questions and disturbances.

From my point of view, I am observing PCT as a challenge, as you and others have noted, to traditional social and behavioral science. I appreciate the new directions in thinking very much. But I think that sometimes the claims about how useless, futile and infantile other approaches to understanding human behavior are, are a bit hyperbolic. Best, KEN

Date: Tue Mar 23, 1993 9:18 pm PST
Subject: Hacker's Tussles and Mary's Ruffles

From Ken Hacker [930323] To Mary Powers

Gee, thanks for informing me about my "half-assed" beliefs. It's truly amazing *&))#(#*&()@)(*#((#*_@*#(_#&(*. But, I will stop the retaliation right there, despite my temptations from surging error signals... :)

You confirmed some interesting points for me and also contributed to OUR learning more about PCT and social behavior, i.e, human behavior:

* Neither of us are CAUSING messages to be sent, but our messages act as INPUT which must be processed and must be compared to reference signals or levels.

* Yes, we have actions that are directly related to our comparisons of input to reference signals.

* You hit the nail on the head by stating that the variation of messages effects in explained by PCT. I think this is a critical point and maybe one of the PCT-comm theory linkages I am looking for.

* You say that my messages are "simply input" and how you interpret them is a "result of a lifetime of social interactions ..." Yes, that is one my main arguments.

In other words, we are doing all of the input-reference comparisons that PCT describes, but there are social sources and reasons for those things we call reference signals.

I think that the comments here about information and signals are correct in so far as they state that information is internally produced more than carried inside the signals themselves. However, I also think that signals, signs, symbols, messages carry enough socially learned signification to which the control system identifies in ways that error signals are present in whatever degree.

Mary, you can continue throwing darts if you want. It's ok. I really dont mind. I have done enough e-mailing and computer conferencing to become flame-proof. But I want you to know that I appreciate your intellect and greatly respect what you and Bill are doing. Please contextualize my comments and questions in the spirit in which they are created -- inquisitiveness and critical thinking.

Best, Ken Hacker

Date: Tue Mar 23, 1993 9:50 pm PST
Subject: Arrogance/Ideology

[From Rick Marken (930323.2100)] Ken Hacker [932303]--

>I perceive arrogance in any form of intellectual discussion
>which begins to sound like ideology instead of exploration.

What, specifically, has any PCTer said that sounds more like ideology than exploration? I am really curious about this because I get this all the time (I've been called a "true believer" on more than one occasion, implying, I presume, some ideological, as opposed to scientific, committment to PCT). Where's the ideology?

>But I think that sometimes
>the claims about how useless, futile and infantile other approaches to
>understanding human behavior are, are a bit hyperbolic. Best, KEN

I will admit to claiming that conventional approaches to understanding human behavior are useless and futile -- but certainly not infantile. I have also explained precisely WHY I claim that such approaches are misguided (and, hence, useless and fultile -- and, I might add, misleading); one explanation is contained in the "Blind men and the elephant" paper. These claims are not religious chants; they are scientific conclusions (many of the experiments that support these conclusions are described in "Mind Readings"). It is not enough to just claim that PCT claims are "hyperbolic" -- that is ideology. What you should do (which none of the "non-ideological" critics of PCT seem to feel they need to do) is demonstrate -- mathematically and experimentally, as we do in PCT -- exactly how PCT claims are "hyperbolic". PCTers tend to be MUCH more interested in working models and experimental demonstrations than sermons about the sins of ideology.

So, please, be specific. Pick one PCT claim that seems "hyperbolic" and explain (or better, tells us how we can demonstrate to ourselves) what is hyperbolic about it.

Thanks Rick

Date: Wed Mar 24, 1993 6:43 am PST
Subject: you CAN tell

[From: Bruce Nevin (Wed 930324 08:44:38)]

The application of PCT to our own ongoing discourse is a Good Thing, and I hope it continues until I am able to participate more.

It is of special interest to me because it involves how language differs from arm movements or getting a beer in that language is socially standardized, else it could not serve our purposes.

It is possible to know what someone is doing by watching what they are doing, when they are doing something that is socially standardized. When Rick writes "this is a colloquial meaning of disturbance" I know that he is producing certain socially standardized words in certain socially standardized relationships to one another and to other words assumed known. (Assumed known, in many cases, because they occur nearby.) We can know the words and the relationships in a determinate way<*> because they are

____

* Up to structural ambiguity for the relationships, and then each branch of the ambiguity is itself also determinately knowable.

____

socially standardized. We can't tell what precisely they mean by what they are saying, and therefore may disagree about that, but what they say in the more limited, literal sense is known and determinate. This is because it must be socially available in a reliable, determinate way for language to be useful to us.

The meanings (nonverbal perceptions) that I associate with the words and words-in-relation may differ from those that you associate with them, and probably do. Nonetheless, we assume that the differences should be matters of detail. (This is the fundamental assumption of science: that the universe is knowable, consistent from one part of it to another and from differing points of view.) So we seek agreements, and in seeking agreements we talk some more, and some of this subsequent talk is words and word-relations not mentioned earlier, words and word-relations associated for each of us with some perceptions that we associated (differently from each other) with the first round of words.

The proposal has been that there is a great deal of recognition-by-production in the processes of associating meanings with heard or read discourse. What would I be meaning if I said that? At an even lower level, what might I want to say next if I had just said that? And even for the word dependencies themselves, e.g. having said "this is a colloquial" what might I want to say next? Given the socially standardized patterning of English, only certain kinds of words are likely to come next. From this process of re-creating what is heard as what is said internally arise expectations as to what is coming, and expectations in some cases so strong that the words need not be fully there. (So, instead of "the meaning of `statements which appear to transform PCT into an S-R like formulation are disturbing' is a colloquial meaning of disturbance", Rick can say "This is a colloquial meaning of disturbance.")

And, as you might have guessed, the message from which I have been quoting one point of many at which I could bring this in:

(Chuck Tucker off-line to Rick Marken):

>   statements which appear to you to transform PCT into an
>S-R like formulation are disturbing to you

[Rick Marken (930322.2100)] --

> This is a colloquial meaning of disturbance. The statements
> which "appear to me to transform PCT into an S-R like formulation"
> are just perceptions (or the cause of perceptions, if you like);
> They are not necessarily disturbances in the sense that they tend to
> push a controlled perception away from a reference. The perceptions
> resulting from these statements are probably uncontrolled, in the
> sense that they depend ONLY on the statements, and not on any of
> my own output effects. That is, the resulting perceptions might
> be described as p(t) = d(t) -- where d(t) is the statements -- but
> not p(t) = d(t) + o(t). If there is no effect of the organism on a
> perception VIA THE EXTERNAL ENVIRONMENT -- that is, if there is no
> contribution of o(t) to a perception, then that perception CANNOT be
> controlled.

> I think the perceptions resulting from the statements you describe
> are uncontrolled -- my outputs have no effect on them. This
> just means that there is no output I can produce that will change
> the way I perceive what Martin, Allan or anyone else is saying;
> I just perceive what they say (or have perceptions based on what they
> say).

When you are using language, you are producing outputs as part of the process of understanding what was said. There is no need to invoke references for uncontrolled perceptions to account for your discomfort at someone saying things about a cherished concept that you would regard as error if you said them yourself. And I hardly think it is the case that your concept of PCT is an uncontrolled perception for you!

Gotta do some work for BBN

Bruce   bn@bbn.com

Date: Wed Mar 24, 1993 9:39 am PST
Subject: A PCT Model of a PCT Net Discussion

[From Rick Marken (930324.0800)]

Bruce Nevin (Wed 930324 08:44:38) --

>The application of PCT to our own ongoing discourse is a Good
>Thing, and I hope it continues until I am able to participate more.

I'm glad you think so because I went to the trouble of developing a diagram of my PCT model of our PCT discourse. I hope it will show you what I mean when I talk about controlled and uncontrolled perceptions.

Here is the diagram:

```
 System Boundary
        ||_____________________________________________________________
"boss   || "me"
reality"||      reference for
        ||   uncontrolled perception
        ||             r
        ||             |
        ||             v
        ||            |C| ---->pure disturbance  --> awareness of
        ||             ^           error               error
        ||             |
        ||             o------> copy of perception ----> awareness of
Other   ||             |          of disturbance          disturbance
People's||             |
Posts   ||             |
        || ____        |       __________
d(t) -->|f(d(t)| ------o------>|          |
        |______| meaning[d(t)] |          |     ___
        ||                     | h[m(d(t) | p(t) |   |  r   reference for
My      ||                     |    +     |----->| C |<---   controlled
Posts   ||                     |  m(o(t)] |      |___|        variable
        ||_____                |          |        |           p(t)
o(t) -->|f(o(t)| ------------->|__________|        |
 ^      |______|  meaning[o(t)]                    |
 |      ||                                         |
 |      ||                                         |
 |      ||_____           e(t)= r-p(t)             |
  ------|g(e(t)|<---------------------------------
        |______|
        ||
        ||_________________________________________________________________
```

The column of double vertical lines is the boundary between the control system (me) and the environment (boss reality). Other people's posts, d(t), and my posts, o(t), are variables in boss reality. They are converted into perceptual signals by the perceptual functions, f(), in the boxes at the interface to boss reality (actually, these boxes probably represent the SAME function operating at different times; I use the same perceptual function to extract the same perceptual variable from my posts and those of others; but I think it is OK to represent this temporal difference as a spacial difference in the model). These perceptual functions are pretty fancy boxes because they turn the posts into a signal (meaning(o(t) or meaning (d(t)) whose magnitude is the "meaning" of the posts; let's imagine that these boxes transform the words in the post into a signal whose magnitude indicates the degree of "cause-effect PCT imagery" in the post. [Of course, I imagine that there are many parallel perceptual functions that produce signals whose magnitude is proportional to other "meanings" in the post; I'm just focusing on one "meaning" aspect of the post in this first level]

The "meaning" signals are the inputs to a "second level" perceptual function box that uses the h() function to convert the two inputs into a higher level perceptual signal. This higher level perception, p(t) is the perception I am controlling in the network dialog about PCT. A name for this perception might be "average amount of cause-effect PCT imagery occuring on the net". (Again, this is just ONE perceptual variable that I might be controlling; many higher order perceptions of the discourse are surely being computed in parallel).

To simplify the model I have drawn the perceptions of BOTH o(t) and d(t) [that is, meaning (o(t)) and meaning (d(t))] as UNCONTROLLED; in fact, only my perception of d(t) is uncontrolled. All this means is that there is no feedback link from me to d(t) that can influence my perception of the "meaning" of d(t) [of course, there is a feedback link to other aspects of the perception of d(t) -- such as whether it exists or not; I can affect the latter pereption by either reading or not reading my e-mail]. There IS a feedback link from my outputs to the aspects of o(t) on which my meaning perception of o(t) depend -- but, for the sake of (relative) simplicity, I have not drawn them in.

One more aspect of the diagram and then I'll throw it out to the net for discussion. Note at the top of the graph that I have shown the perception of d(t) branching off of its path to the second level perceptual input. The perception of d(t) is uncontrolled; it simply exists as my perception of the degree of "cause-effect PCT imagery" in the other person's post. I can become "aware" of that perception (this is one aspect of consciousness in the PCT model; perception of a perception); this is indicated to the signal branch ending at "awareness of disturbance". If I have a reference level for this perception then the perception will be compared to this reference and an error signal will result; this error signal might also become the object of awareness -- so you perceive the disturbance "as an error" -- ie. "that was a silly post -- it has too much "cause effect PCT imagery" in it. But all I can do is experience this error -- since the perception of d(t) (at this level) cannot be affected by my outputs; it is just an uncontrolled perception -- and the error just happens.

Hopefully, this diagram shows the two main points I was trying to make in earlier posts. First, the system that is controlling p(t) -- by varying o(t) -- has no information about what was in the perceptions of the component posts from which this perception is derived; all this systems knows about and controls is p(t) -- the "average amount of cause-effect PCT imagery occuring on the net" -- regardless of the lower level perceptions that contribute to this average. At the same time, however, my consciousness might be aware of the perception of d(t) and o(t) that happen to contribute to p(t). I may have a vague idea that the perceptions of o(t) and d(t) do contribute to the higher level perception I am controlling -- but, with my awareness sitting at the first level of the model, I am most likely to see my own outputs, o(t) as responses to the disturbances, d(t) that I perceive.

This little exercise just made me realize that the first illusion described in the "Blind men and the elephant" paper applies not only to observers of other people's behavior -- it applies to ourselves when we look at our OWN behavior too. Since perceptions of d(t) and o(t) are sure to be at a lower level than the perception that is being controlled by o(t), we are very likely to conclude that our own behavior (o(t)) is a response to input (d(t)). No wonder people believe that inputs (information, perception, "feedback", etc) guide behavior; it not only looks that way when we look at other people -- it also looks that way when we look at our OWN behavior -- from the wrong level.

Best   Rick

Date: Wed Mar 24, 1993 10:28 am PST
Subject: Looking for linkages

[From Rick Marken (930324.0930)]   Ken Hacker [930323] to Mary Powers:

>   * You hit the nail on the head by stating that the variation of
>   of messages effects in explained by PCT. I think this is a
>   critical point and maybe one of the PCT-comm theory linkages
>   I am looking for.

It seems nice and reasonable to try to find "linkages" between one theory (especially one that seems to have worked very well) and another (especially if it is a newcomer). But it can also lead to problems -- and the belief that the newcomer theory must be ideology or heresy. For example, how much linkage could there have been between Ptolemaic and Copernican models of the "universe". On some points there certainly was "linkage" (like the fact that there are planets, a sun, and that these objects move in orbits). But on a fundemental point there was no linkage -- and it turns out that that point prevented any serious linkage between the models

I think the same thing is happening with PCT. On some points there is linkage between PCT and other theories of behavior (like the fact that the muscles exert forces on the skeleton, resulting in outputs that produce events called behavior, etc). But on one fundemental point there is no linkage; in PCT, behavior is controlled input, not caused output. This little point means perception (not output) is the appropriate dependent variable to study in psychological research; it also means that studies of the relationship between output (dependent variable) and stimulation (independent variable) reveal nothing other than (at best) the inverse of the physical law relating these outputs to controlled perceptual inputs.

Looking for linkages might be nice -- but what happens when there are no linkages? What happens when the new theory just flat out contradicts the old one? My guess is that the new theory (or, more likely, the advocates thereof) will be castigated for their impudence and heresy (castigated but never shown WHY their conclusions are wrong -- unless you count pointing to bibical passages).

I think I'll go out and get a bust of Galileo.

Best   Rick

Date: Wed Mar 24, 1993 8:32 pm PST
Subject: Give Galileo a break!

From Ken Hacker [932403]   In response to Rick Marken --

You perceive, therefore you are.   You are, therefore you perceive. That's ideology, my friend.

Arrogance in what is called scientific activity is closing doors to inquiry and citing mathematics or "working models" or experiments as the only acceptable form of reasoning about knowledge claims. I understand the PCT is making solid claims and shaking sacred ground; I like that and find it invigorating. But I prefer reasoning based on data as opposed to reasoning based on someone in PCT says it's so and therefor you must accept it or as Mary implies, you are an idiot if you do not accept this perspective lock, stock, and barrel. Science is based on data and religion is based on faith. I want more data and less faith in regard to PCT.

Of course, Rick, if any of this bothers you, remember that such a thing is impossible. You are bothering yourself; my words hold no information -- any feelings or reactions you have now are produced by yourself. Therefore, if you ticked, you had better castigate yourself, right?

Alfred Whitehead argued that all things are always in process, so that knowledge about anything requires movement back and forth among a variety of perspectives. I think this is needed with theories of human behavior. Various perspectives offer knowledge about different aspects of behavior. Thus, anthropology is no less valid or more valid than PCT, just different in what it chooses to study. Now comparing PCT to behaviorism or cognitivism is certainly different and I can see that PCT could offer some pushes toward a paradigm shift in the study of internal self-regulation and production of behavior. I think that is the central strength of PCT as I have observed it so far. I see no contributions of PCT, thus far, to any claims about social interaction, societies, social systems, organizations, cultures, or language. It may be that PCT does not wish to explain any of this stuff, but all of it is essential to human behaviors at some level. Maybe, as Klaus Krippendorff told me, control theory is best suited to explaining physical movements of human beings more than anything else.

I believe that humans behavior is most wholly a matter of self-organization, and hence, I remain interested in PCT, in part because I see psychology as limited in describing the processes versus variables of human thinking and adapting. Each human regulates his.her behaviors to maintain certain goal states. I believe that when a sensor provides feedback to the comparator and the comparator guides the activator, that the behaviors produced by the activator generate external feedforward-feedback loops which are what I study as social interaction. In other words, I think there is a social loop to human behavior and control which does not deny, negate, replace, or override the internal loops and levels described by PCT, and which exists when the control system is in contact with other control systems.

All conscious, purposeful, and deliberate human behavior requires feedback. Feedback is not a simple or singular structure/process in a complex control system such as a human. A guided missile adjusts its course and hits its target in adapting its motion to the received feedback of radio waves it has directed and had sent back to it from its target. A human receives messages that are sent back in relation to messages that have been sent out. Whether internal or external, or both, the control system still counteracts or amplifies deviation. I am not arguing that feedback from other humans comes orginally from then; my points is that we include other humans in loops of feedback which we create. Wiener said, "This matter of social feedback is of very great sociological and anthropological interest." I agree, but I dont think that he or anyone else has demonstrated why YET.

Cheers, Ken Hacker

Date: Wed Mar 24, 1993 10:10 pm PST
Subject: From the horse's mouth

[From Rick Marken (930324.2100)] Ken Hacker [932403] --

Why all the excitement?

This all started (for me) when you said:

>But I think that sometimes
>the claims about how useless, futile and infantile other approaches to
>understanding human behavior are, are a bit hyperbolic.

And I said:

>So, please, be specific. Pick one PCT claim that seems "hyperbolic" and
>explain (or better, tells us how we can demonstrate to ourselves) what
>is hyperbolic about it.

Now you say:

>Arrogance in what is called scientific activity is closing doors to
>inquiry and citing mathematics or "working models" or experiments as
>the only acceptable form of reasoning about knowledge claims.

So I'm arrogant because I ask for mathematics, working models or experiments to demonstrate that a PCT claim is hyperbolic? What are the alternatives you suggest?

> But I prefer reasoning based on data as opposed to reasoning based on
>someone in PCT says it's so and therefor you must accept it

I am a really big fan of exactly the same thing; reasoning based on data. But you said that citing mathematics, working models and experiments (all of which I count as sources of data on which to base my reasoning) is arrogant (closing the door to scientific inquiry). If you have some other source of data that would help us reason about what you consider to be the hyperbolic claims of PCT I would certainly consider them.

> Science is based on data and religion is based on faith.
>I want more data and less faith in regard to PCT.

Me too.

>Of course, Rick, if any of this bothers you, remember that such a thing
>is impossible.

No bother. I was not trying to be mean or smart ass. What you said about the hyperbolic claims of PCT simply aroused my curiosity -- really. I'm not trying to call your bluff. I'm not interested in defending the "faith". I really just want to know what it is about PCT that you find hyperbolic. My aim is simple -- if you can give me some concrete data (or demonstration or even a verbal description) of what you consider a hyperbolic claim of PCT then maybe -- MAYBE -- I can come up with some way of convincing you that the claim is actually supported by data. If not, I will happily admit that the claim is hyperbolic and I will stop making it. I am interested in doing this for several reasons: 1) yours is a common reaction to PCT, we have you available on the net and maybe -- if we can figure out how to answer your concerns -- we can answer the same concerns among the general population of life scientists and 2) I want to try to convince you that PCTers are more than willing to put the PCT model to the test; PCT is not a religion; the model has been subjected to VERY rigorous tests -- and we want to continue to test it. I want you to see that I am more than happy to subject PCT to ANY and ALL tests -- the more the merrier -- and if it fails even one, the model will be cheerfully exchanged or you can receive a complete refund.

> I see no contributions of PCT, thus far, to any claims about social
>interaction, societies, social systems, organizations, cultures, or language.

Now you don't like PCT because it doesn't claim enough -- yesterday it was because it made hyperbolic claims. You don't do this with your kids, do you?

>I believe that when a sensor provides feedback to the comparator
>and the comparator guides the activator, that the behaviors produced by the
>activator generate external feedforward-feedback loops which are what I
>study as social interaction. In other words, I think there is a social loop
>to human behavior and control which does not deny, negate, replace, or
>override the internal loops and levels described by PCT, and which exists
>when the control system is in contact with other control systems.

That's a lot of believing and thinking. If that is your alternative to evidence from mathematics, working models and experiments then YOU WIN. Everything you say must be correct. On what basis could I possibly say it is not -- other than by saying "you're wrong"? But you could easily counter that by saying "no I'm not" or "I believe such and such".

When you're ready to look into the horse's mouth, I'm ready to help you count teeth.

Best   Rick

Date: Thu Mar 25, 1993 6:40 am PST
Subject: Rick's diagram

[From: Bruce Nevin (Thu 930325 08:36:48)]

[Rick Marken (930324.0800)]

I am glad you read the first sentence of my (Wed 930324 08:44:38).
Any agreements or disagreements with the rest of it?

Because it is agreements for which we control when we use
language, and failure of agreement is a disturbance.

Put that in your diagram thus:

```
 System Boundary
        ||_____________________________________________________________
"boss   || "me"
reality"||      reference for
        ||   uncontrolled perception
        ||             r
        ||             |
        ||             v
        ||            |C| ---->pure disturbance  --> awareness of
        ||             ^           error               error
        ||             |
        ||             o-------> copy of perception ----> awareness of
Other   ||              \         of disturbance          disturbance
People's||                \
Posts   ||                  \
        ||      ______        \         __________
d(t) -->||     |f(d(t)| -----o------->|          |
        ||     |______| meaning[d(t)] |          |      ___
        ||                            | h[m(d(t)  | p(t) |   |  r reference for
What I  ||                            |    +     |----->| C |<---   controlled
would   ||                            |  m(o(t)] |      |___|       variable
say     ||      ______                |          |        |          p(t)
next/   || o(t)-->|f(o(t)| ------------->|__________|        |
mean    ||  ^     |______|  meaning[o(t)]                    |
if I had||  |                                                |
said    ||  |                                                |
that    ||  |      ______            e(t)= r-p(t)            |
so far  ||   -----|g(e(t)|<-----------------------------------
(imagi- ||        |______|
 nation)||
        ||_________________________________________________________________
```

What is indicated here is a "shadowing", in imagination, of the other's utterance in the environment. Many levels and much complexity intervene between the perception of sounds (or graphic shapes) and the association of meanings with the utterance.

We do not do anything like this "shadowing" when perceiving someone's arm movements, except in professional dance, acrobatics, or the like. This is for the very good reason that it is possible only when we have in ourselves (settings of) reference perceptions which we can reasonably expect the other person also to be maintaining.

My discussions of social norms concern how it comes about that people come to have the same settings of the same reference perceptions for things like language. It doesn't happen as a byproduct of interaction with other people who fortuitously share those settings of those references already. As with dance and acrobatics, people have to work very hard at it, and they have to have help from others. These teachers tell the learner how well their outputs approximate the norm, so that the learner can then develop settings of her own perceptions of effort, position, etc. for control such that others will perceive her outputs as conforming to the norm. This is because the learner can in general monitor her own outputs only in a limited way.

A motivation for going to all this trouble (for ordinary social norms, not for acrobatics!) is that it enables cooperative action. I won't elaborate that point again, unless it's not already abundantly clear.

There is a problem indicating the relation between f[o(t)] and meaning[o(t)] with a labelled arrow if that means, as it usually does, a relation between two levels of the perceptual hierarchy. This is because meanings are by Bill's hypothesis any and all perceptions, at any level of the hierarchy. The hypothesis was that a linguistic perception of a word is combined with one or more nonlinguistic perceptions ("meanings") in the input function of an ECS controlling a category perception to match the reference signal in the comparator of that ECS. There isn't a level of the hierarchy at which words are perceived, and a higher level at which meanings are perceived. Linguistic perceptions are controlled from the intensity level at least to the program level, and probably higher, in parallel with nonlinguistic perceptions. Non-linguistic perceptions are associated with linguistic perceptions, as our perception of what the latter mean for us, by some relation other than the familiar arrow connecting one level of the hierarchy to the next.

This association of meanings with words and with combinations of words (and with controlled linguistic perceptions of still more complex structures in utterances) is not standardized. Bill imagines stick figures when he reads a sentence about scientists digging under an old tower. I imagine stereotyped but unmistakably ordinary flesh-and-blood human beings. All that is identical in these perceptions is the word "scientist" and certain category perceptions of features that human beings physically have in common--two arms, two legs, head, etc. But these are category perceptions, so we may well say that all that the nonlinguistic "meaning" perceptions have in common is the words with which they are associated. In other words, the social standardization of language helps us to structure our perceptions categorically in the same way as do others with whom we will need to cooperate.

Does ANYBODY get it yet, or will this be ignored too?

        Bruce   bn@bbn.com

Date: Thu Mar 25, 1993 6:55 am PST
Subject: Re: Rick's diagram

>[re: Bruce Nevin (Thu 930325 08:36:48)]
>
  In other words, the social
>standardization of language helps us to structure our perceptions
>categorically in the same way as do others with whom we will need
>to cooperate.
>
>Does ANYBODY get it yet, or will this be ignored too?
>
I get it. Thanks for stating it so well.

best regards,

Clark

Date: Thu Mar 25, 1993 9:23 am PST
Subject: Re: p(t) = d(t) + o(t)

[Martin Taylor 930325 11:30] Rick Marken 930323.1000

I apologize in advance for an inadequate response. While I was away yesterday I did a count, and realized that I have 19 workdays here between now and June 15. PCT is not what I am paid to do, although I am incorporating it in the "real" work so far as possible. So I can spend little time on PCT discussion, unless the particular item seems likely to further the integration of PCT with my work.

Having said that, and given the CSG-L community the good news that the volume of my posting will be drastically reduced over the next 3 months, I can comment on Rick's posting.

There are several reasons why the discussion of information in PCT has explored so many blind alleys, but apart from the differences among participants in the definitions of words, the most significant one is in a general non-realization that we have been taking multiple views simultaneously.

The "inside view": The result of applying the perceptual input function to the set of sensory inputs--i.e. the perceptual signal. Other, but different inside views could be at the ouput of the reference signal collector function, the comparator, or even the output function. Any one of these inside views is (not "represents") a scalar signal that has a single momentary value that is a number. It exists, in and of itself, representing, from the inside view, nothing at all. An inside viewer recording the signal can see only a waveform, which according to Fourier and subsequent mathematicians, can be represented exactly by a discrete set of numbers. These numbers represent the waveform, and are the most extreme "representative" view that can be taken from inside the ECS.

The "outside view": from outside, any aspect of the universe can be seen and incorporated into some analysis. This includes the signals within the ECS, the state of the CEV that is defined by the Perceptual Input Function (PIF) of the ECS, and the actions of any disturbing variables. An outside view can incorporate all the structure of the hierarchy. It is undetermined what is and what is not in a generic "outside view," though restricted outside views can be postulated for the purposes of specific discussions.

Now, using those quasi-definitions, we can perhaps see why there has been an argument. It is transparently obvious that from an inside view the perceptual signal, being scalar, cannot distinguish between two independent effects upon it. There is nothing there but a unidimensional waveform (ignoring for the moment the fact that the successive independent samples do define a multidimensional space). Nothing in any sample can indicate more than that it has "that" value. From this viewpoint, it is straightforwardly correct that "there is no information about the disturbance (or disturbing variable) in the perceptual signal."

Inside the protagonist (subject), at one level considered alone, no ECS has any information other than "about" its CEV, and even that "aboutness" is not visible from the inside viewpoint.

In the hierarchy, the various PIFs define a mirror hierarchy of CEV's in the world, as we have drawn in the mirror diagram many months ago. A higher PIF represents a CEV that includes among its effects (as seen from outside) a "disturbance" to a lower CEV. Bill Powers reinforced this view a week or two ago when he pointed out that a higher PIF could be formed to represent a CEV that was "the" disturbing variable for a given one and that in this way a disturbance could be extracted out of the perceptual signal of the lower ECS. Apart from some scepticism that one could in this way detect all the external influences on a CEV, this simply is a description of the phase of reorganization that builds new levels in the hierarchy.

The situation is different if we take a full-blooded outside view of the action of a CEV. It is from this kind of view that we argue that the disturbance provides information that passes through the perceptual signal to the output signal. From the outside we can see the disturbing variable do whatever it does to affect the CEV, and we can see the ECS modifying its output to bring the perceptual signal back to its controlled value. From outside we can see the reference signal of the ECS changing, and the ouput changing to move the CEV so that the perceptual signal comes to its new controlled value. From outside, the arguments about there being no information from the disturbance in the perceptual signal lose their force. The use I made of Bill's example was from this point of view, and still seems to me to be correct.

>I have no idea how much the level of "degree of correctness of the
>description of the PCT model" is determined by the contents of your
>posts or mine. All I know is my perception of "degree of correctness
>of the description of the PCT model" and I generate outputs (posts)
>to keep this perception at my reference level.
>
>Now this may seem screwy, I know. What I am saying is that, at the
>level of the perception I have described as the controlled variable
>I have no idea what was "said" in your posts OR mine.

No, it doesn't sound screwy. You are expressing a one-level inside view.

>But I would
>argue that this seems screwy only because we can both move our
>consciousness to the level of our own perceptual systems that perceive
>what was "said" in the individual posts; at that level I can
>see that it was obviously your post that made my perception of
>the "degree of correctness of the description of the PCT model"
>become lower (let's say) -- or that my post obviously made that
>perception become higher.

I'd modify that (and the rest of your posting) to identify where the "degree of correctness of the description of the PCT model" exists. Is it in your perception of my belief, your perception of your own belief, your perception of the change your posting is likely to make in my belief, or what? Otherwise, I agree. The change in level changes the representativeness (as seen from outside) of the perceptual signal.

Enough. 46 messages await my attention. I either don't read them and continue this response, or quit and hope for the best.

Martin

Date: Thu Mar 25, 1993 9:30 am PST
Subject: Re: Bobbing Heads as communication

[Martin Taylor 930325 12:00] Gary Cziko 930323.1700

>You don't need an elevator and don't need to travel alone. Just watch
>people's head bob back and forth in a car (or bus or train or whatever) in
>"reponse" to changes in acceleration. As driver when I slow down to make a
>stop I see my passengers' heads move forward then back again to where it
>should be. As I accelerate from the stop, the heads move backwards before
>coming to rest again a la verticale. No problem with explaining this using PCT.
>
>But what is particularly intriguing is that the driver's head does not seem
>to move nearly as much.

What (to me) is even more interesting is that this is a well-known conventionalized communicative signal. In the days before automatic transmissions, one sign of a good driver was the ability to make smooth gearshifts. If a passenger wished to comment on the lack of skill of the driver, he would allow his (usually male) head to bob quite visible. In PCT terms, he would either increase his transport lag or reduce his gain (I suspect the former, since the head would return quickly to its normal position.) The driver might well feel insulted by this behaviour (I don't know whether to use behaviour or action in this instance).

Martin

Date: Thu Mar 25, 1993 9:33 am PST
Subject: Arrogance vs. visions

[From Bill Powers (930325.0850)] Ken Hacker (930324) --

>Arrogance in what is called scientific activity is closing
>doors to inquiry and citing mathematics or "working models" or
>experiments as the only acceptable form of reasoning about knowledge claims.

Is it arrogance to set standards for accepting knowledge claims so high that even one's own attempts to make claims are more likely to fail than succeed?

I agree with you that there are other approaches to knowledge, but in my view that is all they are: approaches. We have to spend a lot of time mulling over imperfectly stated propositions and guessing how we might investigate them, but I don't think that this kind of preliminary activity belongs in the category of "knowledge claims." It's the sort of thing you do while you're looking for an idea that has enough solidity to be put forth as a candidate for a knowledge claim. When you do finally clear your throat and rap on the table and announce that you have an idea that might actually check out, you should expect to have the idea subjected to the most severe tests that can be devised -- which, preferably, you have already thought up and done yourself.

The criterion that the physical sciences adopted was that a theoretical prediction should fit the data ALL of the time and under ALL circumstances, with an error no greater than the error of measurement. That is a pretty stiff requirement even when some practical latitude is allowed. It has probably caused 99% of the plausible ideas about how physical reality works to be rejected. But look at what it produced. The one percent (or even less) of the ideas that survived have remained undisputable for centuries. Even Newtonian mechanics is still the undisputed theory of choice in very nearly all macroscopic situations -- and has been for over three centuries. This is what you get when you cross the boundary between verbal reasoning and precise analysis -- when you're able to do it.

I don't know why you put "working models" in quotes. A working model is simply a proposition put to the ultimate test: it must generate detailed predictions of real behavior through time, with every deviation from the actual behavior under a critical microscope. This is not merely a case of fitting a theoretical curve to a scatter diagram. It is a committment to predict where EACH POINT on the diagram will fall and to treat each deviation of the model's prediction of a point's position from the actual position as an indication of a deficiency in the model. With that kind of goal in mind, one does not just accept the deviations as natural variability, but keeps going back to experiment, and back again, trying to find out why the model didn't predict correctly. One is never satisfied that a model is as good as it can be. And one certainly doesn't publish the first plausible explanation that fits 2/3 of the observations. Not, at least, while claiming that it constitutes knowledge.

>Various perspectives offer knowledge about different aspects of
>behavior. Thus, anthropology is no less valid or more valid
>than PCT, just different in what it chooses to study.

This sort of eclecticism is generous, but I think misguided. The generalizations that are found in anthropology and other such sciences do not generate predictions of any usefulness. They generate descriptions that more or less fit some of the data. There would be far more progress in these fields (and far fewer publications to wade through) if it were required that each description be recast as a prediction of what will happen under specific circumstances, and then that the prediction be tested against new observations. And, of course, that it predict correctly within narrow limits of error.

I think that over decades and centuries of failure to find rigorous principles of behavior, the behavioral sciences have progressively lowered their ambitions until very few behavioral scientists actually expect to explain behavior in clear and convincing terms. It seems to be generally accepted that behavior is too messy, too variable, too complex to be explained in any but the most approximate way. The statistical approach to describing behavior is in itself an abandonment of rigor. Statistical facts are unexplainable; they simply exist. There is a great temptation to give up on the attempt to find clear and correct explanations, and to assume that nature itself contains random factors that will forever obscure the view. After continual failure to penetrate the fog, it must be a great relief to conclude that the fault is not yours, but that of the universe.

It is probably hard for many people to understand why PCTers wax so enthusiastic over dull and uninteresting tracking experiments. The reason is not so much that tracking behavior itself is interesting. It's that tracking behavior LOOKS as if it contains large random components, but under the aegis of the PCT model, those random components turn out to be highly orderly. People accustomed to the normal appearance of behavioral data and the normal kinds of fits of models to the data experience the modeling of tracking behavior using PCT as a relevation, an experience that reveals possibilities in the study of behavior not even dreamed of before.

This becomes even more apparent when a person ventures to extend this sort of modeling to new situations where a new model must be constructed. Look at what Tom Bourbon did. He got his courage together and asked, "Would this work if I set up models of TWO people and had them do a cooperative tracking task, where the outcome depended on BOTH of them?" He abandoned caution and simply set up two control systems model exactly like the ones used for a single-person tracking task, plus the interactions, and tried it. It worked with the same degree of predictivity: the two models interacted in the experiment almost precisely as the two real people did. The first time.

This kind of result is enough to open a person's eyes and to make all other experimental approaches to behavior look very, very inadequate. It doesn't lead to understanding of the more complex aspects of human behavior, but it reveals a path which goes in that direction. It's clear that by going down this path step by step, learning how to extend the scope of the model while maintaining its essentially perfect fit to the data, we could gradually progress to a more and more complete understanding of behavior, in terms that will stand up for as long as Newton's Laws have withstood the test of time. So what, if others are already exploring those regions of complexity? They're not getting this kind of result or anything even in the same universe of discourse with it. PCT will eventually get where they are, and when it does (it, or whatever it has become by then), it will blow away the fog and show what is actually happening, just as it now can do with simple tracking experiments.

When we finally get to those highest levels of organization, PCT may not be recognizeable in current terms. But what will be recognizeable is the expectation and the confidence that a proper approach conducted under high enough standards will produce knowledge of a kind unfamiliar to practitioners of behavioral science as it is known today. Then what will have been the point of all the groping around in the fog that is going on today? Most of what is accepted as knowledge today will probably simply disappear, having led nowhere. That is what has happened to most of the so-called knowledge that has ever been published in the behavioral sciences, even 20 years ago, even 10 years ago. You could probably pick any study done 20 years ago and replicate it (more or less) today -- but who would care? Who will care, 20 years from now, about most of the findings of today?

Is this arrogance, Ken? Or could it be something more interesting than that?

Best, Bill P.

Date: Thu Mar 25, 1993 9:54 am PST
Subject: Re: defining information

Allan Randall. Rick Marken (930317.1000) on definitions:

> >"perceived disturbance," if I may call
>
> I don't like that term -- it gives the impression that the
> disturbance itself CAN be perceived...

> Calling this the "perceived disturbance" may be the simple
> linguistic basis for all the misunderstandings.

I'm not sure if we have a fundamental disagreement on this point or not. I do not have a particular attachment to the phrase "perceived disturbance," but on the other hand I do not think I understand the extent of your reaction against it.

I was not trying to introduce a new term into the lingo. I was simply saying that both of these interpretations meet the everyday use of the word disturbance in some sense. So we must be clear which is meant by our technical term "disturbance." I was pointing out that variations in the CEV are not what we should mean when we say disturbance, which I think you agree with. I needed a handy phrase that captured why it was possible to think of this as "disturbance." I will be perfectly happy to call it "variations in the CEV."

> >I think we can agree that both are reasonable uses of the
> >everyday word "disturbance."
>
> Absolutely not! Disturbance refers to something that "interferes
> with" something else;
> ...But variations in CEV(t) are NOT interfering
> with anything; they just EXIST.

This is just your technical term disturbance. Why can I not call variations in the CEV "disturbances" to a fixed CEV? What I called "perceived disturbance" interferes with a fixed CEV. While I do not see why you are as opposed to this term as you are, I have no problem with dropping it.

> >But "disturbance" in this discussion
> >should refer to the *external* environmental influences, completely
> >separated out from the output of the control system itself. Agreed?
>
> You betcha! I agree.

Okay, so I think we are agreed on the definition of disturbance.

> >Are we also agreed that this disturbance, while defined in this
> >external point of view, is nonetheless defined in terms of the
> >CEV, which is defined according to the internal point of view?
>
> Say what? Why not just say CEV(t) = d(t) + o(t). If that's what
> the above sentence means then I agree with it.

The point is that the disturbance d(t), if separated out from o(t), is not a meaningful quantity to the ECS. It is meaningful only to the external observer. By drawing an arrow marked d(t) you are talking about something the ECS has no direct access to. From the perspective of the ECS, only the variation in the CEV matters. It cannot separate out its own output from the disturbance. On the other hand, this disturbance is defined in terms of the CEV, since only things in the world that affect the CEV can be said to be disturbance. The only point I was trying to make was that there is a mixing of points-of-view here. The external observer who sees the "disturbance" as separate from the output nonetheless needs to know about what the control system is controlling for to define disturbance. Disturbance is defined partially from an internal and partially from an external point of view. This is fine, but it must be recognized. If you are going to mix points of view like this, you need to play by those rules. It is not contradictory because it *is* possible to have an external observer who takes the point of view of the organism - someone who understands what the organism is controlling for. But this observer is not solely taking the organism's point of view if he talks about disturbance - this is his own definition external to the organism.

> In the formula CEV(t) = d(t) + o(t) it is
> always the current value of the output (occuring at time t) that
> is combined with the current value of the disturbance. The fact
> that o(t) might be the result of processes occuring earlier in time
> is of absolutely no consequence.

I was thinking more of the disturbance being a result of past outputs. You are suggesting, if I read you correctly, that we include all past effects of the output into the disturbance. You are considering the instantaneous output of the control net as separate from the disturbance, but that is all. All previous effects of the output are open to inclusion as part of the disturbance. I have no problem with this. I tend to think it is a largely arbitrary decision.

> >Because of quantum effects, it will at some point
> >become impossible, even in principle, for our external observer
> >to separate the disturbing variables from the past output of the
> >control system.
>
> It's really not necessary to fly before you can walk; we're
> not even close to the point were we need to worry about such
> esoteric phenomena.

I'm *not* suggesting that we need to take quantum effects into account at this point. I was simply pointing out that there is a theoretical limit to our ability to separate past output from disturbance. I'm not saying that we should actually draw the line according to quantum uncertainties - merely that it is not possible to get away without drawing a line somewhere. Taking quantum effects into account would be one end of the extreme. The other end is what I think you are sugggesting - to draw the line at the instantaneous, or near-instantaneous, moment of the output's effect on the CEV.

> >Now we need to agree on a working definition of "information." Can
> >everyone agree that if, by making use of B, it is possible to describe
> >A with fewer bits, then B contains information about A?
> ...
> Why not just measure information in the old fashioned way --

I have no problem with the old-fashioned (Shannon-style) definition. I just thought the algorithmic form might be more intuitive than the probabilistic form, especially for the uninitiated. But both forms turn out to be mathematically equivalent, so it doesn't really matter which we use.

> This is a bit fuzzy for me. How do I know how many bits I need to
> describe D without P? ...
> Would P improve my ability to compress D? What are the
> coding rules that I can use? I guess the answer to the above question
> is "no" from me; I need it to be a bit simpler.

Sorry, I know this kind of language is not always crystal clear. I will try to be more precise. The following is not an introductory course in Information Theory, but I hope it gives enough detail for our purposes.

B contains information about A if

H(A | B,l) < H(A | l)

where

H(x | l) = -log P(x | l)   PROBABILISTIC FORM

OR

H(x | l) = the length of the minimal computer program written in language l that outputs x.   ALGORITHMIC FORM

Without the perceptual input, the control system has a certain H(D | h), where h is the hierarchical structure of the control system. The algorithmic form requires us to consider the hierarchy as a computer language capable of a variety of outputs. The "language" part of the system is normally something that does not change, and hence the "l" or "h" in the above equations will often be dropped for brevity. The "program" part is variable, and different programs produce different outputs. So, H(D | h) is the minimum number of bits needed to program the hierarchy to output D. Without P as an a priori part of the language, this is a fair number of bits, but probably far fewer than would be required to get the hierarchy to output some particular random sequence of equal length. Now, if we preprogram P into the hierarchy and consider it to be part of the language and not the program, then we have a minimal program length of H(D | P,h). With P, the disturbance can be generated at the output with far fewer additional program bits than without P. I.e.: H(D | h,P) < H(D | h). The percept contains information about the disturbance.

This relates to the "old-fashioned" probabilistic interpretation. Saying that fewer bits are required to program the hierarchy to produce D turns out to be the same as saying that D is a more probable output over the ensemble of possible programs.

> H = log2 (variance of D).
>
> Then we can measure the information transmitted by P about D
> as the proportion of variance of D accounted for by variance in
> P (with the appropriate log2 transformations). How about that?

Not if you are talking about a standard statistical variance measure. This might be a possible basis for an entropy calculation, but it is not the basis that corresponds to the point of view of the organism. The standard definition uses probability, not variance. If this is basically what you meant, then I'm okay. If you insist on using variance and not probability, maybe you could define exactly what you mean by "variance"?

> If D is digitized and there are 100 8 bit
> samples then do I need 800 bits to characterize D?

You *may* be able to describe it in fewer bits. Information deals not with how many bits are actually used in a digital channel during transmission, but the shift in uncertainty (or entropy) that occurs in the receiver. If the receiver already has some basis upon which he can describe D in fewer than 800 bits, then that is the amount of information received, not 800 bits. If the message is completely random, and the receiver has no a priori structure that would allow a description in fewer than 800 bits, then (and only then) do you need the full 800 bits for the description.

> Without knowing P I could probably think of some compression schemes all on my
> own.

The compression scheme corresponds to the "point of view" taken. Even an apparently random string of digits can be compressed down to a single bit with the right compression scheme (1 means the message is *this* particular string, 0 means its something else). That is why information depends on a point of view, or frame of reference. Some arbitrary compression scheme you come up with yourself may have some very nice mathematical properties, but it is not a measure of information from the point of view of the organism.

> all I've been saying is that the information (however it is measured)
> about d(t) that exists in p(t) is very close to zero as long as you
> have no information about o(t).

It is the function of the control system to protect its perceptual signal from disturbances. So "very close to zero" is exactly right. But not "exactly zero." However, "exactly zero" is exactly the claim that you have been making. This is impossible for an error-controller because, as Ashby said, error-control is inherently imperfect. Large disturbances are protected against by detection of small disturbances. But these small disturbances *are* there, and the system *must* detect them in order to control against them.

> >Are we also agreed that the reference signal can be considered, for
> >the purposes of this discussion, to be constant?
>
> No. It can be a constant or a variable. This should not influence
> the our conclusions about the infornation in p(t) -- or CEV(t).

I think this is what I meant. It doesn't matter for our purposes whether the reference is constant or a variable. So we might as well assume it is constant for now, since that makes things simpler.

> >Are we also agreed that the *output*, if not the percept, contains
> >information about D...?
>
> Given my definition of information transmittion (variance in one
> signal accounted for by variance in another signal) the *output*
> (o(t)) contains nearly 100% of the information about D. So, yes,
> we are definitley agreed on this one.

As I've said before, this is where I really fail to grasp where you're coming from. Since the system only gets information about the world through the perceptual signal, if the organism can output "nearly 100%" of the information about D, then it follows naturally that the resulting information flow from the closed loop must be to bring information about the disturbance in via the perceptual signal. There is no other physical connection coming into the control system. If the control system can output the disturbance almost perfectly, where is it getting this information from? Do you see why it sounds like you are invoking magic? (Maybe you aren't, but this is how it sounds from my perspective.)

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc. Toronto, ON

Date: Thu Mar 25, 1993 10:02 am PST
Subject: heads "up"

[From: Bruce Nevin (Thu 930325 12:25:05)]

I hadn't meant to sound so querulous when I complained about being ignored. (And thanks, Clark!) I realize there are many Important Perceptions being controlled with fairly high gain on the list. It is not the lack of response that was getting to me, but the responses that ignored most of what I was saying, and the feeling that somewhere down the road I would get from Bill more of the same "I don't know about this social norms stuff, it seems like you're just doing that conventional statistical study stuff and dressing it up in PCT jargon with no real change" and zippo from anyone else. As if I had never said anything to show the basis in PCT for what I am concerned with. So the complaint was an intemperate reaction of frustration, and this is to focus it lest it cause more trouble than that which occasioned it.

[Gary Cziko 930323.1700 GMT] --

The driver of the car "participates" in the movement of the car in various ways, for example by leaning into curves. The driver tilts her head on curves, and probably also in other acceleration (starts and stops). Passengers don't. This is probably why people who experience motion sickness as passengers do not as drivers. I suspect the controlled perception is the attitude of the inner ear with respect to the center of gravity.

Bruce bn@bbn.com

Date: Thu Mar 25, 1993 10:43 am PST
Subject: Re: you CAN tell, diagram

[From Rick Marken (930324.0900)] Bruce Nevin (Wed 930324 08:44:38) --

>It is possible to know what someone is doing by watching what
>they are doing, when they are doing something that is socially standardized.

When I say that you cannot tell what a person is doing by watching what they are doing, I mean that you cannot tell what perception a person is controlling by watching the outputs (or observable side effects of the outputs) that are involved in the control of those perceptions. This applies whether or not these outputs (or side effects) are considered "socially standardized" events or not. A concrete example: can we agree that "stopping at a red light" is a socially standardized event? Your statement above suggests that if I see someone stop their car at a red light I know that this event is "what they are doing" -- by which I mean, the intended result of their actions or (in PCT) the perception that is under control.

But it is certainly possible that the person has not noticed the red light but, instead, has seen a robbery in progress in the store behind the light. Coincidentally, just as the light turns red, driver stops the car -- not to produce the socially standard- ized result "stop at red light" but to produce the socially quite non-standard result; try stop the robbery.

>The meanings (nonverbal perceptions) that I associate with the
>words and words-in-relation may differ from those that you
> associate with them, and probably do. Nonetheless, we assume
>that the differences should be matters of detail.

This is a big assumption. You might make it; but it seems clear to me (especially from these net discussions) that the differences in the meanings (non-verbal perceptions) evoked in different people by the very same words can be quite a bit more than a matter of detail.

>(This is the fundamental assumption of science: that the universe is knowable,
>consistent from one part of it to another and from differing
>points of view.)

I don't see what this assumption has to do with language. I think the success of science results from its non-dependence on language; agreement comes from experimental demonstrations -- finding that you can reliably produce particular perceptions by operating on the world in a particular way. We do have to try to communicate (using language, to some extent) how to operate on the world in order to produce particular perceptions reliably -- anyone who has taken a science lab knows how difficult this communication process can be. But once one is able to translate the linguistic descriptions into appropriate (non-verbal) control systems in oneself (if one is ever able to do this, and this is not guaranteed -- like me in chemistry lab) then one can stop talking and start doing science. This works not only in science but in others areas where language is used (at least in the beginning) to try to communicate control skills -- like in teaching a kid how to drive or do math. The short- comings of language become quite apparent in these situations -- and we sometimes resort to physical demonstrations and what not.

>When you are using language, you are producing outputs as part of
>the process of understanding what was said.

I would say "I'm producing outputs as part of the process of producing perceptions that match my references for those perceptions (perceptions corresponding to "understanding" among them, perhaps); the references are the embodiment of 'what I want to say'"

> There is no need to
>invoke references for uncontrolled perceptions to account for
>your discomfort at someone saying things about a cherished
>concept that you would regard as error if you said them yourself.

I hope my diagram cleared up why they are necessary. It is also worth recalling that, if a perception as under a reasonable level of control, there is not much error. I can't control my perceptions of what others say and these are the perceptions that tend to be "disturbing" -- in the Tuckerian sense, that is, they result in error. I imagine that this is because I have a reference for these uncontrolled perceptions.

>And I hardly think it is the case that your concept of PCT is an
>uncontrolled perception for you!

I can only control what I can affect via the environment. I can affect what I perceive myself saying about PCT; so I can control my concept of PCT as it results from my own actions; but I can't control what I perceive to be your concept of PCT. Having a reference for a perception is not enough to make it a controlled perception. It must also be true that the perception is a joint function of the environment (d(t)) and my effects on that environment (o(t)). [After reading your next post I have had to amend this and admit that, it is possible to control when o(t) is imaginary].

Bruce Nevin (Thu 930325 08:36:48)--

>I am glad you read the first sentence of my (Wed 930324 08:44:38).
>Any agreements or disagreements with the rest of it?

Both. See above.

On your revision of my diagram; I see is that you have removed the effects of my outputs on the environment.

>What is indicated here is a "shadowing", in imagination, of the
>other's utterance in the environment.

I think you would also have to change the controlled variable, p(t), to something like "g(o(t) = g(d(t))"-- that is, the higher level variable is controlling for "shadowing" by generating an imaginary perception (o(t)) to match what is being said, d(t). I suppose this is a good exmaple of controlling completely through the use of imagination. So I guess I am wrong about control only occuring when the system has an effect on the environment. Here is a case where p(t) = imagined(o(t)) + d(t). So I take back what I said in the paragraph above -- I was wrong (how's that for non-arrogance). You can control using only imagination.

>There is a problem indicating the relation between f[o(t)] and
>meaning[o(t)] with a labelled arrow if that means, as it usually
>does, a relation between two levels of the perceptual hierarchy.

I was obviously skipping a lot of steps is the hierarchy. The point is that somehow the nervous system converts strings of written or spoken words into perceptual signals; at one level these signals a preceptions of the word patterns themselves. At other levels they are perceptions of what is represented by the word. My little perceptual function boxes collapsed many layers of perceptual processing into a simple input -- output function; words in -- meaning (in the form of a variable magnitude perceptual signal) out. What that signal "means" is embodied in the function that converts words into perceptual signals.

>Does ANYBODY get it yet, or will this be ignored too?

I don't know what I was supposed to get? But I think I now agree that people can control variables using imaginary output. I just think that if they do it alot they might run afoul of their intrinsic references.

Best Rick

Date: Thu Mar 25, 1993 12:27 pm PST
Subject: Re: definiing information

[From Rick Marken (930325.1100)] Allan Randall (939325) --

>By drawing an arrow marked d(t) you are talking
>about something the ECS has no direct access to. From the perspective
>of the ECS, only the variation in the CEV matters. It cannot separate
>out its own output from the disturbance.

YES!!!!

>On the other hand, this
>disturbance is defined in terms of the CEV, since only things in the
>world that affect the CEV can be said to be disturbance.

I am assuming (and in the tracking experiments it is true) that I know the variable that the person is controlling -- the CEV. But this is a good point -- an event in the world is only a "disturbance" with respect to a controlled variable.

>Sorry, I know this kind of language is not always crystal clear. I will
>try to be more precise. The following is not an introductory course in
>Information Theory, but I hope it gives enough detail for our purposes.
>
>B contains information about A if

>H(A | B,l) < H(A | l)

>where

>H(x | l) = -log P(x | l) PROBABILISTIC FORM

>OR

>H(x | l) = the length of the minimal computer program

>Without the perceptual input, the control system has a certain H(D | h),
>where h is the hierarchical structure of the control system.

>So, H(D | h) is the minimum number of bits needed to program the
>hierarchy to output D.

>Now, if we preprogram P into the hierarchy and consider
>it to be part of the language and not the program, then we have a
>minimal program length of H(D | P,h). With P, the disturbance can be
>generated at the output with far fewer additional program bits than
>without P. I.e.: H(D | h,P) < H(D | h). The percept contains information
>about the disturbance.

NOW WE ARE GETTING SOMEWHERE.

All you have to do now is show that your statements above are true. Show me the minimal program required to generate D and the minimal program required to generate D|P (are there some rules about what constitutes a program step?). I predict that you will need EXACTLY the same size program to generate D (given nothing) and to generate D|P. That is, I predict that H(D|h) = H(D|P,h). If H(D|h) - H(D|P,h) is a measure of the information about D in P then I claim (and you can now prove me wrong) that the information about D in P is precisely zero. Note that in all of your discussion above you quite correctly did not include the output, O, in your calculation of D. Of course, you can easily compute D given P and O (H(D|P,O,h) in one program step -- D = P - O.

>As I've said before, this is where I really fail to grasp where
>you're coming from. Since the system only gets information about
>the world through the perceptual signal, if the organism can output
>"nearly 100%" of the information about D, then it follows naturally that
>the resulting information flow from the closed loop must be to
>bring information about the disturbance in via the perceptual signal.
>There is no other physical connection coming into the control system.
>If the control system can output the disturbance almost perfectly,
>where is it getting this information from? Do you see why it sounds
>like you are invoking magic?

Believe me, I KNOW it sounds like I'm invoking magic. I KNOW that what I am saying must sound irritating and exasperating. Your descrip- tion above makes it very clear WHY my claim (that there is no information about D in P) must seem down right mysticism. It does seem like information about the disturbance MUST be brought into the closed loop by the perceptual signal. The problem is -- that's not the way it works. The perception itself contains NO information about the disturbance -- no usable information anyway --as you will see if you carry out your proposed information measures (H(D|h) and H(D|P,h)) on real (or simulated) data from a simple compensatory tracking task. The fact is that the output mirrors the disturbance -- NOT because the closed loop get's any information about the disturbance in perception. The mirroring of D by O is a SIDE EFFECT; it is the result of the closed negative feeback loop producing O values that continuously keep P nearly equal to R.

But PLEASE DON'T TAKE MY WORD FOR THIS: DO THE INFORMATION CALCULATIONS AND REPORT YOUR RESULTS ON THE NET. Then we can go from there.

Best Rick

Date: Thu Mar 25, 1993 12:37 pm PST
Subject: red light

[From: Bruce Nevin (Thu 930325 15:06:56)] Rick Marken (930324.0900)

I guess I asked for it. I won't be able to respond right away to all the errors (e.g. discrepancies between what I said and what I perceive you saying I said, and discrepancies between what I perceive you saying and how I would say it). Just one point for today: stopping at a red light is legislated; the socially institutionalized norms that constitute language are not (despite the best efforts of Safire et al.).

The behavioral outputs involved in stopping at a red light are not controlled for conformity to social norms in the way that the utterance of a word, a phrase, or a sentence is. At most, stopping at a red light may be controlled as expressive gesture, like Martin's exaggerated nodding of the head to complain about poor driving: being careful to stop smoothly (or, if a teenager, perhaps being careful to squeal the tires). But in that case, stopping for any other reason is likewise subject to the same control of the perception of having expressed some affect or of not expressing affect so as not to draw attention to oneself, etc. (whatever it may be).

The socially institutionalized norms for language should not be confused with rules of conduct and legislation. This confusion has to be cleared away before you can know what I am talking about. But it's hard for you to get it because the former are invisible to you (to most of us most of the time), I think for reasons that I have suggested in the past. Tell me about your experience learning another language, say, Spanish.

hasta man~ana Bruce bn@bbn.com

Date: Thu Mar 25, 1993 12:48 pm PST
Subject: Someone may want to comment

[Martin Taylor 930325 15:30]

The following is an invitation from BBS for open commentary that might appeal to some CSG-L readers (others may suffer from the phenomenon qua fact known as "learned helplessness").
================

Below is the abstract of a forthcoming target article by MARC JEANNEROD, on MOTOR INTENTION, IMAGERY AND REPRESENTATION, that has been accepted for publication in Behavioral and Brain Sciences (BBS), an international, interdisciplinary journal providing Open Peer Commentary on important and controversial current research in the biobehavioral and cognitive sciences. Commentators must be current BBS Associates or nominated by a current BBS Associate. To be considered as a commentator for this article, to suggest other appropriate commentators, or for information about how to become a BBS Associate, please send email to:

harnad@clarity.princeton.edu or harnad@pucc.bitnet or write to: BBS, 20 Nassau Street, #240, Princeton NJ 08542 [tel: 609-921-7771]

To help us put together a balanced list of commentators, please give some indication of the aspects of the topic on which you would bring your areas of expertise to bear if you were selected as a commentator. An electronic draft of the full text is available for inspection by anonymous ftp according to the instructions that follow after the abstract.

THE REPRESENTING BRAIN: NEURAL CORRELATES OF MOTOR INTENTION AND IMAGERY

Marc Jeannerod
Vision et Motricite
INSERM Unite 94
16 avenue du Doyen Lepine
69500 Bron
France

KEYWORDS: affordances, goals, intention, motor imagery, motor schemata, neural codes, object manipulation, planning, posterior parietal cortex, premotor cortex, representation.

ABSTRACT: This target article concerns how motor actions are neurally represented and coded. Action planning and motor preparation can be studied using motor imagery. A close functional equivalence between motor imagery and motor preparation is suggested by the positive effects of imagining movements on motor learning, the similarity between the neural structures involved, and the similar physiological correlates observed in both imagining and preparing. The content of motor representations can be inferred from motor images at a macroscopic level: from global aspects of the action (the duration and amount of effort involved) and from the motor rules and constraints which predict the spatial path and kinematics of movements. A microscopic neural account of the represenation of object-oriented action is described. Object attributes are processed in different neural pathways depending on the kind of task the subject is performing. During object-oriented action, a pragmatic representation is activated in which object affordances are transformed into specific motor schemata independently of other tasks such as object recognition. Animal as well as clinical data implicate posterior parietal and premotor cortical areas in schema instantiation. A mechanism is proposed that is able to encode the desired goal of the action and is applicable to different levels of representational organization.

To help you decide whether you would be an appropriate commentator for this article, an electronic draft is retrievable by anonymous ftp from princeton.edu according to the instructions below (the filename is bbs.jeannerod). Please do not prepare a commentary on this draft. Just let us know, after having inspected it, what relevant expertise you feel you would bring to bear on what aspect of the article.

To retrieve a file by ftp from a Unix/Internet site, type either:
ftp princeton.edu
or
ftp 128.112.128.1
When you are asked for your login, type:
anonymous
Enter password as per instructions (make sure to include the specified @),
and then change directories with:
cd /pub/harnad/BBS
To show the available files, type:
ls
Next, retrieve the file you want with (for example):
get bbs.jeannerod
When you have the file(s) you want, type:
quit
In case of doubt or difficulty, consult your system manager.
A more elaborate version of these instructions for the U.K. is
available on request (thanks to Brian Josephson)>

Where the above procedures are not available (e.g. from Bitnet or other networks), there are two fileservers:
ftpmail@decwrl.dec.com
and

bitftp@pucc.bitnet
that will do the transfer for you. To one or the other of them, send the following one line message:

help

for instructions (which will be similar to the above, but will be in the form of a series of lines in an email message that ftpmail or bitftp will then execute for you).

----- End Forwarded Message -----

Date: Thu Mar 25, 1993 1:12 pm PST
Subject: talking red light blues

[From Rick Marken (930325.1245)]   Bruce Nevin (Wed 930324 08:44:38) said

>It is possible to know what someone is doing by watching what
>they are doing, when they are doing something that is socially standardized.

I gave the "stop at red light" examle of a socially standardized case where you could not tell what someone was doing by watching what they are doing.

Bruce Nevin (Thu 930325 15:06:56) roundly rejected this example and said:

>The behavioral outputs involved in stopping at a red light are
>not controlled for conformity to social norms in the way that the
>utterance of a word, a phrase, or a sentence is.

So let me see if I get this right: you CAN tell what someone is doing by watching what they are doing if what they are doing is uttering a word, phrase or sentence?

Best   Rick

Date: Thu Mar 25, 1993 4:13 pm PST
Subject: BBS article

[From Rick Marken (930325.1530)]   Martin Taylor (930325 15:30)--

>The following is an invitation from BBS for open commentary that might
>appeal to some CSG-L readers

>THE REPRESENTING BRAIN: NEURAL CORRELATES OF MOTOR INTENTION AND IMAGERY

Yeah. I wanted to do that one, too, but I already signed up to do the Killeen article on a mathematical model of reinforcement. The Jeannerod article on motor intention looks like even more fun. I think some CSGnetter should do a review of it.

Best   Rick

Date: Thu Mar 25, 1993 6:44 pm PST

Subject: On disturbances

[From Bill Powers (930325.1530)]   Rick Marken (930317.1000 ff) --

Your excellent diagram makes it clear that we have to redefine the disturbing variable. To call an external variable that affects (potentially) a CEV a "disturbing" variable is to assert that the organism has a reference level for the state of the CEV. A noncommittal way to speak of such cases is to speak merely of independent variables, or distal variables.

In your diagram there is one information path which shows the "disturbing" variable being perceived but without any reference signal or comparator. In that case, the perception is simply that the CEV exists. A perceptual function creates an CEV, but only a complete control system creates a controlled CEV.

A statement made by someone can therefore be perceived and translated into a meaning without constituting a disturbance. It is just a perception. Only when there is a preferred state for the meaning of that perception and a means for comparing the actual state with the preferred state can there be any error signal, and only when the error signal is translated into an effect on the source of the perception can we talk about a disturbance.

Martin Taylor, some time ago, proposed an "alerting system," a proposal that I drew back from, not seeing how it would operate. Now perhaps I understand a little more of what he meant. If we say that an alerting system is simply a control system without its output connected, then we can imagine reference signals defining preferred states of the environment, and error signals that indicate a departure of the environment from that preferred state, but which in fact result in no action. If something does occur as a consequence of such error signals, we can only conclude that the error signals themselves are being monitored by some other system, which acts by changing the complement of active control systems, perhaps bringing some of the "alerting" systems into action by connecting their outputs as appropriate, and disconnecting others. This is a possible mode of hierarchical control that needs further clarification so we can try to model it.

Allan Randall (930325) --

RE: perceived disturbances

Rick's objection to this term, and mine, is that it is too loose. Do you mean that there is a perceptual signal that specifically represents the amount of a change in a CEV, rather than the total magnitude of the CEV (or its perceptual representtion)? If there is a perceptual signal with a magnitude of 100 units, and an independent variable in the environment changes so as to make the perceptual signal become 110 units, a "perception of the disturbance" would indicate 10 units, not 100 or 110.

>Why can I not call variations in the CEV "disturbances" to a fixed CEV?

You can, using an informal meaning of disturbance (the effect). But if you're going to propose that these "variations" are themselves perceptions, you have to propose a perceptual function that specifically reports variations. If you have a control system that works in terms of the total magnitude of the CEV, it is not reporting the CEV in terms of changes, but of magnitude. To get a perception of changes, you have to bring in a different kind of input function, one sensitive to first derivatives or that continually compares values of the CEV at time t with values at time t-tau. That yields a perception of change -- but would be useless for controlling magnitude.

What you can't do, in the spirit of modeling, is to speak of an unexplained or arbitrary change in the CEV. If the CEV changes, something must have changed it; otherwise it would not have changed. This is precisely the same principle that Newton expressed as his first law of motion. This is why I keep insisting on the external independent variable and its physical link to the CEV. A CEV doesn't just change. It changes as a result of something acting on it. When you include an independent environmental variable as one contributor to the change, you also provide a handle by which an experimenter can apply influences to the control system without breaking the control loop. In a model, every change in every variable must be accounted for in some way, either as a function of other dependent variables or as a function of independent variables -- which themselves are eventually accounted for as experimenter actions or as the actions of higher control systems.

In modeling, you can't say "let variable b have the value of 10" if variable b depends on the states of other variables. You have to look for the independent variables and propose influences on the system through manipulating those variables. Otherwise you're simply violating the definition of the model. Variable b, if dependent on other variables, is an unknown in the system equation, for which you can solve if you know or specify the states of the independent variables.

The state of a CEV depends on two variables: the output of the control system, and a representative external independent variable that I call the disturbing variable.

>While I do not see why you are as opposed to this term as you
>are, I have no problem with dropping it.

I would much rather you understood the opposition. I hope the above explanation helps you understand.

>The point is that the disturbance d(t), if separated out from
>o(t), is not a meaningful quantity to the ECS.

Whenever you see d(t) in one of Rick's equations, or D in one of my diagrams, you should think of it as an environmental variable outside the feedback loop, connected to the CEV as one end of the rubber bands is connected to the knot. It does not EVER mean "an arbitrary change in the CEV." d(t) could be varying between plus and minus 100 units, while the CEV remains within 1 unit of its reference level and is varying in a waveform that has no resemblance to the waveform of d(t). All of this would be so much easier to understand if you would just run one of our simple tracking experiments and look carefully at the data!

d(t) is ALWAYS "separated out from o(t)." It's a physically distinct variable. And d(t) and o(t) are physically distinct from the controlled variable.

>By drawing an arrow marked d(t) you are talking about something
>the ECS has no direct access to.

Yes, this is correct. But d(t) is not the arrow; it is the variable at the start of the arrow. The arrow merely indicates the connection that gives d(t) an influence on the variable at the head of the arrow. We would associate with the arrow, perhaps, a constant of proportionality -- but not a variable. The constant of proportionality written over the middle of the arrow, times the value of the variable at the start of the arrow, is the magnitude of the influence on the variable at the head of the arrow.

Unfortunately, this convention holds in the environmental part of a control-system diagram but not inside the control system. Inside the control system, the boxes labeled "input function" and so on contain the transformations, while the arrows indicate the variables -- signals traveling from one place to another, carrying the magnitudes of the outputs of the boxes. I have thought many times of making the conventions uniform inside and outside, but have never done so because in the environment, it seems to me, the variables are associated with physical places where we could measure them, while the arrows express the invisible physical laws that connect the visible variables. Inside the nervous system, the variables are associated with the signals that go from one place to another, while the functions are performed in specific physical locations (distributed or not). It would be very nice to be able to express processes inside and outside of the system in the same way. But maybe there is a fundamental difference that the present convention acknowledges.

>The external observer who sees the "disturbance" as separate
>from the output nonetheless needs to know about what the
>control system is controlling for to define disturbance.
>Disturbance is defined partially from an internal and partially
>from an external point of view. This is fine, but it must be
>recognized.

I agree, and it's good to have this pointed out explicitly. The external observer here is not a naturalist or a behaviorist, but a modeler. The variable affected by the disturbing variable and by the output of the system is, in the eyes of the modeler, more than just one variable caught between two others. It is the visible entry-point to a control system, and the output is the visible exit point. What goes on between entry and exit is imagined, but carefully. Without that model of the control system, there would be no reason to characterize one variable as a disturbing variable and the other as an output. They would both just be physical variables with effects on many other variables, including the one between them.

>I was thinking more of the disturbance being a result of past outputs.

Do you mean in general? In general this isn't true, although it might be true (the transient following complete removal of all external sources of disturbance).

In general a disturbing variable is like the wind affecting the path of a car. The disturbing variable, the wind, changes completely independently of the path of the car and the steering efforts of the driver.

The wind velocity (the disturbing variable) is converted by the laws of aerodynamics (the arrow) into an influence on the path of the car (at the head of the arrow).

>You are suggesting, if I read you correctly, that we include
>all past effects of the output into the disturbance. You are

>considering the instantaneous output of the control net as
>separate from the disturbance, but that is all.

No, this is not what Rick or I would mean. Is the reason becoming clearer? By disturbance we do not mean the change in the controlled variable itself, although there's many a slip in trying to stick strictly to the proper meaning. We mean a variable physically distinct from both the output and the controlled variable, which acts through some environmental link on the controlled variable at the same time that the output is acting, through a different link, on the same variable.

Past effects of the output are still part of the closed negative feedback loop. The disturbing variable is not in the negative feedback loop.

I sure hope this is getting across. When you see exactly what we mean here, I believe that all this confusion will melt away.

I'm not going to get into the probabilistic stuff. I'm having a hard enough time sorting it out in private conversations with Martin.

Bruce Nevin (930325.1215) --

I don't mean to ignore your communiques. Despite my quibbles I like what you're doing and believe that you're working your way toward a real PCT understanding of linguistics -- individual and social. But I've run out of brain cells -- there are too many conversations going on, and I feel that you and I have no argument over basic principles. I'm gradually working my way back into the area where my competence lies: modeling. I'm trying to drag others in the same direction, but basically I just want to get back into it myself and do something real.

Best to all,   Bill P.

Date: Fri Mar 26, 1993 5:35 am PST
Subject: Arrogant values?

From Greg Williams (930326)   Bill Powers (930325.0850)

>Is it arrogance to set standards for accepting knowledge claims
>so high that even one's own attempts to make claims are more
>likely to fail than succeed?

It is arrogance to suppose that one's own goals are the only, or even the most, important ones. Such a view then can lead to the belief that the methods useful in achieving one's own goals are the best methods for everyone. Just try selling an insurance exec on the method of modeling -- he/she does fine with statistical descriptions of populations and doesn't care about predicting individual behaviors. In fact, the method of modeling individuals would be unwieldy here, to say the least, not to mention the fact that it is still in its infancy and cannot be applied to "high-level" behavior now (or, I predict, for decades to come).

>I agree with you that there are other approaches to knowledge,
>but in my view that is all they are: approaches. We have to spend
>a lot of time mulling over imperfectly stated propositions and
>guessing how we might investigate them, but I don't think that
>this kind of preliminary activity belongs in the category of

>"knowledge claims." It's the sort of thing you do while you're
>looking for an idea that has enough solidity to be put forth as a
>candidate for a knowledge claim.

As Phil Runkel pointed out in his book, there is a lot which can be accomplished by the "Grand Method" (his term) of descriptive statistics of populations. I would suggest that there are some human goals which can be met more efficiently by using descriptive statistics than by using the (again, "Grand Method") of individual models. Yes, the user of statistics must be careful in the ways which have been abundantly discussed on this net. But I believe that statistics of populations IS a kind of knowledge. It is not as detailed a knowledge as models of the population individuals; this confers both disadvantages AND advantages. Surely you don't expect epidemiologists sometime in the (far?!?!) future to replace statistical description with individual models. Some physicists still work with thermodynamics and statistical mechanics while others model individual molecules (SMALL ones!) -- why? because they have varying goals, for which varying approaches are appropriate, mainly in terms of efficiency.

>Even Newtonian mechanics is still the undisputed theory of choice
>in very nearly all macroscopic situations -- and has been for
>over three centuries. This is what you get when you cross the
>boundary between verbal reasoning and precise analysis -- when
>you're able to do it.

Ah, there's the rub... when you're able to do it. When you're able, I'd be happy to submit to a tracking experiment (say, at a CSG meeting however many years from now when you are ready); you tell me what the task is supposed to be and then try to predict what I actually do. When will your modeling and associated methodolgy be sufficiently sophisticated to predict when I will COOPERATE and when I won't? Not for a long time, of course. And when PCT has leapt that hurdle, the next one is to make models for THOUSANDS of individuals and combine them some way to predict population measures. Lots of luck. No, it is better to follow the example of physics and stick with descriptive statistics for generating SOME kinds of knowledge.

As you know, I am very sympathetic to the method of modeling in general and to PCT models in particular. But I am not arrogant enough to think that modeling individuals is the only path to knowledge. It IS the only path to SOME kinds of knowledge. But some people don't need that kind of knowledge.

As ever, Greg

Date: Fri Mar 26, 1993 8:43 am PST
Subject: Re: talking red light blues

[Martin Taylor 930326 11:15] Rick Marken 930325.1245

>So let me see if I get this right: you CAN tell what someone is
>doing by watching what they are doing if what they are doing is
>uttering a word, phrase or sentence?

Consider what one is likely to be doing (PCT sense) when one utters a word, phrase, or sentence. Colloquially, it is called "communicating," which might be translated as "attempting to let another know what one is doing (PCT sense)."

(Avoiding for the moment lies and other deceptions). So, if one is controlling well, your communicative partner CAN tell what you are doing in the specific case of language (construed to include body language, etc.). The conventionalization of language has evolved just so that this can happen.

When one is interacting with non-linguistic entities, there is no evolutionary benefit in an outsider's possibly being able to see what one is doing. It is enough that one does it. When one is interacting with a linguistic entity (person or now computer) it is very important that one controls for the partner to know what one is doing, whether to enhance that knowledge (cooperative communication) or to defeat it (deception).

Language IS special that way.

Martin

Date: Fri Mar 26, 1993 8:59 am PST
Subject: Arrogance?

[From Rick Marken (930326.0800)] Bill Powers (930325.1530)--

>I'm gradually working my way back
>into the area where my competence lies: modeling. I'm trying to
>drag others in the same direction, but basically I just want to
>get back into it myself and do something real.

I'll be coming along soon. One of the "real" things that I hope might come out of this "information in perception" is an idea for models, demos or experiments that would illustrate some of the points we are trying to make about the behavior of a control system. I think your modelling efforts are important because they show what a control system model can do in terms of producing observable behavior that is considered "interesting" or "complex". That is why the ARM demo is so important. What I want to try to do is develop demos of the "implications" of a control organization for conceptualizing behavior. Your ARM demo, for example, manages to produce complex pointing behavior even though there is no information in the perceptual inputs to the model that allows it to do this. But this fact about the operation of the model is not obvious; one could view the ARM demo as another clever approach to producing observable behavior instead of what it actually is; a fundementally differnet conception not only of the processes that generate behavior but of the nature of behavior itself. I want to try to develop models (demos, whatever) that show the remarkable implications of the control model (such as the fact that the model works without any information about "what to do" coming from its percpetual inputs). I hope that by pushing on this topic (to the extent that others are willing to push back) ideas for new demos-models may be ignited by the sparks of friendly conflict.

While Bill P. is off looking for something real, I'll take the liberty of responding to:

Greg Williams (930326) >>Bill Powers (930325.0850)

>>Is it arrogance to set standards for accepting knowledge claims
>>so high that even one's own attempts to make claims are more
>>likely to fail than succeed?

>It is arrogance to suppose that one's own goals are the only, or even
>the most, important ones. Such a view then can lead to the belief that
>the methods useful in achieving one's own goals are the best methods
>for everyone. Just try selling an insurance exec on the method of
>modeling -- he/she does fine with statistical descriptions of
>populations and doesn't care about predicting individual behaviors.

I was assuming (and Bill was too) that Ken Hacker shares many of the same goals of PCTers -- such as, the goal of gaining some knowledge about the processes underlying the phenomenon of human behavior, both individual and collective. I don't see where Bill, Mary or I (we've been the only participants in this discussion) have supposed that our goals are the only or the most important ones. What gave you that impression? I really am curious. I'm willing to believe you; I just don't know what "goals" you are talking about and, because of that, your reply to Bill seems like a non-sequiter. Let me try to be a bit more specific. I see this "arrogance" discussion this way:

1) Ken Hacker said that PCTer are arrogant because they make "hyperbolic" claims. I agree that making hyperbolic claims is arrogant so PCTers who make such claims are arrogant.

2) I said: Please tell me what hyperboloic claims have been made and please explain why they are hyperbolic, using mathematical, working model or experimental evidence.

3) Ken Hacker said that my asking for such evidence was "arrogant".

4) That was followed by Bill P's delicious question:

>Is it arrogance to set standards for accepting knowledge claims
>so high that even one's own attempts to make claims are more
>likely to fail than succeed?

I don't see how you got from this dialog to:

>It is arrogance to suppose that one's own goals are the only, or even
>the most, important ones.

Help me out here.

By the way, I think its fine for people studying population level phenomena (insurance companies, polling agencies, etc) to use sampling statistics. Their goals are different from mine (in PCT) and that's fine with me. I have no interest in trying to derive actuarial phenomena from individual behavioral models.

Best Rick

Date: Fri Mar 26, 1993 9:13 am PST
Subject: talking blue flashing lights

Me (Thu 930325 15:06:56):

>The behavioral outputs involved in stopping at a red light are

>not controlled for conformity to social norms in the way that the
>utterance of a word, a phrase, or a sentence is.

[Rick Marken (930325.1245)]

>So let me see if I get this right: you CAN tell what someone is
>doing by watching what they are doing if what they are doing is
>uttering a word, phrase or sentence?

Yes. Specifically, you can tell that they are uttering that word, that phrase, that sentence.

Because of the socially standardized word dependencies in language, you can tell that they produced an utterance with a certain structure. This structure of their utterance is the linguistic information in the utterance. It is in the utterance, that is, socially available and not just the private perceptions of one party or another, because people have worked very hard to learn to control linguistic structure in a socially standardized way. As to why it is called information, refs on request if you have lost them.

You CAN'T tell what they are saying it FOR, which is part of what they mean by it. You can tell what you imagine you would be saying it for if you were saying it. And you imagine that is their motivation too. (Or perhaps you construct a model of them and their motivations in imagination, but that comes back ultimately to projection [in imagination] of your own motivations.)

And there is ambiguity. A given utterance can have more than one structure concurrently, e.g. Harris' example (used later by Chomsky) "Flying planes can be dangerous." If you are willing and able to display the ambiguous possibilities to yourself in imagination, you can tell that the other person means at least one of a set of category perceptions, but you can't tell which one.

But usually you are unaware of possible ambiguity. Usually you have imagined a version of meanings and motivations consistent with your perceptions of the subject matter domain and of the social relations between you, and don't notice that the utterance is structurally ambiguous (that the linguistic information in it might be different from what you had perceived). You don't try out the alternative branches of the ambiguity.

An example of ambiguity: substitute "one" for "you" in the preceding, especially in case you think I am making generalizations about you personally. This is because "you" as an indefinite noun used interchangeably with "one, someone" is a different word from "you" the person I am talking to. They're pronounced and spelled the same (and have a common ancestor in earlier English), but they are now different words, like "beet" and "beat". Another example: "socially standardized" does not mean society standardizes people or their references. People set their reference perceptions in a way that they perceive (and are willingly helped to perceive) as being the norm. This helps them to control a perception of being able to cooperate with others.

Ambiguity is quite apart from the differences in nonverbal perceptions (non-category perceptions) that we associated with utterances in an idiosyncratic way that is not socially standardized. You can't call that ambiguity, because ambiguity is a choice between structurally defined alternatives.

I started to say more about ambiguity and motivations and social relations, but I've ripped it out and stuffed it in a file. Not too many steps at a time.

So when you see someone stopping at a red light you don't know what they are doing. Your perception that they are interrupting their travel to let other traffic cross might accord with theirs, even probably does (robberies are not *that* common, and passersby stopping to interfere in one even less so). But the conforming to the social arrangement about red lights and traffic is motivated by avoidance of direct hazards to life and limb, or perhaps only to one's driver's license and insurance rate if there's no traffic. And that's it. It is isolated; it is not part of a system of socially standardized interdependencies. If I hear a sound like "pu" it might be my little daughter blowing out the candle. (You recall the story about how Poo Bear might have got his name.) It might be any number of other things, a piece of paper stuck to her lip, anything. If I hear it in context I have no doubt that she has produced the word "put" in the sentence "put the candle out", and that the blowing out of the candle was an accidental byproduct, or that pretending that it was is intended for a joke. And whatever perceptual consequences might be associated with "put" in this case are not the same in general (as for the hazards associated with running a red light). They are different in "I won't put up with that", and different yet again in "up with which I will not put," not to mention "shot put." These differences are structurally determined by the linguistic contexts of words and word dependencies in which the different occurrences of the sound "put" are heard.

Too much? Are you still with me? And can you tell me about your experience learning Spanish?

Bruce bn@bbn.com

Date: Fri Mar 26, 1993 9:26 am PST
Subject: Arrogance and other opinions

[From Bill Powers (930326.0830)] Greg Williams (930325.0850) --

Back on the net in fine fettle, I see.

>It is arrogance to suppose that one's own goals are the only,
>or even the most, important ones.

Oh, heck. Well, is it arrogance to suppose that one's own definition of arrogance is the only, or even the most important one? What is the most important goal, then? And where can I look out there to find it?

>Just try selling an insurance exec on the method of
>modeling -- he/she does fine with statistical descriptions of
>populations and doesn't care about predicting individual behaviors.

Seems to me that I've commented a number of times on the fact that people who deal with populations do fine (for themselves) with statistics, because statistics concerns the properties of whole populations, not individuals. When that's your concern, who cares if a few of the ants get stepped on?

Insurance executives do use a model to predict behavior. It's called the actuarial tables. It predicts the behavior of populations.

>In fact, the method of modeling individuals would be unwieldy
>here, to say the least, not to mention the fact that it is
>still in its infancy and cannot be applied to "high-level"
>behavior now (or, I predict, for decades to come).

If enough people don't adopt the method of modeling and try to apply it to understanding human nature, you can change "decades" to "centuries."

In the 19 March _Science_, p.1773, there's an interesting review of _Testing testing_ by F. Allan Hanson. "Hanson goes on to argue that 'institutional analysis' reveals that history can be 'read' as the gradual perfection of the manipulation and subordination of human beings to serve institutional ends." One of the primary tools for achieving this end is statistical analysis of population tests.

>As Phil Runkel pointed out in his book, there is a lot which
>can be accomplished by the "Grand Method" (his term) of
>descriptive statistics of populations.

I agree. I don't agree with much of what it is used to accomplish. And Phil agrees that casting nets is no way to understand specimens.

>I would suggest that there are some human goals which can be
>met more efficiently by using descriptive statistics than by
>using the (again, "Grand Method") of individual models.

Whose goals are you talking about? If you own an insurance company, it's most efficient to use statistics to set the rates, to raise the rates on people who make claims, and to reject applicants who belong to high-risk populations. That's the most efficient way to satisfy your personal human goal of making money out of selling insurance to a population.

>But I believe that statistics of populations IS a kind of knowledge..

Yes, it is. It is knowledge about populations.

>It is not as detailed a knowledge as models of the population individuals ...

It is not detailed at all. It is a mass measure, applying only to the entire mass.

>Surely you don't expect epidemiologists sometime in the (far?!?!) future
>to replace statistical description with individual models.

It might help them if they could say why person A catches the disease while persons B ... Z do not. Identifying individual carriers has occasionally been important.

>Some physicists still work with thermodynamics and statistical mechanics ...

Yes, and who cares if a molecule or two drops through the cracks? No problem here.

>When will your modeling and associated methodolgy be sufficiently
>sophisticated to predict when I will COOPERATE and when I won't?

Probably never. Predicting behavior is not what PCT is about. PCT is about understanding behavior, and what it is being used to control. If I ask you to do a tracking task and you do something else, it will be clear that your goal is not cooperation, but something else. I might be able to find out what the something else is by interacting with you long enough. I'll never find out by studying "people like you" who show "behavior like yours" in "circumstances like these."

>And when PCT has leapt that hurdle, the next one is to make
>models for THOUSANDS of individuals and combine them some way
>to predict population measures. Lots of luck.

Why would I want to do that? I'm not trying to make a living by selling insurance or proving that I am -- on the average-- a successful doctor or educator or politician. My interest is in understanding the next individual I meet, by some means that doesn't involve formalized prejudice.

>No, it is better to follow the example of physics and stick
>with descriptive statistics for generating SOME kinds of knowledge.

Better for everyone? Is this one of those "more important human goals?"

>But I am not arrogant enough to think that modeling individuals
>is the only path to knowledge. It IS the only path to SOME
>kinds of knowledge. But some people don't need that kind of knowledge.

Not for their professions, I agree. What about for getting along with their families and friends and the salesperson and the waitperson and themselves? It seems to me that by relying on statistical generalization in such person-to-person circumstances, people create more problems for each other than they solve.

I think what you're saying basically makes sense: use statistics for appropriate purposes, and models for other appropriate purposes. Partly this is just a practical matter of what we currently know how to do. Weather modeling doesn't work very well, and perhaps can't, so weathermen also use statistical data. Even when a model would in principle be the best tool, if you don't have a model developed well enough to use you fall back to relying on generalizing from experience just as people have always done.

It's really a question of how well you need to understand things. In principle the statistical approach, given huge populations and easily-measured mass effects (as in thermodynamics) can yield very precise predictions for the population, with small error spreads. The laws relating gas pressure, volume, and temperature do extremely well in predicting the behavior of large enough packets of gas. They say nothing about the behavior of molecules, of course. But the precision of the mass measures is quite good enough to serve almost any practical purpose.

This is not true of statistical measures of human populations. I don't know what the smallest packet of people is that would provide an accurate characterization of a population, but it must be at least as large as largest studies so far done. You probably need hundreds of thousands, perhaps millions, of people as subjects, and even then, human circumstances vary so much that in a population of millions you may still have subpopulations that are too small to provide accurate data.

What people really want is a way to predict the behavior of others within the very small subpopulations of people with whom they are likely to interact. With how many people will you have an important face-to-face interaction during your lifetime? A few hundred? A few thousand? That, for you, is the relevant subpopulation. It is far too small to characterize reliably by any statistical means. Whatever generalizations might be made about the behavior of that number of people, you are likely to find more negative than positive instances of it. To do better than that, you need a model that says true things about every single person you could possibly meet, or pretty close to it.

The PCT model is that kind of model. At present, it can say such universally true things only in simple circumstances like the rubber-band experiment. Even then, the kind of knowledge it gives you isn't the sort that psychologists look for: what specific actions people will perform under certain circumstances. The knowledge that PCT gives is about relationships and processes, not events. It's more like relaizing that the other person is controlling the knot connecting the rubber bands. Knowing that, you then know that the other person will do anything necessary with that end of the rubber bands, in order to keep the knot where that person wants it. When you have figured out where the person wants the knot, or how the person wants the knot to behave, you can then do some very accurate predicting over a limited time span, and you can predict the result of doing things to your end of the rubber band that have never occurred before. This has nothing to do with what any other person might do: you're talking strictly about this person.

You might say that where we can't use the PCT model, why not use the statistical approach, because it's all we have? I've heard this a lot: what else can we do? My answer is always the same: nothing. You just do what you do, and get the results you get, as has always been the case. This doesn't make the results any better than they have ever been. Where you can't figure out how to use the PCT model, you're stuck with life as it was before PCT. That's doesn't mean you have to be happy with it.

Best, Bill

Date: Fri Mar 26, 1993 1:11 pm PST
Subject: Aaaargance

From Greg Williams (930326 - 2) Rick Marken (930326.0800)

>I was assuming (and Bill was too) that Ken Hacker shares many of
>the same goals of PCTers -- such as, the goal of gaining some
>knowledge about the processes underlying the phenomenon of human
>behavior, both individual and collective.

I believe that assumption is correct (maybe we could ask Ken -- but, uh, that's another thread!), but my concern was Ken's explicit complaint that some other scientists' aims (and hence their methods) were being deprecated by some "arrogant" folks on this net. I could go back and get the quote if you want.

>I don't see where Bill, Mary or I (we've been the only participants
>in this discussion) have supposed that our goals are the only or the
>most important ones. What gave you that impression?

I got the impression from such comments as (paraphrasing here) non-PCT behavioral scientists and/or pseudo-PCTers don't know what a real model looks like. This sort of line, TO ME, implies that "real models" are where it's at, and anything else is, well, to paraphrase again, not actual knowledge. I guess I tend to read between the lines some, and what I get is the "flavor" that modeling individuals is much more important for Bill, Mary, and you than other methods of science (especially statistical). Now, that's fine, as long as it is more important because it is the appropriate (i.e, efficient) technique for reaching (or attempting to reach) YOUR OWN goals. The "arrogance" which I believe Ken was concerned about is in hinting (if not explicitly saying) that OTHERS' goals SHOULD BE less important FOR THEM than goals similar to yours. It is not arrogance to point out that certain techniques will fail or work poorly for achieving someone's aims; it is arrogance to claim that those aims are not as "important" as your own. And the latter is, I think, what Ken was perceiving.

>By the way, I think its fine for people studying population
>level phenomena (insurance companies, polling agencies, etc) to
>use sampling statistics. Their goals are different from mine (in
>PCT) and that's fine with me. I have no interest in trying to
>derive actuarial phenomena from individual behavioral models.

Then you aren't arrogant, even though you might sound that way to some netters. (:>)

>Bill Powers (930326.0830)

>Oh, heck. Well, is it arrogance to suppose that one's own
>definition of arrogance is the only, or even the most important
>one? What is the most important goal, then? And where can I look
>out there to find it?

PCT itself claims a radical individualistic stance for the notion of "importance." We each decide what views, goals, and definitions are most important TO US. Others certainly might have differences of opinion about the importance of any of these. One manifestation FOR ME of arrogance is assuming that others should ascribe the SAME importance to one of these as you do. As long as you say something like "if you want to do so-and-so, I think you'll find that a good way to do it is via thus-and-thus," then you are in the nonarrogant realm. When you say something like "this is what you should want to do, and doing something else isn't as good," you are simply saying (arrogantly) that your own goals are better (i.e., "more important" or "real knowledge") than the goals somebody else prefers.

>Seems to me that I've commented a number of times on the fact
>that people who deal with populations do fine (for themselves)
>with statistics, because statistics concerns the properties of
>whole populations, not individuals. When that's your concern, who
>cares if a few of the ants get stepped on?

Are you suggesting that the concerns of people who deal with populations are in some sense unseemly, because "a few" individuals generally get stepped on? Are you suggesting that these people could use PCT to do a better (to you) job in some way? If so, how? I said in my last post that it appeared extremely unwieldy to develop population measures from individual models. Perhaps your preference is for the population dealers to simply quit dealing?

>>I would suggest that there are some human goals which can be
>>met more efficiently by using descriptive statistics than by
>>using the (again, "Grand Method") of individual models.

>Whose goals are you talking about?

Insurers, epidemiologists, some government workers, statistical physicists, many agricultural researchers and engineers, etc.

>If you own an insurance company, it's most efficient to use
>statistics to set the rates, to raise the rates on people who make
>claims, and to reject applicants who belong to high-risk populations.
>That's the most efficient way to satisfy your personal human goal of
>making money out of selling insurance to a population.

Personal human goals are the only ones we've got. You seem to be implying sleaziness or selfishness again here. Do you really know that much about the people who use population statistics to say that what they (some? many?) are doing is "bad" ethically?

>>When will your modeling and associated methodolgy be
>>sufficiently sophisticated to predict when I will COOPERATE and when I won't?

>Probably never. Predicting behavior is not what PCT is about. PCT
>is about understanding behavior, and what it is being used to control.

OK, the challenge becomes: show me you understand my behavior and what it is being used to control. But first, YOU tell ME what a passing grade at such "understanding" would consist in. You speak of highly accurate PREDICTIONS in tracking experiments; why not also in more complex experiments? It sounds to me as if you are going back on the promise of Glory Days (sometime in the future) when PCT experimenters can predict .99+ correlations in more than tracking trials. At any rate, I'd like to hear more about "understanding" as opposed to predicting.

>If I ask you to do a tracking task and you do something
>else, it will be clear that your goal is not cooperation, but
>something else. I might be able to find out what the something
>else is by interacting with you long enough.

I'm ready to try the experiment when you are.

>>And when PCT has leapt that hurdle, the next one is to make
>>models for THOUSANDS of individuals and combine them some way
>>to predict population measures. Lots of luck.

>Why would I want to do that? I'm not trying to make a living by
>selling insurance or proving that I am -- on the average-- a
>successful doctor or educator or politician. My interest is in
>understanding the next individual I meet, by some means that
>doesn't involve formalized prejudice.

YOU don't want to do it, but MANY OTHERS would, IF PCT could provide more reliable population measures. The big question is: can it? When? If you don't think it should (or can) be used that way, then how is PCT going to help these folks? Perhaps you think it SHOULDN'T help them, because they are misguided?

>>No, it is better to follow the example of physics and stick
>>with descriptive statistics for generating SOME kinds of knowledge.

>Better for everyone?

Obviously not. That's why I said "SOME"! You can ask individuals what they want to do. (And you can tell them that they shouldn't want to do THAT, but that's not likely to get you very far. You'll probably just be labeled "arrogant.")

>>But I am not arrogant enough to think that modeling individuals
>>is the only path to knowledge. It IS the only path to SOME
>>kinds of knowledge. But some people don't need that kind of knowledge.

>Not for their professions, I agree. What about for getting along
>with their families and friends and the salesperson and the
>waitperson and themselves? It seems to me that by relying on
>statistical generalization in such person-to-person
>circumstances, people create more problems for each other than
>they solve.

That is quite possible. But it remains to be seen whether or not the individual modeling method can do any better. This is because it has not been attempted to the degree where a verdict can be given. I'm game -- have been for several years -- to try it out, but if I were to prejudge the outcome and say that (sometime in the future) PCT will allow solution of interpersonal problems which are now exacerbated by statistical generalization (yes, even the fallacy of applying population measures to individuals, since "good" results can sometimes come from mistakes!), I would be arrogantly predicting the future to match my own preferences.

>I think what you're saying basically makes sense: use statistics
>for appropriate purposes, and models for other appropriate
>purposes. Partly this is just a practical matter of what we
>currently know how to do. Weather modeling doesn't work very
>well, and perhaps can't, so weathermen also use statistical data.
>Even when a model would in principle be the best tool, if you
>don't have a model developed well enough to use you fall back to
>relying on generalizing from experience just as people have
>always done.

Yes. And I'm also saying that it is still to be determined whether PCT models will ever be efficient for understanding some population effects.

>What people really want is a way to predict the behavior of others within the
>very small subpopulations of people with whom they are likely to interact.

WHICH people want this? ALL people? I think YOU want it, and are extrapolating to what you think others SHOULD want. I don't think all others "really want" this. Where's your evidence?

>The PCT model is that kind of model. At present, it can say such universally
>true things only in simple circumstances like the rubber-band experiment.

Where's your evidence that the PCT approach is going to be sufficiently generalizable? That's the kind of evidence people want in order to not characterize a statement like "Here's the way you should be doing it" as arrogant.

>You might say that where we can't use the PCT model, why not use
>the statistical approach, because it's all we have?

The problem is convincing folks that the PCT model will do a better job FOR THEM.

As ever,   Greg

Date: Fri Mar 26, 1993 1:33 pm PST
Subject: Linguistic Behaviorism

[From Rick Marken (930326.1000)]

When I say "you can't tell what a person is doing just by looking at what they are doing" I don't mean to imply that we (as observers) are always confused about what people are doing (what variables they are controlling). I just mean that, IN PRINCIPLE, because people control perceptions, not output, you cannot be SURE what perceptions a person is controlling (what they are "doing" in the PCT sense) by simply observing their outputs (what they are doing in a behavioristic sense). I'll call this the "PCT perspective on behavior". All this perspective says is: in order to be SURE of what a person is doing (PCT sense), you must do some version of the Test for the Controlled Variable. Such a Test must be a sine qua non of any scientific approach to understanding behavior; I don't see how one can successfully model behavior with- out knowing what is being a modelled. That is why the "PCT perspective on behavior" is important; it is important if one wants to achieve a systematic, scientific understanding of behavior .But in our everyday interactions with each other we rarely take a SYSTEMATIC approach understanding behavior. Our approach to determining what people are doing (what they are controlling) is rather informal; we either take outputs at face value or we apply the Test unsystematically (by poking and proding -- applying disturbaces --as we do on the net) to come up with some guess about what is being controlled (what a person is trying to accomplish with their outputs). I imagine that we are usually pretty close to being right about what other people are doing (PCT sense) just by looking at what they are doing (behaviorist sense).

Martin Taylor and Bruce Nevin seem to think that linguistic behavior is somehow exempt from what I call the "PCT perspective on behavior". I'll dub their perspective "linguistic behaviorism" because it assigns to language alone the status that is assigned to all behavior in behaviorism. In behaviorism, behavior is whatever an organism does (its observable outputs). Of course, in PCT, behavior is only what an organisms intends to perceive.

Bruce Nevin (930326) is quite clear about his belief in linguistic behaviorism (as I've defined it).

I said --

>So let me see if I get this right: you CAN tell what someone is
>doing by watching what they are doing if what they are doing is
>uttering a word, phrase or sentence?

and Bruce replies:

>Yes. Specifically, you can tell that they are uttering that
>word, that phrase, that sentence.

Boy, it's sure refreshing to a get nice, straight-forward answer like that for a change.

Martin Taylor (930326 11:15) says --

>So, if one is controlling well, your communicative partner CAN tell what
>you are doing in the specific case of language (construed to include
>body language, etc.). The conventionalization of language has evolved
>just so that this can happen.

Which is also a nice, straight-forward answer.

Let me try to nail this down a bit more. Are you guys saying that, if I hear someone say "Hi there Rick", then I know for sure that this sound pattern was an intended result of the speaker's outputs (efferent neural impulses)? However, if I see someone stop at a stop light I cannot be sure that this visual pattern is the intended result of the actor's outputs? Martin implies that the observer's ability to tell what the speaker is doing depends on how well the speaker is controlling the controlled variable. Let's say that one variable being controlled is, indeed, the sound configuration that we recognize as "Hi there Rick". How do I know that that sound pattern is well controlled just by hearing it? What if the speaker meant to say "Hi there Prick" and said a perfect version of "Hi there Rick" by mistake. How can I tell from just what I hear (a perfect rendition of "Hi there Rick") that what I am hearing is NOT a perfectlly controlled (intended) version of "Hi there Rick"?

Bruce says:

>And there is ambiguity. A given utterance can have more than one
>structure concurrently, e.g. Harris' example (used later by
>Chomsky) "Flying planes can be dangerous." If you are willing and
>able to display the ambiguous possibilities to yourself in
>imagination, you can tell that the other person means at least
>one of a set of category perceptions, but you can't tell which one.

Well, if I can't tell which meaning the speaker meant (intended) to communicate, doesn't that satisfy my claim that "you can't tell what a person is doing by looking at what they are doing"? Assuming that the speaker intends to produce one or the two possible meanings (he may mean to produce neither) all I can tell from the output (the sentence "Flying planes can be dangerous") is that one of the two meanings was intended. So I can't tell (meaning, know for sure) what the person intended. Your arguments about ambiguity seems like the points I should be making. Remember, you're the one arguing that you CAN tell what a person is doing by just looking -- when what they are doing is language.

>But usually you are unaware of possible ambiguity. Usually you
>have imagined a version of meanings and motivations consistent
>with your perceptions of the subject matter domain and of the
>social relations between you, and don't notice that the utterance
>is structurally ambiguous (that the linguistic information in it
>might be different from what you had perceived). You don't try
>out the alternative branches of the ambiguity.

Yes! So you're sitting there, certain that your buddy intended to say something about the hazards of being a pilot when, in fact, s/he was actually making a joke -- planes aren't really dangerous unless you start flying around in them.

>Too much? Are you still with me? And can you tell me about your
>experience learning Spanish?

Just right. I'm still listening. El experiencio estere muy dificile. Yo creo que yo nunca appredendre lo muy bueno. Pero, yo puedo comprar si que yo quieres en Mexico.

Best Ricardo

Date: Fri Mar 26, 1993 1:39 pm PST
Subject: Subjective probability intro

[Martin Taylor 930326 16:20]

The following was drafted some months ago as an introduction to the "Information leads to PCT" paper I am trying to draft. But that has been stalled by higher priority (i.e. paid for) work, so I thought it would be a reasonable idea to post this part. There has been some discussion involving probability, and will be more. In this, I try to describe my approach to probability, and why I think something like it must be necessary. This document, or something like it, will be incorporated into the final paper.

Martin

Information, Perception and Control

M. M. Taylor, DCIEM, Box 2000, North York, Ontario,
Canada, M3M 3B9

(Epigraph)
The only justification for our concepts and system of concepts is that they serve to represent the complex of our experiences; beyond this they have no legitimacy. I am convinced that the philosophers have had a harmful effect upon the progress of scientific thinking in removing certain fundamental concepts from the domain of empiricism, where they are under our control, to the intangible heights of the a priori. For even if it should appear that the universe of ideas cannot be deduced from experience by logical means, but is, in a sense, a creation of the human mind, without which no science is possible, nevertheless this universe of ideas is just as little independent of the nature of human experience as clothes are of the form of the human body. (A. Einstein, 1922, The Meaning of Relativity, 4th Edition, 1950, London: Methuen, p2).

Prologue: Probability and Perception

Newton and Einstein

Before the middle of the 17th century, Physics consisted of a great many aphorism and folk truths, and a few numerical descriptions of experimental or observational results. The motions of the planets were well described by Ptolemaic epicycles, and if a description failed in some minor detail, another epicycle could always be added to correct the error.

Then came Newton. Newton imagined a consistent physical world, in which the interactions among the parts could be described by a few simple rules that would apply whether the objects were very large, like the sun, or very small, like a grain of sand. We might not be able to use these rules to predict all the motions of the universe, but our limitation was only our inability to observe all of the elements in the universe and make the necessary calculations. Perhaps there were a few rules that had not yet been discovered, but an omniscient observer who knew all the rules and could see the state of all the interacting elements would be able to determine the fate of the universe for evermore.

Newton's simple laws worked very well, provided that they were applied to elements that were not too big or small, and didn't move too fast. But there were some nagging problems that they could not account for, such as the advance of the perihelion of Mercury, or the shape of the radiation spectrum of a black body. It took some 250 years before a new advance occurred in the way we look at the physical universe; in fact two advances, both eventually based on the same core concept, and in apparent contradiction with each other. Einstein and Heisenberg both pointed out that the universe is intrinsically unknowable to any single observer, and developed the consequences of that truth.

Einstein considered the consequences of the finite speed of signal transmission, and developed the Theory of Relativity out of the fact that there cannot be an omniscient observer who can see all of the Universe at the same moment. Things that happen in one place cannot have any effect on things that happen at another until signals have traversed the intervening space, and this takes a finite time. At point A, an event EA happens, and it is observed at point B some time after an event EB happens at point B. To an observer at B, the sequence is EB then EA. But it can happen that an observer at A does not get the signal from EB until after EA has happened. For the observer at A, the sequence is EA then EB. The two observers, who may later exchange communications, might disagree about which event came "first." If the two observers are stationary relative to each other, they can resolve the disagreement by factoring in the speed of the light signal, and they will then agree on which event "really" came first. The reason this agreement is possible is that the two observers share a common "frame of reference." They can get away with pretending that their own frame of reference is the real, absolute point of view--the point of view that would be held by a God-like observer.

But what if the observers are in significant motion relative to each other? A might think that EB preceded EA, while B might equally legitimately claim that EA came first, even when the speed of light effect is factored out. Neither would be wrong. Two observers in relative motion can no longer settle the argument. They will disagree about the temporal ordering of events, just as we on Earth can legitimately disagree about spatial ordering. Two people standing at different locations on the shore might disagree on whether one ship in a harbour is to the left or right of another ship, since they observe the ships from different angles. Since there is no absolute spatial perspective from which to view the ships, there is no absolute answer as to which is to the right or left of the other. Likewise, since A and B view the timing of events from different angles in space-time, they will disagree on which came first. In such a universe, there can be no simultaneity, except as defined by some particular observer. There is no one true frame of reference, no absolute God-like point of view.

Einstein's way of solving this puzzle was to consider as legitimate only those aspects of the Universe that were potentially accessible to any single all-powerful observer limited only by the signal velocity. The result was a unifying view of the interactions in the world that provided the same numerical results as Newton's laws for middle-scale slow-moving objects and that corrected the errors of Newton's predictions for interactions involving very large or very fast objects. The central core of the theory was that even if there were a God-like observer not limited by the speed of signal transmission, no real observer could use those observations, and real observers would see a world different from the one the God-like observer would see. We speed-limited people would see a relativistic world.

Einstein's world is simpler than Newton's. It contains fewer arbitrary laws. The really essential law is the one that limits the access of an observer to information from distant places. All the rest follows with few essential added assumptions.

Heisenberg also developed a theory that we now see as based on the idea that there is a limit on the amount of information we can acquire from the world. In his case, the limitations come when the interacting elements of the world are very small. If we obtain a very accurate measure of some parameter of an element, we cannot obtain an accurate measure of a dual parameter of that element at the same time. It is not possible, for example, to determine simultaneously the position and the momentum of a particle. Time itself is a member of such a pair of dual parameters; we cannot determine both the time and the frequency of an oscillatory event.

There is no need to dwell on the theories of Einstein and Heisenberg. Their detail is irrelevant to the theme of this paper. The moral to be drawn is that the great advances in science and technology of the 20th century are based on consideration of one simple truism, that we can work with only what we can observe, not with what God might observe. And that is the theme of this paper. There are two ways of looking at the world: a fantasy way, in which we imagine we are God, or a realistic way, in which we imagine what we might be able to observe. The former is Newtonian, the latter Einsteinian.

Probability: Frequentist or Subjective?

What does the phrase "the probability of event E" mean? Most people are taught that if there are an infinite number of opportunities for E to happen, and E does happen on a certain fraction of them, then that fraction is the probability of the event. If a coin is tossed an infinite number of times, it should fall heads half the time, so the probability of a head is 0.5. It does not take long to realize that such a definition cannot be used to measure the probability of event E, so this "ideal" definition is taken as a target for more practical measurement techniquesQ"an infinite number" is replaced by "many," for example, or a mechanism is described that would inherently result in E occurring on some specified fraction of the opportunities. An ideal coin would fall heads exactly half the time in an infinite number of tosses, because that is the definition of an ideal coin. But it doesn't help in determining the probability that the next toss of this coin will be a head.

Even to ask the question "What is the probability that the next toss will result in a head" is to deny the validity of the definition of probability most people are taught. What does one toss have to do with infinity? It comes out heads or it does not. There must be some other way of looking at the notion of probability, because we certainly feel that there is some value in thinking about whether one event is more probable than another. Is it more probable that the next time we see Joe he will be wearing brown shoes than that it will be raining in Toronto at noon on July 12, 1997? One feels that there is some sense in asking such questions, even though neither event has more than one opportunity of occurring.

The notion that probability has to do with the fraction of opportunities for an event on which it actually happens can be called "frequentist." It is a Newtonian view of a world in which situations are indefinitely repeatable, observations can be carried on for infinite time, and can be infinitely precise. The Newtonian world is not a world accessible to ordinary mortals, and thinking about a frequentist probability that can occur only in a Newtonian world can lead one into great confusion and paradox.

A real observer can only go on what is observable. The real, observable world, we can call Einsteinian. In this world, probability depends only on what has been observed. If Joe's friend has just telephoned to say that Joe bought new brown shoes and is coming round to show them off, we might say that there is a high probability that the next time we see Joe he will be wearing brown shoes, even though we have never before seen him wear brown shoes. Observations include all sorts of things related to the uncertain event in question, not simply observations of the critical factor as it occurred or did not occur on past occasions we deem to have been similar. If we believe that a coin has been made with fair balance and evenly milled edges, we will judge it to have a probability 0.5 of landing heads, even though we have never tossed it before. But if we believe that the coin tosser has a special skill in how high and with what spin to toss the coin, we may alter that judgment depending on our belief as to which result the tosser wants to see.

Probability is a subjective matter. The only consistent and reliable way to deal with probability is to treat it as a property of an observer, not of the world. There may be probabilities in the world, but real observers can no more detect them than they can assert a correct, universal sequence of events in the world.

To be subjective does not mean to be arbitrary. If one believes that two events are mutually exclusive and that one of the two has to happen, then to be consistent, the subjective probabilities of the two events must sum to unity. One cannot arbitrarily say that p(A) = 0.8 and p(B) = 0.8 and p(A xor B) = 1. Subjective probability has constraints, if it is to be dignified with the name of probability rather than wantedness or hope, or something such word. For example, the subjective probabilities of all mutually exclusive events that could happen in particular circumstances must sum to unity. If one does not believe that A can happen unless B does, then the subjective probability of A cannot be greater than that of B. All of the usual arithmetic applied to frequentist probability is appropriate in dealing with subjective probability.

One of the constraints is that all probability is conditional. It makes no sense simply to say of an event that its probability is P. As with the frequentist kind of probability, one must think of what might be the alternatives to the eventQwhat counts as the event not happening, what occasions might be considered as opportunities for the event to happen, on which the probability is based. Would Joe appearing barefoot count against the event "Joe wearing brown shoes the next time," or would he have to be wearing some non-brown shoes? We notate a conditional probability with a vertical bar: P(A|B) is the probability that A will be true, given that B is true. B represents the occasion, A the event. B, like A, must be something observable. The observer must be able to determine whether B is true, whether this is an occasion on which it is interesting to observe whether A is true. We may, for example, be interested in whether Joe is wearing brown shoes only on condition that he is wearing some kind of shoes. Or, conversely, Joe may usually wear no shoes, but when he does, they are usually brown, so we are interested in whether he is wearing brown shoes given that we see him, shod or no.

How one defines the condition has a great impact on the subjective probability of an event. Given that one has observed a coin thjat has nearly come to rest after being tossed, one can put a high value on the probability that it will lie "head" (or not, as the case may be), when it finally does stop. The probability will be even higher if the condition is added that there be no earthquake or other disturbance of the floor before the final observation, and higher yet if another condition is added that no person disturb the coin's movement. This example may seem extreme, but such conditions apply to every judgment of probability, and in many cases they are both less obvious and more important than these.

Probability depends on knowledge. If one had not observed the coin since determining that it had a head side and a tail side just before it was tossed, one would probably judge that it had a probability of 0.5 of landing head or tail, regardless of the earthquake or interference conditions. After all, we model those events as being equally likely to end up with it lying either way, and that is the same as our judgment of the result of the undisturbed toss, so those conditions have no effect of the subjective probability judgment. But notice that it is our model that allows us to make that assessment. If we had knowledge that a certain unscrupulous person with skill in magical tricks had a vested interest in seeing the coin land heads, we might change our subjective probability that it would do so. A condition that we did not observe interference would not be sufficient to bring the probability back to 0.5. In our Einsteinian universe, we cannot impose a condition that no interference occur, only that we do not observe it to occur.

In everything that follows, whenever the notation P(x) occurs, it must always be remembered that in the background there is a condition Y, so that the correct notation would have been P(x|Y). The omission of Y may be justified on the grounds that Y is obvious, but sometimes it is not so obvious. The notation P(x) can sometimes be seriously misleading, if the background conditions are not intuitively apparent.

Probability and replication

Whatever the condition Y, an observer can assess P(x|Y) for any event x. Of course, for most cases of x and Y the observer will have no reason to do so. Y may be very unlikely to happen, given the current condition C, or the observer may deem Y to be irrelevant to x, given C (i.e. P(x|Y,C) = P(x|C)). But in other cases, the observer has some reason for being interested in P(x|Y). Let us suppose that the observer knows nothing about the relation between x and Y, other than that it might be interesting. To find out something of the relation, the observer has to see what happens when Y occurs.

When Y occurs, there is an opportunity for x to occur. If it does, the observer can make a note "Yes," and if not, a note "No." After N occurrences of Y (replications of an experimental observation), there will be a certain number X of "Yes" notations. Since we hypothesized that the observer knew nothing beforehand about the relationship of x and Y, what the observer now knows is that x occurred on X/N of the times that Y was true. A rational observer would use that knowledge to set a value for $P(x|Y)$ close to X/N. If X/N differed appreciably from $P(x|C)$ averaged over the various current conditions C prevailing when Y happened to be true in addition to C, then the observer is likely to say that Y probably affects the probability of x. The observer can also assign a value to P (Y affects x, given C), which we can notate as $P(M|C)$, where M is the model (Y affects x) . If X/N was not very different from $P(x|C)$, then $P(M|C)$ would be low; the observer would not believe very strongly that Y affects x.

X/N is a proportion, not a probability. If in future more occasions occur in which condition Y is true, X and N will change. A frequentist view of probability asserts that the "true" probability of X given Y is the value of X/N that would be observed after an infinite number of occurrences of Y (replications). A practical approach to the frequentist view is less dramatic: X/N will approach an unknown but true ideal limit ever more closely as the number of occurrences of Y increases. There exist theories, which we have no reason to dispute, about the probability that X/N would take on any particular value after N occurences of Y, if the ideal has a certain value Z. A rational observer would be likely to incorporate these theories into the subjective judgment of $P(M|C)$, by, for example, comparing the probability $P(M1|C)$, where M1 is that X "Yes" events would be observed in N opportunities if Y had no effect, with $P(M2|C)$, where M2 is that $P(x|Y,C) = X/N$.

There is a hidden assumption in the foregoing: that condition Y can be repeated (that replications are possible). What does it mean to say that a condition recurs? If everything in the observable universe is taken into account, no condition can ever recur. If nothing else, the universe has aged since the first occurrence of the condition. Its original age can never be recovered. But the observer may well think that this doesn't matter, given that the difference in age is probably only a few parts in 10^10. Such an observer will accept that two occurrences in which Y is the same except for a small increase in the age of the universe can be considered together in evaluating $P(x|Y)$. But the age of the universe is not the only difference in conditions between the two occasions. The position of the sun, moon and planets in the sky will have changed as well. If their arrangement ever recurs, it is at intervals long enough that the positions of the stars will have changed significantly. But if our observer is not astrologically inclined, this difference may not matter either. So the observer may say that the condition Y recurs even though the age of the universe is different, and the planets have moved in the sky. That is strictly a personal judgment by the observer who is trying to get data that will allow a reasonable value to be developed for the subjective $P(x|Y)$, for subjective probabilities are greatly affected by observation. The value determined with the aid of observation, of course, applies to some future occasion of condition Y, not to any that has already occurred, because for those occurrences of condition Y the observer knows whether the answer was "Yes" or "No."

Usually, what the observer considers to be an occurrence of condition Y is specified not by listing all the irrelevant conditions, but by listing some of the conditions that would change the situation into something other than Y if they were varied. All events that happen when the condition does not include those falsifying conditions ought to be counted as opportunities for x to occurQas replications. Of course, in practice, they are not. The observer notices that when (Y but not Z) x tends to occur less often than when (Y including Z), and so changes the definition of Y so that it necessarily includes Z and a failure for Z to be true makes the situation not suited to an observation relevant to determining $P(x|Y)$.

The important point about this discussion is that what constitutes a replication of an observation is a matter totally subjective to the observer. For example, in an experiment in psychophysics, one observer may say that the presentation to a trained listener of a particular waveform to which is added noise from a well-controlled noise generator constitutes a replication of an observation to determine how well people can detect that signal in that noise. Another observer may note that the detection probability depends on whether the trial is presented early or late in a series of trials, and that therefore early trials should be considered separately from late ones. Yet another observer may note that it depends on whether the observer has just eaten lunch, or on the particular waveform obtained from the noise generator, and so on and so forth. It is never possible to specify objectively what constitutes a replication of an observation.

So far, I have tried to make the point that the term "probability" refers only to an attribute of an observer, who should apply it only to unique events. But the observer must have some reason to evaluate the probability of the event, and this reason may come from anywhere. It could be based on being told something by a trusted (or untrustworthy) source, it could come from a belief that some known mechanism causes the event to occur or not to occur, or it could come from observations of whether a similar event occurred under apparently similar circumstances on one or more other occasions. It is up to the observer to determine how to relate the various sources of information so as to adjust the subjective probability of the event. And most importantly, it is up to the observer to determine what constitutes a replication of an observation in case the occurrence of similar events under similar circumstances is a contributor to that probability estimate.

Probability and measurement

What is the probability that the width of this page is 8.5 inches (assuming you are reading this on North American letter-size paper)? What is the probability that the width is 0.4 inches? Technically, the answer to each question is "zero." The page may be roughly 8.5000001 or 8.4999999 inches wide, but that is not 8.5. Nevertheless, one feels that the probability of it being 8.5 inches should be higher than of being 0.4 inches. We "know" that the width is not 0.4 inches, but we do not "know" in the same way that it is not 8.5 inches. Indeed, nominally, the width is 8.5 inches, so if someone asks "Are you reading something on 8.5 inch paper or 6 inch paper?" we could confidently answer "8.5 inch," and would not have to say "neither" as we would if the question were "Are you reading something on 1 inch paper or 6 inch paper." So, there is a range of width over which we are satisfied that this is 8.5 inch paper, more or less.

How well must we measure, to agree that this paper is 8.5 inches wide? That depends on what else we "know." If we are accustomed to North American standard sizes, we need only see well enough to say "this is letter paper, not some other size," because paper does not come in other widths near 8.5 inches. But if the paper might have come from some place more in tune with world standards, then it might be A4, and thus a little taller and narrower than 8.5 x 11. What is the probability that the paper might be A4? Is it enough to require us to measure the paper more precisely than by a quick glance?

A quick glance suffices to assure us that the paper is 8.5 x 11 rather than, say, 11 x 14 inches. But a quick glance is insufficient to assure us that the paper is not A4. A closer look to get more information about the page is needed for that judgment. But if the question is whether the paper-cutting apparatus that was supposed to cut the sheet to be 8.5 x 11 is properly adjusted, no amount of looking will be enough. The paper must be carefully compared with some standard, such as a rod that is asserted to be 8.5 inches long. Ignoring for the moment the trust that has to be placed in the assertion about the rod length, what is the next step in the determination? One looks to see whether the rod is longer than the paper is wide. Maybe there is a clear difference, and one can say something like "On the condition that I am not drunk or hallucinating or something like that, I am quite sure this paper is less than 8.5 inches wide." But maybe the difference is very small, and the best one can say is "I'd bet 3 to 1 that this paper is less than 8.5 inches wide," or even "I have no idea whether this paper is more or less than 8.5 inches wide."

What is the next step? Get a microscope? That may answer the question with respect to a specific piece of paper for one observer, but for another piece or a different observer, it may not. The same set of possibilities will exist. There will always be some width of paper for which you cannot say whether it is greater or less than the specified width, no matter how precise the measuring instrument. Your measurement is always a probability distribution. Using the naked eye and the measuring rod, you may say that the probability is essentially zero that the paper is less than 8.48 inches wide (on the condition that you are not hallucinating) or that it is wider than 8.52 inches. Perhaps you would bet even odds that it is between 8.495 and 8.505 inches. But this is a lot better than you could have done before you got the measuring rod. Without it, you might have given even odds that the paper was between 8 and 9 inches wide. The measuring rod has reduced your uncertainty 100-fold, and if you used a microscope you might be able to reduce it another 100- or 1000-fold. You cannot reduce it to zero, no matter what measuring instrument you use. All you can do is to get more and more information by closer observation.

Any number that is the result of a measurement is a discrete approximations to the "real" value of the thing being measured. Indeed, all numbers that we can use in any way at all are either rational with a finite number of digits in their fractional numerator and denominator, or are describable by an algorithm of finite length. Such numbers are useful, since they are the only means we have of representing values in the world, one must always be aware that they are not reality. Any such number represents only a distribution of subjective probability about reality.

(More to come later -- Much later).

Date: Fri Mar 26, 1993 2:06 pm PST
Subject: Arrowgance

From Rick Marken (930326.1330)]

I said:

>I don't see where Bill, Mary or I (we've been the only participants
>in this discussion) have supposed that our goals are the only or the
>most important ones. What gave you that impression?

Greg Williams (930326 - 2) replies --

>I got the impression from such comments as (paraphrasing here) non-PCT
>behavioral scientists and/or pseudo-PCTers don't know what a real
>model looks like.

Yes. Thanks Greg. I can see how that would, indeed, sound like what I would call "arrogance". I know where that sentiment comes from; it is difficult to make a PCT point (always based on the operation of a working model) when the opposition just dismisses it with words (instead of an alternative model). This is particulary exasperating when dealing with reinforcement theorists; since they have no working models it is impossible to convince them that their theories (as stated) don't account for what they observe. But still, I think we could try to make our point about modelling more diplomatically

So I would like to apologize for any contribution that I might have made to this perception of arrogance; while I do think it is important to have working models in order to be able to deal coherently with the behavior of living systems, non-modellers can certainly make important and worthwhile contributions to our efforts by observing phenomena and applying models (I'm thinking of the clinicians here, for example).

>Then you aren't arrogant, even though you might sound that way to some
>netters. (:>)

Oh, Greg, you sweety. You know I'm an asshole in many ways; but I'm not an arrogant asshole.

And I'm still waiting for Ken Hacker to point out the hyperbolic claims of PCT.

Best Rick

Date: Fri Mar 26, 1993 2:08 pm PST
Subject: disturbance

[from Wayne Hershberger 930324]

I just got around to reading my E-mail or I would have had my oar in the disturbance controversy before this. I hope the dialectic will continue until thesis and antithesis begets a constructive synthesis. Imagine what Claude Shannon and Harold Black could have discovered about the nature of information and control by arguing with each other in the way M & M have been doing here on the net-- it reminds me of Wilbur and Orvile Wright's heated arguments about the nature of flight: very productive.

Recognizing the fact that a control system's disturbance can be mirrored in its output without being represented in its input is a matter of the first importance in understanding the nature of closed-loop control. Rick Marken's steadfast defense of this fact, both as fact (no loose canon, this) and as the essence of HPCT has been marvelous.

As Bill Powers has noted from time to time, a control system may be viewed as an analog computer that determines the magnitude of a variable, d, not by sensing it directly, but, rather, by controlling the value of an alternate, sensed variable, p, that is disturbed by variable d whose magnitude is being computed; the system's output, o, is the system's estimate of d. This is the fact that Rick is insisting that we not fudge: control systems compute an estimate of d while sensing only p. It is magical, but true.

However, since this magical fact is natural, not supernatural, it should be possible to explain the "trick." The trick, of course, is negative feedback--feeding the system's output back on itself so that it is self-limiting. That is to say, the control loop's output is at once error driven and error reducing. Yet, saying this, it seems to me that the question Martin is asking remains unanswered. That is, how exactly does the control system compute its estimate of d from p?

I believe the answer is implicit in M & M's observation that the better the control, the less p is attributable to d, with the limit being zero--not >>0 as Ashby supposed. The control system's estimate of d depends upon this limit. That is, the control system's irreducible error, at the limit of control, is attributable exclusively to d, so the estimate of d is a function of this error and the system's gain--which determines the error's limit.

Martin, I believe you overstated the case when you acknowledged ([Martin Tailor 930319 14:30) that:

> With infinite precision perceptual signals and zero
> transport lag around the loop, the perceptual signal is
> always completely under control and the disturbance
> is never represented there.

Only if the gain is infinite and the bandwidth of the disturbance is not too great. Suppose that the system is perfectly stable but the gain is not infinite. Suppose that the forward gain is 990, and the reference value is a constant .11 units; further,

> (Bill Powers 930320.2100) Suppose the disturbing
> variable is a constant 10 units, and the output is a
> constant -9.9 units, both measured in terms of effect
> on the CEO when acting alone. The perception,
> referred to the environment, is 0.1 units.

The irreducible error is .01 units. Since this error is irreducible at the limit of control it is an amount of p that is attributable exclusively to d. Therefore, O = .01 * 990, is a good estimate of d.

(Gary Cziko 930323)

> When I'm driving, I control the acceleration of the car
> and so I seem to use this advance knowledge of
> impending accelerations to minimize my head bobbing.
> This phenomenon is what I'm having some trouble
> understanding as in PCT terms since it appears to be
> a good example of what a "normal" psychologist would
> probably refer to as FEEDFORWARD.

Yes, or classical conditioning, or both, as I did in my chapter " Control theory and learning theory" in Rick's special issue of American Behavioral Scientist. You may think of any such feedforward (or conditional reflex) as an endogenous disturbance added to the output and timed so that it coincides with an anticipated exogenous disturbance thereby mutually canceling each other. You may also think of it as a pulse added to the error signal (i.e., added to the output before it is amplified), in which case the pulse is effectively being added to the reference signal; that is, r - p + pulse = r + pulse - p (this is what Tom Bourbon was describing several months ago on the net). This bumps the matter up a level in the hierarchy where the pulse may be either anticipatory (feedforward; i.e., output added after amplification) or error driven (feedback; i.e., added before amplification). If it is the latter, then the anticipatory pulse is added to the reference signal at level 2--which bumps the question of whether it is ultimately feedforward or feedback up to the 3rd level, etc., etc. Bill Powers insists that ultimately if is feedback. I'm not convinced.

Gary, the paper you sent me to read was incomplete, comprising only the odd-numbered pages. Or was it the even-numbered ones? I forget. I apologize for not telling you this earlier, but the matter is academic because I don't know when I'll be able to get around to it. It seems, these days, that the hurrieder I go, the behinder I get.

Warm regards, Wayne

Date: Fri Mar 26, 1993 4:01 pm PST
Subject: Re: Disturbances that don't disturb

[Allan Randall (930326.1616 EST)]

Bill Powers (930317.1200) on definitions:

```
> >disturbance: the total sum environmental influence on the CEV.
>
> No, and the reason we don't agree, once understood, should clear
> everything up. The word "influence" is ambiguous...
>
>             A --- /////////// --- B --- \\\\\\\\\\ --- C
>
> Now let's get the language straight. What are "the two influences
> on B?" They are A and C, the positions of the two objects. So
> here we are defining an "influence" as "something capable of
> affecting something else." ...
> Influence, however, has another sense. We can ask, "How much
> influence do the positions of A and C have on B?" Now the problem
> becomes clearer. If A moves 1 cm to the left while C
> simultaneously moves one cm to the right, and the springs are
> identical, the answer is NONE. Neither A nor C has, in fact, any
> influence on B _in terms of an effect on B_. The total force
> acting on B will remain zero, and B will not move.
```

Yes, exactly. However, I don't think this is the source of our disagreement, as I think we are basically in agreement on this point (correct me if I'm wrong). By "influence" in my definition I am talking about your first sense of the word, not the latter: the effect of the output from the control system is excluded from what we mean by "disturbance". So A is said to truly have an "influence" on B, since C is considered a separate influence. In other words, influences can cancel out. Perhaps the word "force" would be better? I used the phrase "environmental influence" in the definition so that only influences caused by the external environment are included. If a disturbance is 100% canceled out by the output of the control system, then the disturbing variables have an "influence" (exert a force) on the CEV, but fail to actually affect it, and the perceptual signal, representing as it does the CEV, will carry NO information about the disturbance.

```
>                        **************************
>
> Disturbing variable
>
> *********************
>
> -------------------------------------------------------------
>                        *
>                         *
> CEV                    *
>                           *
> ***********************      *           ************
>                                   *  *
>
> -------------------------------------------------------------
>
>                                   *  *
>                               *          *************
>                            *
> (Opposing) Output        *
>                         *
>                        *
> **********************
```

Note that the change in the disturbing variable is reflected in the CEV and in the output. If I similarly created a disturbance over and over again in a pattern to send a Morse-encoded message containing the complete works of Shakespeare, the message could be read off the CEV, the perceptual signal, or the output. Note that the controlling organism might deny any knowledge of having trasmitted the complete works of Shakespeare. It will try to convince us it was just playing with some rubber bands, or whatever. Although the control system works as a closed loop, it provides a transmission channel from disturbance to perceptual signal which we can tap into, unbeknownst to the control system, in order to send a more traditional open-loop straight-through message containing all of Willy's plays.

> It's clear that the behavior of the CEV is not like the behavior
> of the disturbing variable.
> ...the CEV itself doesn't reflect the state of
> the disturbing variable. Knowing that the CEV has a value of zero
> tells you nothing about the value of the disturbing variable.

The compressed binary file of a computer image does not look much like the original image either: it appears to be a largely random collection of numbers. But I doubt anyone would claim it had no information about the original image. The value of bit 1254 of the compressed file may not correspond in a simple way to any particular pixel value. The point is that the image can be reconstructed using this compressed file, so the file must be said to contain information about the image.

> This is where we have been sliding past each other. It makes no
> sense to say that the control system's perceptual signal contains
> no information about the CEV. That is why it has seemed so self-
> evidently true to you that a disturbance (meaning an actual
> change in the CEV) conveys information to the control system --
> and so stupid of us to claim that it does not.

I really don't think this has been the problem. I'm not just saying that the perceptual signal has information about the CEV. I really *am* saying there is information about the disturbance itself - as a variable separate from the output of the ECS.

In any case, are we at least settled on the definition of disturbance?

disturbance: the total sum environmental force on the CEV.

This is as opposed to "the total sum effect of the environmental forces on the CEV," which is not what we mean.

Anyway, I think we are getting somewhere. If we can even just agree on the terms of this debate, we will have accomplished something.

Allan Randall

Date: Sat Mar 27, 1993 12:16 pm PST
Subject: disturbance:0.00001, information, 0

[From Rick Marken (930327.0900)]

Wayne Hershberger (930324) made some excellent points relevant to the "information in perception" discussion. In particular, the following comment made me realize that I was taking one of Martin's claims for granted -- when, in fact, it is not true. Wayne says:

> Martin, I believe you overstated the case when you
>acknowledged ([Martin Tailor 930319 14:30) that:

> With infinite precision perceptual signals and zero
> transport lag around the loop, the perceptual signal is
> always completely under control and the disturbance
> is never represented there.

Martin's claim here turns out to be FALSE, a fact that can easily be determined from the equations for a closed negative feedback control loop. In particular, I point you to the closed loop solution for p which is given as equation 6 on page 276 of "Behavior: The control of perception".

6) p = (k.e k.o r + k.d d)/(1+k.e k.o)

where k.e is the environment function relating output to input, k.o is the output or amplification function that transforms error into output and k.d is the disturbance function that transforms variations in a distal environmental variable (the disturbance, d) into effect on the input.

The product k.e k.o is the loop amplification factor or "loop gain"; the greater the loop gain, the "tighter" control. This means that, as k.e k.o approaches infinity, equation 6 approaches p = r. But even when k.o k.e becomes infinity, p never actually equals r and there is never zero contribution of the disturbance. This can be seen if we write out equation 6) substituting infinity for k.e k.o:

6.i) p = (infinity/(1+infinity))r + (k.d/(1+infinity))d
What this shows is that, when the loop gain is infinite, the perceptual signal will (in technical math terms) be DAMN CLOSE to being equal to r, but it will be NOT QUITE equal to r (thanks to that pesky 1 in the denominator of the multiplier of r). The remainder of p (that is not equal to r) is k.d/(1+infinity) of d; which (again in technical math terms) is a little, teensey, weensey (LTW) amount of d, but still, it more than zero.

So even with infinite gain (the best you can get) there is still a LTW amount of d in the perceptual input -- not zero. This means that there is NEVER absolutely NO disturbance contributing to the perception. Yet, there is still NEVER any information about the disturbance present in the perception. The amount of disturbance in p may be .0000000000000000001 -- but the amount of information in p is always, precisely 0.0. This can be seen from equation 6. The perception, p, is ALWAYS equal to a value that is proportional to the reference input, r, and the disturbance, d. There is no information about d in p because the system has NO WAY OF KNOWING which component of p is attrtibutable to r and which to d -- all the system sees is p, always.

Equation 6, by the way, just shows the consequence of the output ALMOST completely cancelling the disturbance. p is still equal to o + d. Equation 6 just shows that when o is approximately equal to -d, p is proportional to r (with a LTW amount of uncancelled d -- what Wayne referred to as the uncancellable error).

Again, the point is that some amount of disturbance is ALWAYS present in the perceptual signal, even when control is perfect (infinite loop gain). Martin says that the perfect control situation is the only one in which it is appropriate to say that there is no information about the distrubance in perception because the disturbance is completely cancelled. But the disturbance is NOT completely cancelled -- even when control is perfect (which Martin correctly says is an impossible condition anyway). My point is that the perfection of control is irrelevant to the question of whether there is information about the disturbance in perception. If you believe that the information in perception is communicated by the effect of the disturbance on the perception, then even when control is perfect this effect DOES NOT DISAPPEAR -- it is (k.d/(1+infinity))d. I am arguing, even though the effect of the disturbance present with perfect control, their is STILL no information about the disturbance in perception because the perception, p, is still the SUM of two influences -- distrubance AND output.

The trick of how output mirrors disturbance even when there is no information about the disturbance in perception has been described over and over again; it is so simple that maybe it just doesn't seem like muck of a trick. Here is it again:

The mirroring of disturbance by output is a SIDE EFFECT of the fact that output is proportional to r-p. The time integral of this difference mirrors the output. There is nothing in p that tells the output "how much it should be" at any instant in order to equal -d at any instant. Whatever output is generated at any instant, it is added to the current d to produce p -- which, because of the negative feedback loop, TENDS to be (but is not always) a bit closer to r than it was an instant before.

On another topic:

As I have mentioned, I have volunteered to review a BBS paper on reinforcement theory. I don't know if I'll really follow through on this -- or whether they'll chose me as a reviewer -- but I have looked over the paper and it is very revealing about the current state of MODERN reinforcment theory. I'm interested in modern reinforcement theory because I submitted a paper to Science several years ago describing the e. coli experiment ("Mind Readings" ch 4, paper 1) as a challenge to the idea that reinforcement shapes behavior by strengthening responses. The paper was rejected (of course) and one of the more common reasons for rejection was that my experiment challenged a "straw man" version of reinforcement theory. Bill P. and I tried another approach to challenging the reinforcment theory establishment -- using modelling. But again the paper was usually rejected (when it was presented as a challenge to reinforcement theory) because we were attacking a "straw man"; the straw man being our claim that reinforcement theory says that reinforcement strenthen's responses (we eventually got that one published by billing it as a model of control with random consequences -- we just mentioned at the end that the results might be interesting to people who think that consequences shape behavior).

Well it's 1993 now and here is what an article targetted for BBS, one of the leading journals in the behavioral sciences, has to say about reinforment theory:

"What does reinforcement strengthen? Responses, to be sure ...

These are the first words of the article. I guess straw men for PCT are state of the art for conventional psychology.

SIGH.   Best   Rick

Date: Sat Mar 27, 1993 12:16 pm PST
Subject: Arrgnce; opinions;challenge II

[From Bill Powers (930327.0700)]

Greg Williams (930326 - 2) --

>I guess I tend to read between the lines some, and what I get
>is the "flavor" that modeling individuals is much more
>important for Bill, Mary, and you than other methods of science
>(especially statistical).

I think that's true; it's true that you tend to read between the lines some, quite often correctly, and that I at least prefer modeling to statistical methods. I don't prefer modeling, however, just because I like modeling, and I don't hesitate to use statistics when there is a use for it just because it's statistics. My reason for preferring modeling is that I like explanations that fit the observations as much of the time as possible, preferably all of the time. This means that when an explanation (of whatever kind) fails to fit any valid observation, I want to know why and I generally attribute the reason to a failure in the explanation, not to some innate barrier to discovery such as an allegedly inherent variability of natural systems. if modeling didn't help to produce explanations that work, I wouldn't like modeling much, either.

This is what leads to the aim of explaining individual behavior. The exceptions to generalizations about behavior always turn out to be some individual acting in a way that doesn't fit the generalization. If one believes, as I do, that we should reserve the term "knowledge" to mean explanations with no known exceptions, then even a single individual case that contradicts an explanation means that the explanation is still inadequate. We can live with inadequate explanations, but doing so doesn't mean we should pretend that they are more adequate than they are.

I think that most scientists actually have the same attitude. Nobody likes to have exceptions staring one in the face. One way to avoid exceptions is to back off and make more general statements, statements about populations instead of individuals. You can't say that a certain piece of advertising will influence Greg Williams of Gravel Switch, KY to rush to the store and buy a widget, but perhaps you can say that in a population of a million people reached by the ad, at least 100 people are all but certain to rush to the store. You can't prove that your reasoning about why these people are influenced is correct, and you don't know which 100 people will respond, but at least you can say that events haven't proven your reasoning wrong. After some experience with making this sort of prediction, you can become pretty confident that your advertising theory will produce at least 100 customers per million, with no exceptions.

The difficulty arises when the client says "That's all very fine, but what I want is a return of 1000 customers per million, and furthermore, I want to be sure that these 1000 people will actually pay their bills." If you then give in to temptation and promise what the client wants, you're going to start experiencing failures of the theory at a significant rate, because all you really _know_ is that at least 100 persons per million will buy the product. You start promising more than your knowledge can justify.

This, it seems to me, is the bind in which people who use statistics in lieu of knowledge about individuals find themselves, more often than not. In the medical profession, for example, we have all kinds of claims about what pills, diets, and physical regimines will do for you, all based on statistical studies of populations. One result is that hundreds of billions of dollars are spent each year by people for pills, diets, and medical attention that do almost all of them no good at all. Yet the medical profession can point to the statistics and show that in fact, all these treatments have had a clear statistical effect on the population. They point out that taking aspirin every day reduces the chance of getting a heart attack by 40 percent. They don't say what the chance was initially, or how small a chance there is that you would have had a heart attack without taking aspirin. Aspirin is cheap. But what about the effects of taking expensive drugs that produce even smaller effects in a population, on problems with even lower incidences?

We have the same problems almost everywhere that statistics is used in a way that affects people's lives, whether it be in psychotherapy or scholastic tests or job screening tests. The knowledge is used in situations where it is actually invalid and in ways that are actually harmful. That I what I object to, loudly and strongly. You can call that arrogance if you wish.

There are plenty of situations in which statistical knowledge is used correctly, in such a way that it could easily pass my criterion that there be no, or insignificantly few, exceptions. Those are the situations in which the knowledge is applied to the same population from which it is derived, and is used only to make predictions about the population.

Even when the knowledge is limited, it can still be used honestly. The advertiser could easily meet my criteria by continuing to claim that using his knowledge of responses to advertising claims, at least 100 people per million will respond to his ads, but also making it clear that the chances of getting 1000 respondents are 100-to-1 against. The doctor could meet my criterion by telling his patient that these $20-each pills, taken every day for a year, have a 0.2 percent chance of alleviating whatever is causing the patient's current complaint, and a 25% chance of not causing something worse. The psychotherapist could meet it by telling the parent that this test for Attention Deficit Disorder in the child is 85% accurate, but that the incidence of ADD is only about 2%, so the chances are that a positive indication of ADD has only about one chance in 7 of being correct -- and Ritalin, which is indicated for the condition, not only seems to help about half (or whatever) of children taking it, but also causes side-effects of dullness and other things that children don't like.

Unfortunately, avarice, self-importance -- and yes, arrogance -- usually triumph over honesty. One must, after all, make a living, and what's wrong with making a profit?.

>One manifestation FOR ME of arrogance is assuming that others >should ascribe the SAME importance to one of these as you do.

Does this prevent you from trying to persuade others to your point of view? You can acknowledge that each person selects goals privately for private reasons, and has every natural right to do so, without passing a law that prevents one person from trying, by nonviolent means, to get others to see an advantage to themselves in adopting a different goal. If people believe that certain goals and ways of achieving them will impress others, is there anything wrong with letting them know that you, at least, are not impressed? If someone expresses an opinion about something of general importance, is there anything wrong with stating that you have a definitely different opinion? I get the impression that you interpret PCT as saying that people should adopt a strictly hands-off attitude toward other people's goals, attitudes, and opinions. I don't read the principles of PCT that way at all. I haven't noticed others reading them that way when it comes to arguing with me.

>When you say something like "this is what you should want to >do, and doing something else isn't as good," you are simply >saying (arrogantly) that your own goals are better (i.e., "more >important" or "real knowledge") than the goals somebody else >prefers.

Finding the right attitude isn't always easy, and one doesn't always manage to live up to the attitudes that seem right. I try to put the "if"s in my arguments as appropriate, but sometimes I forget to say "If you want knowledge that fits the data..." Sometimes I make the mistake of just assuming that everyone would prefer their knowledge to be supported by experience instead of being contradicted by it. It still seems rather silly to have to say that all the time. Why would anyone who prefers theories that are the worst at making predictions be listening to or participating in this conversation?

>Are you suggesting that the concerns of people who deal with
>populations are in some sense unseemly, because "a few"
>individuals generally get stepped on?

Oh, yes. Definitely. I think they are doing harm. Unless they are willing to take real responsibility for the harm they cause, and take steps to prevent it, they should be doing something else for a living. That is my opinion and I will defend it.

>Are you suggesting that these people could use PCT to do a
>better (to you) job in some way?

Perhaps, if they would actually try to use it. It's also possible that the job doesn't have to be done, or should be done in an entirely different way that doesn't cause harm.

>Perhaps your preference is for the population dealers to simply quit dealing?

For the most part, yes. That is what I would prefer, if they're unable to make the effects they have on others commensurate with their actual knowledge. And before you ask if I would prefer that doctors stop treating illnesses, please think more deeply about what I mean. I ask only that doctors live up to their own credo, which begins, "First, do no harm." And I hope you are keeping your own (apparent) principle in mind: that it is arrogant to tell someone else what to prefer.

>Personal human goals are the only ones we've got. You seem to
>be implying sleaziness or selfishness again here. Do you really
>know that much about the people who use population statistics
>to say that what they (some? many?) are doing is "bad" ethically?

Yes. Many of them. The ethical "badness" is, of course, in terms of my own ethics. Many of them do things which, if I did them, would make me feel ashamed; which I would want to present to others in a better light so they would not know that I am ashamed, or look down on me. I would be ashamed to deny medical insurance to a person who is suffering from effects of a previous car accident, and then turn around and sell the same person medical insurance at three times the cost, as was done to Mary by Blue Cross, which happens also to run the uninsurable pool in Colorado. I would be ashamed at dictating the future course of another person's life on the basis of a test that I knew had a significant chance of misjudging that person. I would be ashamed at doing most of the things that I have seen done by the people who rely on poor statistics to make important decisions about other people's lives. I wouldn't do those things, and don't. I am disgusted about such practices and see no reason to keep quiet about it. Perhaps people who have a good opinion of my judgments in other matters will take my attitude as something to be considered when they form their own opinions. Perhaps not. That's beyond my control.

>OK, the challenge becomes: show me you understand my behavior
>and what it is being used to control. But first, YOU tell ME
>what a passing grade at such "understanding" would consist in.

My own ethics tells me that if I go around issuing challenges, I can't very well turn one down.

Passing grade? Why not make it as high as possible? Let's say that I will have met your challenge if you agree that I understand some behavior of yours and what it is being used to control. I'm not claiming that I can get a passing grade, but why settle for a lesser goal

A couple of questions first. Will you be trying to help me understand a behavior of yours and what it is being used to control? The easiest way to find out what a person is controlling for is, of course, to ask. If the person's intention is to reveal the controlled variable, then the problem shouldn't be too hard to solve. On the other hand, if the intention is to make it difficult, that intention shouldn't be too hard to discern, either, nor the means of carrying it out. The only really difficult case would be the one in which a person uses deception, trying to give the impression of controlling for something when there is no actual control of it, with the intent not only of concealing a real controlled variable, but making it appear that something else is being controlled.

That, however, would be difficult to achieve, because to give the impression of controlling for something without actually controlling it becomes impossible when the spurious controlled variable is disturbed. Either one must control it, thus revealing it, or let it change under the disturbance, in which case it is clearly not controlled. You should be certain that you want to issue this challenge, because I see no way for the Test for the Controlled Variable to fail. I trust that you don't mean the challenge as saying that I must lay out the entire structure of all your controlled variables from bottom to top and from now to the end of your life, a project that would require rather a long time and rather intimate interactions on a continuing basis, not just through words but through extended personal contact. I think I would have to admit that I couldn't meet that challenge. I assume, however, that you have something more practical in mind.

So:

You have issued a challenge to me to explain what variable a behavior of yours is being used to control. Issuing the challenge is clearly a behavior of yours; what is not immediately apparent is the perception in you that is being controlled by issuing this challenge. So let me try the easiest way first:

What is the effect you intend to produce by issuing this challenge?

Best, Bill P.

Date: Sat Mar 27, 1993 12:21 pm PST
Subject: JOINING THE CONTROL SYSTEMS E-MAIL NETWORK

I would like to join the control systems e-mail network. I would appreciate any help avialible in making this come about. thanks.

Bryan Ennis

Date: Sat Mar 27, 1993 3:59 pm PST
Subject: Information about disturbance

[From Bill Powers (930327.1600)] Rick Marken (930327.0900) --

Rick, there's a mathematician's way of deciding whether a function approaches a value of zero. It goes, "Let epsilon be some small number. Prove that for any epsilon, the value of the function is not greater than epsilon. If this is true for all epsilon within some small distance from zero, then it is true for epsilon equals zero." You can show that 1/(1+g) is less than or equal to any small epsilon for g greater than some amount. Therefore the limit of the function as g goes to infinity is zero.

Actually, your point is made just as well with g = 10. I'm beginning to think that our claim is best made by referring to attribution. To get information from a signal that depends on a different variable, you have to know what variations in the signal to attribute to the variable. If there is more than one contributor, this becomes impossible -- or rather, there are infinitely many ways to attribute the signal partly to one variable and partly to the other.

We're probably indulging in overkill, trying to prove an elementary point. If I tell you that the value of the perceptual signal is 3, and ask you to deduce from that the values of the disturbing variable and the output variable, the task is clearly impossible. But this is a proof from the world of continuous variables. Let me try one using the terminology of information theory, about which I now know dangerously little (I've been reading a book by Garner, recommended by Martin Taylor).

Uncertainty depends on the number of possible outcomes, n. It is formally defines as the log to the base 2 of n, but we can just use n -- it's only a matter of the kind of graph paper you use for plotting uncertainty.

When we are given the perceptual signal only, its magnitude might come from any combination of values from the range of the output variable and the disturbing variable that have that sum. If we know the range of the perceptual signal, we can say that the MINIMUM range of the output and disturbing variables must be equal to the range of the perceptual signal. Dividing that range into n discrete values, then, there are n ways of achieving the observed value of the perceptual signal by picking one value from the minimum range of the disturbing variable and another value from the minimum range of the output variable. A given observation of the perceptual signal might be accounted for in any of those n ways.

This means that if the initial uncertainty about the value of the disturbance was n before we looked at the value of the perceptual signal, it is still n after we look at the perceptual signal. The uncertainty about the value of the disturbance has not been reduced at all by knowing the perceptual signal's value (whether you calculate it in logs or linearly). Therefore no information has passed from the disturbance to the perceptual signal.

I believe the proof is general -- but I'll leave it to a real mathematican to say yea or nay to that.

Best, Bill P.

Date: Sat Mar 27, 1993 6:53 pm PST
Subject: DME & THEORIST - RKC

[From Bob Clark (930327.1945 EST)]

The following remarks are not derived from the posts of the past week or so. They bring up another subject, one that can lead, I think, to some very interesting and helpful results.

Over the last ten days I have tried to write this material from several different viewpoints. Each is pertinent and interesting, but tends to become too long and complicated for a reasonably short (2 - 3 pages) item. This viewpoint appears to offer a framework that can be used to explore additional important (useful) subjects, including some of those recently on the Net.

The Decision Making Entity (DME) [see Bob Clark 12052.] can be considered from several Viewpoints. Each is interesting, but the Theorist's is the most general, and may be the most useful.

THEORIST'S VIEWPOINT
This Viewpoint is defined here by paraphrasing and quoting from BCP p 18.

The HCPCT Theorist proposes to construct a "model of the brain's internal organization" where "observed behavior is deduced ... from the way in which these internal entities interact with each other and the external world."... These entities have been chosen not only to "behave properly," but also to fit anatomical hints about the nervous system, physical models of the organism and its environment, subjective experience, and elementary mathematical logic.

I. PRIMARY CONCEPTS -- greatly condensed summaries of BCP
A. "Behavior is the Control of Perception" -- "Perceptual Variables"
B. The Negative Feedback Control System and its intrinsic properties.
C. The Hierarchical Structure of Negative Feedback Control Systems.
D. Problem-solving Programs -- fixed instructions with choice points;
E. Intrinsic Variables -- genetically determined;
F. Reorganization -- change in the properties or number of components
G. Memory -- recording and playback switches;

These concepts, together with their analysis and development, cover a remarkably large range of human (and other) activities. However, this structure is largely fixed in form, changing only by the addition of new Recordings ("Memories") or Reorganization. Problem-Solving Programs, including associated Choice Points, are composed of Recordings. They are derived from combinations of Recordings and/or Reorganizations. New Programs result only from new, and re-arranged, Recordings and Reorganization. This results in limited flexibility leading to several problems.

II. POSSIBLE PROBLEMS
1. Minor Changes in behavior may be needed because of inadequate or "incorrect" Problem-Solving Programs. Reorganization is unnecessary and not initiated.
2. Minor Changes in behavior may be needed because of inadequate, or "incorrect" Recordings. Reorganization is unnecessary and not initiated.
3. An Operator is needed to control the Recording-Playback Switches.
4. A Source of Reference Levels is needed at the top of the Hierarchy.
5. Arbitrary Action is observed in the absence of Intrinsic Error.
6. Initiative is observed, but not explained in present PCT.

7. Anticipation of unexpected events is observed, but not explained in present PCT.
8. Errors, accidents, misdeeds are observed but Assignment of Responsibility is not provided in present PCT.
9. Subjective Reports ("User's View") of the process of Selecting among alternatives is not described in present PCT.
10. An "Observer's View" of Subjects' unexpected actions is not described in present PCT.

A "Decision Making Entity," "DME," is proposed as a partial solution of these questions. The concept seems to be generally taken for granted and accepted by many people -- including most (if not all) of those participating in the CSG Net. Such acceptance is demonstrated by the frequent use of the First Person Singular. "DME" is proposed as a name for this concept when personal associations are removed, leaving nothing but the process of selecting from among alternatives for action. It offers a straight-forward way to solve some of the above problems, and possibly others, by the addition of a single element with its associated capabilities and characteristics. This concept is consistent with several others discussed in BCP, and helps to clarify the operations and relations within HPCT as summarized above.

III. OPERATION OF THE DME -- Summary
i. Reacts to attention-getting events;
ii. Searches for relevant memories (by association and/or content);
iii. Compares their anticipated results;
iv. Selects those preferred on the basis of selected guidelines;
v. Puts them into effect by using them as Reference Levels for selected Orders within the Hierarchy.

These and other topics can be discussed separately. Enough for now.

Regards, Bob Clark

Date: Sat Mar 27, 1993 9:00 pm PST
Subject: I was wrong.

[From Rick Marken (930327.2030)]

I hate to do this but I have to eat crow; I WAS WRONG. I said (Rick Marken (930327.0900))

>Martin's claim here [that perfect control means zero disturbance in the
>perceptual signal] turns out to be FALSE

This statement was based on a mathematically incorrect analysis of the following equation:

p = (k.e k.o r + k.d d)/(1+k.e k.o)

I thought that even if k.e k.o (loop gain) went to infinity there was still a little, teensy weensey amount of disturbance left over. But Bill Powers (930327.1600) set me straight:

>You can show that 1/(1+g) is less than or

>equal to any small epsilon for g greater than some amount.
>Therefore the limit of the function as g goes to infinity is zero.

So infinity was a bad choice of values on my part -- and when the loop gain goes to infinity, p = r and there is no disturbance represented in the perceptual signal. So MARTIN WAS RIGHT: I WAS WRONG; if control were perfect (loop gain = infinity) there would unquestionably be no information about the disturbance in p because there would be NO effect of the disturbance AT ALL on p. So we are back to where we were: Martin claiming that there is no information about the disturbance ONLY when loop gain is infinity; me claiming that there is no information about the disturbance EVER.

I tell you; admitting that you are wrong is NOT EASY. I hate this. But I did it, not only because my error was made in public so I was caught "read" -handed, but also to show that I am willing to admit when I am wrong (NB. Ken Hacker -- is this arrogance?). So Martin and Allan can be assured that if they can persuade me (with evidence) that there IS information about the disturbance in the perceptual input to a control system then I will (eventually) admit that they are right and I am wrong. I'll hate doing that; just as I imagine they would hate admitting that Bill and I are right. I know admitting error is really tough -- but it's not the end of the world. Now I know that infinite loop gain does produce perfect control -- ie. p = r. Now I just have to figure out whether such a system is physically possible (it probably isn't); if it is, then I'm going to have to do some figuring just like Martin and Allan might have to do some figuring if they are able to convince themselves that there actually is no information about disturbances in perception.

Anyway, it is a far far better thing I do ....

Back to the fray. Best Rick

Date: Sun Mar 28, 1993 5:38 am PST
Subject: Building a better challenge

From Greg Williams (930328) Bill Powers (930327.0700)

>My reason for preferring modeling is that I like
>explanations that fit the observations as much of the time as
>possible, preferably all of the time.

This appears to be a profession of FAITH, in the following way: you see that the method of modeling individuals in some of the "hard" sciences (especially physics) does indeed result in "explanations that fit the observations" often with high precision, then you conjecture that similarly high precision will be achievable by applying the method to high-level human behaviors "in the field" (extending far beyond laboratory tracking experiments). But such precision is NOT currently possible.

>if modeling didn't help to produce explanations
>that work, I wouldn't like modeling much, either.

With regard to modeling high-level human behavior, as I said in my last post, the verdict is still out about whether the explanations produced (some time in the future) will "work." By comparing what actually works rather poorly RIGHT NOW (statistical description of human behavior) with what does NOT work RIGHT NOW (individual modeling of complex human behavior), you open the door to charges of arrogance. Building a better mousetrap and showing it off is certainly not arrogance, but the temptation for an observer to see arrogance is great if someone touts the "better" mousetrap which doesn't yet exist and demeans the existing mousetraps RELATIVE TO the non-existent "better" one.

>We have the same problems almost everywhere that statistics is
>used in a way that affects people's lives, whether it be in
>psychotherapy or scholastic tests or job screening tests. The
>knowledge is used in situations where it is actually invalid and
>in ways that are actually harmful. That I what I object to,
>loudly and strongly. You can call that arrogance if you wish.

That kind of objection doesn't appear arrogant to me. It is only when it is coupled to unsupportable claims that your approach is better -- especially if it is supposed to be better with respect to the goals of those happily using statistical description (nonfallaciously, as public health workers do when they want to see whether a disease incidence is going down in a population, and don't care WHICH individuals are infected) -- that I say it is arrogant. Arrogant individuals make unsupportable claims of self-importance: they profess to have a better mousetrap than the others, but it turns out that they don't even have a mousetrap. Of course, someday they might actually build a mousetrap and it might actually be better than all others (with respect to their own goals and even the goals of others) and will have a good reason to claim legitimate, non-arrogant selfimportance.

>You can acknowledge that each person selects goals
>privately for private reasons, and has every natural right to do
>so, without passing a law that prevents one person from trying,
>by nonviolent means, to get others to see an advantage to
>themselves in adopting a different goal.

Yes. But some ways of trying to get others to see an advantage appear to be more fruitful than others. That probably sounds pretty statistical to you, but you can try an experiment to see how good the statistics are: stand on a street corner in Durango and (if you can make it that long) say to 100 passersby, in turn: "You asshole! My theory about how you work is better than your theory!!"

>If people believe that certain goals and ways of achieving them will
>impress others, is there anything wrong with letting them know that
>you, at least, are not impressed?

And then what? I expect that several will then ask something to the effect: "Well, what do you propose that is more impressive?" You'd better be prepared with something THEY will find impressive, notwithstanding that YOU find it more impressive than their stuff.

>If someone expresses an opinion about something of general importance,
>is there anything wrong with stating that you have a definitely
>different opinion?

No. But then you have to back up your opinion with evidence (here comes the hard part) IMPRESSIVE TO THE OTHER PERSON. Otherwise, you will sound arrogant.

>I get the impression that you interpret PCT as saying that people

>should adopt a strictly hands-off attitude toward other people's
>goals, attitudes, and opinions.

No. And I don't think that PCT says one "shouldn't" be arrogant (or even deluded) -- PCT isn't ethically prescriptive. But if you are trying to persuade scientists, arrogance (unsupportable claims of "I've got a better way to do it than you") doesn't work very well. Yes, more statistics; some scientists no doubt can be swayed by arrogance. LOTS of rock fans appear to cherish arrogance at certain heavy metal concerts. And delusion also can persuade some folks, too. Witness Waco. It seems that the arrogant and the deluded need to choose their audiences carefully, if they are to be persuasive.

Before we get to the challenge business, I want to express my hope that you will say a bit more about whether or not PCT is (or will ever be) appropriate for predicting behavior (i.e., with .99+ correlations), as well as for understanding (I assume, post hoc only) behavior. Is the most that can be expected from individual modeling of high-level human behavior in the field an after-the-fact analysis? I note that individual modeling in physics is valued (especially by engineers) for its ability to predict (i.e., that this spacecraft will take this path). Why should prediction be downplayed when it comes to high-level human behavior? Are tracking experiment predictions ("one whole minute," even after several months) to be the limits of PCT predictions?

>Passing grade? Why not make it as high as possible? Let's say
>that I will have met your challenge if you agree that I
>understand some behavior of yours and what it is being used to control.

Fine.

>A couple of questions first. Will you be trying to help me
>understand a behavior of yours and what it is being used to control?

Yes.

>You have issued a challenge to me to explain what variable a
>behavior of yours is being used to control. Issuing the challenge
>is clearly a behavior of yours; what is not immediately apparent
>is the perception in you that is being controlled by issuing this
>challenge. So let me try the easiest way first:

>What is the effect you intend to produce by issuing this challenge?

I don't know. Perhaps you have a way to get me to say that I "realize" some particular effect is the one I intended? The question then is: how do we know that was REALLY the effect I intended (yesterday), and not just some sort of after-the-fact fiction than seems to make sense? How do we avoid "just-so" opining? One way I can think of is to analyze my history and current environment and then predict what I'll do next -- that would be an impressive feat, if it could be repeated at will. But post hoc analysis leaves the same nagging doubt in both the theorist's and subject's minds as psychoanalytic "explanations": how much simply SOUNDS reasonable (to both parties)?

As ever, Greg

Date: Sun Mar 28, 1993 11:14 am PST
Subject: Of mice and lint

[From Rick Marken (930328.1000)]

Greg Williams (930328) --

It seems like you are trying to hone in on what makes PCT seem arrogant. I sense that you have deep feelings about this (like Ken Hacker) so I want to jump in here and see if I can get a better feeling for what you think of as PCT arrogance.

>By comparing what
>actually works rather poorly RIGHT NOW (statistical description of
>human behavior) with what does NOT work RIGHT NOW (individual modeling
>of complex human behavior), you open the door to charges of arrogance.

I don't know if this is the comparison being made by Bill (or any other PCTer). This is like saying that any criticism of the methods and theories of astrology is arrogance if I don't have a better method or theory of the phenomena that astrology is trying to explain.

>Building a better mousetrap and showing it off is certainly not
>arrogance, but the temptation for an observer to see arrogance is
>great if someone touts the "better" mousetrap which doesn't yet exist
>and demeans the existing mousetraps RELATIVE TO the non-existent "better" one.

I agree that it would, indeed, be arrogance (a better word might be "cant") to tout a "better" mousetrap which doesn't exist. It is also rather annoying to have someone say that you are doing this, when you are not. Could you please give us one example of where a PCTer has done this. When have we ever said that PCT is a better mousetrap when we did not demonstrate this claim? When dealing with phenomena where 1) we have a good idea that control (pur- poseful behavior) is involved and 2) we also have a good idea what variables are controlled, then we have been able to demonstrate that PCT is better than existing mousetraps. But I don't believe that PCT has claimed to be a better model of the kinds of statistical phenomena that we criticize. In fact, we explicitly say that there is probably nothing to explain (as would be the case if an astrologer asked for your better explanation -- than that fire and water signs don't get along -- of why Aries and Pisces don't make good mates).

So again, I ask you where PCT has claimed to be a better mousetrap without showing that it is. I want a case where PCT says that it is a better mousetrap for catching MICE (purposeful behavior); I don't want examples where PCT has said "we have a better mousetrap but there aren't any mice around here; follow us and we'll show you where to find them; what you've been catching is large clumps of lint (side effects of purposeful behavior)".

>Arrogant individuals make unsupportable claims of self-importance: they profess
>to have a better mousetrap than the others, but it turns out that they
>don't even have a mousetrap.

Again, give me an example of this. I think that you are just confused. PCT is not not saying that it has a better mousetrap and not producing it; it is (usually) saying that that there ain't no mice around here (behavioral science) and what you (behavioral science mouse catchers) have been catching is lint; PCT is a mousetrap to catch mice (purposeful behavior); the linttrap (SR models) that that the behavioral sciences have been using might work to catch lint; we want mice.

Best Rick

Date: Sun Mar 28, 1993 1:00 pm PST
Subject: Output to Disturbances

From Ken Hacker [930328]

Some output in relation to some disturbances:

[1] I am trying to understand PCT
in two ways: first, as a explanatory perspective or theory of human perception, and second, as a behavioral theory related in some ways to other theories. To understand PCT, I ask questions, probe, and query what does not make sense to me. If I ask questions or raise objections that offend someone, I apologize for the personal affront, but do not apologize for the intellectual challenge. In science, there should be fewer articles of faith than testable and tested propositions. I guess I am PCT-challenged and Theoretically Incorrect at this point, but c'est la vie.

[2] Mr. Marken may want to debate the arrogance or hyperbole which I argue are evident in PCT postings (by some), but he cannot win (or lose) the debate for the simple reason that we are tussling over my perceptions -- not empirical truths.

I really doubt that even Mr. Marken can tell a human being what his or her perceptions are more accurately than the person themself. Good reasons to me are not going to be good reasons for someone else unless we agree on some ground rules. I believe that scientific discussion should be grounded in the norms of education; multiple points of view are discussed and evaluated in a climate of tolerance. I do not believe that scientific discussion should be ground in norms of propaganda, whereby only one point of view is assumed correct and others are dismissed.

[3] There are claims made on this net which I do, in fact, think are problematic (arrogant or hyperbolic). These include the following (if Mr.ta, he should print out the notebooks and start counting -- I don't have time to do it for him):

* All social science is wrong.
* All behavioral science is wrong.
* All social scientists and behavioral scientists believe what is wrong.
* All social scientists and behavioral scientists do not know what they are doing with research.

* Only PCT has something useful to say about how humans regulate their behaviors.

* Communication theory is wrong.

* PCT equations (confirmed by PCT theorists) PROVE that PCT is correct.

* PCT is real science. Other approaches to human behavior are pseudo-scientific or "half-assed."

* People who challenge PCT are misguided, ignorant, and not yet fully developed intellectually.

The claims above are simply the ones I could recall at the keyboard. There are more, but these are enough to make my point: PCT theorists need to be more receptive to INQUIRY, which necessitates consideration of questions by those who are new to the ideas, concepts, models, etc.

[4] There are several points which have been made that I think are useful and insightful (from my view):

a. People are influenced by others, but only individuals can control their own behavior.

b. While some views of behavior, such as common notions in traditional psychology, assume that behaviors are done in response to stimuli, PCT assumes that behaviors are the independent variables, not dependent ones.

c. There are hierarchical orders of input-comparison in human control.

d. Much of behavioral science is rooted in the belief that operant conditioning still guides much of those behaviors which appear automatic.

These are just a handful of points I have picked up and recalled. There are many more I have not listed.

[5] There are many questions which I have about PCT, including:

a. The scope of PCT. Is it limited to describing and explaining physical behaviors of human beings such as motor actions?

b. If there is no information inside of input, what prevents one from asserting behavioristic-like internal responses to internal stimuli or a kind of reverse behaviorism?

c. What happens to Ashby's Law of Requisite Variety?

d. How does PTC account for stored knowledge such as schemata (or whatever other term you choose)?

d. What makes PCT more than intensive descriptions and explanations of neural pathways and signals?

e. Why must PCT be understood as some sort of revolution as opposed to some sort of new and productive thinking?

f. What are essential differences between PCT and cybernetics?

g. What do axioms or propositions from PCT contribute to everday human adapting, living, changing, improving, that other forms of analysis do not? Where is the proof of human behavioral successes with PCT, as opposed to success in the abstract? Is too soon to ask this question?

h. Most scientists are concerned with not only describing and explaining phenomena, but also wish to predict and sometimes "control" them. Why does PCT deny the desirability of prediction? I can predict one PCTer; I have already deleted some expletives because of that fact! Is prediction not always part of human anticipation, whether in ordinary life or in science?

i. Where do reference levels or signal originate? In other words, if my control system has error signals, what constitutes the sources of the error signals?
What is stored that makes me think that I do not like or wish to accept certain input? How conscious are the processes of matching input to the reference signals?

[6] So far, it seems to me that PCT is doing very productive work in describing and explaining individual control and the complex mechanisms or processes which are involved with physiological or neural signalling. This is valuable and informative. It does not explain social behavior, social systems, or interactions between control systems, but it cannot be faulted with this, since its focus is on single individual self-regulation. As PCT should not be criticized by me or anyone else for not explaining social behavior, social sciences which do focus on social behaviors should not be excoriated for not explaining individual human control mechanisms and processes.

[7] Finally, my interests in PCT are something like this:

a. Communication studies, unlike traditional psychology and maybe other behavior/social sciences, no longer assumes that people are simply affected by stimuli in varying conditions and ways. Since the 1970s, my discipline has said that each human TAKES from mass media, from interpersonal interactions, from printed words, from any messages, and recodes what is decoded. We threw out Shannon and Weaver's model in the 1970s as anything useful to explain human communication. Both Aristotelian and electrical engineering models of message sending and receiving say little of importance to understanding the complexities of human communication, mainly because people do not simply

take in messages. Nor do
messages do contain meaning. Meaning
is created internally by receivers of messages and the messages
themselves are only one part of meaning construction. Along
with this, I have thought
that humans, as Vygotksy and others note,
are always adapting and when they interact with others, they
attempt to adapt more by their communication behaviors.

b. I believe that each human being has needs, goals, and reference levels and essential variables. When people interact, they are attempting to accomplish their goals, minimize instability in their own life, and regulate themselves to gain what is important to them. I think PCT's descriptions of how people control are very useful for understanding what each communicant is doing endemically as a human being, so that we describe social ineractions, we understand the behaviors that each communicator naturally engages in while responding and adjusting to the other human being.

c. I also acknowledge that social science methods have been problematic and need to mature into studying more integrated and detailed processes of social behaviors and structures. As Chomsky points out, and Bill Powers does also, social science needs to adopt more of the rigor of the physical sciences and to produce clear explanatory principles of phenomena. I think that today, much of social science consists of scattered topics, perspectives, and mountains of findings. I think much of it is interesting and useful, but that there are strong needs to explain how things work, as opposed to how various variables are related to each other within a context of general topics.

I have to stop here for interests of brevity and to get back to grading and other such exciting things. I truly enjoy the debate here and am learning much from it. I will have more responses and comments later as I catch up with about 4 days of messages I printed out. My goal is not to win any arguments here, but to state my perceptions, check my own thinking, and make sense out of what I see PCT doing and how it may relate do what I study regarding human communication. For those who are patient with me, I salute you.

Best, Ken Hacker

Date: Sun Mar 28, 1993 1:29 pm PST
Subject: Where is the proof? Information of disturbance

******** FROM CHUCK TUCKER 930328 ********

Some time ago Rick (930315.1500) stated that his research in "Mind Readings" Chapter 3 ("The Cause of Control Movements in a Tracking Task" proved "... THERE IS NO INFORMATION ABOUT DISTURBANCE IN THE PERCEPTUAL SIGNAL CONTROL SYSTEM. This means that THE PERCEPTUAL INPUT TO A CONTROL SYSTEM CANNOT BE WHAT CAUSES THE OUTPUR OF THE SYSTEMN (sic) TO MIRROR THE DISTURBANCE." Reading his article again I don't believe that he proves that at all; in fact, the subjects did have an awareness of a disturbance - they were told about it and through several test runs of the task knew that a disturbance was influncing the cursor - it is the case that they did now know WHAT is was (neither did Marken - it was random). In the study "Subjects were tested individually. Each subject was seated before the videodisplay and asked to keep the cursor aligned with the target by turning the game paddle appropriately. After several practice sessions subjects were tested in 20 experimental runs...(63)." I would have refused to published this study since WHAT the subjects were told is not to be found in this description BUT having done this type of work I know that they had to be told something (unless all of the subjects had done this many times) BUT IT SEEMS OBVIOUS TO ME THAT THE SUBJECTS WHERE AWARE THAT A DISTURBANCE WAS INFLUENCING THE CURSOR POSITION ALONG WITH THEIR MOVEMENT OF THE GAME PADDLE. If this investigator believes otherwise let him present evidence that they were ignorant of a disturbance. He should has ASKED THEM or used THE TEST.

Date: Sun Mar 28, 1993 4:02 pm PST
Subject: Granting Rick's wishes

From Greg Williams (930328 - 2) Rick Marken (930328.1000)

>>Building a better mousetrap and showing it off is certainly not
>>arrogance, but the temptation for an observer to see arrogance is
>>great if someone touts the "better" mousetrap which doesn't yet exist
>>and demeans the existing mousetraps RELATIVE TO the non-existent
>>"better" one.

>I agree that it would, indeed, be arrogance (a better word might be
>"cant") to tout a "better" mousetrap which doesn't exist. It is also
>rather annoying to have someone say that you are doing this, when you
>are not. Could you please give us one example of where a PCTer has
>done this. When have we ever said that PCT is a better mousetrap when
>we did not demonstrate this claim?

When you claim that all of psychology needs to be rebuilt from the ground up on a PCT foundation, I think that sounds like a claim for having a "better" trap to catch WHATEVER psychologists want to catch. (See examples below.) To date, no one has demonstrated that after the rebuilding, the resultant edifice will actually suffice to meet psychologists' (including your own) goals for "understanding" or "predicting" (take your pick) high-level human behavior generally. Or do you claim to know in advance that it will suffice? If so, upon what basis do you make the extrapolation from rubber-banding and tracking experiments?

>When dealing with phenomena where
>1) we have a good idea that control (pur-poseful behavior) is involved
>and 2) we also have a good idea what variables are controlled, then we
>have been able to demonstrate that PCT is better than existing
>mousetraps.

I agree, subject to (3) in some CIRCUMSCRIBED situations. Where does any extrapolatory faith come from?

>But I don't believe that PCT has claimed to be a better
>model of the kinds of statistical phenomena that we criticize.

I think the closest a PCTer has come to this claim is when Bill Powers suggested some time back that the best way to handle social behavior is by first handling individual behavior. I suppose Bill will deny that he meant using individual models to build up predictive social models (from which could be derived population measures), so I guess you are correct in your belief.

>In fact, we explicitly say that there is probably nothing to explain (as
>would be the case if an astrologer asked for your better explanation -
>- than that fire and water signs don't get along -- of why Aries and
>Pisces don't make good mates).

Many users of statistical data on human behavior understand already that they are NOT explaining. And at least in some cases they are happy about not needing to try any explaining. But they do try to predict sometimes, and if you think prediction should be an important aim for psychologists (as you sometimes seem to, see examples below), perhaps you think that PCT (someday) can make better predictions of population measures than correlation studies? If so, how is that claim supported by the accomplishments to date of PCTers?

>>Arrogant individuals make unsupportable claims of self-importance: they profess
>>to have a better mousetrap than the others, but it turns out that they
>>don't even have a mousetrap.

>Again, give me an example of this. I think that you are just
>confused. PCT is not not saying that it has a better mousetrap and
>not producing it; it is (usually) saying that that there ain't no
>mice around here (behavioral science) and what you (behavioral science
>mouse catchers) have been catching is lint; PCT is a mousetrap to
>catch mice (purposeful behavior); the linttrap (SR models) that that
>the behavioral sciences have been using might work to catch lint; we
>want mice.

It is entirely possible that I am just confused. Maybe even more than I usually am! Be that as it is, it seems to me (see below) that you claim that non-PCT behavioral scientists can NEVER understand/predict high-level purposeful behavior of individuals, yet PCT can... SOMEDAY, you believe. You say that the non-PCT linttrap cannot catch mice... but I think you'll have to admit that the PCT mousetrap has so far not caught any really big mice. A nice big juicy one would be predicting my behavior for "one whole minute" -- Bill Powers already passed on that challenge; perhaps you would like to try (someday), since you think a tool "isn't much of a tool" if it "can't help us predict things" (see below).

Below are some examples (from Rick's posts during the first three weeks of March 1993) of what I believe some non-PCT behavioral scientists would consider "arrogant" (and a few other things of tangential interest). {} delimit Rick's remarks about which I make comments delimited by {{}}.

>[From Rick Marken (930301.1000)] Ed Ford (930228:0920) --

>>A close friend needs good references in the current literature
>>(books and articles) on the best explanation of cognitive theory.
>>I would appreciate anything you could offer.

>I don't have easy access to the current literature on cognitive
>theory. {And it would be hard for me to evaluate what might
>constitute the "best" explanation of this theory -- since I
>think they are all equally ridiculous.}

{{Here you don't explicitly say that "ridiculous" refers only to trying to handle purposive behavior; it would be easy for a reader to accuse you of arrogantly claiming that cognitive theory is "ridiculous" with respect to its own goals, which are not necessarily limited to handling purposive behavior.}}

>Allan Randall (930226.1730) --

>> You do not like information theory.

>It's not a matter of like or dislike. It's a matter of "SO WHAT"?
>Information theory contributes zilch to our understanding of
>living control systems (though, I'm sure, Martin disagrees). But
>there is no need to argue; {WITHOUT information theory, Bill Powers
>has been able to build a simulation of a system that can produce
>realistic complex behavior in a realistic environment; WITH
>information theory the life sciences (with thousands of bright
>researchers and decades of research) have been barely able build
>simulations of systems doing unrealitic, simple behavior in
>unrealistically simple environments.}

{{This dichotomy could easily be construed as an arrogant judgment of the "failure" of non-PCT behavioral science SOLELY with respect to YOUR goals.}}

>[From Rick Marken (930304.1400)]

>There are two kinds of psychologists who embrace Perceptual Control
>Theory (PCT). {One kind (including myself) believes that we should
>start psychology over from scratch based on the principles of PCT.}

{{This could be construed as an arrogant assertion that PCT principles, and ONLY PCT principles, can (sometime?) suffice to meet EVERYONE'S goals for psychology, which might include generating population measures.}}

>[From Rick Marken (930310.1400)]

>{Well, IT [Information Theory] isn't much of a tool if it can't help
>us predict things; looks like PCT WITH IT is no better off than
>Isaac's brother Phil WITHOUT calculus.}

{{Really just a side issue: How important do you think prediction should be in PCT? Isn't the current state-of-the art of PCT prediction rather limited?}}

>[From Rick Marken (930312.1300)]

>I don't think Bill or I know what the heck you folks think IT
>is good for. Feel free to use IT along with anything you like
>to make a prediction. {All I want to see is how IT (calculus-like)
>can improve what we do in PCT -- which is try to discover controlled
>variables and model now they are controlled.}

{{Ditto on the side issue, as immediately above.}}

>[From Rick Marken (930317.0800)]

>Avery Andrews (930317.1514) --

>>People just don't care about the input-output model of behavior as
>>much as Rick thinks they do.

>Multi-millions of dollars are spent in the US (and probably
>Australia too) in support of behavioral science research (psych,
>sociology, econ, poli sci, etc) where the data is collected
>and analyzed in the context of the general linear model; multiple
>regression, ANOVA, etc. I bet few of these people would consciously
>say "I assume that the basic model underlying behavior is a
>cause-effect model" but they sure ACT like this is what they
>assume-- and big bucks are being spent in tacit support of this
>assumption. I believe it is important to know that this is the
>model that behaviooral scientists are "controlling for" -- consciously
>or not -- because I am sure that it is the reason why PCT -- after,
>what, 30 or so years on the scene -- has made virtually NO headway
>in the behavioral sciences. {Its either that PCT is just a stupid
>model and all the behavioral scientists have been smart enough to
>notice that (but there are some other obviously stupid "models"
>running around in the behavioral sciences and, nevertheless, they
>get a lot of attention) or it is because of active resistence. I
>think PCT has made no headway because there IS active resistence
>and I think the underlying reason for this resistence is that PCT
>is a disturbance to the assumption in the behavioral sciences of a
>cause-effect model.}

{{Maybe not arrogant in content (but so in flavor, to me!), but just mistaken. One of the first things one learns in statistics is that correlation doesn't mean cause-effect. The "Grand Method" of population description is DIFFERENT IN KIND from a generative model, and to claim that all who work with statistical descriptions instead of generative models THEREFORE are committed to a lineal generative model of behavior is simply a nonsequitur.}}

>>Indeed, I suspect that one of the reasons
>>linguists don't spend much time in the psych lab is that the
>>input-ouput, IV-DV stuff just seems stupid and boring to them, w.r.t.
>>the things they are interested in, but that's the only thing that people
>>seem to know how to do in labs.

>Well, I don't know if I would brag about not having any model
>at all. {Just observing is very genteel and all -- but unless you
>try to predict and explain what you see with a model, what have you
>learned?}

{{Better ask Bill Powers, who recently said that prediction isn't the point for PCT (though he earlier bragged about "one whole minute" of accurate prediction of tracking movements).}}

>I think that linguists do have implicit models -- cause
>effect models. If they don't want to test them then that's there

>problem. I don't care if they test them in labs or in the real world;
>but I don't think you have much of a science unless you test models.
>{If linguists don't know how to do anything other than IV-DV research
>when they do it it's because their basic (unconscious model) is
>cause-effect. If they don't go into the lab to do IV-DV research it
>must be because they don't like the implications of their own models.
>If their models were not cause-effect -- if, for example, they were
>control of input models -- I'm sure these bright folks would have
>noticed very quickly that there is an alternative to the IV-DV
>approach to research and they would have very quickly understood
>research based on testing for controlled variables.}

{{This raises the nonsequitur to Freudian heights!}}

>[From Rick Marken (930319.1300)]

>So, the reason this topic is controversial is because it seems to me
>that what you are saying about the operation of a control system is incorrect.

>Why is this alarming? Because it seems to me that you are trying to
>squeeze the beautiful concept of control of perception into the
>Procrustean bed of conventional linear cause effect thinking. {PCT
>demands that we start psychology all over again -- psychology beautifully
>reborn.} But, if we keep trying to see control in linear cause effect
>terms (instead of in terms of the circular causality that is actually
>involved) this renaissance will be continually be delayed.

{{Here is arrogance if the "demand" assumes that "we" includes everyone --
including those who don't have goals to which PCT is appropriate.}

>[From Rick Marken (930321.1200)]

>I think we have different "enemies". You have problems with SR
>theories of behavior. {I have problems with those ideas too -- but,
>more importantly, I have problems with the whole conventional approach
>to studying behavior -- both its goals (finding relationships
>between IVs and DVs) and it's assumptions (that such relationship
>reveal something about how behavior works).}

{{Do you have problems with users of statistical descriptions who DON'T assume
that such descriptions "reveal something about how behavior works"? Are such folks
not "studying behavior"? If you say they aren't, I think they would have good
reason to label you as arrogant.}}

As ever, Greg

Date: Sun Mar 28, 1993 5:41 pm PST
Subject: Re: DME & THEORIST - RKC

[Martin Taylor 930328 20:30] Bob Clark 930327.1945

>A "Decision Making Entity," "DME," is proposed as a partial solution of
>these questions. The concept seems to be generally taken for granted and
>accepted by many people -- including most (if not all) of those

>participating in the CSG Net. Such acceptance is demonstrated by the
>frequent use of the First Person Singular.

Count me among the non-acceptors. Your DME sounds very much like the old homunculus who sits behind the sensors and effectors, manipulating. How does he work? Does he have his own little hierarchy?

I give up. I've spent 10 minutes trying to edit one line. Bye.

Martin

Date: Sun Mar 28, 1993 8:11 pm PST
Subject: Where is the proof? Information of disturbance

[From Rick Marken (930328.1500) CHUCK TUCKER (930328)

>Some time ago Rick (930315.1500) stated that his research in "Mind Readings"
>Chapter 3 ("The Cause of Control Movements in a Tracking Task" proved
>"... THERE IS NO INFORMATION ABOUT DISTURBANCE IN THE PERCEPTUAL SIGNAL
>CONTROL SYSTEM. This means that THE PERCEPTUAL INPUT TO A CONTROL SYSTEM
>CANNOT BE WHAT CAUSES THE OUTPUR OF THE SYSTEMN (sic) TO MIRROR THE
>DISTURBANCE." Reading his article again I don't believe that he
>proves that al all

Good. This is all I want; look at the research. (By the way, did I use the word "prove" -- it's not in quotes above? If I did, I'm sure I meant it to be synonymous with "test" rather than "deductive proof"). So why doesn't my experiment "prove" this to your satisfaction? You say it's because:

> the subjects did have an awareness of
>a disturbance - they were told about it and through several test runs
>of the task knew that a disturbance was influncing the cursor - it is
>the case that they did now know WHAT is was (neither did Marken - it was
>random).

Two problems here: 1) they were not told about the disturbance -- though they could surely tell that it was present and 2) it was a sine wave, not random (though it works just as well with a filtered random disturbance).

But I don't see what either of there points (subject's being aware of the disturbance and the disturbance being random) has to do with the point of the paper: that there is no information in the cursor movement that can be (or is being) used by the subject to determine how to make their responses? I assume "information about the disturbance" means that there is some something about cursor movement (p(t)) that allows the subject to compute o(t) -- ie. o(t) = f(p(t)) -- so that o(t) = -kd(t). The experiment you read showed, I think, that there is no f() that could be turning p(t) into o(t) because you get virtually the same o(t) twice when you use the same d(t) but p(t) is totally different on each occassion.

>I would have refused to published this study since WHAT the subjects
>were told is not to be found in this description BUT having done
>this type of work I know that they had to be told something (unless all
>of the subjects had done this many times) BUT IT SEEMS OBVIOUS TO ME

>THAT THE SUBJECTS WHERE AWARE THAT A DISTURBANCE WAS INFLUENCING THE
>CURSOR POSITION ALONG WITH THEIR MOVEMENT OF THE GAME PADDLE. If this
>investigator believes otherwise let him present evidence that they
>were ignorant of a disturbance. He should has ASKED THEM or used THE TEST.

Does this study upset you, Chuck?

Why would the subject's being aware or not aware of the disturbace make any difference? The subject can certainly tell that there is a disturbance operating and I would have reported that fact if I (or anyone) could have seen why it would be germain. What is far more important is "did the subject have any idea about the nature if the function d(t)". In fact, one clever reviewer (Tom Bourbon, it was, before he fell into the PCT all the way) pointed out that, because I used sine waves, the subject could tell from their own hand movements what the disturbance function was. Tom suggested that the subjects could be producing the sinusoidal outputs from memory. This was a far-fetched, but at least pertinent, criticism of my study. I pointed out that the phase and frequency would have been impossible to predict on the repetition trials (which were non-continguous and the subject had no way of knowing when the d(t) on a particular trial was a repetition of the d(t) on a previous one-- but, still, I should have done the study with random disturbances. In fact, I have done the study (just the other day) with random disturbances and it works just the same.

So it looks like your only complaint about the study is that the subjects were aware of the disturbance; in fact, they would be aware of the distrubance whether I told them about it or not; they know that they have to do something to keep the cursor on target. I think it would be hard to keep the subject "unaware" of the existence of the disturbance. So I guess I have to ask how this awareness might affect the results -- ie, the fact that the correlation between output traces on two (very separate) trials with the same disturbance was typically .99+ while the correlation between p(t) (the only thing the subject could see on those trials) was typically less than .2, once as low as .0032?

Best Rick

Date: Sun Mar 28, 1993 9:22 pm PST
Subject: Video, second edition

[From Dag Forssell (930328 21.30)]

On March 11, Christine and I presented an introduction to PCT to a second Deming Users Group. We now have an edited video ready which we think is twice as educational as the first one was. We have finally settled on the equipment for editing and have learned a few things about what to do and not to do.

As before, I will mail the video with supplemental booklet to anyone who asks. I'll ask for $10 for the tape and booklet with additional for postage as follows:

Postage to	Surface, 4th class	Air
USA	$2.-	$3.-
Canada	$3.-	$4.-
Europe	$4.-	$9.-

Pacific $4.- $11.-

I will honor E-mail, fax or letter request and trust that you add up and send U.S. funds by snail mail. State surface or air.

Note: The video is 120 minutes, 1/2 inch VHS, NTSC (U.S. video signal). Overhead slides are shown in closeups.

Dag & Christine Forssell
23903 Via Flamenco
Valencia, Ca 91355-2808   Phone (805) 254-1195   Fax (805) 254-7956

Date: Sun Mar 28, 1993 10:02 pm PST
Subject: Monday's lookin' good

[From Rick Marken (930328.2100)]

Oy vay. Such a rough weekend. First I have to admit I'm wrong about infinite loop gain (it really does eliminate the disturbance); then my contrition is rewarded with my own arrogant posts being thrown right back in my face (graphically by Greg Williams (930328 - 2) and impression- istically by Ken Hacker [930328]). A weekend like this makes monday look good -- thanks fellas.

Ken Hacker [930328] says --

>[3] There are claims made on this net which I do, in fact, think are
>problematic (arrogant or hyperbolic). These include the following (if
>Mr. Marken wants more data, he should print out the notebooks and
>start counting -- I don't have time to do it for him):

```
>         *   All social science is wrong.
>         *   All behavioral science is wrong.
>         *   All social scientists and behavioral scientists
>             believe what is wrong.
>         *   All social scientists and behavioral scientists do not
>             know what they are doing with research.
>         *   Only PCT has something useful to say about how humans
>             regulate their behaviors.
>         *   Communication theory is wrong.
>         *   PCT equations (confirmed by PCT theorists) PROVE
>             that PCT is correct.
>         *   PCT is real science.  Other approaches to human behavior
>             are pseudo-scientific or "half-assed."
>         *   People who challenge PCT are misguided, ignorant,
>             and not yet fully developed intellectually.
```

Well, that's quite an impressive list. I have to assume that these state- ments really reflect your perception of what we have been saying on CSG-L. I, of course, would not have stated them this way -- these sentences don't evoke in me a perception equivalent to what I think we've been saying on CSG-L. But there are a lot of sentences here to dispute. Maybe we could make this a bit easier by picking ONE and seeing what we really said (or tried to say). I suggest we start with the forth one -- which seems to come fairly close to describing something I might have said. But I don't think I would have said that social and behavioral scientists "don't know what they are doing with research". I think they know very well what they are doing -- I just don't think they are learning what they think they are learning by doing this research. What I think I might have been trying to say was this:

"Virtually all social and behavioral science research is based on the assumption that stimulus inputs cause behavioral outputs -- o = f(i). This assumption can be seen in the prediliction for doing experimental research in which an independent (stimulus) variable is manipulated and its effect on a dependent variable (behavior) measured under controlled conditions. The goal of this research is to learn something about the characteristics of the system that turn inputs into outputs -- ie. the function, f(). If, however, the systems being studied are control systems, then this type of research reveals little or nothing about f()."

Is this a fair statement of what you consider to be a hyperbolic PCT claim?

Greg WIlliams (930328 - 2) --

>When you claim that all of psychology needs to be rebuilt from the
>ground up on a PCT foundation, I think that sounds like a claim for
>having a "better" trap to catch WHATEVER psychologists want to catch.

But that is only if they want to catch mice, not lint. What this "sounds like" depends on how well people are listening.

>To date, no one has demonstrated that after the
>rebuilding, the resultant edifice will actually suffice to meet
>psychologists' (including your own) goals for "understanding" or
>"predicting" (take your pick) high-level human behavior generally.

Hey, it's a risk. All we do is present evidence; whether anyone thinks it's worth taking the plunge, based on this evidence, is their decision.

>it seems to me (see below) that you claim that non-PCT behavioral scientists
>can NEVER understand/predict high-level purposeful behavior of individuals,
>yet PCT can... SOMEDAY, you believe.

Behavioral scientists who operate based on an input-output model of behavior will never understand/predict purposeful behavior at ANY level; I'll stick with that one.

> You say that the non-PCT linttrap cannot catch mice... but I think you'll
>have to admit that the PCT mousetrap has so far not caught any really big mice.

They look pretty big to me. But their size doesn't matter. The non-PCT lint trap has caught NO mice of any size. If people keep working with the lint trap (input output model) they'll NEVER catch any mice (purposeful behavior) -- big or small.

>Below are some examples (from Rick's posts during the first three
>weeks of March 1993) of what I believe some non-PCT behavioral
>scientists would consider "arrogant".

Some people consider the Reagan/Bush era "the wonder years". OK, I can be considered arrogant. Mea culpa.

I stand by the comments I made in the posts that you posted. Some may be construed as arrogant -- that's fine. Actually, I am trying to be a bit of a strong source of distrubance because I want to provoke public rebuttals (outputs). I think I have been fairly successful in provoking and maintaining the "information in perception" conversation. I believe this has been very worthwhile because this fact about control system operation (that output does not depend on information in perception about the disturbance) is of fundemental significance to understanding the operation of living control systems.

I am trying to provoke some reorganization -- the only way to really learn something. And I see some reorganization occuring -- though most of what I get is what I expect -- defense. But some of that defense may actually be substantive (like Allan Randall's last couple of excellent posts) and informative for observer's for whom PCT is NOT part of a controlled perception.

As far as how little PCT has to show compared to conventional behavioral science: could you give one example of one of the big accomplishments of the latter?

Thanks Rick

Date: Mon Mar 29, 1993 1:46 am PST
Subject: My Parenting Program

From Ed Ford (930328.2007) originally sent 1000

The following is a copy of an article which was published in this month's Journal of the National Center For Outcome Based Education. I've had a number of private contacts about the program and what I do, so I thought this would be helpful.

At the several schools where this course has been taught, it has had very exciting results. As you might guess, it is very much a process or hands-on course, full of skills building and very specific techniques. What was most exciting is how well the administrators, counselors, teachers, and parents worked together, and how proficient they had all become at the end of the course. Many are now teaching other parents and teachers within the school district using the same tools with which I taught them. I never dreamed how efficient this course would be, especially at enabling the participants to teach others the same skills and techniques I taught them. It is all very exciting.

Administrators of one of the districts has ordered copies of all the PCT books for their library. There seems to be a growing desire to understand more about PCT and what it is really all about.

Building an ODDM Family
by Edward E. Ford

Abstract

Focusing on school district parents, this article briefly describes a parenting program design specifically to fit the ODDM model, on making them partners and part of the ODDM process. Selected teachers and parents are taught to work with parents wanting help. It explains how perceptual control theory is the psychological basis for teaching children responsibility. Parents are taken through a three- step process. First they are taught the key to building a strong relationship so they can access their children when difficulties arise. Next

they're taught how to set standards and ask for choices and consequences, followed by intensive training on the skills of teaching responsible thinking to their children.

\- - - - - -

ODDM is a process for school improvement where the success of all children is the major goal or outcome. This process demands a knowledge of the latest research and the use of a psychological basis, the how and why of human behavior, as a way to increase effectiveness when dealing with children.

Parents should not only be knowledgeable of this process, but they should be a part of it. They should be dealing with their children in a way compatible with the theoretical basis of ODDM. This is especially true with regard to how they teach responsible thinking to their children at home, thus becoming equally effective with the ODDM program at school.

An Outcome Based Parenting Program has been developed which trains both parents and teachers how to work with parents who want help in dealing with their children. The most important element to children's success, namely their parents, can now become an integral part of the ODDM program. This means that parents and teachers can work together as partners, using the same method when dealing with their children.

More importantly, parents will have a better idea of the ODDM model because they will now be part of the process. Also, with a more complete understanding of the creative changes taking place in their child's school, they are more likely to be accepting and cooperative. The school district will now be perceived as a partnership, a place where children can succeed and parents are a welcome part of the what's happening rather than outsiders looking in.

Consequently, the parent's understanding of the psychological basis of ODDM and the basic strategy for teaching responsibility is critical to the success of children. The parenting program will provide both.

Psychological Basis For The Parenting Program

Perceptual Control Theory

The psychological basis for this parenting program is Perceptual Control Theory (PCT), a theoretical model of how people think, based on the creative thinking and writings of William T. Powers, author of Behavior: The Control Of Perception. It is also based on the research and work of the Control Systems Group, an organization including scientists, psychologists, sociologists, engineers, educators, economists, social workers, and those in business, all working on models, ideas, and practical applications pertaining to the development of this theory. (See selected references at end of article)

PCT states that the actions of people are not caused by some stimulus from either the environment or inside the human system. Rather what we do is caused by a comparison of what a person wants with what represents that want which is called feedback.

For example, if I wanted to drive 50 MPH, I would compare that goal with my speedometer, the feedback representing what I want. My actions, including my foot on the accelerator, would be my attempt to work at reconciling the difference between what I want with the reading I perceive on the speedometer.

If while writing this article I felt thirsty and wanted a glass of water, I would do a number of things, including getting up, getting a glass from the shelf, turning the faucet, filling the glass, and lifting the glass to my lips. The sensing of my thirst being quenched would be the feedback. All the various actions would be done without much thought to what I was doing but only to satisfy my desire to match the feedback to what I wanted.

In both the above examples, I never think about what "my actions" are because my attention is centered on the comparison of what I want with my present situation. While most people try to change what others are doing, control theory suggests dealing with what they're thinking, namely their various wants and how they perceive their present situation. The important point here is that whereas most teachers and therapists attempt to deal with the actions of their students or clients, the key to helping children is to teach them to deal effectively with what causes those behaviors, namely, how their children think.

With PCT as a basis, much of what we do with and for others can be made more efficient. For example, one thing PCT makes clear is the need for measurable goals and proper feedback when helping a child to think through the creation of an effective plan. Supposing a child decides to "study more at home". Obviously, the plan is vague.

Now supposing she changes that to a measurable goal such as "study 30 minutes a day".   Further, she wants to set up a chart that reflects her goal. The chart would look like this:

```
        50-|
minutes 40-|
spent   30-|- - - - x - - - -x- - - -x- - - goal
doing   20-|  x        x           x
activity10-|              x
            --|--x--|--|--|--|--x--|--|----days of week
              M  T  W  T  F  S  S  M  T
```

Note: dot marks are the x's, you connect x's with line.

Notice how the chart resembles PCT. The goal is the time she wants to spend studying. The dot marks are the times she actually spends every day or the feedback, and the line connecting the dots shows the historical progress. The distance between any dot and the goal gives a picture of how big the difference is between what she wants and her feedback. Thus, the chart actually is a visual picture of the PCT model in action.

Building Strong Relationships:

The Most Important Step

As the ODDM process does with school personnel, this parenting process of change involves the parents in a very creative and effective way of thinking and working with their children. Parents first have to go through the experience of learning what it is that helps them to gain access to dealing with their children reasonably and rationally, especially when differences arise.

Once they've begun to establish a closer relationship with their children, they then have to learn how to set standards. Finally, they have to learn the difficult art of teaching their children to think responsibly.

The first lesson PCT teaches is that you can't control another person. Since people create their own goals as well as their own perceptions, and since their actions are their attempt to reconcile the difference between their goals and their present situation, obviously any attempt to get parents and their children to change is going to be met with failure unless they're willing to work together toward a common goal.

What gives a teacher or a parent access to working with children is when the youngsters have a belief that who ever is working with them cares about them and believes in their ability to resolve their problems.

Thus, the most important step when teaching responsible thinking involves spending the kind of time which is going to create on the part of children a belief that someone cares about them. I call this quality time. This belief creates a willingness to not only resolve problems, but to do so with their parents. Over the past 25 years, I have dealt with couples and parents with children looking for help at building sufficient strength in their marriages and families such that they could resolve their problems in a reasonable and rational way. This quality time is the most effective program I've found.

From PCT, I learned that we create our perception of others from the mutual, interactive type of experiences people have with each other. The criteria for these experiences are as follows: first, what ever you do, you must be aware of each other; second, you must create the enjoyment, rather than passively watch the TV or movie; third, you must spend this time alone together with whom ever you are trying to build a relationship; and fourth, this activity should happen on a daily or regular basis.

Set Standards: Ask For Choices & Consequences

Making Responsibility Possible

The strengthening of the relationship then allows the teacher or parent to access the child more easily. It also creates a greater likelihood of the child's willingness to accept rules or standards. This leads to the second step in the parenting process.

Responsibility is the willingness and ability of people to follow standards and rules and ultimately to set their own, without infringing on the rights of others. Thus, parents must establish reasonable rules in the home and children should be taught to obey them.

In order for parents to deal with a child, they have to set specific and reasonable rules and standards that must be consistently applied over time and enforced fairly with each person. On those occasions when children are not willing to follow these standards or obey rules, the parents will then have some already established standards with which to help them deal with the child.

When children refuse to obey rules, the parent should then ask them to name the various choices they have and to explain the consequences that result from making those choices. The consequences should include the loss of the privilege which is related to the responsible choice the child refuses to make.

In the PCT model, standards and rules are the basis for how people make choices. Problems often arise when parents or children consider choices that are in conflict with the standards at home.

Teaching Responsible Thinking

A Skill-Learning Process

The third step in the process is the way we deal with children when differences arise, and this is the heart of teaching children responsible thinking. And probably the single most difficult aspect of this process is to overcome the desire to "preach" to children, to tell them "what they're doing wrong or what they should be doing".

One thing PCT teaches is that there is no sense trying to change those things over which you have no control. Human beings cannot be directly controlled, unless overwhelming force is used. And schools and homes are no exception. Thus, in the ODDM process, parents as well as teachers are taught that if you want to be effective with children, you have to get them to deal with their world, with what they want, and the consequences of what they are doing to get what they want.

The first thing to learn is to deal with children by asking questions, not by telling them what you think. The questioning process forces the parent or teacher to stay in the child's world. In PCT terms, you're dealing with the children as individually and internally-driven control systems, which is what they are. By asking the right questions, you're teaching children to deal with themselves in such a way as to satisfy their own internal goals.

And, you're teaching them to do this in the most effective way possible, which is to deal with the key elements of responsibility, namely, what they want, which is their goal; what the standards or rules are, which is what they base their choices on; and what their specific choices and actions have been.

Ultimately, you get to the most crucial area which is the responsibility questions: first, you get them to compare (again, using PCT) what they want to what they are doing, or what they're doing as compared to what the rule or standards are; and secondly, once they've accepted responsibility for their thinking in that area, then you find out if they are willing to set a goal to work at resolving their problem. In PCT terms, how strong is the internal signal that represents what they want as compared to other wants.

Now the children are ready to work out a plan to satisfy what they want with measurable goals and charts, which were discussed above.

As you can see, this program not only provides uniformity of standards in both the home and school and a way for parents to learn how to teach responsible thinking to their children, but it gives parents a real partnership in the ODDM process.

To Find Out More About

Perceptual Control Theory

A comprehensive checklist of articles and books on Perceptual Control Theory, updated annually, is available for $5 postpaid from Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328.

CSGnet, the electronic mail network for individuals interested in PCT as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. No sign-up or connect-time charges. The Bitnet address is CSG- L@uiucvmd. CSG-L@VMD.CSO.UIUC.EDU is the Internet address. To join, send note to network manager, Gary Cziko, at the Internt address G-CZIKO@UIUC.EDU or Phone 217 333-8527.

A partical list of writings on PCT is as follows:

Powers, William T. (1973) Behavior: The Control Of Perception. Hawthorne, NY: Aldine DeGruyter - Basic Text

Robertson, R.J. and Powers, W.T. (Eds). (1990) Introduction To Modern Psychology: The Control Theory View. Gravel Switch, KY: CSG Book. - College Level Text

Powers, W.T. (1989). Living Control Systems: Selected papers. Gravel Switch, KY: CSG Books. Previously published papers, 1960-1988.

Powers, W.T. (1992). Living Control Systems II: Selected Papers. Gravel Swtich, KY: CSG Books. Previously unpublished Papers, 1959-1990.

Marken, Richard S. (Ed.) (1990) Purposeful Behavior: The Control Theory Approach. American Behavioral Scientist. (Thousand Oaks, CA: Sage Publications) 11 articles on control theory.

Ford, Edward E. (1989). Freedom From Stress, Scottsdale, AZ: Brandt Publishing. A self-help book. PCT in a counseling framework.

Runkel, Philip J. (1990) Casting Nets and Testing Specimens. New York: Praeger. When statistics are appropriate; when models are required.

McPhail, Clark (1990). The Myth Of The Madding Crowd. New York: Aldine de Gruyter. Introduces Control Theory To Explain Group Behavior.

Richardson, George P. (1991). Feedback thought in social science and systems theory. Philadelphia:University of Pennsylvania Press.

Cziko, Gary A. (1992). Purposeful behavior as the control of perception: Implications for educational research. Educational Researcher, 21 (9), 10-18.

Petrie, Hugh G. (1981) The Dilemma Of Inquiry and Learning. Chicago: University of Chicago Press.

About The Author

Ed Ford, MSW, has authored eight books, his two latest, Love Guaranteed and Freedom From Stress, reflect PCT concepts. Besides his family counseling practice, he has worked and consulted in schools, social services centers, correctional, residential, mental health facilities, and with businesses. He has taught at various universities in the Phoenix area and for over 10 years, was a faculty and certifying member of the Institute for Reality Therapy. He is a founding member and past president of the Control Systems Group. He and Hester, married 43 years, live in Phoenix, have eight children, and a growing number of grandchildren.

Ed can be reached by contacting Ed Ford & Associates, 10209 N. 56th St., Scottsdale, Arizona 85253. Phone 602 991- 4860.

Date: Mon Mar 29, 1993 4:54 am PST
Subject: Challenge 2

[From Bill Powers (930328.0930)]   Greg Williams (930327.1945 EST) ---

In your 930326-2) you said

>...show me you understand my behavior and what it is being used
>to control. But first, YOU tell ME what a passing grade at such
>"understanding" would consist in.

I did respond to your request that I tell you what passing grade we should use. Your request was a behavior of yours. When you read my answer, you replied "Fine." Was that an indication that my answer gave you the information you wanted? If so, I would explain that behavior by saying that you issued the request for information because you wanted to see my answer (explanation 1), and that at another level, my answer met whatever criteria you had set for an acceptable answer (explanation 2). Is either of those explanations 99% accurate?

As to the main part of the challenge statement, I asked you what effect you intended to produce by issuing the challenge. Your reply was

>I don't know. Perhaps you have a way to get me to say that I
>"realize" some particular effect is the one I intended? The
>question then is: how do we know that was REALLY the effect I
>intended (yesterday), and not just some sort of after-the-fact
>fiction than seems to make sense? How do we avoid "just-so" opining?

Is the challenge still in effect, or did it end with your issuing of it? If it's still in effect, then unless you consider that I've successfully met it, there must still be some sort of error signal in existence, and the reference signal must still be in effect, right now. It shouldn't be necessary to remember yesterday's motive. Even if the nature of the reference signal changes in your mind, it's only the goal as it exists right now that you or I can try to meet, because that is the one you will compare the results against.

Your words indicate that if I suggested what was "really" your intent, you would tend to reject it as guesswork. I agree. I'm not a psychoanalyst who will interpret your own words to yourself in some surprising way. If there is some hoped-for result in your issuing the challenge, it is still here and in effect. Otherwise, we're finished, aren't we? You would have withdrawn the challenge.

Assuming you haven't withdrawn it, perhaps we can work out the intended effects by considering it at different levels. I presume, for instance, that by maintaining the challenge, you expect some sign that I am accepting it (or not). Was (and is) this one of the immediate goals in issuing the challenge? That is, did you want to get back an indication of my acceptance or rejection of it? It would be a reasonable interpretation on my part to assume that you did (and do) want to know whether I (then and now) accept or reject it, and I might even guess that your preference was and is that I accept it. But you are the only one who knows the truth of the matter. I'm out here. You're in there where you can see what's going on right now.

So, do you still want to know whether I still accept the challenge? And is your preference right now that I do or do not accept it?

Best,   Bill P.

Date: Mon Mar 29, 1993 4:57 am PST
Subject: Mary on communications loops

[Mary Powers 9303.28]   Ken Hacker

The problem I see with with the external feedback loop between people interacting is that what looks like one loop is really two.

Here is how I visualize interaction in your terms:

```
          feedfoward
Control  --------------> Control
          feedback
System 1 <-------------- System 2
```

I think of it like this

```
Control  .-.-.-.-.-.-.-.> Control
System 1 <-.-.-.-.-.-.-.- System 2
```

which dissects out as

```
         output
Control  ---------------> Control
         input
         <--------------
         output
System 1 <..............  System 2
         input
         ...............>
```

System 1 outputs and inputs being dashes, and System 2 being dots.

The point is that the first diagram makes it look as though 1 & 2, in conversation or whatever, are a single system. I think they are two seperate systems using the same external events in two different ways - much like the rubber band demo, in which both parties are using the same rubber band, one to keep the knot on the dot, and the other to move the first one's hand to a mark.

Each system should also have a loop like this:

```
CS 1  ---*----------->
         *
          *
         *
      <--*-----------
```

Now when CS 1 speaks, his own output is fed back, which is how he knows what he's saying. He also perceives the answer from CS 2, which is also feedback from what he's saying. But CS 1's feedback is also CS 2's output. This makes it look like CS 2 is "giving" feedback to CS 1. I think not. CS 2's output is input to CS 1, and part of that input is CS 1's own feedback, a consequence of his own actions.

Feedforward has nothing to do with it.

I agree with you and Wiener that social interaction is indeed very interesting. I also think it can be dealt with in PCT terms, using the PCT model of individuals.

I'm glad Klaus Krippendorf thinks control theory explains something. But does he understand that movements are the ultimate means by which control systems at every level control perception?

Mary Powers

Date: Mon Mar 29, 1993 6:34 am PST
Subject: insulting behavior

[From: Bruce Nevin (Mon 930329 08:42:58)]

I can't contribute much this week. I played fast and loose with my deadlines a bit too much last week. Just one comment regarding ambiguity in language.

[Rick Marken (930326.1000)]

> Well, if I can't tell which meaning the speaker meant (intended) to
> communicate, doesn't that satisfy my claim that "you can't tell what a
> person is doing by looking at what they are doing"? Assuming that the
> speaker intends to produce one or the two possible meanings (he may mean
> to produce neither) all I can tell from the output (the sentence "Flying
> planes can be dangerous") is that one of the two meanings was intended.
> So I can't tell (meaning, know for sure) what the person intended. Your
> arguments about ambiguity seems like the points I should be making.
> Remember, you're the one arguing that you CAN tell what a person is
> doing by just looking -- when what they are doing is language.

I did not say that you can tell everything that they are doing. For example, I disclaimed that you could know why they were saying what they said. Their motivation might be quite indirectly related to the ostensive meaning of the utterance--they might be devious, manipulative, or outright lying, f'rinstance. They might even be concealing some of their motivations from their own awareness (happens all the time).

I also noted that their motivations, whatever they might be, are part of the meaning of the utterance. "Meaning" is far richer, and differs far more from one individual to another, than that aspect of meaning that is socially available in the structure of the utterance, which I have called by Harris' term of linguistic information.

If an utterance is ambiguous, it is because the speaker is presuming that the hearer understands that certain unspoken parts of the utterance are in the mutually understood context of the part that is spoken. "Flying planes can be dangerous" is ambiguous only in isolation.

"Hi there, Rick!" is ambiguous if the speaker is concealing his desire to insult you. He must be concealing it from you, or else he would have pronounced the P before rick unmistakeably loud and clear. Like the man I met in high school who got some sort of passive-aggressive (sic) kick out of closing certain telephone conversations with the phrase "f!ck you very much!", using the intonation of "thank you very much" and not that of "f!ck you!" There are other creative uses of ambiguity, such as equivocation and prevarication.

Ambiguity is a degeneracy that is introduced by reductions. If the context enabling the reduction is not recognized as such by the hearer (as the speaker is presuming it to be), the hearer may choose a branch of an ambiguity other than that intended by the speaker. The intended linguistic information is in the utterance. Unintended linguistic information is also there.

I don't see how this constitutes linguistic behaviorism. Please try to avoid insults. Or did you mean something different from psychological behaviorism? :-)

A second point: without language, science is impossible. Also mathematics. No language, no math and no science. The development and conduct of mathematics and science depend upon linguistic information. I won't try to prove this to you by email. For supporting argument and demonstration, I refer you to _The Form of Information in Science_ and _A Theory of Language and Information_.

Bruce bn@bbn.com

Date: Mon Mar 29, 1993 6:35 am PST
Subject: Ah, Monday!

From Greg Williams (930329) Rick Marken (930328.2100)

>As far as how little PCT has to show compared to conventional
>behavioral science: could you give one example of one of the
>big accomplishments of the latter?

As you (correctly) say, conventional behavioral science CANNOT catch mice (explain purposive behavior), while PCT HAS caught some (at least little) mice. Likewise, conventional behavioral scientists have caught some lint (population measures) -- perceived size depending on one's goals; some CBS's appear to think the size pretty big, but you appear to think it completely unimportant -- which PCT will (perhaps someday), I predict (based on the inefficiency of using the method of individual models in physics and other hard science), be able to do only in very special cases and with great difficulty. So the situation is almost symmetrical, in my view. CBS CANNOT explain purposive behavior AT ALL and PCT CANNOT (today) and WILL NOT AS A PRACTICAL MATTER be used to generate population measures, in general. OK, so it looks like somebody who wants to catch lint should use CBS and someone who wants to catch mice (at least small ones) should use PCT. But then you call for a complete rebuilding of psychology based on PCT; I guess the lint-catchers are out of luck. I haven't heard CBSers saying (arrogantly) that there is NO place in psychology for explaining purposeful behavior.

As ever, Greg

Date: Mon Mar 29, 1993 6:51 am PST
Subject: Prediction vs. Understanding; Arrogance, pooey.

[From Bill Powers (930329.0700)]

Ken Hacker, Greg Williams (930328 etc.) --

I'd like to clarify some things about PCT and prediction. I've made some contradictory statements here: PCT is about understanding, not prediction; PCT can predict behavior better than other approaches can. This sounds a lot like talking out of both sides of my mouth, if you haven't been paying close attention.

The basic PCT model says that the actions a person performs arise from two independent variables: external disturbances of a controlled variable, and the setting of the reference signal. In order to describe behavior using a PCT model, it's necessary to deduce what the controlled variable is, how the reference signal is set, and what the external disturbances are. In situations where this has actually been done satisfactorily, the PCT model predicts behavior very accurately. The accuracy of the prediction is an indication of how well the requirements have been satisfied because, according to the model, accurate prediction would not be possible if they weren't satisfied.

To establish the circumstances where such predictions can succeed, it's necessary to persuade the human participant to agree to maintain a constant goal during an experiment (and actually to do so), and for the experimenter to know the quantitative sum of all important disturbances acting on the controlled variable. Either of these requirements can be difficult to meet under field conditions where the experimenter can't predict all disturbances and has not deduced the nature of the relevant controlled variable and its reference level. This is very much like the S-R theorist's excuse that one can't control all stimuli. So what's the difference?

One difference is that under PCT, the point of the analysis is not to predict the visible (audible) actions of the person, but to predict relationships between actions and disturbances that maintain controlled variables stable against the disturbances. In order to predict actions, EVEN WHEN ONE'S MODEL OF THE CONTROL SYSTEM IS PERFECT, it is necessary to predict disturbances -- that is, to predict future states of the environment. I may have a terrific model of your umbrella-using control system for keeping dry, but I can't predict whether it's going to rain, nor where your umbrella is going to be when you need it. So when PCTers claim to be able to predict behavior, it must be understood that the behavior is almost totally dependent on the environment. Even knowing the complete structure of a person's control systems at all levels would not permit predicting behavior unless the disturbances that are going to occur are known.

However, a correct PCT model would be able to predict the states of certain environmental variables with great accuracy even without any ability to predict disturbances. Those are the controlled variables that correspond closely to controlled perceptions. With a one-level model, you can measure the characteristics of the control system and derive a model that will predict the controlled variable quite well as long as the reference signal remains constant. With a two-level model, you can predict the second-level controlled variable too, over a somewhat longer period of time, until the second-level reference signal changes. There is reason to think that as the model is applied to higher and higher levels, the reference signals and hence the controlled variables remain constant for longer and longer times, so that the prediction of the controlled variable can remain correct for longer times.

However, it continues to be true that behavior -- output -- is only as predictable as disturbances. For higher-level variables, disturbances tend to be more predictable, too, because the effective disturbance is only what can get past all the lower- level systems and disturb a higher-level variable. So it is possible to predict higher-level behavior for longer periods of time. This observation, I believe, is common to most studies of human characteristics, from whatever viewpoint.

This is why I put the emphasis on understanding more than on prediction, even though prediction can be very impressive in laboratory experiments. The more completely you understand the organization of a person's control systems, the more clearly you will know what is in principle predictable and what is not, under realistic circumstances. You will also be able to understand the relationship between actions and environmental processes when unanticipated disturbances arise; in effect, every action that arises when a new disturbance occurs tells you more about what the person is controlling for. The actions bring disturbances to your attention; the disturbances test your concept of what the person is controlling for. The concept of a control system explains what you are seeing even when you didn't predict it, and couldn't possibly have predicted it. This helps you to see how to interact with other people, and to understand how they interact with others.

One of the frustrating aspects of our conversations on PCT and other approaches is that right now there is only a tiny handful of people actually trying to develop ways of exploring behavior using the principles and methods of PCT. The rest of you guys, however justified your objections and criticisms and comments, are sitting on your butts and defending your personal status quo and not doing a lick of experimental work to extend the scope of PCT -- except from your armchairs.

Every person who has ever thought of a new idea and claimed that it is better than old ideas has been accused of arrogance, of not showing proper respect for the great achievements of others, of making overblown claims, of not understanding what is good about the old ways. Maybe those accusations are correct, but so what, if the new idea is in fact better than the old ones? People who love the status quo have always got their knickers in a twist over upstarts who dare to get cocky about their own achievements and go around pointing out the Emperor's lack of modest attire. That's the problem of the complainers. It has always been their problem. They're responsible for it. They can very easily make it go away: get into the trenches and start showing that the new idea is wrong, or get into the trenches and start working to make it even better. All the wailing and gnashing of teeth and complaints about violating the comfortable standards of polite deference to the ideas of others is, in my mind, the sound of oxen being gored and signifies mostly that the point is being missed.

If I'm arrogant or Rick is arrogant, we have a right. We're doing the work, not watching a TV show.

With sweetness and light,   Bill P.

Date: Mon Mar 29, 1993 6:59 am PST
Subject: Re: Where is the proof? Information of disturbance

***************** FROM CHUCK TUCKER 930329 ***************

So, the subjects, contrary to your claim that they didn't, DID have information that a disturbance was going to influence the cursor and randonness was invloved; you stated: "The phase and frequency of the disturbance were determined randomly for each run. The same disturbance was repeated on pairs of nonconsective runs (63, in Reader)." That the actual form of the disturbance was not known by either you or the subjects (BTW, your description of the subjects leaves open the possibility that one is you - a real problem if so) BUT the point is that you previous stated that the subject had NO INFORMATION about a disturbance; that simply is not true if your own description (as poor as it is) is an accurate representation of what was done in the study.

You can translate all the words into other symbols (sy=symbols) and put ='s, +'s, or any other sign you care to between them and they still turn out for me to be words; I still have to read them to have some idea of what you might be stating; actually I find your use of such symbols a pretention on your part to appear to be systematic and scientific when your research does not have such features for me.

YES, this study (and all others that I have seen you do) does bother me. It does so because it is so poorly done that I would not consider it worthy of inclusion in a list of scientific work. The claims that you make for such work are usually not supported for me by your study. You never indicate to the reader WHAT you have told the subjects to do; this for me is basic in all research - what procedures were used in the study!!! (e.g., I just found out in your answer to my post that you did not TELL the subjects that there was a disturbance - did you use deception in this study? - if you did this creates not only empirical problems for me it creates some serious ethical ones - I hope you aren't lying to your subjects in your research !?!). I suggest that you begin to consider language and words more seriously in your research.

BTW, I found some more Markenisms that were not mentioned by Greg but may have served as data for Ken:

> Rick (930307.1500)
>
> It's why we say PCT is revolutionary. I know that it seems impossible that all the old revered theories in psychology are invalidated by the work of a nice engineer from Chicago who doesn't even have a PhD in psychology - but that's the fact Jack."
>
> Rick (930312.1300) to Randall (930312.1200)
>
> "There is no information about the disturbance in the stumulus. Goodby behavioral science as usual, hello looney bin."

Rick (930315.0900)

"Actually this is not intentional; Bill P. is just a truly wonderful person and I am a schmuk."

I have a sense that you and I don't use a similar criterion to make a judgment on whether a statement is categorized as "hyperbolic" but I do judge the first two statements that way; the last statement is just a half-truth.

Best regards, Chuck

Date: Mon Mar 29, 1993 7:03 am PST
Subject: Re: Challenge 2

From Greg Williams (930329 - 2) Bill Powers (930328.0930)

>In your 930326-2) you said

>>...show me you understand my behavior and what it is being used
>>to control. But first, YOU tell ME what a passing grade at such
>>"understanding" would consist in.

>I did respond to your request that I tell you what passing grade
>we should use. Your request was a behavior of yours. When you
>read my answer, you replied "Fine." Was that an indication that
>my answer gave you the information you wanted? If so, I would
>explain that behavior by saying that you issued the request for
>information because you wanted to see my answer (explanation 1),
>and that at another level, my answer met whatever criteria you
>had set for an acceptable answer (explanation 2). Is either of
>those explanations 99% accurate?

I don't know. I don't have direct access to my motivational control structure, so I can't say what explanation is "99% accurate." I can CONJECTURE that both of your explanations (and a lot of other possible explanations) make some sense. But I think "from the outside" we'd need to apply disturbances to see which error signal(s) I was REALLY trying to keep minimized. As they say, hindsight tends to appear 20/20, even if it really (on the basis of more solid evidence) isn't.

>As to the main part of the challenge statement, I asked you what
>effect you intended to produce by issuing the challenge. Your reply was

>>I don't know. Perhaps you have a way to get me to say that I
>>"realize" some particular effect is the one I intended? The
>>question then is: how do we know that was REALLY the effect I
>>intended (yesterday), and not just some sort of after-the-fact
>>fiction than seems to make sense? How do we avoid "just-so"
>>opining?

>Is the challenge still in effect, or did it end with your issuing of it?

Still in effect.

>If it's still in effect, then unless you consider that I've

>successfully met it, there must still be some sort of error signal in
>existence, and the reference signal must still be in effect, right now.

Yes, but not necessarily the same one(s) in effect when I originally issued the challenge.

>It shouldn't be necessary to remember yesterday's motive. Even if
>the nature of the reference signal changes in your mind, it's only the
>goal as it exists right now that you or I can try to meet, because
>that is the one you will compare the results against.

I thought you were going to explain why I issued the challenge in the first place, not why I continue to say you haven't succeeded in meeting it.

>Your words indicate that if I suggested what was "really" your
>intent, you would tend to reject it as guesswork. I agree. I'm
>not a psychoanalyst who will interpret your own words to yourself
>in some surprising way. If there is some hoped-for result in your
>issuing the challenge, it is still here and in effect. Otherwise,
>we're finished, aren't we? You would have withdrawn the challenge.

You are overlooking the possibility that an output can remain (at least look) the same while the intended outcome is changing. My motivations could be changing continuously for not withdrawing the challenge. So what about an explanation for my originally issuing the challenge?

>Assuming you haven't withdrawn it, perhaps we can work out the
>intended effects by considering it at different levels. I
>presume, for instance, that by maintaining the challenge, you
>expect some sign that I am accepting it (or not). Was (and is)
>this one of the immediate goals in issuing the challenge?

Again, I don't know. It sounds reasonable, but how would I gather any additional evidence beyond (my own) hearsay?

>That is, did you want to get back an indication of my acceptance or
>rejection of it?

It is possible. I'm not sure.

>It would be a reasonable interpretation on my part to assume that you
>did (and do) want to know whether I (then and now) accept or reject
>it, and I might even guess that your preference was and is that I
>accept it. But you are the only one who knows the truth of the matter.

You have quite a bit more confidence in my "truth"-seeing than I do. How do I KNOW -- not just GUESS -- what motivated me two days ago?

>So, do you still want to know whether I still accept the
>challenge? And is your preference right now that I do or do not accept it?

Right now, I CLAIM that I want to know and that I prefer that you accept it. So, what about explaining my original challenging?

P.S. Are you going to address "understanding" vs. prediction?
_____

As ever, Greg

Date: Mon Mar 29, 1993 10:00 am PST
Subject: More Monday sweetness and light

From Greg Williams (930329 - 3) Bill Powers (930329.0700)

>I'd like to clarify some things about PCT and prediction.

That you did. Thanks very much.

>To establish the circumstances where such predictions can
>succeed, it's necessary to persuade the human participant to
>agree to maintain a constant goal during an experiment (and
>actually to do so), and for the experimenter to know the
>quantitative sum of all important disturbances acting on the
>controlled variable. Either of these requirements can be
>difficult to meet under field conditions where the experimenter
>can't predict all disturbances and has not deduced the nature of
>the relevant controlled variable and its reference level.

Then how can behavioral scientists interested in predicting human behavior in the field (especially if they are interested in predicting population measures) be expected to get very excited about PCT? The requirements for accurate predictions are very stringent and in many cases simply impossible to meet -- yet the (blind to explanation) statistical approach can sometimes produce precise predictions of population measures.

>One of the frustrating aspects of our conversations on PCT and
>other approaches is that right now there is only a tiny handful
>of people actually trying to develop ways of exploring behavior
>using the principles and methods of PCT. The rest of you guys,
>however justified your objections and criticisms and comments,
>are sitting on your butts and defending your personal status quo
>and not doing a lick of experimental work to extend the scope of
>PCT -- except from your armchairs.

I think I'm sitting on my butt and trying to test the limits of PCT notions as strinently as possible, through critique and discussion. Of course, it is only my opinion that tough criticism is essential in science -- just as essential as doing experimental work.

>Every person who has ever thought of a new idea and claimed that it is
>better than old ideas has been accused of arrogance, of not showing proper
>respect for the great achievements of others, of making overblown claims,
>of not understanding what is good about the old ways.

I doubt that. I don't recall reading of accusations of arrogance directed toward Einstein, though I've read much on his life and work (such as Clark's biography). I suspect it is true that MANY creators of new ideas have been called arrogant. I also suspect that acceptance of their good ideas was slowed by the accusations, so why give an opportunity for such accusations to be made?

Maybe those accusations are correct, but so what, if the new idea is in fact better than the old ones?

The so what, as noted above, is that acceptance will probably come slower if you're perceived as arrogant. And that results in at least one result which you obviously don't like: fewer investigators working on your new idea for a longer time.

>People who love the status quo have always got their knickers in a twist
>over upstarts who dare to get cocky about their own achievements and go
>around pointing out the Emperor's lack of modest attire. That's the problem
>of the complainers. It has always been their problem. They're responsible
>for it. They can very easily make it go away: get into the trenches and
>start showing that the new idea is wrong, or get into the trenches and
>start working to make it even better.

I consider myself in the trenches working to make the new idea better.

>All the wailing and gnashing of teeth and complaints about violating the
>comfortable standards of polite deference to the ideas of others is, in
>my mind, the sound of oxen being gored and signifies mostly that the
>point is being missed.

Maybe I'm missing some points, but one I am NOT missing is that scientific ideas deserve the closest scrutiny and the most extreme criticism.

>If I'm arrogant or Rick is arrogant, we have a right. We're doing the work,
>not watching a TV show.

I don't see how it follows that "doing the work" gives someone a right to be arrogant. Nevertheless, if you exercise that "right," I predict that your frustration about few folks contributing to PCT will last for a longer time than if you yourself moderate your own wailing and gnashing of teeth and get back to modeling and experiments. But I'm not about to arrogantly suggest that you better spend less time arguing on the net and more time in the lab. That choice is yours.

As ever, Greg

Date: Mon Mar 29, 1993 10:20 am PST
Subject: Movement

[From: Dennis Delprato (930329)] Mary Powers 9303.28

>movements are the ultimate means by which control systems at every level
>control perception[.]

You said that you wonder if someone understands the above. I suspect that your point is frequently overlooked, both in and out of PCT literature and thinking.

I'll leave it to the reader to ponder some consequences of failing to take into account your point.

Date: Mon Mar 29, 1993 10:30 am PST
Subject: Reply to Tucker's Review

[From Rick Marken (930329.0900)] CHUCK TUCKER (930329) --

Gee Chuck. I feel like I just got read the riot act by Innocent XII.

>So, the subjects, contrary to your claim that they didn't, DID have
>information that a disturbance was going to influence the cursor

Bill P. suggested that this might be the misconception under which Martin and Allan were laboring. I never said that subjects had no information THAT a disturbance was going to influence the cursor; the debate is about whether there is information in the movements of the cursor itself that can be used to determine how to move the handle so that handle movements compensate for the disturbance.

>(BTW, your description of the subjects leaves open
>the possibility that one is you - a real problem if so)

Why?

>BUT the point is that you previous stated that the subject had NO
>INFORMATION about a disturbance;

The ambiguity here can be helped by the mathematical notation that you seem to dislike. When I say "disturbance" I mean the time series, d(t), of values that is added to the cursor. I am claiming that there is no information about the values of this time series in the time series of cursor movements to which the values of the disturbance time series contribute. This is amazing because the time series of cursor movements is all that the subject can see.

>that simply is not true if your own description (as poor
>as it is) is an accurate representation of what was done in the study.

I think the meaning of my claim was clear from the context of the paper; but apparently you think the fact that the subject's knew THAT there was a disturbance is enough to disprove my claim that there is no information in perception that will allow the subject to make the appropriate outputs. I hope you will excuse me for not being convinced by your disproof.

>YES, this study (and all others that I have seen you do) does bother
>me. It does so because it is so poorly done that I would not consider
>it worthy of inclusion in a list of scientific work.

So you are offended by the poor quality of the study, not the results. Perhaps you could do a high quality version of this experiment (or describe what an acceptable version of the experiment might be) so that I can see the right way to do it.

>The claims that you make for such work are usually not supported for me by your
>study.

Well, could you suggest a study that might convince you, one way or the other, on this "information in perception" question? Of course, if you already know the answer then I don't need to waste any more time on these sloppy, biased experiments.

>You never indicate to the reader WHAT you have told the subjects
>to do; this for me is basic in all research - what procedures were used
>in the study!!! (e.g., I just found out in your answer to my post that
>you did not TELL the subjects that there was a disturbance - did you
>use deception in this study? - if you did this creates not only empirical
>problems for me it creates some serious ethical ones - I hope you aren't
>lying to your subjects in your research !?!). I suggest that you begin to
>consider language and words more seriously in your research.

I don't know how to "consider" it? Would you be happier if I did the same study with a white rat -- then language would not be an issue, would it?

My evaluation of the results of this (and most other studies) is ultimately based on matching the behavior of the subjects to that of a control model in the same situation. As I mention at the end of the paper, the result I consistently obtained with people (perfect correlation between outputs on two occasions, very low correlations between the inputs on the same occasions) was also found for a working control model (which didn't talk at all and which had no knowledge "that" a disturbance was influencing its input). Given the fact that people and model behaved in exactly the same way, I conclude that the characteristics of the study that you mention as being its shortcomings (how I instructed the subjects, their probable knowledge that there was a distrubance, the fact that I was one of the subjects) are irrelevant to the results. If you think they are relevant, I think it would be helpful if you would explain why they are and how the results would be different if these "shortcomings" were changed.

Markenisms:

>I have a sense that you and I don't use a similar criterion to make a
>judgment on whether a statement is categorized as "hyperbolic" but I
>do judge the first two statements that way;

Here they are again:

> It's why we say PCT is revolutionary. I know that it seems impossible that all the old revered theories in psychology are invalidated by the work of a noce engineer from Chicago who doesn't even have a PhD in psychology - but that's the fact Jack."
>
> "There is no information about the disturbance in the stumulus. Goodby behavioal science as usual, hello looney bin."

I agree that the first one could be considered hyperbolic; I stand by it, though; at least when "old revered theories" is taken as I intended it -- as working model implementations of these theories. The second statement is not hyperbolic at all -- it's just true. Perhaps you misread the "looney bin" comment. What I meant was that people (like me) who claim that there is no information in perception (which is what the quoted comment implies) would be considered (by conventional behavioral scientists) as candidates for the looney bin. I know that the claim that there is no information in perception is a pretty strong one; it goes against the most basic assumption of conventional behavioral science. I'm sure it must seems like mysticism. Why do you think I'm leaning on it so strongly. As Bill P. has said, it's a pretty obvious point (because p = o + d -- sorry about the symbols); but, as you can see, obvious or not, there is VERY strong resistance to this fundemental fact about the operation of a control system -- as evidenced so marvelously by your posts.

Greg Williams (930329)--

>But then you call for a complete rebuilding of psychology based on PCT; I guess
>the lint-catchers are out of luck. I haven't heard CBSers saying (arrogantly)
>that there is NO place in psychology for explaining purposeful behavior.

Whenever I have suggested that conventional behavioral science (CBS) will have to be rebuilt, I always meant to imply that this is only true if behavioral scientists want to understand purposeful behavior (catch mice). No rebuilding is necessary if CBS is only interested in estimating population parameters or input -- output relationships from sample data (catch lint).

Best Rick

Date: Mon Mar 29, 1993 11:14 am PST
Subject: Ken Hacker's !s and ?s

[From Bill Powers (930329.1000 MST)] Ken Hacker (930328) --

A reasonable set of comments and questions. I'll try to deal with some of them.

Your list of our criticisms of conventional behavior science is generally on the mark, but without any reasons being supplied they sound more extreme than they are.

Consider:

* All social science is wrong.

* All behavioral science is wrong.

* All social scientists and behavioral scientists believe what is wrong.

* All social scientists and behavioral scientists do not know what they are doing with research.

In the first place, I always try to put somewhere in the discussion the proviso "if the PCT model actually explains behavior...", which make this somewhat less arbitrary and ad hominum.

In the second place, it makes a difference just what claim of wrongness is made. Wrong about what? The fundamental disparity between PCT and conventional theories of behavior is the treatment of cause and effect. Under PCT, the actions of an organism are not caused by stimuli with the organism simply mediating between cause and effect. In the PCT model, but not in any conventional model, there is a controlled variable between the apparent cause and its apparent effect as usually observed. The behavior is directed toward control of the controlled variable; it is not simply a response to the stimulus or a dependent variable being determined by an independent variable. As a result, all theories in the conventional sciences that rely on some cause-effect or IV-DV paradigm, including methods of data analysis that assume such a relationship, are fundamentally contradicted by PCT. It isn't necessary to investigate every detail of any theory that is proposed; all that is needed is to see whether the internal validity of the theory depends on assuming that behavior is a dependent variable and inputs to the organism are independent variables. If that is true, then no more need be known: the theory is contradicted by PCT and in terms of PCT the conclusions are false.

I leave it to you to decide how much of the social and behavioral sciences would remain intact if it were to prove true that all behavior is organized around controlled variables, and none is open-loop.

* Only PCT has something useful to say about how humans regulate their behaviors.

As stated, this is contradictory to PCT on the face of it. Human beings, under PCT, do not regulate their behaviors at all, as least as behavior is conventionally understood. PCT says nothing at all about the regulation of behavior; it is concerned with the regulation of perceptions, of inputs. Behavior, or action, becomes what it must become to prevent disturbances from having an important effect on the controlled inputs. So it is true that only PCT has something useful to say about the regulation of input, if that is what is meant by behavior. But that is not what conventional sciences have meant by behavior.

* Communication theory is wrong.

Communication theory, as a mathematical treatment of certain ways of representing variables and signals, is completely right in terms of its mathematical operations. Where I have difficulties with it is in the initial definitions and assumptions, which do not seem to me to have any necessity to them. But that's me, not PCT.

* PCT equations (confirmed by PCT theorists) PROVE that PCT is correct.

Fortunately, confirmation of the PCT equations by comparison with experiment is done by public methods easily replicable by anyone, with high reliability. No faith or special knowledge is required. The validity of these equations is put at risk every time they are used to fit a model to real behavior. The equations don't prove that PCT is correct; as far as they go, they ARE PCT.

* PCT is real science. Other approaches to human behavior are pseudo-scientific or "half-assed."

PCT is real science because it risks everything with every application to data. It isn't simply assumed to be correct and twisted to make the data seem to fit it. It can't live with serious exceptions or counterexamples. Its premises are themselves testable through experiment, and are tested every time they are used.

I would accept as real science any other approach to behavior that had the same characteristics. We're not talking about control theory here: we're discussing what science is about.

> * People who challenge PCT are misguided, ignorant, and not yet fully developed intellectually.

Depends on how they go about challenging it. I see no shame in being any of the things you mention -- is there anyone who doesn't fit the description? I object to only one kind of challenge to PCT: the kind that is made without knowing what PCT is about, and is based only on a difference in conclusions.

The questions:

> a. The scope of PCT. Is it limited to describing and explaining physical behaviors of human beings such as motor actions?

In that all visible behaviors are physical and entail motor actions, including debating about quantum physics and women's rights as well as tracking behavior, yes. One could say, given a few precautions, that there are "really" only four kinds of behavior: push, pull, twist, and squeeze. All the rest is controlled perception, consequences of behavior. And even those four are controlled perceptions. PCT can be applied to behavior at any level of organization you please. HPCT is meant to apply at all levels at once.

> b. If there is no information inside of input, what prevents one from asserting behavioristic-like internal responses to internal stimuli or a kind of reverse behaviorism?

Nothing. You can propose that behavior is caused by the phases of the moon if you like, but you'll have to defend the proposal. See my earlier post to Rick on this subject. I'm not going to generalize about information and input any more until I'm sure what's being claimed or asserted.

> c. What happens to Ashby's Law of Requisite Variety?

I don't use it, although it's implicit in control theory. Variety is defined by Ashby as follows:

"The word _variety_, in relation to a set of distinguishable elements, will be used to mean either (1) the number of distinct elements, or (2) the logarithm to the base 2 of the number, the context indicating the sense used." (p.126 of _introduction to cybernetics_).

The Law of Requisite Variety states that the output of a regulator must have at least as much variety as the disturbance, if the result is to be regulation of some variable. That is to say, the output must have at least as many discriminable states as the disturbance has.

In terms of control theory, we would say that if a controlled variable is to be maintained exactly at a reference level, the output of the system must be at all times quantitatively equal and opposite to the magnitude of the disturbing variable, both measured in terms of effect on the controlled variable.

Note that the control-theoretic statement goes much farther than the LRV goes. It says that not only must the output have AS MANY discriminable states as the disturbance, but that these states must have the correct quantitative magnitudes, and they must occur in pairs: one specific output state for each state of the disturbing variable. The output and disturbing states must be quantitatively equal, and of opposite signs. So the LRV, while it states a weak necessary condition, by no means states a sufficient condition for control or regulation to exist.

d. How does PCT account for stored knowledge such as schemata (or whatever other term you choose)?

PCT itself doesn't; that's a question of fact for neurophysiology or neurochemistry. I have proposed some possible relationships of stored information to the operation of control systems in the hierarchical model. As Greg Williams puts it, those proposals are embellishments on the basic model.

"Schemata", as I have heard the term used, could have various meanings, corresponding to reference signals (when they relate to goals), to perceptual input functions (when they determine how lower-level perceptions are interpreted) or memories (in various circumstances). "Knowledge" has similar usages. Because these words are used in relation to functionally very different aspects of a brain model, it's hard to pin them down to any one meaning. The problem is somewhat similar to the uses of "want." If you say "I want some ice cream" you could mean to point to a reference signal, the definition of what it is that you want, or you could mean to indicate an error signal, emphasizing the fact that you don't have the ice cream. In my opinion, terms like these are too loose to be used in a model, although we may use them informally when context supplies the missing discriminations.

d. What makes PCT more than intensive descriptions and explanations of neural pathways and signals?

The fact that it makes almost no attempt to describe neural pathways and signals. The primary uses of PCT are in modeling externally-visible behavior in real environments, showing that relationships expected under the hypothesis of control do in fact occur and can be predicted. The models propose that certain functions are carried out inside the control system: perception, comparison, and conversion of error to output. Nothing is said about what neural circuits are involved in implementing these functions, although in some cases we know quite a bit about specific neural signals and pathways involved in specific control systems.

e. Why must PCT be understood as some sort of revolution as opposed to some sort of new and productive thinking?

Because PCT directly contradicts most traditional conceptions of how behavior itself works. It also predicts phenomena of a kind that conventional sciences have known nothing about, or that under conventional assumptions go in a direction opposite to the direction that control theory would predict and that experiment supports. For example, under control theory, the prediction (and the observation) is that doubling the sensitivity of a perceptual function to stimuli will result in halving the input to the perceptual function, not doubling its output signal.

Control theory is to conventional theory as Newtonian celestial mechanics was to Ptolemaic epicycles. Many of the observations may be the same, but the explanation is radically different.

f. What are essential differences between PCT and cybernetics?

Cybernetics abandoned control theory at about the time of the 5th Macy conference, when most of the main misconceptions about control systems were laid in.

g. What do axioms or propositions from PCT contribute to everyday human adapting, living, changing, improving, that other forms of analysis do not? Where is the proof of human behavioral successes with PCT, as opposed to success in the abstract?

I'd say you should ask people who are applying it. Try Ed Ford, Dag Forssell, and David Goldstein. Or ask people in the CSG who have been interacting with each other for 9 years with an understanding of PCT. In a lot of cases, PCT vindicates commonsense ideas that science has rejected, such as the inportance of goal-setting, perception, and intention. If you see concepts like that creeping back into polite scientific society, I think you can credit PCT with having inspired at least some of those changes in thinking. Even when an idea is rejected, something of it sticks.

h. Most scientists are concerned with not only describing and explaining phenomena, but also wish to predict and sometimes "control" them. Why does PCT deny the desirability of prediction?

PCT does not deny the desirability of prediction. I, personally, deny the desirability of lousy predictions, particularly when they're used as if they were good predictions.

Is prediction not always part of human anticipation, whether in ordinary life or in science?

Yes, I think it is. It's explained, more or less, in HPCT as a phenomenon of imagination, which is part of the story of mental modeling. People do it, so it belongs in the model. But I can't say a lot, theoretically, about HOW they do it. Most of my objections to prediction involve pointing out that it's not being done very well, and shouldn't be relied on as a method of control in most situations -- not if good control is important.

i. Where do reference levels or signal originate? In other words, if my control system has error signals, what constitutes the sources of the error signals?

I guess you haven't read anything about the hierarchical model. When you say "my control system" I wonder if you have read anything about the model at all. People have hundreds, thousands of control systems, all active at the same time and at different levels. Higher-level systems act not by producing motor outputs but by adjusting reference signals for lower systems. If you haven't been aware of this aspect of the model, you must have found a lot of the conversations on the net pretty confusing. Next you'll ask "But what about the highest levels of control system? Where do THEY get their reference signals?" And I will answer, your guess is as good as mine, but asking the question shows that you get the picture.

What is stored that makes me think that I do not like or wish to accept certain input?

You're assuming that it's something stored that "makes you think" that you don't like etc. A reference signal specifies a certain amount of a given perception, anywhere between the maximum possible and zero. That becomes the amount that you prefer; when you say you prefer that state of that perception, control theory explains this by saying that there is a reference signal in some control system set to that value, probably as part of controlling for some higher-level perception. This may or may not involve "storage" of something. There are lots of proposals on this subject in BCP.

> How conscious are the processes of matching input to the reference signals?

It depends. The same control system can operate consciously or automatically, at any level. A perceptual signal can occur with or without consciousness of its presence. The nature of consciousness is not explained in PCT or HPCT (or as far as I know, in any T).

> As PCT should not be criticized by me or anyone else for not explaining social behavior, social sciences which do focus on social behaviors should not be excoriated for not explaining individual human control mechanisms and processes.

Clark McPhail, Chuck Tucker, and Kent McClelland (the CSG's sociologists) and Tom Bourbon (who has modeled simple social interactions) ought to comment on this. My view is that PCT, by explaining the interactions of individuals with their surroundings, lays the groundwork for explaining what happens when groups of independent control systems interact with other control systems. Social "laws" emerge from the properties of interacting individuals and the shared environment. While naturalistic observation is needed to determine the existence of regularities in social behavior, PCT, I claim, is needed to explain these regularities.

> a. Communication studies, unlike traditional psychology
>and maybe other behavior/social sciences, no longer assumes
>that people are simply affected by stimuli in varying
>conditions and ways. Since the 1970s, my discipline has said
>that each human TAKES from mass media, from interpersonal
>interactions, from printed words, from any messages, and
>recodes what is decoded. We threw out Shannon and Weaver's
>model in the 1970s as anything useful to explain human
>communication. Both Aristotelian and electrical engineering
>models of message sending and receiving say little of
>importance to understanding the complexities of human
>communication, mainly because people do not simply take in
>messages. Nor do messages do contain meaning.

Good. Give my regards to Klaus. And what is the reason for all these changes in thought? Is there any theory from which you could have deduced all these new ideas? Or are they simply ideas that were proposed, and that others found acceptable for unstated reasons? Everything in this paragraph could be deduced from PCT, and has been familiar to PCTers for many years (except for the remaining buzzwords like "coding" and "recoding"). You're talking about observable phenomena, and that's what PCT is for: to bring observable phenomena into a single common framework and make some sort of sense of them -- not just to accept that they happen, but to explain, in terms of a unified structure of theory, why they must happen. The same theory that explains tracking behavior explains why people give their own meanings to their inputs, and why that is the only way it can happen.

Hope you had fun with the grading....

Best,   Bill P.

Date: Mon Mar 29, 1993 11:43 am PST
Subject: Re: Reply to Tucker's Review

     *************   FROM CHUCK TUCKER 930329-2 *************

Yes, doing this science stuff is tough especially if you are claiming results from your research that can not be supported by a reader who is quite familar with your model and your research procedures.

You are the one who made the claim and said "chapter and verse" (remember that phrase) that your study supported the claim. It now up to you to to the study properly by doing at least three things: (1) have written statements to either to be read by your subjects or you say verbatim to them, (2) make as certain as you possible can that the subjects do not have any awareness in any manner that a disturbance is involved even remotely in the study, and (3) ask each subject AFTER the ALL the "runs" questions which they can answer which will indicate if they were or were not aware of any disturbances while doing the study.

I would be willing to review your research proposal BEFORE you do the study so we don't get into the problems having me to tell your after it is done that you made some errors which made it impossible for you to support your claims.

Regards, Chuck

Date: Mon Mar 29, 1993 11:49 am PST
Subject: Monday Sweetness and Light

[From Rick Marken (930329.1100)]   Greg Williams (930329 - 3) --

> I don't recall reading of accusations of arrogance
>directed toward Einstein, though I've read much on his life and work
>(such as Clark's biography).

Both sides had very precise working models and methods for testing them. So instead of yelling "arrogant claims" the establishment TESTED Albert's claims; oops, they turned out to be correct.

>I don't see how it follows that "doing the work" gives someone a right
>to be arrogant.

The things we are saying that are being called "arrogant" are usually just testable assertions. The "no information in perception" claim is one. We are called arrogant for making it and saying that it invalidates the IV-DV approach to studying behavior -- even though we have presented evidence for this position over and over. I think the only way that we can be assured of not being seen as arrogant is to say something like this:

"Of course there really is information in the perceptual input to a control system about how to behave; we were just being silly when we said there wasn't. So just keep studying behavior by manipulating independent variables and measuring dependent variables -- if you keep doing this you will keep learning more and more about the processes that generate behavior".

Does that sound a bit less arrogant?

Oh, and pay no attention to that man behind the curtain.

Best Rick

Date: Mon Mar 29, 1993 12:18 pm PST
Subject: info about disturbance

[From: Bruce Nevin (Mon 930329 14:44:46)] Chuck,

In my opinion Rick is right on this point and you are barking up the wrong tree. Whether or not the subject knows or figures out that there IS something other than her mouse movements that is making the cursor move is irrelevant to the point being claimed.

In fact, even if you carefully avoid mention of any disturbance, the subject will presume that there is one, though perhaps not using that word for it. As a matter of conversational pragmatics it would be impossible to give a tracking instruction, like "keep the cursor on this mark", unless it were stated to the subject or presupposed by the subject that something was liable to separate the cursor and the mark, otherwise what's the point? If there is no disturbance to change that relationship, the subject would justifiably be confused and say "what do you want me to do, sit and watch as it continues not to change?" The subject is ostensively presented with a task. The task description calls her attention to a perception, prescribes a reference value for that perception, and tells what to do to change the perception just in case it should move from the reference value. The obvious necessary presumption, even if the subject is not explicitly told there will be a disturbance, is that something is going to move the perception away from the reference value. No disturbance, no task.

I think you tripped up on an ambiguity in the phrase "information about the disturbance." You appear to be reading it as "information that there is a disturbance present." The intended meaning is "information as to what the disturbance is", that is, information about the changing moment-by-moment values of the disturbance as a change of position of the cursor that is added to the changes due to mouse movements. (In another situation it could be pressure, light intensity, pitch, etc.)

Bruce bn@bbn.com

Date: Mon Mar 29, 1993 12:43 pm PST
Subject: Still more Mon. s + l

From Greg Williams (930329 - 4)   Rick Marken (930329.1100)

>> I don't recall reading of accusations of arrogance
>>directed toward Einstein, though I've read much on his life and work
>>(such as Clark's biography).

>Both sides had very precise working models and methods for testing them.
>So instead of yelling "arrogant claims" the establishment TESTED Albert's
>claims; oops, they turned out to be correct.

I suspect that Einstein's non-arrogant demeanor played a part in that.

>>I don't see how it follows that "doing the work" gives someone a right
>>to be arrogant.

>The things we are saying that are being called "arrogant" are usually
>just testable assertions. The "no information in perception" claim is one.

How can you test, today, whether rebuilding psychology on a PCT foundation will actually meet your own and/or other psychologists' goals? You are currently just ASSUMING it will, not having tried, I think, to apply the foundational method of testing for controlled variables to "hard" cases in the field. Want to join Bill Powers and try it on me?

>We are called arrogant for making it and saying that it
>invalidates the IV-DV approach to studying behavior -- even though
>we have presented evidence for this position over and over.

An IV-DV basis for generative behavioral modeling IS contradicted by PCT, but the "IV-DV approach to studying behavior" also includes, as I've been trying to say recently, deriving statistical measures of population characteristics, with no commitment to an IV-DV generative model. When you don't make it clear that you don't include the later in "studying behavior," you sound arrogant to me. When you are more careful, you don't (as I've said before). Saying that ALL of psychology needs to be REVOLUTIONIZED by PCT is not being very careful.

As ever,   Greg

Date: Mon Mar 29, 1993 12:45 pm PST
Subject: From arrogance to pretension, yet

[From Bill Powers (930329.1230 MST)]   Chuck Tucker (930329) --

>You can translate all the words into other symbols (sy=symbols)
>and put ='s, +'s, or any other sign you care to between them
>and they still turn out for me to be words; I still have to
>read them to have some idea of what you might be stating;
>actually I find your use of such symbols a pretention on your
>part to appear to be systematic and scientific when your
>research does not have such features for me.

Hmm. If they're a pretension on Rick's part, they must be a pretension on my part too. Judging from referees' comments on various papers, you probably aren't alone in dismissing the simple little equations as just a pretentious way of saying the same thing you could say in words. There's certainly enough of that sort of pretension in the literature: (m + o)^2 meaning something like motive and opportunity considered twice, and so on.

But the equations are real equations in this case. When the position of the cursor c is said to be equal to the disturbance d plus the output o, and this is expressed as c = d + o, the meaning is just what the algebra says. It is not that the presence of a disturbance sort of affects the cursor, and so does the output. It is that the cursor position in pixels measured from the center of the screen is numerically equal to the measure of the handle position plus the numerical value of the disturbance. If the disturbance is 100 units and the handle output is -99 units, the cursor will be exactly at the position 1, not 2 and not 0. The accuracies of the apparatus and the computer are many times greater than needed to make this true within one pixel on the screen. Saying that c = o + d pins down the meaning exactly; there are no ordinary-language words that can do this without adding many, many sentences -- and sentences explaining the sentences. As above.

If you just say in words that the disturbance and the output affect the cursor, you're leaving all kinds of slop in the meaning. Words are blunt instruments for conveying quantitative meanings. There is no pretension in using equations that actually mean something. The symbols are just words from a language with which you're not familiar, indicating meanings with a precision that you're not accustomed to.

It may also help you to know that ordinarily, the disturbances I use in my experiments are generated by a random-number generator and smoothed by a fixed algorithm. For convenience I usually record 10 or 20 tables of disturbances for later use, but often I just generate and record them as the experiment proceeds. In any case, I never know beforehand what is in those tables unless I write a plotting program to check on the range and variability of the disturbances, which I do only when adjusting the smoothing to affect the general difficulty of the task. And then I usually regenerate the whole set. The disturbances are a complete surprise to me and to the simulation, and when they're being generated in real time they have a different pattern on every experimental run. So it makes no difference at all who the subject is -- myself or a neophyte. There's no way to cheat.

As to instructions: about the minimum instructions for a tracking experiment are "Keep that (point to the cursor on the screen) between those marks (point to the target marks) using this (point to the handle)." The subject will immediately begin tracking when the experiment begins. You can convey this information in a hundred ways, and after 10 seconds of tracking you can't tell which instructions were used. Of course you can use incomprehensible instructions, but even then the subject is likely to say, shortly after the run begins, "Oh, I see -- is this all you meant?"

I'm not recommending poor instructions. But I think that you sometimes go overboard in instructing people what to do, reading exact wordings from cards, as if this somehow guaranteed that each person would get the same meanings from them -- or as if any slight deviation would totally screw up the experiment. Perhaps this is a holdover from the kinds of experiments you're used to in sociology, where the results are so ambiguous to begin with that any variation in the instructions could spoil the results altogether. Maybe -- just guessing -- this is why having the experimenter be a subject seems to be a such a no-no for you. It's probably true in normal-science experiments that if the experimenter knows what's going on, he or she can fudge the results. But that's not true in control-system experiments. You can't control any better than you can control, and the disturbance is simply not knowable, even to the person who wrote the program. Unless you're some sort of eerie prodigy who can predict the next 1800 outputs from a random number generator and mentally run them through a 3-stage smoothing filter.

I do have to stand up for Rick when you say

>YES, this study (and all others that I have seen you do) does
>bother me. It does so because it is so poorly done that I
>would not consider it worthy of inclusion in a list of
>scientific work. The claims that you make for such work are
>usually not supported for me by your study.

Rick's claims are always exactly upheld by his data in the published experiments. The only reason you can't see this is that you can't read the algebra or the program steps, so you have to rely on the verbal approximations to make your judgments. That leaves too much room for you to bend his meanings to fit your preconceptions, and where he makes statements that conflict with what you believe, to see plenty of room for him to be wrong. This is a general problem with trying to convey the results of modeling to people who are used to working with words alone. It's exacerbated by the fact that the models we use, and the programs, are so simple-looking. It's hard to believe that the author isn't giving them a lot of interpretive help in order to make them fit the data so well. Most people who are used to normal experimental results in the behavioral sciences just can't believe that such simple expressions could achieve such accuracy of prediction. And they haven't the experience with quantitative modeling to see that there are no tricks or hidden assumptions.

I think you have to blame your reactions to Rick's experiments on your own lack of experience with this sort of modeling process. You're putting form before substance.

I do agree with you that Rick gets a little wild in his claims on the net. But so would you, for a while, if you understood what he understands. It will pass.

Best, Bill P.

Date: Mon Mar 29, 1993 1:16 pm PST
Subject: Mary on Hacker's comments and questions

[Mary Powers 9303.29] Ken Hacker:

I hope that a number of people try answering your questions - various approaches may yield one that's right on the money for you.

>a. The scope of PCT. Is it limited to describing and explaining
>physical behaviors of human beings such as motor actions?

The problem with this question is an implication (that I perceive, and think you intend, consciously or not) that movement is the outcome or end result of whatever is going on in the organism. In PCT, motor actions are the means by which perceptions are controlled - all perceptions, at every level.

The idea that PCT is all very well for explaining mere movement has been around in cybernetics for years - Klaus probably heard it the same time I did, from Varela, in November, 1984. The idea was that it did not explain having high thoughts about great cybernetic ideas. How one went about speaking to an audience about those great thoughts didn't enter into it.

>b. if there is no information inside of input, what prevents one
>from asserting behavioristic-like internal responses to internal
>stimuli or a kind of reverse behaviorism?

I'm not sure I understand this question. Do you mean input as a stimulus causing behavior? The point is joint causation - there has to be sensory input AND a reference input to a comparator. The difference drives the output. If both signals have the same value, nothing happens, no matter how stimulating the input appears to an outsider.

>c. What happens to Ashby's Law of Requisite Variety.

I'll leave this to Bill.

>d. How does PCT account for stored knowledge such as schemata
>(or whatever term you choose).

I'll leave this to neurophysiology. It suffices for the moment that knowledge is stored and can be accessed.

>d. (again) What makes PCT more than intensive descriptions and
>explanations of neural pathways and signals?

PCT is an explanation of the functional organization of neural pathways and signals. There are a lot of pathways and signals whose functions are unknown unless they are seen as components of control systems. PCT is consistent with what is known to exist in the nervous system, but a lot of what exists is not explained by those unfamiliar with PCT.

>e. Why must PCT be understood as some sort of revolution as
>opposed to some sort of new and productive thinking?

Because of the active resistence and consistent misinterpretations with which it has been received over the past 33 years. PCT was offered initially as simply something new and productive. The idea that it is revolutionary was a gradual development of an explanation for other people's reaction to it.

>f. What are essential differences between PCT and cybernetics?

This would be easier to answer if cyberneticists were willing to define cybernetics. Those I have asked have been unbearably coy.

I think the answer is that PCT IS cybernetics. In The Science of Control and Communication in the Animal and the Machine, PCT is the science of control. No one in cybernetics has even come close to PCT in developing it. PCT began when Bill P. read Wiener. While waiting for further dvelopments of these marvellous new ideas, he began working them out for himself. No one else seems to have bothered over the past 40 years. So by default, PCT is it.

>g. What do axioms or propositions from PCT contribute to
>everyday human adapting, living, changing, improving, that other
>forms of analysis do not? Where is the proof of human behavioral
>successes with PCT, as opposed to success in the abstract? Is
>[it] too soon to ask this question?

PCT affirms and explains some principles that have been around a long time. For instance, if you push on people, they push back. Why? Because they are also control systems, resisting disturbance. Terry Brazelton advises parents to let their children be autonomous as much as possible, and only insist and clamp down when it is really necessary, in order not to establish a pattern of resistance to everything. This is empirical wisdom. PCT explains it.

PCT explains why Rick has escalated over the years into the high emotional pitch you see now - he keeps meeting resistance and pushes back harder. Maybe some day he will reorganize and find a new way to go about getting PCT across to people who don't want to know about it. He's tried quite a few, and I hope keeps on thinking up more demos and experiments. Meanwhile, ranting and raving on the net keeps that error from getting too large. (Incidentally, what's wrong with arrogant statements on the net? This isn't formal publication, it's a conversation. Mostly polite, sometimes not. Pretty mild compared to some, we are told).

Control systems are all around us, and include us. Our proposition is that people are doing their everyday adapting, etc etc AS CONTROL SYSTEMS. Once you start analysing them this way, it is obvious. PCT contributes a point of view.

It is probably too soon to ask about the practical value of PCT, by which I assume you mean changing people. There are a few therapists, a few people involved in education, a few in management. What they are doing, primarily, is teaching people that they, and others, are control systems. Adopting this point of view seems to help people understand themselves, others, and why some things they do don't and can't work in getting along with people, resolving conflicts, achieving goals, etc. But so far the evidence of success is anecdotal.

>h. Most scientists are concerned with not only describing and
>explaining phenomena, but also wish to predict and sometimes
>"control" them. Why does PCT deny the desirability of
>prediction...is prediction not always a part of human
>anticipation, whether in ordinary life or science?

Sure. People want to predict and control. We act to prevent anticipated errors all the time. I control for not running out of milk, toilet paper and cigarettes by buying enough to last at a predicted rate of consumption until the next time I plan to drive into town.

Scientists and others who want to control behavior are not likely to succeed unless they understand what behavior is FOR. The behavior they want to control is now viewed as a bunch of undesirable outputs. Ultimately, the only way to control those is by physically preventing them - lock up the bank robbers, shoot the abortionists, kick the trouble-makers out of school.

Everyone has goals. What their behavior is for is to achieve those goals - that is, to make their perceptions match their reference levels. If the behavior they are using to achieve those goals is unacceptable, in some cases simply pointing this out and asking for different behavior is enough. Sometimes teaching new ways of behaving that hadn't occurred to the person is enough. Sometimes nothing is enough, and then you have to get into the particular structure of reference levels in the individual person - what they want. Much of what they want may be unconscious, conflicting, and impossible. I can't get into how a psychologist can deal with that, but the PCT approach of going up levels and developing awareness can work - and I think provides a rationale for how therapy works. If therapy isn't available or successful, control comes back down to physical force.

>i. Where do reference levels or signals originate? In other
>words, if my control system has error signals, what constitutes
>the source of the error signals?
>What is stored that makes me think that I do not like or wish to
>accept certain input? How conscious are the processes of matching
>input to the reference signals?

Reference signals originate in the brain. Where in the brain depends on which reference signals you are talking about - for breathing, walking, speaking, painting a picture, etc. See BCP.

The sources of error signals are comparators, which receive reference signals and input signals, compare them, and generate a signal representing the difference. That isn't the question you thought you were asking, is it?

Not liking an input means that there is an emotional component to error. I don't know where reference signals hide out. Not the liver or the spleen.

The process of matching input to reference signals may be totally unconscious or totally conscious or in between. One problem with the self-regulatory folks is a preconception that it is all conscious - thus Karoly mis-cites Bill as claiming that humans are unique in the ability to achieve consistent ends by variable means (so there, e. coli!). I personally have no idea how my body weight stays the same month after month. I am conscious of hunger periodically, and of buying and cooking and eating food, and I'm certainly aware of the need to finish the digestive process, though I don't consciously control intestinal contractions, or any of the digestive processes that preceed them, or cellular metabolism. Lots of control systems doing their thing, all me, some conscious some of the time.

Mary Powers

Date: Mon Mar 29, 1993 1:38 pm PST
Subject: IV-DV; challengs

[From Bill Powers (939329.1415 MST)] Greg Williams (930329-4) --

>... the "IV-DV approach to studying behavior" also includes, as
>I've been trying to say recently, deriving statistical measures
>of population characteristics, with no commitment to an IV-DV
>generative model.

I'm not sure it doesn't. If a population study finds a correlation between taking aspirin and a decrease in heart attacks, doesn't this carry the implication that if you take aspirin, you'll have a decreased risk of a heart attack? It's certainly publicized that way.

It's hard for me to think of a circumstance in which, once an IV- DV relationship has been found, the IV isn't then used as a predictor of the DV. Maybe the prediction is just "yes-no", but that's a generative model. Not one of your GREAT generative models, you understand.

You to Rick:

>You are currently just ASSUMING it will, not having tried, I
>think, to apply the foundational method of testing for
>controlled variables to "hard" cases in the field. Want to join
>Bill Powers and try it on me?

You should have said that to me. Are you trying to present yourself to me as a "hard" case? You're not a hard case.

Best,   Bill P.

FROM MARY:

[from Mary Powers 9303.29]   Greg Williams

What's going on? You want PCT to explain your behavior when you yourself don't have access to your own motivational control structure? If Bill gives you an analysis, correct or incorrect, you're saying you aren't equipped to confirm or deny. So why are you asking?

Mary

Date: Mon Mar 29, 1993 1:43 pm PST
Subject: insulting behavior

[From Rick Marken (930329.1230)]   Bruce Nevin (Mon 930329 08:42:58) --

>If an utterance is ambiguous, it is because the speaker is
>presuming that the hearer understands that certain unspoken
>parts of the utterance are in the mutually understood context of
>the part that is spoken. "Flying planes can be dangerous" is
>ambiguous only in isolation.

I agree. But I think that the same is true of ambiguous non-linguistic behavior too (I'm saying that all behavior, linguistic and non-linguistic, is ambiguous inasmuch as we can never be positive (without doing the test) what controlled variable might be involved). "Context" lets us do a partial "test for the controlled variable". When I see a person running across the street at top speed I might think the person is "running from an assailant" or "trying to make it to an appointment" until I see the bus pull up just as the person reaches the bus stop to get on. The context reveals one of the person's controlled variables -- "catch the bus".

>I don't see how this constitutes linguistic behaviorism. Please
>try to avoid insults. Or did you mean something different from
>psychological behaviorism? :-)

No insult intended. As I said, I use the word "behaviorism" to describe the tendency to take observable outputs at face value -- as "behavior". You seem to be saying that language (which is certainly an observable output) can be taken at face value -- ie. if someone says the word "hello" then that is their "behavior" in both the conventional and PCT sense (the PCT sense being that behavior is a controlled perceptual consequence of output). I'm saying that language is like all other behavior -- you can't tell (for sure) what constitutes a person's behavior (an intended consequence of their outputs) without doing "The Test" in order to determine what perceptual inputs they are controlling.

Bruce Nevin (Mon 930329 14:44:46) --

>Chuck,

>In my opinion Rick is right on this point and you are barking up the wrong tree.

I'll never call you a linguistic behaviorist again, Bruce.

Best Rick

Date: Mon Mar 29, 1993 2:06 pm PST
Subject: Rick and Albert

[From Rick Marken (930329.1330)] I said:

>Both sides had very precise working models and methods for testing them.
>So instead of yelling "arrogant claims" the establishment TESTED Albert's
>claims; oops, they turned out to be correct.

Greg Williams (930329 - 4) replied:

>I suspect that Einstein's non-arrogant demeanor played a part in that.

Einstein's non-arrogant demeanor is known to us only after relativity was pretty well recognized as a major accomplishment. I seem to recall reading somewhere that he was not exactly Mr. Humble pie when he was developing his ideas at the ol' Polytechnique. Watch how nice and laid back I get when all the media in the country acknowledge PCT as the great new theory of living sysems. I might even grow a mustache, long hair and a cute accent.

>How can you test, today, whether rebuilding psychology on a PCT
>foundation will actually meet your own and/or other psychologists'
>goals? You are currently just ASSUMING it will,

True. Based on what I know, PCT seems to offer a better possibility for building a science of living systems; conventional behavioral science demonstrably provides no such possibility. But, in fact, the best approach might not have been developed yet. How did (dare I say it) Galileo know that his silly little ball rolling experiments were the best way to go about developing a science of physics?

> not having tried, I
> think, to apply the foundational method of testing for controlled
> variables to "hard" cases in the field.

And Galileo never tried "hard cases" like launching satellites to
test his theories.

> Want to join Bill Powers and try it on me?

No way. I've got my hands full the "easy" cases. Witness the Tucker-Marken
debates.

Best   Rick

Date: Mon Mar 29, 1993 2:08 pm PST
Subject: Challenge 2: ??

[From Bill Powers (930329.1330 MST)]   Greg Williams (930329-2) --

>>Is either of those explanations 99% accurate?

>I don't know. I don't have direct access to my motivational
>control structure, so I can't say what explanation is "99%
>accurate."

This is going to make it difficult for you to know whether I've met the challenge.
Well, I did say you are the judge, so maybe you will win by default.

>I thought you were going to explain why I issued the challenge in the
>first place, not why I continue to say you haven't succeeded in meeting it.

It's hard to experiment with a previous condition of a control system. I think
it's better to assume that the same control system continues in existence, and see
what can be found out about it as time progresses. If it's really gone, so some
completely different process is going on right now, I don't think I'll have much
luck in guessing how it was organized then. It didn't occur to me at the time I
chose the challenge-issuing behavior as the test object that the control process
in existence just then might disappear immediately and be unrecoverable.

If that happens to be true (I'm not admitting it yet), what conclusion would you
draw about the possibility of understanding human behavior by ANY means?

>You are overlooking the possibility that an output can remain
>(at least look) the same while the intended outcome is
>changing. My motivations could be changing continuously for not
>withdrawing the challenge.

That's a possibility. Is that what's happening?

>So what about an explanation for my originally issuing the challenge?

If that's your condition, I may well admit defeat. Any reaction to that?

>>Was (and is) this one of the immediate goals in issuing the challenge?

>Again, I don't know. It sounds reasonable, but how would I
>gather any additional evidence beyond (my own) hearsay?

Try trusting your own "hearsay" (I thought that referred to repeating what someone else said). I don't think you need to look outside to find out what you want. If you want it, you still want it. You can imagine what outcome would be satisfying to you, and know whether it seems OK or if there's something wrong with it.

>>That is, did you want to get back an indication of my
>>acceptance or rejection of it?

>It is possible. I'm not sure.

Well, in that case I'm not going to tell you whether I intend to continue with this challenge. OK with you?
Best,   Bill P.

Date: Mon Mar 29, 1993 3:03 pm PST
Subject: On and on
From Greg Williams (930329 - 5)   Bill Powers (939329.1415 MST)

>>... the "IV-DV approach to studying behavior" also includes, as
>>I've been trying to say recently, deriving statistical measures
>>of population characteristics, with no commitment to an IV-DV
>>generative model.

>I'm not sure it doesn't. If a population study finds a
>correlation between taking aspirin and a decrease in heart
>attacks, doesn't this carry the implication that if you take
>aspirin, you'll have a decreased risk of a heart attack?

I say no. So do at least some epidemiologists.

>It's certainly publicized that way.

Yes, at least by some aspirin advertisers and probably even some doctors. What does that fact have to do with my point, which was about behavioral scientists, not aspirin advertisers or doctors (or even epidemiologists)?

>It's hard for me to think of a circumstance in which, once an IV-
>DV relationship has been found, the IV isn't then used as a
>predictor of the DV. Maybe the prediction is just "yes-no", but
>that's a generative model. Not one of your GREAT generative
>models, you understand.

I think if we perused the refereed behavioral science literature together, we'd find a lot of cases where correlations are found between two population measures, and no generative model for IV-DV variables of individuals is supposed.

>You to Rick:

>>You are currently just ASSUMING it will, not having tried, I
>>think, to apply the foundational method of testing for
>>controlled variables to "hard" cases in the field. Want to join

>>Bill Powers and try it on me?

>You should have said that to me. Are you trying to present
>yourself to me as a "hard" case? You're not a hard case.

I seems that any case in the field, where ongoing disturbances cannot be known by the experimenter, might be "hard" for succeeding with The Test. At least, that's one thing I conclude from your post on prediction earlier today.

>Mary Powers 9303.29

>What's going on? You want PCT to explain your behavior when you
>yourself don't have access to your own motivational control
>structure? If Bill gives you an analysis, correct or incorrect,
>you're saying you aren't equipped to confirm or deny. So why are
>you asking?

Maybe Bill can convince me that there aren't any problems with the kind of introspection he wants me to use. Or maybe he will come up with a different way of attempting to meet my challenge without relying on my own impressions. I don't think he has given up yet. I'm waiting to see what happens next. Also, when I first asked, I didn't know he would need my access to my own motivational control structure. Shouldn't the Test work even on non-verbal creatures? In that case, such access would be moot.

>Bill Powers (930329.1330 MST)

>This is going to make it difficult for you to know whether I've
>met the challenge. Well, I did say you are the judge, so maybe
>you will win by default.

Maybe it would be better for YOU to be the judge?

>It didn't occur to me at the time I chose the challenge-issuing
>behavior as the test object that the control process in existence
>just then might disappear immediately and be unrecoverable.

>If that happens to be true (I'm not admitting it yet), what
>conclusion would you draw about the possibility of understanding
>human behavior by ANY means?

It all depends on the dynamics of disappearing and reappearing, I guess. Having access to the detailed dynamics some way (electrical probes?) might in principle be able to reconstruct individual motivations backward in time, if chaos isn't problematic (and it might not be, if the dynamics were governed by limit cycles). And for predicting POPULATION behavioral measures, statistical description might work (predicting "backwards."

>>You are overlooking the possibility that an output can remain
>>(at least look) the same while the intended outcome is

>>changing. My motivations could be changing continuously for not
>>withdrawing the challenge.

>That's a possibility. Is that what's happening?

I don't know.

>>So what about an explanation for my originally issuing the challenge?

>If that's your condition, I may well admit defeat. Any reaction to that?

Why do you say you may well admit defeat?

>>>Was (and is) this one of the immediate goals in issuing the challenge?

>>Again, I don't know. It sounds reasonable, but how would I
>>gather any additional evidence beyond (my own) hearsay?

>Try trusting your own "hearsay" (I thought that referred to
>repeating what someone else said). I don't think you need to look
>outside to find out what you want. If you want it, you still want
>it. You can imagine what outcome would be satisfying to you, and
>know whether it seems OK or if there's something wrong with it.

Why should I trust my own "hearsay"?

>>>That is, did you want to get back an indication of my
>>>acceptance or rejection of it?

>>It is possible. I'm not sure.

>Well, in that case I'm not going to tell you whether I intend to
>continue with this challenge. OK with you?

OK. As ever, Greg

Date: Mon Mar 29, 1993 3:31 pm PST
Subject: Fear and loathing in scienceland

[From Rick Marken (930329.1500)] CHUCK TUCKER (930329-2) --

>Yes, doing this science stuff is tough

Not really. But understanding it can be quite painful, apparently.

>You are the one who made the claim and said "chapter and verse"
>(remember that phrase) that your study supported the claim.

I presume you are referring to the claim that "there is no information about the disturbance in perception"? It would help if we could agree on what we mean by that claim. If you insist that the claim means "a person can't tell that a disturbance is acting in a compensatory tracking task" then you are correct; not only does the study not support that claim, it's not even relevant to it. [Why would I even do the repetition of tracking with the same disturbance and measure the correlation between inputs if all I wanted to show was that people don't know that a distrubance is present? Did that question ever occur to you while reading the reserach report? If you really thought I was trying to see if people could detect the presence of the disturbance, why wouldn't it occur to you that the logic of the experiment is not relevant to that question.] If, however, the statement means "it is not possible to reconstruct the disturbance given the perceptual input in a tracking task" then your criticisms of my experiment are not germain and the results of the experiment do provide strong evidence that there is, indeed, no information about the disturbance in the perceptual input.

>I would be willing to review your research proposal BEFORE you do the
>study so we don't get into the problems having me to tell your after
>it is done that you made some errors which made it impossible for you
>to support your claims.

Thanks for the non-arrogant offer.

Since there are so many errors in my study, can I take it, then, that you don't buy my claim that there is no information about the disturbance in the perceptual input to a control system? I acknowledge that there is information THAT there is a disturbance in that input; no experiment was ever needed to test that. I mean, do you understand and reject my claim that "it is not possible to reconstruct the disturbance given the perceptual input in a tracking task"?

Just curious Rick

Date: Mon Mar 29, 1993 4:03 pm PST
Subject: Re: defining information

[Allan Randall (930329.1700 EST)] Rick Marken (930325.1100) writes:

> >...H(D | h,P) < H(D | h). The percept contains information
> >about the disturbance.
>
> NOW WE ARE GETTING SOMEWHERE.
>
> All you have to do now is show that your statements above are true...
> ...I predict that you will need EXACTLY the same size
> program to generate D (given nothing) and to generate D|P. That is,
> I predict that H(D|h) = H(D|P,h).
> ...
> But PLEASE DON'T TAKE MY WORD FOR THIS: DO THE INFORMATION CALCULATIONS
> AND REPORT YOUR RESULTS ON THE NET. Then we can go from there.

First, let's be clear what this would entail. I will describe a possible experiment, predict the results, and provide some analysis. You can respond as to which of these three, if any, you have a problem with. Let's take a computer simulation of a simple control system (one of Bill Powers' Primer examples might be a good choice). Our algorithmic definition of information divides the problem into language, programs and output. We will take the hierarchy itself as the computer language. The perceptual input will be the program, and the control system's output is, simply enough, the output of the program. Since the hierarchy is the only language we will be concerned with, all the entropy measures will include (... | h). So we'll just drop the 'h' and assume it throughout.

We would first run the control system under normal closed-loop conditions, recording disturbance and output, and noting that the disturbance is almost perfectly countered at the output. The control system's perceptual line would then be cut, and superceded by a new experimental line. All possible perceptual inputs would then be presented, one at a time, starting with the shortest and working up. This would be repeated until the disturbance recorded from the first experiment appears on the output. Let's assume that the output in the first experiment contained 100% of the information about D, since we all seem to agree that this is very nearly the case. Now we need only look for a replication of the initial experiment, and H(D) will simply be the length of the successful percept.

Let's label the possible percepts {Pi} = {P0,P1,P2, ...}. P0 will be the null message, which is equivalent to simply cutting the perceptual line and not providing any replacement input at all. Let's call the first Pi to successfully replicate the initial output Pk. Pk might very well be equal to the original P, but it could be shorter. We are guarenteed that it will not be longer. The entropy of the disturbance is:

   H(D) = |Pk| <= |P|

We then perform the same procedure to compute H(D | P). This time, P is a given, so it is always provided to the perceptual line "for free" and not counted in |Pi|. This time we will be successful on the very first attempt. Since this setup is equivalent to our initial experiment, without any extra input to the perceptual line, the first successful Pi will be P0, the null message. The entropy of the disturbance given the percept therefore is:

   H(D | P) = |P0| = 0

I claim:

  H(D | P) < H(D), which gives |Pk| > 0

You claim:

  H(D | P) = H(D), which gives |Pk| = 0

So your claim can only be correct if the disturbance can be reproduced "nearly 100% perfectly" at the output with a completely cut-off perceptual line. This *is* theoretically possible, if the environment is completely predictable. But this is obviously not the case in the real world. As far as I can see, Rick, you are forced by your claim to conclude that the system can control blind - with no input whatsoever. I think we can all agree that this, if theoretically possible, is not generally the case in the real world.

My question is: what is wrong with the above reasoning? Do you actually support the idea that the system can control blind? If not, for your claim to be right, I must have goofed up in my reasoning. Do you disagree about the results that I have assumed we would get from this experiment? Do you disagree with the way the problem was divided into program, language and output? Do you disagree with the derivations? Definitions? The discrete nature of the computer simulation? There are many places in the above argument where our disagreement might lie. If we can agree as to where we disagree, we just might be able to come to an understanding.

Allan Randall

Date: Mon Mar 29, 1993 5:49 pm PST
Subject: info in d(t)

[From Jeff Hunter (930329-A)]

Well hello again CSG. This information discussion has tempted me back out of my lurker status.

Rick's claim is that p(t) contains no information about d(t), however p(t) and o(t) together contain complete information about d(t).

Avery makes the useful observation that p(t) can be considered an encoded message about d(t) where o(t) is the key. He seems to agree with Rick that p(t) is completely encrypted.

[Avery Andrews 930317.1404]
>I'd also like to point out that the situation with d(t), o(t) and p(t) is
>basically the same as encryption with a one-time pad: the pad (o(t))
>certainly has no content at all, the message (p(t)) has no accessible content
>for people who don't have the pad, but if you have the message and the pad
>then you get the content. Maybe the philosophy of cryptography, if there is
>such a thing, has some application here.

I've just been reading "Spycatcher", and have been working up to a posting using the metaphor of cryptography. This is a perfect invitation. (Thanks Avery!)

The short answer is that p(t) is not very well encrypted.

The claim that p(t) contains *no* information about d(t) is a strong one. Thus a single counter-example will disprove it. This gives me a fair bit of leeway to make simplifying assumptions. (After this is done we can haggle over how much info is in p(t) in more normal circumstances.)

Putting on my test-setter-upper hat I'll choose the ECS from Bill Powers' primer (Part II). I'll also choose a reference r(t) that stays constant for long periods, and only changes in abrupt steps that are much greater than the average internal noise of the ECS. Likewise I'll choose a d(t) that also changes only in large steps at infrequent intervals.

Now putting on my codebreaker's hat I'll forget what I know about r(t) and d(t), but remember the ECS (as shown below):

```
                         r(t)
                          |
                          v
            -----p(t)--> (-)
            |             |
            |            *ko
            |             |
           *ki     leaky integrator [qo = qo(1-dt/ks) + error*dt/ks]
            |             |
       ................................
            |
    d(t)--->+<---o(t)-----|
```

From here on down "I" am the codebreaker, staring at p(t).

Most of the time p(t) will just jitter slightly from noise internal to the ECS. If I average out the noise I can find the current value of r(t). Pretty soon I'll find that r(t) is constant.

Now I wait to see if d(t) or r(t) changes. When this happens I'll see one of the two events below (thanks to Bill P. for the diagram):

```
                       *
                        *
                         *
                           *
***********************       *          ************
                                  *  *
                       ^d(t) changes

 - or -
                                  *  *
                              *          ************
                           *
                         *
                        *
                       *
***********************
                       ^r(t) changes
```

Either will show the convergence time of the ECS. This gives me "ks", the leak factor of the "leaky integrator".

At any change in the disturbance there will be a brief period of time before o(t) changes. This means that I can measure the peak of the change in p(t). This is just "ki" times the change in d(t). From here on I can sum the changes in d(t) to get ki*d(t) (plus some unknown value of disturbance before I started watching, plus a growing error term).

Now I exploit the fact that the leaky integrator is non-linear. The integrator reacts slightly differently depending on the current magnitude of o(t). This is a very small effect, and I have to watch a large number of transitions to cancel out the effects of noise. This can be used to sharpen my estimate of ki*d(t).

The conclusion is that in this special case I can find almost complete information about d(t) solely from p(t), the form of the ECS, and the assumption that no-one is manipulating d(t) and r(t) and the internal noise in an attempt to spoof me.

Thus the statement "there is NO information about d(t) in p(t)" is disproved. ... Jeff

Date: Mon Mar 29, 1993 6:05 pm PST
Subject: IV-DV; Challenge 2

[From Bill Powers (930329.1830)] What a day!

Greg Williams (930329-5) --
> I think if we perused the refereed behavioral science literature together, we'd
>find a lot of cases where correlations are found between two population
>measures, and no generative model for IV-DV variables of individuals is
>supposed.

You have better access than I, although I'll look, too. Let me know what you find, when you have time.

> I seems that any case in the field, where ongoing disturbances cannot be known
>by the experimenter, might be "hard" for succeeding with The Test.

True. When the experimenter isn't there to ask questions and observe what's happening, it's pretty difficult. However, with enough verbal transaction and cooperation, it's not impossible. I think that you and I will manage to find some of your controlled variables even under the present circumstances. You're not that hard to figure out.

You to Mary:
> Maybe Bill can convince me that there aren't any problems with the kind of
>introspection he wants me to use.

It sounds as though you don't want to try a method of introspection that you see as having some problems. Is this true?

> Or maybe he will come up with a different way of attempting to meet my
>challenge without relying on my own impressions.

I trust your impressions. Don't you?

> Shouldn't the Test work even on non-verbal creatures? In that case, such access
>would be moot.

Pretty tough to apply the Test for control of nonverbal perceptions over the net.

to me:
> Maybe it would be better for YOU to be the judge?

Oh no, I would have won already. That would be too easy. I prefer hard cases, if I can find any.

>... what conclusion would you draw about the possibility of >understanding human behavior by ANY means?

> It all depends on the dynamics of disappearing and reappearing, I guess. Having
>access to the detailed dynamics some way (electrical probes?) might in principle
>be able to reconstruct individual motivations backward in time, if chaos isn't
>problematic (and it might not be, if the dynamics were governed by limit
>cycles).

An informative answer. Let me try some test questions now, for which I think I know what the answers will and will not be.

1. Do you want me to go on questioning you in this same way?

2. Do you want this process to end up with a valid understanding on my part of one of your controlled variables?

3. Do you sense that I will fail or succeed in meeting the challenge?

I predict that your answers will be, very approximately:

1. That's up to you, I don't know what I want to happen, there's
no evidence of any desire one way or the other.

2. I have no idea how it will end up; I can't see evidence for any preference.

3. My senses don't lead to any prediction; I'm just waiting to see how this will turn out.

I predict that the answers will NOT be

1. Either Yes or No.

2. Either Yes or No.

3. Either Yes or No.

Having revealed my predictions ahead of time, I expect you to play fair. Note that I still haven't said what I think you're controlling for. I'm still gathering evidence, without a lot of help from you.

Best, Bill P.

Date: Mon Mar 29, 1993 7:03 pm PST
Subject: Allan Randall's proposed experiment

[From Bill Powers (930329.1930 MST)] This day is never going to end.

Allan Randall (930329.1700 EST) --
> We would first run the control system under normal closed-loop conditions,
>recording disturbance and output, and noting that the disturbance is almost
>perfectly countered at the output. The control system's perceptual line would
>then be cut, and superceded by a new experimental line. All possible perceptual
>inputs would then be presented, one at a time, starting with the shortest and

>working up. This would be repeated until the disturbance recorded from the first
>experiment appears on the output.

This is not very practical. The disturbances we use range in value from about -350 to 350, or roughly a 10-bit number. The outputs have, of course, about the same range. The range of deviations of the controlled quantity, the input, is about 5% of this range, with a moderately difficult disturbance (medium bandwidth). So the record of inputs will be representable by a string of 5-bit numbers.

During a one-minute run on a VGA screen, we record 1800 consecutive 5-bit numbers for the input deviations. In order to present all possible inputs, you would have to start with a single 5-bit number, cycling through all 32 possibilities, then do the same for the next number, and so forth until you reached a string of 1800 5-bit numbers. To reproduce an arbitrary disturbance waveform in this way you would have to try, on the average, half of the possible strings of 1800 numbers, which amounts to something like (32^1800)/2 possibilities. To say the least, the chances are small that you would ever run across the input that matches the waveform seen with the loop closed. Even if every proton in the universe were a Cray computer. If I remember correctly, there are only an estimated 10^72 protons in the universe.

The other problem is that if you cut the input line and start feeding the subject aribtrary inputs, you will cease immediately to get tracking behavior -- just as soon as the subject realizes that the control handle is no longer affecting the input.

So I don't think your proposed experiment can be done, unless I have misunderstood it.

Best,   Bill P.

Date: Mon Mar 29, 1993 7:21 pm PST
Subject: It's Monday

From Ken Hacker [930329]:   Marken quotes Marken (930329) --

"Virtually all social and behavioral science research is based on the assumption that stimulus inputs cause behavioral outputs -- o = f(i). This assumption can be seen in the prediliction for doing experimental research in which an independent (stimulus) variable is manipulated and its effect on a dependent variable (behavior) measured under controlled conditions. The goal of this research is to learn something about the characteristics of the system that turn inputs into outputs -- ie. the function, f(). If, however, the systems being studied are control systems, then this type of research reveals little or nothing about f()."

Marken then asks me:
Is this a fair statement of what you consider to be a hyperbolic PCT claim?

The answer is no, it is not hyperbole, it may be arrogant, but what it definitely IS, is a mix of half truths.

How do you define "virtually all" of the social and behavioral science research? 99%? 95%? 80%? Whatever sounds good?

My point is that many of the paper, articles, and presentations done by social and behavioral scientists reject the equation o = f(i).

As a member of that community, I can say that there are repeated and growing claims against S-R assumptions and the protests against those assumptions has been going on for years.

Why is a predilection for doing experimental research a problem? Are you not doing experiments?

Of course, you do not accept the goals of some experimenters which include describing causal relationships between variables. I understand and appreciate the point about rejecting the IV-DV view for all human behaviors, but are there not some questions about human behavior where they are useful -- if we take causality out of the assumptions? For example, if I test 2 groups (which I will be doing in the fall) of students, one with one type of learning program and one with another, and see what differences there are in knowledge retention, recall, etc., what is wrong with what I am doing?   The answer is NOTHING is wrong with it if I am simply comparing retention and recall differences by program differences. Can I go deeper into student behaviors, perceptions, control?   YES, of course, but I may not need to in order to answer my questions.

Ah, and finally, I found something we can agree on: Yes, social science and behavioral science, as I see them, do not go after control systems as an area of study. I fully support your assertion that they are not contributing to knowledge about control.   I believe that control is essential to human being and that PCT is creating knowledge about control.

I do not believe that I have to attack any perspective to prove my own. Thus, I am willing to try understanding PCT while I stand neck-deep in other views of human behavior. The reason is that I don't see the contradictions that you posit, although I do see that PCT gets at control where the others neglect it.

As for lumping all social and behavioral scientists into one large army, I have to say the references remind me of people complaining about how biased "the media" are.   The Media?   Anyone who has ever looked at media knows that they are disparate, competitive, and hardly consonant on what they say and do at all times. I think the same is true with behavioral and social scientists. We are odds much of the time and are continuosly jousting with perspectives, data, and epistemology.

ken hacker

Date: Mon Mar 29, 1993 7:45 pm PST
Subject: Re: Mary on Hacker's comments and questions

Ken Hacker [930329]:

Mary and Bill,   Thanks for some rich food for thought. I am now going to control for circuit overload and simply tell you that I am reading your comments and will respond Wednesday. Only one other question for now (I knew something would sneak through):   If we are sets of control systems with muliplexed reference and error signals, what happens to imagery, ideas, feelings, sensations, and all of the things like mental models?   Best, Ken

Date: Mon Mar 29, 1993 8:43 pm PST
Subject: Detecting d(t); ideas and imagery

[From Bill Powers (930329.2300)]

This is positively the last for today.   Jeff Hunter (930929-A) --

If you know the output of the output function, o(t), and you know that p(t) = d(t) + o(t), then you know right away that d(t) = p(t) - o(t). Noise aside, you can completely reconstruct the waveform of the disturbance knowing p(t) AND o(t).

But that's you: you can separately perceive o(t), but the control system can't. To say that there is information about d(t) usable by the control system itself (not some external observer), you'd have to do your reconstruction of d(t) strictly on the basis of p(t), without any knowledge of o(t). And that's what we're claiming is impossible.

In your block diagram I think you have a minor error in the equation for the leaky integrator. It's

o(t) = o(t) + (ko*error - o(t)) * dt/ks

I don't think that's equivalent to what you wrote.

In the diagrams, the upper diagram is right (for the controlled variable), but the lower diagram is shown as if the change in output were due to a change in the reference signal. As I drew the diagrams originally, the lower diagram should also have "d(t) changes" instead of "r(t) changes" where the curve starts to rise. The lower curve shows how the output changes AT THE SAME TIME that the controlled variable is changing, after the disturbance steps suddenly from one value to another. If you changed the reference signal, the controlled variable (upper plot) would rise toward a new value and stay there, not return to zero. Check it out with Simcon.

Ken Hacker (930329) --
> If we are sets of control systems with muliplexed reference and error signals,
>what happens to imagery, ideas, feelings, sensations, and all of the things like
>mental models?

"Multiplexing" means putting several independent signals through one transmission channel, normally by time-sharing (assigning time slices to samples from each channel). A synchronized demultiplexer is required to sort the signals out again at the other end. This is not the HPCT model.

In the HPCT model, all the control systems operate independently and simultaneously. For example, all of the 600 to 800 control systems that operate individual muscles work at the same time, independently of each other. At each level above the first, the perceptual signals of a given control system are sets of perceptual signals from a lower level, put through an input function that determines what aspect of the signals is perceived. One lower-level signal can contribute simultaneously to many different higher-level perceptual functions. All the perceptual signals in a given level and at all the levels, as well, are present at the same time in their own channels. All of this creates a perceptual world in which many sensations, objects, transitions, events, relationships, and so on are present at the same time, totally filling the world of experience.

Imagery and ideas -- imaginary experiences in general -- are, hypothetically, created when a control system of a given level receives a copy of its own output signal at its own perceptual function, instead of sending that signal to serve as a reference signal at lower levels. The result is just as if the lower system had perfectly controlled its input, bringing its perceptual signal exactly to the required reference level and passing a copy, as usual, to the higher system's input function. The difference is that there can be no environmental disturbances of a purely imagined perceptual signal, and there is no limitation on speed of control due to lags in lower systems and delays of actual physical movements. See Ch. 15 of BCP in which some detailed proposals about memory and imagination are laid out.

"Sensations" are level 2 in the hierarchical model.

"Feelings" is one of those terms with different meanings. If you mean emotions, HPCT explains those as a combination of a hierarchical goal and a set of sensations from inside the body, reflecting the biochemical state. If you mean hunches, probably the imagination mode at the program level (level 9) would fit most instances.

How about actually reading BCP? It would answer most of these questions. Your university library should have a copy, or can get one via interlibrary loan.

Best, Bill P.

Date: Mon Mar 29, 1993 9:24 pm PST
Subject: information in controlled perception

[From Rick Marken (930329.2000)] Allan Randall (930329.1700 EST)--

>I will describe a possible
>experiment, predict the results, and provide some analysis. You can
>respond as to which of these three, if any, you have a problem with.

Wonderful!

>Let's take a computer simulation of a simple control system

Excellent idea. Do this information measurement with a simulation of a control system.

>We would first run the control system under normal closed-loop conditions,
>recording disturbance and output, and noting that the disturbance is
>almost perfectly countered at the output. The control system's perceptual
>line would then be cut, and superceded by a new experimental line. All
>possible perceptual inputs would then be presented, one at a time,
>starting with the shortest and working up.

Let me see if I have this right. I assume that cutting the "perceptual line" means that the system is now operating open loop. Perceptual signals are injected directly into the system as Pi, the system does its usual job of transforming Pi values into output -- we use a simple integrator in many of our models so O := O + k*(R-P) -- but now O has no effect on the Pi values. Is this right? The equation for the open loop output, O, is the same as the one used in the closed loop version of the experiment. Is my interpretation correct so far? Also, what is the length of a perceptual input? Are you referring to the length of the vector of values that make up Pi?

> This would be repeated until
>the disturbance recorded from the first experiment appears on the output.
>Let's assume that the output in the first experiment contained 100% of
>the information about D, since we all seem to agree that this is very
>nearly the case. Now we need only look for a replication of the initial
>experiment, and H(D) will simply be the length of the successful
>percept.

I really have to know what you mean by "length"? If it's the number of samples in the Pi vector, finding the first Pi value that maps into the disturbance could take a hell of a lot of iterations.

>Let's label the possible percepts {Pi} = {P0,P1,P2, ...}. P0 will be
>the null message, which is equivalent to simply cutting the perceptual
>line and not providing any replacement input at all. Let's call the
>first Pi to successfully replicate the initial output Pk. Pk might
>very well be equal to the original P, but it could be shorter. We are
>guarenteed that it will not be longer. The entropy of the disturbance is:

>   H(D) = |Pk| <= |P|

Assuming that "length" is the number of samples in Pi, then what you are saying is that H(D) is proprtional to the length of an input perceptual signal that is put into the open loop model and from which O = D is recovered. So if D had 1000 samples and I could recover all 1000 values by plugging in a Pi with 500 samples, then the entropy of D is 500? Is that right?

>We then perform the same procedure to compute H(D | P). This time,
>P is a given, so it is always provided to the perceptual line "for free"
>and not counted in |Pi|.

Now I really have to be clear about this. To find H(D|P) you do the SAME THING you did to find H(D)? So you inject a perceptual signal into the open loop (as you did before) and find the resulting disturbance predictions? Is this right?

> This time we will be successful on the very
>first attempt. Since this setup is equivalent to our initial experiment,
>without any extra input to the perceptual line, the first successful Pi
>will be P0, the null message. The entropy of the disturbance given the
>percept therefore is:

Now I'm confused. Is P coming in via the usual closed loop perceptual input or "for free" as an open loop input? This is a VERY important point. If P in coming in closed loop then you are not computing the distrubance in the same way as you were when you injected Pi.

In the first part of the experiment (if I understand you correctly) you are computing O = O + (R-P) open loop and seeing if the resulting O matches D. The first (and sortest) P that gives you this match is the measure of the entropy of the disturbance (measured as the number of samples needed to reduce the uncertainty). Now you want to find the entropy GIVEN that you know the P that "works" -- call this P'. I think the appropriate way to measure this is the same as you measured H(D) -- ie. compute O := O + (R - (P'+P0)) open loop as you did before. Since P' presumably produces the values of O that match D, adding P0 to P' should just screw things up -- which is why you imagine the length of P0 should be zero so that

> H(D | P) = |P0| = 0

Is the above a fair description of what you propose? If so, I approve and would recomend that you do the experiment.

But let me know if I have it right before you start programming.

>My question is: what is wrong with the above reasoning?

Nothing, if my understanding of it is correct.

>Do you actually support the idea that the system can control blind?

Of course not. Control is the control of perception.

>If not, for your claim to be right, I must have goofed up in my reasoning.

If I have interpreted your reasoning correctly then there is nothing wrong with it.

>Do you disagree about the results that I have assumed we would get
>from this experiment?

Yes. But you'll see when you actually do the experiment.

>Do you disagree with the way the problem was divided into program,
>language and output?

I don't think so. Looks OK.

>Do you disagree with the
>derivations? Definitions? The discrete nature of the computer simulation?

Nope. Looks good.

>There are many places in the above argument where our disagreement might
>lie. If we can agree as to where we disagree, we just might be able to
>come to an understanding.

There is only one possible place where we might disagree. This is in how P' is injected into the system. If it comes in via the usual operation of the loop, then I will not accept the proposal. There are two reasons for this. The first is logical; in the computation of H(D|P) you are proposing to measure the length of the P0 that is added to P' to increase the ability of P' to produce D. In order to try more than one added P0 string, you must have the SAME P' available each time. If P' enters through the operation of the closed loop it is not true that P' will be EXACTLY the same on each repetition of the operation of the loop. So you must have P' available (as an independent variable, so to speak) in order to test for improved prediction with several added P strings. I know you expect to need only one P0 value -- that any a P0 value of length 0 can be added to P' and maintain the ability of P' to produce O=D. But you do have to try adding at least ONE P0 string (other than the zero length one) to show that this is true.

The other reason also seems logical (to me) but I don't think it is necessarily as strong. The reason is "fairness" really -- you are measuring H(D) by injecting P strings open loop; it seems reasonable to measure the change in entropy with P present (H(D|P)) in the same way -- that is P, is injected open loop to see how well it predicts the disturbance.

Hope we are in agreement.

Best Rick

Date: Mon Mar 29, 1993 10:29 pm PST
Subject: Re: Detecting d(t); ideas and imagery

From Ken Hacker [930329]:

Bill, I have BCP. In fact, our libary has two copies. I confess that I need to do some re-reading, but most of my questions are spontaneous.

Ken

Date: Tue Mar 30, 1993 3:44 am PST
Subject: TGIT?

From Greg Williams (930330) Bill Powers (930329.1830)

>>I think if we perused the refereed behavioral science
>>literature together, we'd find a lot of cases where
>>correlations are found between two population measures, and no
>>generative model for IV-DV variables of individuals is supposed.

>You have better access than I, although I'll look, too. Let me
>know what you find, when you have time.

Will do.

>I think that you and I will manage to find some of your controlled
>variables even under the present circumstances. You're not that
>hard to figure out.

Do you mean "putative" controlled variables that SEEM to make sense, like certain psychoanalytic "explanations" of individual behavior SEEM to make sense?

>>Maybe Bill can convince me that there aren't any problems with
>>the kind of introspection he wants me to use.

>It sounds as though you don't want to try a method of

>introspection that you see as having some problems. Is this true?

YOU are the one who is trying your method. I just work here (as the judge of whether you've understood my behavior).

>>Or maybe he will come up with a different way of attempting to
>>meet my challenge without relying on my own impressions.

>I trust your impressions. Don't you?

Sometimes, but not always. It remains to be seen what credence I'll assign to my impressions related to the challenge.

>>Shouldn't the Test work even on non-verbal creatures? In that
>>case, such access would be moot.

>Pretty tough to apply the Test for control of nonverbal perceptions over the
>net.

Cute, but beside the point I was making.

>>Maybe it would be better for YOU to be the judge?

>Oh no, I would have won already. That would be too easy.

And too unconvincing to observers, especially if I protested your judgment?

>>... what conclusion would you draw about the possibility of
>>understanding human behavior by ANY means?

>>It all depends on the dynamics of disappearing and reappearing,
>>I guess. Having access to the detailed dynamics some way
>>(electrical probes?) might in principle be able to reconstruct
>>individual motivations backward in time, if chaos isn't
>>problematic (and it might not be, if the dynamics were governed
>>by limit cycles).

>An informative answer. Let me try some test questions now, for
>which I think I know what the answers will and will not be.

>1. Do you want me to go on questioning you in this same way?

Yes.

>2. Do you want this process to end up with a valid understanding
>on my part of one of your controlled variables?

Yes.

>3. Do you sense that I will fail or succeed in meeting the challenge?

Yes. (Unless there's something in-between a binary possibility!)

>I predict that your answers will be, very approximately:

>1. That's up to you, I don't know what I want to happen, there's
>no evidence of any desire one way or the other.

Wrong.

>2. I have no idea how it will end up; I can't see evidence for
>any preference.

Wrong.

>3. My senses don't lead to any prediction; I'm just waiting to
>see how this will turn out.

Wrong, sort of.

>I predict that the answers will NOT be

>1. Either Yes or No.

>2. Either Yes or No.

>3. Either Yes or No.

Wrong.

>Having revealed my predictions ahead of time, I expect you to
>play fair. Note that I still haven't said what I think you're
>controlling for. I'm still gathering evidence, without a lot of
>help from you.

Playing fair,

Greg P.S. Thanks for the SIMCON stuff.

Date: Tue Mar 30, 1993 6:37 am PST
Subject: language is not behavior

[From: Bruce Nevin (Tue 930330 08:36:29)] Rick Marken (930329.1230)

> > "Flying planes can be dangerous" is
> >ambiguous only in isolation.

> I agree. But I think that the same is true of ambiguous
> non-linguistic behavior too (I'm saying that all behavior,
> linguistic and non-linguistic, is ambiguous inasmuch as we
> can never be positive (without doing the Test) what controlled
> variable might be involved).

I'll repeat something from my (Fri, 26 Mar 1993 11:56:11):

> Ambiguity is quite apart from the differences in nonverbal
> perceptions (non-category perceptions) that we associated with
> utterances in an idiosyncratic way that is not socially
> standardized. You can't call that ambiguity, because ambiguity

> is a choice between structurally defined alternatives.

Now, I was wrong to say "you can't call that ambiguity". You can call anything you like anything else you like, of course. And in particular there is a common sense of "ambiguous" meaning that something is ill defined. So let's use "ill defined" for the indeterminacy of observed behavioral outputs w.r.t. the reference perceptions of the observed control system. (Or if you want to keep the word "ambiguous" add some qualifier, "ambiguous because of its indefiniteness" or some such.)

The linguistic ambiguity that we see in "flying planes" is not ill defined at all. The two alternative linguistic structures (linguistic information) that we perceive in that utterance are both quite well defined. Even the fact that there is more than one structure is quite well defined. It is well defined because those structures are socially available, common knowledge of anyone who knows the language, by the very nature of what "knowing the language" is.

Now, when you successfully perform the Test for the controlled variable, the reference perception back of some observed behavioral output becomes well defined in two respects at once: the reference perception becomes socially available, common knowledge of anyone who has observed (and understands) the Test; and there is only one reference perception put in correspondence with the behavioral outputs, as opposed to the indeterminate range of possible reference perceptions (including none, if the behavior was incidental, not a consequence of perceptual control). But notice that the range of possibilities was indeterminate, not socially available prior to the Test. In language, ambiguity is structurally determined, that is, the range of possibilities is socially available prior to disambiguation.

Furthermore, with linguistic ambiguity one need not always perform the Test to determine which linguistic structure (which linguistic information) the speaker intended. Frequently, more careful attention to context suffices. And suffices not in a probable, statistical sort of way, but precisely and determinately. The speaker would have had to change the subject twice in quick succession, once just before the sentence and again back to the original subject just after. A change of subject is signalled by structural cues in the discourse (e.g. "by the way"), as well as by intonation and other gesture in face-to-face interaction. Such a change would violate conversational norms. Most importantly, it would be incoherent, not fitting the informational structure of the ongoing discourse. If the intention of the speaker was in fact to shift topic and use the ambiguous utterance in a non-obvious sense, then she would take steps to make sure that you knew it (structural cues, intonation, etc.). Prevarication is possible, NB, even with the Test: the observed person can manipulatively pretend to control a perception that in fact she doesn't care about. She has to control that perception in fact at the time of the Test, but likewise a person exploiting linguistic ambiguity has to speak as though intending one meaning, while secretly holding an alternative meaning in reserve.

> language is like all other behavior -- you can't tell
> (for sure) what constitutes a person's behavior (an intended
> consequence of their outputs) without doing "The Test" in order
> to determine what perceptual inputs they are controlling.

Language is not behavior. Speech is behavior. Language is socially standardized reference perceptions, the control of which has speech as a byproduct. The reference perceptions are socially available without recourse to the Test because they are learned (and taught--cf. Bruner) as socially standardized perceptions.

The observed speech varies from the socially standardized references in ways that can at best be described statistically. It is common for people to control for different norms with respect to different other perceptual control (am I speaking "correctly" or as a non- alienating member of my group or am I slumming, etc.). This usually adds to the variability of speech, as people alternate from one norm to another depending on their perception of their social situation, but it can result in conflict.

(I think this perspective, familiar to linguists, answers Bill's qualms about Labov's work.)

Bruce bn@bbn.com

Date: Tue Mar 30, 1993 7:00 am PST
Subject: Challenge 2

[From Bill Powers (930330.0700)] Greg Williams (930330) --

>YOU are the one who is trying your method. I just work here (as
>the judge of whether you've understood my behavior).

One of the assumptions I have to make in trying to discern a control process is that there is some control process going on: outputs that affect the external world to produce an effect intended by the outputter. A person who chooses to play the role of passive spectator and commentator is not obviously controlling anything, so there is little to push against. However, it's not hard to see the resistance to being drawn into conjecture, so I assume that you have a low reference level for indulging in conjecture (or introspection). This, I will admit, does make the challenge a little harder (although I submit that the foregoing analysis captures some of the situation -- is that right?).

>It remains to be seen what credence I'll assign to my
>impressions related to the challenge.

This suggests that you're waiting for something to reach a point at which you will experience a definite impression one way or the other. As long as there is no definite impression, you will continue to feel that the challenge has not been resolved. It also suggests a higher level, such that when some impressions actually do occur, they will not be taken at face value, but will be judged against some other criterion and only then given credence or not. Do I interpret your meanings correctly?

>>[You be the judge?] Oh no, I would have won already. That would be too easy.

>And too unconvincing to observers, especially if I protested your judgment?

That's what I meant -- too easy to be convincing, even to me.

Test questions:

My surface hypothesis was evidently wrong, because it led me to wrong predictions of each of the answers I thought you would NOT give, and also wrong predictions as to each of the answers I thought you WOULD give. This is essentially what I expected according to a deeper hypothesis.

Previously, when I asked questions about what you wanted, you said that you had no contact with your inner motivational structure and therefore didn't know, or could only guess. This time, when I asked what you want, you gave a clear YES answer to two of the questions, and when asked about success or failure preferred to predict some degree of success between the extremes.

With regard to the first two questions, do the YES answers still stand today?

If so, we would seem to have a glimpse of two controlled variables: (1) a perception that I am questioning you in a particular way, and a reference signal specifying that this perception should continue rather than cease, and (2) a reference signal specifying that a perception of a valid outcome of this process would be desireable, and (judging from the fact that you have not yet accepted the success of the process), a perception that the outcome is not yet perceived as valid. The implied error signals mean that you are not yet ready to judge that the challenge has been met and will continue participating.

As to the third question, you indicated a preference for an answer somewhere between absolute success and absolute failure. May I correctly interpret this as a reference level regarding the concepts of success and failure, indicating that you prefer a continuous scale over a binary one in this case?

This is, of course, a model-based interpretation as are all my interpretations. You're the only one in a position to say whether the interpretations fit what you observe. How about it?

Best,   Bill P.

Date: Tue Mar 30, 1993 7:23 am PST
Subject: IV-DV

[From Bill Powers (930330.0800)]   Ken Hacker (930329) --

>I understand and appreciate the point about rejecting the IV-DV
>view for all human behaviors, but are there not some questions
>about human behavior where they are useful -- if we take
>causality out of the assumptions? For example, if I test 2
>groups (which I will be doing in the fall) of students, one
>with one type of learning program and one with another, and see
>what differences there are in knowledge retention, recall,
>etc., what is wrong with what I am doing?

I don't want to say there is something wrong with what you're doing, but I would like to know what you ARE doing. As you actually do experiments of the kind you describe, you can probably answer some questions about it (I seem to be in a question-asking mode left over from talking with Greg Williams).

If you do this experiment and find that there is a significant difference in learning between the groups, to what will you attribute this difference?

If one learning program is associated with a significant improvement in learning, would you recommend using it in the future, over the other?

My last question would probably be answered best using the results of some study like this that you've already done. Since I don't know what it is, I'll frame the question generally.

Assuming that there is some treatment that differs between two groups, and that the different treatments yielded a significant difference in performance between the two groups:

1. How many people took part in the study?

For each of the treatments:

2. How many people clearly performed better,

3. how many clearly performed worse,

4. and how many did not clearly perform better or worse?

Best, Bill P.

Date: Tue Mar 30, 1993 9:50 am PST
Subject: What a day!

[From Dag Forssell (930330-09.00)]

What a day indeed. 33 messages, 142,768 bytes. A record for one day? It is becoming very hard to even scan, much less read - unless you live for the net only. A lot of clarification and much worth saving. I am hanging in there. No complaints. Best Dag

Date: Tue Mar 30, 1993 10:17 am PST
Subject: Re: ambiguity

[From Bill Powers (930330.1100 MST)] Bruce Nevin (930330.0836) --

Another possible definition of ambiguity: an input that is perceived in two different ways at the same time. To say that it is perceived means that there are two perceptual structures, well-defined, that already exist, and that the input satisfies the conditions needed for both of them to produce a perceptual signal.

This definition applies outside language, too: If I get in my car, drive it 25 feet, and turn it off again, two perceptions are produced: (1) I have uncovered the part of the road where the car was, and (2) I have covered the part of the road where the car now is. Both results are outcomes of the action; most likely, only one (at most) is an intended outcome. If I were "punning," both could be intended outcomes -- using one action to create two perceptions.

The Test would be needed for a second party to find out which, if either, was the intended outcome.

>The observed speech varies from the socially standardized
>references in ways that can at best be described statistically...

>(I think this perspective, familiar to linguists, answers
>Bill's qualms about Labov's work.)

It does.

Best, Bill P.

Date: Tue Mar 30, 1993 10:33 am PST
Subject: SIMCON fixes

[From Bill Powers (930330.1110)]

RE: SIMCON44

Some of you who downloaded SIMCON may have found that it needed Borland Graphics Interface files when I said it didn't. Others with 8088 processors may have found that it wouldn't run at all. Those problems have been fixed, and I've uploaded a new SIMCONZ.EXE to Bill Silvert's server at biome.bio.ns.ca. It should be available shortly (when Bill says it is). Sorry for the bother.

The new file is bigger because it includes a new primer1 and primer2 that are organized around SIMCON.

To recap: SIMCONZ.EXE is a self-extracting zipped BINARY file, for IBM-compatible DOS PCs only. When you execute it, it expands into a series of files, one of which is SIMCON.EXE. After that works, you can delete SIMCONZ.EXE. SIMCON.DOCiled instructions for using SIMCON and a couple of examples.

Also on The Silvert Server:

SIMCONZ.UUE is an ASCII file, which must be unpacked with uudecode, a program available from several sources via ftp. SIMCONZ.ASH is a "burned" file≈F+~? using Pat and Greg Williams' program; you need UNBURN.EXE to turn it back into binary.

You can also obtain SIMCONZ.EXE from me at

Bill Powers
73 Ridge Place, CR 510
Durango, CO 81301

or from Wolfgang at

Wolfgang Zocher
Hauptstrasse 15
3225 Capellenhagen
Germany

Send a stamped self-addressed disk mailer with a formatted 3.5-
inch or 5.25-inch floppy in it.

Best, Bill P.

Date: Tue Mar 30, 1993 10:55 am PST
Subject: Re: Allan Randall's proposed experiment

[Allan Randall (930330.1330 EST)] Bill Powers (930329.1930 MST) writes:

> >...All possible perceptual
> >inputs would then be presented, one at a time, starting with
> >the shortest and working up...
> ...To reproduce an arbitrary
> disturbance waveform in this way you would have to try, on the
> average, half of the possible strings of 1800 numbers, which
> amounts to something like (32^1800)/2 possibilities. ... If I
> remember correctly, there are only an estimated 10^72 protons in
> the universe.

Yes, to do this experiment for real would require either a very small control experiment, or some very big simplifying assumptions. I was hoping that we could agree on the outcome without the need to actually do the experiment. The number of possible programs for most languages you can think of grows exponentially with the program size. This, however, is a legitimate problem in trying to calculate the "true entropy" of something. If you want to be sure you have the "minimal program" for something, you will have to check all programs that are shorter and make sure none of them produce the same output. Nonetheless, for a sufficiently restricted control task, such an experiment could be done.

However, I do not see why we should need to do the experiment at all (or at least not any of it requiring exponential search). Showing that H(D|P) = 0 is trivial and involves no exponential growth. One of the assumptions was that the original experiment produced output with 100% of the information about D. Thus, computing H(D|P) will *exactly* replicate the output of the first experiment on the first try: P0. So we can conclude that H(D|P) = 0 without actually doing the experiment. Rick claims that H(D) = H(D|P), so it follows mathematically that H(D)=0, which is equivalent to blind control. If we can all agree that blind control is not an interesting case of control, there is no need to compute H(D). Rick's claim has been disproved before the experiment got past the starting gate. We don't actually get around to doing any exponential search of P-space. In other words, actually doing the experiment is mostly redundant. The main point of the experiment is that it shows Rick's claim to be either logically unsound *or* a claim for blind control.

> The other problem is that if you cut the input line and start
> feeding the subject aribtrary inputs, you will cease immediately
> to get tracking behavior -- just as soon as the subject realizes
> that the control handle is no longer affecting the input.

I think you misunderstood on this point. I was talking about a simple control system completely under our control - a small simulation on a computer - not a real world experiment with human subjects. That would be most impractical (although still possible as a thought experiment).

Allan Randall

Date: Tue Mar 30, 1993 10:59 am PST
Subject: Re: ambiguity

[From: Bruce Nevin (Tue 930330 13:16:36)] [Bill Powers (930330.1100 MST)] --

Thanks for the note, Bill, and the ack re Labov.

>The Test would be needed for a second party to find out which, if
>either, was the intended outcome.

In the case of nonverbal perception, the ambiguity is in the observer's perceptions, with no basis for assuming that it is also in the observed's. This is why you must say "which, if either" (or "if any"). Did the person moving the car even perceive the space that the car covered, before or after moving it? What if she was controlling line of sight, or shade? She could move the car a few feet or a car length without looking at the surface before or behind the car.

In the case of language, the ambiguity, the alternatives of structure (information), is a socially established property of the utterance. You need some version of the Test to determine which, but you would never hedge the question with "if any".

(The version of the Test typically is to proceed with conversation as though the interpretation you have assumed-- often without considering that there might be alternatives-- is the correct one. If there's a wrong assumption, error might arise in one or both of you farther on. If you never notice, it didn't matter [yet]. In other words, one's acting on expectations in conversation and cooperative interaction constitute many repetitions of the Test, though not intentionally so, because of which results are often misinterpreted. And so on, yakatayakatayak.)

Bruce bn@bbn.com

Date: Tue Mar 30, 1993 12:17 pm PST
Subject: ... and on....

From Greg Williams (930330 - 2) Bill Powers (930330.0700)

>A person who chooses to play the role
>of passive spectator and commentator is not obviously controlling
>anything, so there is little to push against. However, it's not
>hard to see the resistance to being drawn into conjecture, so I
>assume that you have a low reference level for indulging in
>conjecture (or introspection). This, I will admit, does make the
>challenge a little harder (although I submit that the foregoing
>analysis captures some of the situation -- is that right?).

Gee, Bill, it sounds so REASONABLE! But do you really want to stake a science's reputation on introspectively determined reasonability? I don't have problems with the hypothesis that "everything is perception" and so one cannot escape the subjectivity of individual judgments (including one's own), but I think it would be a lot MORE reasonable to perform tests aimed at understanding (by experimenter and subject alike) one's behavior which are not COMPLETELY dependent on possibly fallacious subjective reports. Appropriate use of the Test for Controlled Variables needn't be so dependent on subjective reports, I think -- or am I wrong? (It seems obvious that the Test could be applied to non-human animals -- not over the net very easily, of course, as you so wittily pointed out yesterday.)

But (a big "but," I think), the Test cannot be applied to behavior which happened in the past (at least not with current technology), and the Test, as you have noted, is difficult to apply in field conditions. The point I think (but don't know!) I'm trying to make NOW (not necessarily when I first issued the challenge) is that understanding (and especially predicting) complex individual human behaviors in the field by using The Test looks to be much more difficult than laboratory tracking experiments. Rick, especially, is prejudging a quite open issue: how significant PCT's contributions to the aims of behavioral scientists in general will ever be. Granted that statistical description of population measures (if nonfallacious) cannot result in understandings of individual behaviors, it is at least possible that, IN PRACTICE, PCT won't be able to generate population measures which some behavioral scientists will continue to desire.

Back to the challenge.

>>It remains to be seen what credence I'll assign to my
>>impressions related to the challenge.

>This suggests that you're waiting for something to reach a point
>at which you will experience a definite impression one way or the
>other. As long as there is no definite impression, you will
>continue to feel that the challenge has not been resolved. It
>also suggests a higher level, such that when some impressions
>actually do occur, they will not be taken at face value, but will
>be judged against some other criterion and only then given
>credence or not. Do I interpret your meanings correctly?

I can't presume to say that your interpretation is "correct" or not. It SOUNDS reasonable. How can either of us know when reasonability equals correctness?

>>That's what I meant -- too easy to be convincing, even to me.

>Previously, when I asked questions about what you wanted, you
>said that you had no contact with your inner motivational
>structure and therefore didn't know, or could only guess. This
>time, when I asked what you want, you gave a clear YES answer to
>two of the questions, and when asked about success or failure
>preferred to predict some degree of success between the extremes.

>With regard to the first two questions, do the YES answers still stand today?

Yes.

>If so, we would seem to have a glimpse of two controlled
>variables: (1) a perception that I am questioning you in a
>particular way, and a reference signal specifying that this
>perception should continue rather than cease, and (2) a reference
>signal specifying that a perception of a valid outcome of this
>process would be desireable, and (judging from the fact that you
>have not yet accepted the success of the process), a perception

>that the outcome is not yet perceived as valid. The implied error
>signals mean that you are not yet ready to judge that the
>challenge has been met and will continue participating.

Again, that all SOUNDS reasonable. Now, assuming that the reference signals and implied error signals are as you say (heck, the reference signals are "implied," too), why are they as they are, and not otherwise?

>As to the third question, you indicated a preference for an
>answer somewhere between absolute success and absolute failure.
>May I correctly interpret this as a reference level regarding the
>concepts of success and failure, indicating that you prefer a
>continuous scale over a binary one in this case?

Ditto.

>This is, of course, a model-based interpretation as are all my
>interpretations. You're the only one in a position to say whether
>the interpretations fit what you observe. How about it?

I don't think I'm in that position. It doesn't feel like I'm introspecting about reference signals and such when I am behaving, except when the behavior is introspecting about reference signals. What if when I think I'm introspecting about reference signals I am deluding myself? Applying the Test "directly" to the behaviors which you want to understand would get around this problem; too bad the Test works well only in circumscribed situations (at least, that's my hypothesis, based on the evidence to date).

As ever, Greg

P.S. Did you e-mail me the NEW Simcon? May I have the new primers, also? Thanks in advance.

Date: Tue Mar 30, 1993 12:24 pm PST
Subject: Re: Allan Randall's proposed experiment

[From Rick Marken (930330.1130)] Allan Randall (930330.1330 EST)--

>I was hoping that we could agree on the outcome without the need to
>actually do the experiment.

A very common sentiment when the cause-effect crowd meets the PCT loonies. Hang in there; you're in for some BIG surprises when you start to run the experiments.

>Showing that H(D|P) = 0 is trivial and involves no exponential growth.

Well, I think you'll find showing this to be quite NON-trivial. Try the experiment!

>Rick's claim has been disproved before the experiment got past the
>starting gate.

That would make things a lot easier for your horse, indeed. But don't scratch me yet. Just fire up the ol' simulator and see what she does. Ol' "Stewball" Marken might surprise you.

Best

Rick ("There's no information about disturbances in controlled perceptions") Marken

Date: Tue Mar 30, 1993 1:11 pm PST
Subject: Re: Mary on communications loops

[From Dick Robertson] 9303.30
Very nice Mary. I wish I had time to comment in more detail. I think I tried to say something similar in my chapter on social psychology in the textbook, but your reference to the rubber band demo really makes it nice and neat.
best, Dick R.

Date: Tue Mar 30, 1993 1:27 pm PST
Subject: They call me Dr. Marken

[From Rick Marken (930330.1300)] Ken Hacker [930329] --

In reply to my statement that "Virtually all social and behavioral science research is based on the assumption that stimulus inputs cause behavioral outputs -- o = f(i)." Ken says:

> How do you define "virtually all" of the
> social and behavioral science research? 99%? 95%? 80%? Whatever sounds good?

I define it as >99%. I was a conventional scientific psychologist for many years. I even wrote a well - received textbook on research methods and statistics in experimental psychology that might be in your library (Methods in experimental psychology, Brooks/Cole, 1981). I think I'm pretty familiar with the methods and assumptions of scientific psychology and I've looked at a hell of a lot of the research articles published in MANY fields -- including perception, physiological, cognitive, social, developmental, operant behavior, clinical, etc, etc. I think I have a pretty good idea of what's out there in my field (psychology). I know that there are many "correlational" studies where the authors celebrate their undergraduate awareness that "correlation does not imply causality"; but all this means is that the researcher admits that a causal relationship between input and output variables is not established by the results; but the clear implication is that such a relationship could have been found if the study were (or could be) done properly -- ie. with the appropriate controls. The causal model is still assumed to characterize behavior -- even when you methods don't let you unambiguously determine the causal variable(s).

In my 12+ years as a conventional psychologist I ran across only two or three experimental research articles that did not assume a lineal causal model of behavior. These articles were published by W. T. Powers. I would estimate that I've read well over 1000 social science research articles. I can think of 3 articles by Powers that described non - causal model based research. So that means 997/1000 causal based articles. So my estimate of "virtually" is 99.7%. This is between the years 1965 to 1979. Since then, Tom Bourbon and I have added some research papers to the non-lineal causa-based research collection. So now, out of 1000 research articles you might find only 996 that are based on a cause effect model; so the "PCT revolution" has brought "virtually" down to about 99.6%. Were movin' right along.

>My point is that many of the paper, articles, and presentations done by
>social and behavioral scientists reject the equation o = f(i).

I'm sure they do; but they still do research as though that assumption were true. The rejections of the o = f(i) mdel that I have read are nothing more than verbal shenanigans. Just look at the Karolyans -- a group that ostensibly understands and accepts the basic principles of the PCT model of behavior. Once they get into the lab it's IV-DV all the way.

But I would love to have a reference to one or two experimental research articles that do not assume an underlying lineal causal model -- other than those done by PCT people, that is.

>Why is a predilection for doing experimental research a problem?
>Are you not doing experiments?

Experimental research is essential; the IV-DV approach to doing experiments (which assumes an underlying causal model of behavior) is NOT the only way to do experiments. Read about "The Test for Controlled Variables" in BCP. But first read BCP and do some of the "simple" demo experiments that illustrate what is going on. You won't understand the Test unless you understand PCT.

>I understand and appreciate the point about rejecting the IV-DV
>view for all human behaviors, but are there not some questions about
>human behavior where they are useful -- if we take causality out of
>the assumptions? For example, if I test 2 groups (which I will be
>doing in the fall) of students, one with one type of learning
>program and one with another, and see what differences there are in
>knowledge retention, recall, etc., what is wrong with what I am
>doing? The answer is NOTHING is wrong with it if I am simply
>comparing retention and recall differences by program differences.

There is certainly nothing morally wrong with it (until you start applying these statistical results to individual human beings). I just find studies of this sort (and there are thousands of them being done all the time) a tragic waste of time, money and intellectual talent.

>I do not believe that I have to attack any perspective to prove my own.
>Thus, I am willing to try understanding PCT while I stand neck-deep in
>other views of human behavior. The reason is that I don't see the
>contradictions that you posit, although I do see that PCT gets at control
>where the others neglect it.

This is very humane and decent sounding of you. It's difficult, however, to say things like "there is no information in controlled perceptions" without at least implicitly stepping on some toes. But what's wrong with showing that a perspective is wrong, anyway? Was it wrong for Galileo to show that Ptolemy was wrong? Was it wrong for Darwin to show that God was wrong? What's wrong with showing that an idea is wrong if the idea IS wrong? If you think there is some value to a "perspective" that we are "attacking" then just defend it; show that we are wrong. This is not a religious war -- this is how science works. In fact, PCT WELCOMES attacks -- ie. disciplined attempts to show that PCT is WRONG. That is why I'm so thrilled that Allan Randall is trying to show that PCT is wrong about the "no information in controlled perception" claim.

I think it's a little over-sensitive to feel like someone is "attacking" you or your ideas when they say that your ideas are wrong -- especially when they try to provide evidence to back up what they say. I don't feel attacked by Allan Randall. I feel (very pleasantly) challenged.

That's the nice thing about science. When you have a theory that happens to be correct (or, at least, more correct than other theories of the same phenomenon), then there is no threat when people say or try to demonstrate that your theory (or some claim you make about it) is wrong. The only people who feel threatened by "attacks" on their theories are people who 1) MUST be right and/or 2) have nothing but words to support their theories -- that is, religious people. The Catholic church felt "attacked" when Galileo SAID that the world was stationary because this contradicted a theory that HAD TO be right -- a theory that was backed up only by a bunch of words in an old book. Well, who gives a flying **** what the stupid book SAYS. The fact is that there is evidence that the world turns; if the church were really confident about their theory, they would have been happy to subject it to critical test. But the only critical test was to see what the book had to say about this. I'm more interested in what it says in the real book -- the book of my own experience.

So don't fret about PCT "attacks" on what you hold dear. Just show us how we're wrong. At the moment there is only one person who seems to be willing to do this with real demonstrations -- good ol' Allan Randall. I hope that he continues to attack with a frenzy those aspects of PCT that he thinks are wrong; it's the only way we can get things done here in the non-religious world.

Best   Rick

Date: Tue Mar 30, 1993 1:50 pm PST
Subject: SIMCON fixes installed

>[From Bill Powers (930330.1110)]   RE: SIMCON44
>
>Some of you who downloaded SIMCON may have found that it needed
>Borland Graphics Interface files when I said it didn't. Others
>with 8088 processors may have found that it wouldn't run at all.
>Those problems have been fixed, and I've uploaded a new
>SIMCONZ.EXE to Bill Silvert's server at biome.bio.ns.ca. It
>should be available shortly (when Bill says it is). Sorry for the
>bother.

The new version was installed at 17:30 AST on March 30.
--
Bill Silvert at the Bedford Institute of Oceanography
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2
InterNet Address: silvert@biome.bio.ns.ca
(the address bill@biome.bio.ns.ca is only for mailing lists)

Date: Tue Mar 30, 1993 1:53 pm PST
Subject: INTRODUCTION AND A COMMENT

(930330)

I suppose that it is time to get involved. My name is Dan Miller, and I am a sociologist at the University of Dayton. When people press I tell them that I am a social psychologist interested in interacting perceptual control systems. I have been a member of CSG for six or seven years, and I have been reading CSGNet daily for eight months. This is my first dip into the pond.

The dialogue between Mary Powers, Ken Hacker, Bill Powers, and Rick Marken on perceived arrogance convinced me to make the plunge. Then the March 29th [930329] barrage of posts almost seemed to smother the topic. Still, here goes...

The arrogance perceived by Ken Hacker is, I believe, what my kids call ATTITUDE, as in "It is better to give than to receive attitude." Like Ken Hacker, there are times that I could do without the attitude from my kids, or colleagues, or whomever. But is this arrogance? I have been called arrogant on a few occasions. Each time I was shocked by the accusation, and tried to make sense of the other's perceptions by re-membering what I had done. More pertinently, I have seen this charge leveled at my teacher while in graduate school as an act of bad faith in order to demean and damage.

The graduate school incidents surrounded the publication of a study that I had been involved with in which Carl Couch was the project director. Couch had given up on standard statistical techniques and experimental laboratory procedures, and was developing the use of the small groups laboratory as a provocative stage on which interacting humans may construct social acts - purposive social acts. We videorecorded (and audiorecorded) the proceedings. I coded, mapped, transcribed, and looked for redundant sequences and patterns. No frequencies, no general categories, no statistics. Couch wanted invariant interaction patterns that specified how two autonomous individuals moved from independent action to purposive social action toward a shared focus. It worked. Couch was on a fantastic rush, making bravado claims of scientific revolution. I saw the looks and heard the comments of arrogance. Few colleagues bothered to look at his work (and mine). They were not controlling for scientific explanation. Their bad faith suggests that they were controlling for tenure, promotion, and prestige - part of which was attained through the dismissal of Couch, his students, and his "kind of research." What had we done? We had explained every (or nearly so, Clark) outcome with some precision and elegance given the time - 1973. Of course, now we know that we had just begun to understand how two interacting perceptual control systems construct social acts. I am still trying to learn the language of PCT.

The point is that those claims of arrogance hurt. Their intent was to stifle talk - to censor. Most often they were made by those who neither understood, nor really wanted to understand the implications of this approach to sociology. It seems to me that Bill Powers, Rick Marken, and others with attitude have a great deal to offer - from neurophysiology, to robotics, to psychology, to sociology, to philosophy, to aesthetics. Let's not let their attitude become the issue. It wastes my time.

Dan Miller   Sociology   University of Dayton   Dayton, Ohio 45469-1442

MILLERD@DAYTON

Date: Tue Mar 30, 1993 2:41 pm PST
Subject: Randall's experiment; Challengs 2

[From Bill Powers (930330.1330 MST)]   Allan Randall (930330.1330 EST) --

Sorry to be so slow, but reading your conversation with Rick, I realized that I don't even understand what experiment you're proposing to do (with or without a real subject).

Let's get it into a diagram so I can tell what's connected to what under what conditions. Here's a start, the basic control system. It's simplest just to use an integrator in the output function:

```
                     | ref sig (constant)
                     |
            p----- comp ---- e= error
             |               |
          inp fuct       integrator
             |               |
        i=input           o=output
          +| |+              |
           |  ---------------
           |
           d
```

for simplicity we assume that p = i; the input function is just a unity multiplier.

As I understand it, you want to do a short run and record p, o, and d. Then you want to cut the input line and insert an artificially generated signal a:

```
                     | ref sig (constant)
                     |
           p ----- comp ---- e = error
             |               |
          inp funct       integrator
       /                     |
     a     (i)              (o)
          +| |+              |
           |  ---------------
           |
           d
```

Then you want to start creating all possible signals a, while comparing the output o with the disturbance d. The search ends when o(t) is the same as it was with the loop closed.

In the successful case, d + o will produce a value of i (= d + o) that has the same waveform as a. So you're really searching for the waveform of a that produces a waveform of i such that i(t) = a(t) for all t

But this is no different from simply solving the closed-system equations for the value of the input variable, which is the same as the perceptual variable. If I understand what you mean, you're proposing to find the solution by systematically going through all possible solutions until you find the one that works.

Is this correct so far? I don't want to try to follow the rest until I'm sure what the actual procedure and method are.

Bruce Nevin (930330.1316) --

>Did the person moving the car even perceive the space that the
>car covered, before or after moving it? What if she was
>controlling line of sight, or shade?

Right, those are the questions that the Test would have to resolve. I didn't pick a wonderful example, but you get the idea: actions can have more than one outcome, and generally only one of them is intended, although more than one may intended.

>In the case of language, the ambiguity, the alternatives of
>structure (information), is a socially established property of
>the utterance. You need some version of the Test to determine
>which, but you would never hedge the question with "if any".

Agreed. The same actually holds for any action, because according to the basis hypothesis, the only reason for which any action is ever produced is to maintain control of some perceptual variable. The proviso "if any" was meant merely to acknowledge what you said: that the actual intended outcome of the observed action may not have been on the offered menu of choices. It's conceivable that the purpose of uttering the words "Flying is dangerous" is merely to check that the microphone at the podium at the linguistics conference is working, with neither possible meaning being intended.

Greg Williams (930330 - 2) --

>I think it would be a lot MORE reasonable to perform tests
>aimed at understanding (by experimenter and subject alike)
>one's behavior which are not COMPLETELY dependent on possibly
>fallacious subjective reports.

Subjective reports can be fallacious if they contradict evidence that can be obtained about the same thing without relying on subjective reports. They can't be fallacious if the question is how the world appears to the observer. The investigation of illusions, for example, depends completely on subjective reports: if a person says that the crossbar of a T looks shorter than the stem, one has to accept that report as truthful in order to establish that the illusion even exists.

I can accept your statements as truthful without a qualm, because what I am doing is trying to determine how your experiences appear to you. I'm not trying to catch you up in a mistake by showing that what you report to me isn't true. Of course I have to trust that you will report, as truthfully as you can, what seems to you to be going on in present time. This is no different from asking you if it seems to you that "if A is greater than B and B is greater than C, A must be greater than C." That is a subjective impression which most human beings who reason experience as the truth, even though sometimes it isn't true. I recognize that your words may mean something different to you than they mean to me, but we don't generally do too badly in spite of that. At least we can come to an understanding that can be checked out in other contexts.

You may not have noticed, but I've been applying Tests all over the place, attempting to disturb variables that I hypothesize that you're controlling for, and noting whether you resist or not. So far I've established that it's difficult to get what would normally be accepted as an answer to a question -- that is, a response cast in the same terms as the question. I've found that the answers I do get tend to invalidate my trial hypothesis with interesting regularity. I'm still trying to learn enough about this phenomenon to make it worth while to offer a real hypothesis.

>Appropriate use of the Test for Controlled Variables needn't be
>so dependent on subjective reports, I think -- or am I wrong?
>(It seems obvious that the Test could be applied to non-human
>animals -- not over the net very easily, of course, as you so
>wittily pointed out yesterday.)

Yes, you apply a disturbance to the variable you hypothesize to be under control, and look for systematic resistance to it. In nonverbal situations, this means direct physical interaction of some sort, which we can't do over the net. In purely verbal situations, like the one we are in, it means uttering words calculated to have a certain effect on the words that another person speaks, under the hypothesis that the other has a certain intent in uttering the words. If there is resistance, one is on the track of something controlled; if not, one has to try another hypothesis. Success will depend largely on how well we can each communicate meanings to each other. Fortunately, we speak much the same language.

When I said you were an easy case, by the way, did you experience anything that you didn't tell me about?

>But (a big "but," I think), the Test cannot be applied to
>behavior which happened in the past (at least not with current technology) ...

Right. But I'm interested in controlled variables and control organizations, not behaviors. There's not a lot anyone can do to understand a behavioral organization that appears and disappears in five minutes. But I don't think that real human control processes come and go so whimsically, unless the changes are intentional (in that case, there is still regularity to be found). If you were controlling for something important to you a week ago, chances are that you're still controlling for it. What other assumption can one go on, when trying to determine the characteristics of any system?

>... and the Test, as you have noted, is difficult to apply in field conditions.

Not as difficult as you think. I would say challenging.

>The point I think (but don't know!) I'm trying to make NOW (not
>necessarily when I first issued the challenge) is that
>understanding (and especially predicting) complex individual
>human behaviors in the field by using The Test looks to be much

>more difficult than laboratory tracking experiments.

Well, I don't necessarily buy that. Outside the laboratory there are far more controlled variables to be found. The chances of finding one that will open the door to another are much greater, and the kind of controlled variables that will be found are much more natural. It's harder to quantify them, to be sure, but at the higher levels that's not our immediate concern.

>... it is at least possible that, IN PRACTICE, PCT won't be
>able to generate population measures which some behavioral
>scientists will continue to desire.

Just curious -- is that all you intend to say about my proposal the other day for predicting population demand curves from models of individual behavior?

Back, as you say, to the challenge.

It now seems to me that you are reluctant to report subjective impressions and offer them as answers to my questions. Is this correct? If so, this might explain some of the failure of my questions to elicit answers cast in the same terms. But I won't know whether this impression fits your experiences unless you tell me.

P.S. primers on the way, ASCII. Yes, it was the new Simcon. Unless I've lost another dozen neurons.

Best to all,   Bill P.

Date: Tue Mar 30, 1993 4:16 pm PST
Subject: Re: Randall's Experiment

[From Rick Marken (930330.1500)]   Bill Powers (930330.1330 MST) --

Thanks for the diagrams. I still have no response from Allan to my questions about his experiment; but your diagram's will help me describe what I think he's proposing.

My understanding was that he would do a simulation run with the standard control system as below:

```
                     | ref sig (constant)
                     |
           p ----- comp ---- e = error
            |                |
          inp funct       integrator
                            |
           (i)              (o)
          +| |+              |
           |  ---------------
           |
           d
```

The disturbance would be, say, 1000 samples of filtered random noise. Allan would then save the 1000 values of d, i and o resulting from the run.

Now, in order to measure H(D) he injects a vector of i values (the length of the vector going from 0 to 1000) in a run consisting of 1000 iterations -- and computes the resulting 1000 outputs from each run in an open loop preparation as below:

```
                     | ref sig (constant)
                     |
           p ----- comp ---- e = error
            |                |
          inp funct       integrator
            |                |
           (a)              (o)
```

The candidate i vectors enter through (a) and Allan compares each resulting set of 1000 o values to the actual disturbance vector. The length of the first candidate i vector that produces o values that match d perfectly is H(D) (There is a problem here; What if you don't get any perfect matches to d, Allan?.Nothing but an infinite loop gain system could produce outputs that match d perfectly anyway: how about doing this just until a candidate i vector produces o values that match the disturbance to the same degree as did the o values generated in the closed loop case. So, if the closed loop correlation between disturbance and output was .99967, pick the first i vector that produces o values that correlate .99967 with the disturbance).

Assuming you can find H(D) using this decidedly peculiar technique (why not just measure the variance of d?) the next step is to find H(D|P). My understanding is that Allan would "play" the original i vector (from the actual tracking run) into the open loop system (injected at point a) ALONG WITH the candidate i vectors that were used in the determination of H(D). Since the shortest candidate i vectors are added in first, Allan is assuming that the output resulting from the original i vector along with the "null" (0 length) candidate i vector will produce an output that perfectly matches the disturbance (or, at least, matches the disturbance as well as the output did in the original run).

Is this a correct description of the experiment Allan?

Best Rick

Date: Tue Mar 30, 1993 5:23 pm PST
Subject: BURN/UNBURN programs: enjoy!

From Pat & Greg Williams (930330)

We hereby release our programs BURN.EXE and UNBURN.EXE into the public domain (both source code and executables). Feel free to use them for whatever purposes you have, but be advised that we offer ZERO warranty.

Bill Powers, maybe you could put them where netters can get them if they want?

Best wishes,

Pat & Greg Williams 460 Black Lick Rd.
Gravel Switch, KY 40328 606-332-7606

P.S. If anyone wants copies of the programs direct from us, they should send $5.00 U.S. for handling/postage.

Date: Tue Mar 30, 1993 5:29 pm PST
Subject: ... and on!

From Greg Williams (930330 - 3, I conjecture)

>Bill Powers (930330.1330 MST) Greg Williams (930330 - 2) --

>>I think it would be a lot MORE reasonable to perform tests
>>aimed at understanding (by experimenter and subject alike)
>>one's behavior which are not COMPLETELY dependent on possibly
>>fallacious subjective reports.

>Subjective reports can be fallacious if they contradict evidence
>that can be obtained about the same thing without relying on
>subjective reports. They can't be fallacious if the question is
>how the world appears to the observer.

My understanding what you are supposed to do to meet the challenge I originally issued is to come up with an "understanding" of a particular sample of my behavior which is satisfactory to me. My judgment on that involves more than just questions of "how the world [including what I perceive to be or have been my motivations] appears" to me, doesn't it? Isn't it possible that I could judge, solely on the basis of "how the world appears" to me at time A, that you have met my challenge, and then at later time B, be shown to be wrong by other evidence?

>I can accept your statements as truthful without a qualm, because
>what I am doing is trying to determine how your experiences appear to you.

But my veracity in reporting "how my experiences appear" to me is independent of whether or not "how my experiences appear" to me can be used to establish the validity of your meeting the challenge or not. What I'm saying in other words is that using introspective reports to make claims about introspective reports themselves (i.e., "The cross- bar on that T looks longer than the upright") could be accepted as reasonable, but using introspective reports alone to decide whether an understanding of a behavior is acceptable could be considered unreasonable, because of the questions it begs.

>When I said you were an easy case, by the way, did you experience
>anything that you didn't tell me about?

I (truthfully!) forget, at this time. It is entirely possible that I did, but maybe I didn't.

>If you were controlling for something important to you a
>week ago, chances are that you're still controlling for it. What
>other assumption can one go on, when trying to determine the
>characteristics of any system?

Maybe you can't determine the characteristics of some (many?) systems? How big of a sample did you derive the statistical "chances are" from? (No, I realize you wouldn't make the fallacious attempt to infer individual characteristics from population measures. Honest!)

>>... and the Test, as you have noted, is difficult to apply in field conditions.

>Not as difficult as you think. I would say challenging.

I thought you used the word "difficult" in your post on prediction. Sorry if I misremembered.

>>The point I think (but don't know!) I'm trying to make NOW (not
>>necessarily when I first issued the challenge) is that
>>understanding (and especially predicting) complex individual
>>human behaviors in the field by using The Test looks to be much
>>more difficult than laboratory tracking experiments.

>Well, I don't necessarily buy that.

Really?

>Outside the laboratory there are far more controlled variables to be
>found. The chances of finding one that will open the door to another
>are much greater, and the kind of controlled variables that will be
>found are much more natural. It's harder to quantify them, to be sure,
>but at the higher levels that's not our immediate concern.

Oh, oh. More "chances." Have you done studies on large populations to be so confident of your projected correlations?

>>... it is at least possible that, IN PRACTICE, PCT won't be
>>able to generate population measures which some behavioral
>>scientists will continue to desire.

>Just curious -- is that all you intend to say about my proposal
>the other day for predicting population demand curves from models
>of individual behavior?

Right now, no. I'll add that, from a PRACTICAL standpoint, it will be a stupendous task to try to understand just one behavior of, say, only 1000 people, to judge by the progress of this challenge, and then use the resultant knowledge to predict population demand curves. Maybe I'll intend to (and maybe even actually) say even more some time in the future. Any guesses as to when I might?

>It now seems to me that you are reluctant to report subjective
>impressions and offer them as answers to my questions. Is this correct?

Right now? No.

>P.S. primers on the way, ASCII.

Thanks again.   As ever,   Greg

Date: Tue Mar 30, 1993 7:27 pm PST
Subject: Re: IV-DV

FROM KEN HACKER [930330]:

Bill, I sense that you want me to say that I will attribute increased recall or retention to the learning program. Aha. BUT, I am not doing that. I am going to attribute whatever significant differences there are to what the subjects appear to need to learn what I am trying to teach them with the program. Essentially, the program is a self-paced tutorial where students learn how to use computer conferencing. I am trying to compare the program method (each student learns by him or herself with a pc) versus a lecture-by-the-expert method of training. You see, I am trying to find the best method of teaching-learning from the point of view that the student has needs and those needs will best be met by one method over the other (at least that's my hypothesis). KEN

Date: Tue Mar 30, 1993 8:43 pm PST
Subject: Allen's model; Challenge 2

[From Bill Powers (930330.2000 MST)] Rick Marken (930330) --

I sort of get it. Actually, it's a mistake to assume that o = d instead of using the actual value of o. With an integrator as an output function, the pattern of o will differ from d by an amount depending on how fast d is changing.

Anyway, it's beginning to seem to me that the point is once again being lost. If you know what o is, you can exactly reconstruct the waveform of d from knowing o and p (assuming p is equivalent to the CEV). You don't need to know anything else about the control system. We've never said you couldn't do that, or that some higher-level system capable of perceiving both o and p couldn't do it. Allan is using information about o in his method, as I vaguely understand it now, with your help. It shouldn't be surprising if he can also reconstruct d.

But what if he were limited to the same circumstances as the control system, which has perceptual information only in the form of p and has no information about o? Suppose I gave him JUST the record of the behavior of the CEV, and asked him to reconstruct d (or for that matter, o) on that basis alone?

By the way, I think this notation H(p|d) is not just an ordinary function, but represents some sort of probability calculation with base-2 logs and all that. Allen?

Greg Williams (930330 -3) --

>My understanding what you are supposed to do to meet the
>challenge I originally issued is to come up with an
>"understanding" of a particular sample of my behavior which is
>satisfactory to me.

Explaining a sample of behavior doesn't make much sense in terms of PCT. According to PCT, the challenge was an output intended to bring about some desired perception. Under somewhat different external circumstances, some other output might have been used toward the same end. The same output, for that matter, might have been used for a different end. The output in itself is of little interest. The task under PCT is not just to explain the output, but to find out what it was being used to control for -- what input was wanted. When that is known, all behaviors that bear on keeping that perception under control will be understood. Control of that particular perception may never entail just that behavior again, and recurrance of that behavior may mean nothing regarding a particular goal. Trying to explain a "sample of behavior" is a waste of time. You might say that this is the essence of my objections to conventional psychology.

If you are unaware of what you expected to get out of issuing the challenge (or are suffering from amnesia), it's a little unreasonable to expect me to deduce (from what?) the nature of your goal as it existed at a time that you no longer remember and when you were in a state of mind which is apparently beyond recapture. And I might add, that is beyond my ability to experiment with. All I can do is to try to discern what some current goal is, by seeing how your outputs vary with disturbances now.

Perhaps there is a simple question I should have asked originally, if I failed to do so: what kind of "understanding" would be satisfactory to you? Am I aiming at a moving target, or any target at all? Is there any substance in your challenge?

It has crossed my mind that you are mainly defending against being understood. That would certainly explain a great deal. I seem to recall a time at one of our miniconferences when you announced that nobody could understand you if you didn't want them to. I recall, I hope correctly, your description of a strategy that would always make the Test fail: simply change your mind continually about the nature of the goal. Of course if this strategy were detected, it would reveal the aim of preventing the Test from working, in which case the Test would succeed.

I don't know if such a guess is near the truth, and I'm uncomfortably aware that even to make such a guess might sound like looking for excuses. However, it's a rather obvious guess, so it might as well be put on the table. A scientist who is embarrassed to ask an obvious question won't get anywhere.

This leads to an interesting impasse. Suppose I were to announce that I have discovered your purpose, which is to show that the Test can't work on you. To demonstrate this, all I have to do is use the Test to search for signs of a controlled variable, and show that every time the Test appears to succeed for a moment, it ceases to work. Is this evidence that the Test isn't working, or that it is? Since you are the judge of success, you have both complete control over the success of the Test in any specific situation (other than the overall strategy) and the ability to say that it has failed even when you can see that it has succeeded -- because that would be honestly consistent with the same strategy.

Another consideration is whether the Test would always seem to fail with any person to whom it was applied in similar circumstances. If the Test were to succeed with everyone but you, that, too, would be evidence in favor of my hypothesis -- but you would still have to be the judge, so that kind of "objective" evidence wouldn't wash.

This raises some other interesting questions in my mind, but I'll leave them for another time.

I'll put BURN and UNBURN on Silvert's Server, but didn't you and Pat have a version that could be transmitted in ASCII and then converted using a debug script, also in ASCII? People who can't use ftp won't have much use for your programs if they can be accessed only through ftp, and if they have access to ftp they can probably download binaries. It would make the most sense to publish the

ASCII-accessible version on the net. Then I could send SIMCONZ1.ASH directly to them, on request.

Best to all, Bill P.

Date: Tue Mar 30, 1993 10:06 pm PST
Subject: Allen's model

[From Rick Marken (930330.2100)] Bill Powers (930330.2000 MST) --

>Allan is using information about o in his method,
>as I vaguely understand it now, with your help. It shouldn't be
>surprising if he can also reconstruct d.

Still nothing from Allan on this: but I think I must not have been clear in MY description of what I THOUGHT Allan was doing. I don't think Allan is going to use information about o (the o generated in the real, closed loop run) at all. In fact, the way I conceive of Allan's study, you don't need to save o at all; all you need is d and p from the closed loop run. Again, here's what I think Allan is proposing:

To get H(D):

Put p's of varying length into the open loop model:

p --> model --> o

Correlate each resulting o (a vector of 1000 values) with d (also a vector of 1000 values). As you pointed out in your first post on this topic, this will involve going through a whole lot of different p's. At some point you get to a p which, when put through the model, yields a correlation between the resulting o and d that is equal to the corerlation between o and d obtained in the closed loop experiment. The length of the successful p vector is H(D). I suspect that the successful p would be the same length as the disturbance vector. So, from my point of view, Allan is proposing a kind of inverse regression analysis , where the regression equation is d = model(p). We are given d and the model (the "regression coefficients") and we are trying to find values for p that satisfy this equation (with the outputs of the open loop model being the predictors of d).

To get H(D|P):

We again put p's of varying length into the open loop model, put this time these p's are added to the p vector that was observed during the closed loop run. Call the p vector from the closed loop run p'. So

(p'+ p) --> model -->o

Again we would correlate the resulting o's with d. But this time, Allan imagines that the correlation between o and d will be maximum when the length of p is 0 -- that is, when the p vector is all 0's.

I think you should be able to see, now, why Allan will not get the results he expects. But I want Allan to explain whether this is a correct description of his plan. If I explain the problem I'm sure he'll say that he meant something different and we'll have to go around and around again. I think we have to get Allan to commit to a precise method of measuring the information about d in p (p' in my notation here). I would like to have the method agreed on in the form of a computer program (I'll write it if he likes). I think we're pretty close to the point where we can do this.

So let's get this precisely nailed down in terms of program steps (I'm all set to run Allan's model as soon as I get his concurrance about what it is) -- are you with us, Allan?

Best Rick

Date: Tue Mar 30, 1993 10:12 pm PST
Subject: They call me Dr. Marken

From Ken Hacker [930330]: Dr. Miller --

Apologies for unwarranted arrogance are just as bad as the arrogance originally identified. BTW, I was being exceptionally NICE by using the word arrogance. Other terms were not used when they could have been.

Rick Marken --

There you go again.

You apply what you know about psychology to all of social science. That is inane. There are many points of view in social and behavioral sciences and you continue the fallacy of monolithic sameness for them. It is not a big lie; it is a HUGE lie!

If you were more secure in what you are doing with PCT, perhaps you would not have to launch daily attacks on what other people with other perspectives are doing. What is the strength of PCT -- the limitations of non-PCT? Or the explanatory value of PCT within itself? I would hope the latter because I am convinced that much, not all, of the former is unmitigated bullshit.

Dare I also suggest that the megalomania about Galileo is a distraction from what PCT is attempting? Yes, I do dare. Rick, you are not Galileo, you are not Einstein. >99% of the world does not know who you are and probably never will. (me too -- but I don't pretend). So can we please get real and drop the delusions?

>But I would love to have a reference to one or two experimental research
>articles that do not assume an underlying lineal causal model -- other
>than those done by PCT people, that is.

Do some literature searching yourself. Look at what social scientists have done since 1950 and stop lumping them all in one bundle.

>.Experimental research is essential; the IV-DV approach to doing
>experiments (which assumes an underlying causal model of behavior)
>is NOT the only way to do experiments. Read about "The Test for
>Controlled Variables" in BCP. But first read BCP and do some of the
>"simple" demo experiments that illustrate what is going on. You
>won't understand the Test unless you understand PCT.

I fully agree.

>There is certainly nothing morally wrong with it (until you start
>applying these statistical results to individual human beings). I
>just find studies of this sort (and there are thousands of them
>being done all the time) a tragic waste of time, money and intellectual talent.

Translation: Only Marken knows who to do research - c'mon!!
Can we please drop this pretentious bullshit?

>This is very humane and decent sounding of you. It's difficult, however,
>to say things like "there is no information in controlled perceptions"
>without at least implicitly stepping on some toes. But what's wrong
>with showing that a perspective is wrong, anyway? Was it wrong for
>Galileo to show that Ptolemy was wrong? Was it wrong for Darwin to
>show that God was wrong? What's wrong with showing that an idea
>is wrong if the idea IS wrong? If you think there is some value to
>a "perspective" that we are "attacking" then just defend it; show
>that we are wrong. This is not a religious war -- this is how science
>works. In fact, PCT WELCOMES attacks -- ie. disciplined attempts to
>show that PCT is WRONG. That is why I'm so thrilled that Allan Randall
>is trying to show that PCT is wrong about the "no information in
>controlled perception" claim.

Henri Poincare (idiot non-PCTer) said, "Science is facts; just as houses are made of stones, so is science made of facts; but a pile of stones is not a house and a collection of facts in not necessarily science."

KEN HACKER

Date: Wed Mar 31, 1993 2:42 am PST
Subject: ... and on!!

From Greg Williams (930331)   Bill Powers (930330.2000 MST)

>Perhaps there is a simple question I should have asked
>originally, if I failed to do so: what kind of "understanding"
>would be satisfactory to you?

A kind which is not based solely on my introspective assessment that you have "correctly" identified the appropriate subset of my motivations. I believe that such a kind is possible in theory, although it might be very difficult (challenging?) to come up with in practice in some (many? most?) cases. If I were really motivated to perceive X, then if you applied disturbances in an attempt to make me NOT perceive X, I would attempt to correct for those disturbances, and such attempted corrections would constitute evidence for my REAL motivations, independent of what I THINK my motivations are.

>Am I aiming at a moving target, or any target at all?

I think the target is moving, if at all, rather slowly. But that's only my (honest) impression.

>Is there any substance in your challenge?

I believe that there is.

>It has crossed my mind that you are mainly defending against
>being understood. That would certainly explain a great deal. I
>seem to recall a time at one of our miniconferences when you
>announced that nobody could understand you if you didn't want
>them to. I recall, I hope correctly, your description of a
>strategy that would always make the Test fail: simply change your
>mind continually about the nature of the goal. Of course if this
>strategy were detected, it would reveal the aim of preventing the
>Test from working, in which case the Test would succeed.

I DON'T have the impression that I am "mainly defending against being understood." And I honestly don't remember anything about such a discussion at a miniconference, though it is entirely possible that you are correct about it.

>This leads to an interesting impasse. Suppose I were to announce
>that I have discovered your purpose, which is to show that the
>Test can't work on you. To demonstrate this, all I have to do is
>use the Test to search for signs of a controlled variable, and
>show that every time the Test appears to succeed for a moment, it
>ceases to work. Is this evidence that the Test isn't working, or that it is?

We are speaking hypothetically (unless I am deluded or intentionally deluding you regarding my introspective assessment that I'm not defending against being understood). If it were true that I am defending against being understood, and you used the Test to show that I was controlling for the perception of your not being able to understand, then I would accept that you had met the challenge. Applying the Test with respect to the PARTICULAR perception of your not being able to understand would not lead to an impasse: the Test would either succeed or fail; applying the Test to (at least some) OTHER kinds of perception could, as you suggest, lead to an impasse.

>I'll put BURN and UNBURN on Silvert's Server, but didn't you and
>Pat have a version that could be transmitted in ASCII and then
>converted using a debug script, also in ASCII? People who can't
>use ftp won't have much use for your programs if they can be
>accessed only through ftp, and if they have access to ftp they
>can probably download binaries. It would make the most sense to
>publish the ASCII-accessible version on the net. Then I could
>send SIMCONZ1.ASH directly to them, on request.

I have a program (from PC MAGAZINE) to convert binary files to ASCII script files which can be used with MS-DOS's DEBUG program (reasonably clear instructions for doing this are embedded in the script files) to convert back to binary. To save money, I'll send you a disk with the script files for BURN and UNBURN; then you can upload them to the net.

As ever, Greg

Date: Wed Mar 31, 1993 5:42 am PST
Subject: Re: tracking task--a story

[Martin Taylor 930330 17:30] (Bruce Nevin 930329 14:44:46 re Tucker-Marken)

>As a matter of conversational pragmatics
>it would be impossible to give a tracking instruction, like "keep
>the cursor on this mark", unless it were stated to the subject or
>presupposed by the subject that something was liable to separate
>the cursor and the mark, otherwise what's the point?

I am reminded of the experience a colleague reported many years ago from the time when he had been an undergraduate subject in a tracking experiment. There was a turntable like that of an old gramophone, and on the turntable a small brass disk. He was told "keep the pen on the disk." The experimenter expected he would track the disk as the turntable rotated. Instead, he assumed that it was an intelligence test, and that only a dumb person would track the disk like that. He dismounted the turntable, laid the pen tip on the disk, and sat back, succeeding perfectly in his assigned task (even if it wasn't the task the experimenter thought had been assigned).

Martin

Date: Wed Mar 31, 1993 5:44 am PST
Subject: Re: Movement

[Martin Taylor 930330 16:15] Dennis Delprato 930329 to Mary Powers 9303.28

>> movements are the ultimate
>> means by which control systems at every level control perception[.]
>
>You said that you wonder if someone understands the above. I suspect
>that your point is frequently overlooked,
>both in and out of PCT literature and thinking.

It tends to be ignored that we have other ways than movement of affecting our environment, and that these ways can also be important in controlling perception -- very, very important. Chemical output, which can be terribly important in mate selection, for example. Other species perhaps make more use of this than we do, but chemical outputs are by no means negligible with ourselves. There may be others besides movement, but I can't think of one for now. We shouldn't ignore the possibility, though.

Martin Only 50 more mail items from yesterday to go !!!

Date: Wed Mar 31, 1993 5:45 am PST
Subject: Re: Where is the proof? Information of disturbance

[Martin Taylor 930330 11:20] Rick Marken various

(930315.1500 as an example)
>"... THERE IS NO INFORMATION ABOUT DISTURBANCE IN THE PERCEPTUAL SIGNAL

>CONTROL SYSTEM. This means that THE PERCEPTUAL INPUT TO A CONTROL SYSTEM
>CANNOT BE WHAT CAUSES THE OUTPUR OF THE SYSTEMN (sic) TO MIRROR THE
>DISTURBANCE."

Let's consider a thought experiment to test this. If I understand the claim, oft repeated, Rick means that no function that takes as input (a) the perceptual signal and (b) any other signal that is agreed to have no information about the disturbance can reconstruct the disturbance, but that nevertheless the disturbance is mirrored in the output.

I'll leave the logical problem with the "mirroring" unaddressed, and assume that Rick accepts as correct what he says, that the output mirrors the disturbance.

In my thought experiment, I will take the ECS, and add a simple function that takes as its input the reference signal to the ECS (which I think we can agree has no information about the disturbance) and the perceptual signal, which Rick CAPITALIZES as having no information abou the disturbance. Let us see whether a function can be constructed that takes these two inputs and produces a signal that matches the disturbance. If so, I would consider it conclusive evidence that information about the disturbance is to be found in the perceptual signal.

```
             ------------> Signal X (which should match the disturbance)
            |
       mystery function M(r, p)
        ^      ^
        |      |
        |      |        V  (reference signal R(t) into ECS)
        |      |        |
        |       <-------|
        |               V
        |---------->comparator------- error = P-R
        |                           |
    perceptual                     output
     signal P(t)                       function O(error)
        ^                           |
        |                           V
        |                         output signal
        |             (accepted as mirroring the disturbance)
```

If Signal X matches the disturbance, the perceptual signal must be the route from which the mystery function M(r, p) gets the information about the disturbance. Right?

Now let the function M be indentical to O(R-P). Signal X will then be the negative of the output signal, which is the disturbance. The only question here is whether O(error) is a function or a magical mystery tourgoodie. I prefer to think we are dealing with physical systems, and that O is a function. Therefore, information about the disturbance is in the perceptual signal, and moreover, it is there in extractable form.

QED.

(Actually, QED is too strong, since I imagine most of us will want to challenge Rick's claim that the output mirrors the disturbance. But that way lies the argument to information rate, which I will pursue whether or not Rick accepts that way out of the Q that ED.)

Martin

Date: Wed Mar 31, 1993 6:25 am PST
Subject: experimental social science

[From: Bruce Nevin (Wed 930331 08:37:59)]

Ken Hacker, I can sympathize with your views (though I would urge that it causes as much trouble to take offense as it does to give it). Bill, Rick, and others here have presumed that linguistics prior to PCT was just another case of S-R IV-DV behavioral science. Avery remarked that it is not. (Refer also to Chomsky's classic review of Skinner's _Verbal Behavior_.) Bill's rejoinder was that this could be, if linguists just happen not to do any lab work.

Now I do agree with the view of some linguists that more experimental work is needed, as opposed to Chomsky's stated aversion to "data bound" linguistics. But linguistics is concerned with a different order of problems, distinct from the problems before any study of behavior. I believe this distinction also applies to all the social sciences. I have been working to bring this different order of problems under the PCT umbrella.

I stated it yesterday (Tue 930330 08:36:29):

>Language is not behavior. Speech is behavior. Language is
>socially standardized reference perceptions, the control of which
>has speech as a byproduct. The reference perceptions are
>socially available without recourse to the Test because they are
>learned (and taught--cf. Bruner) as socially standardized perceptions.

There is good evidence that this distinction between competence and performance (to use Chomsky's terms) applies to the social sciences generally. That is, behavior that is properly studied in the social sciences is a byproduct of control according to socially standardized reference perceptions. The appropriate subject matter of the social sciences, including linguistics, is not the behavioral outputs but rather the socially standardized reference perceptions and the processes by which individuals come to calibrate their private reference perceptions according to a perceived public standard. Dan Miller described some early (1973) investigations of "redundant sequences and patterns". Some years ago there was a William Pierce at UMass Amherst whose study of what he called constitutive rules in social interaction still seems to me to be very fruitful. Perhaps I should dig out some papers of his that I got in the 1970s.

To the point here, the study of socially standardized reference perceptions is not subject to the same experimental techniques as the study of behavioral outputs as byproducts of experimentally determined reference perceptions. The former has a "meta-" relation to the latter. Other sorts of questions must be addressed, questions about how people establish and change reference perceptions, how the learning process (or "acquisition" process, the question-begging term preferred in linguistics) may relate to higher-level perceptions such as the perception of being in a cooperative relationship, being able to rely on another, being perceived as reliable by others, whether or not some of these perceptions are "intrinsic", and so on.

It is not clear to me how experiments regarding these matters might be designed. Part of the problem is my lack of experience and training in experimental methodology, but part of it is that there is no obviously direct transfer of methodology from PCT modelling of behavioral outputs to PCT modelling of calibrating one's reference perceptions to a social standard. My time and energy are severely limited, and my imagination probably as well. I would greatly appreciate help in this from anyone.

        Bruce   bn@bbn.com

Date: Wed Mar 31, 1993 8:23 am PST
Subject: Study of teaching methodsX-

[From Bill Powers (930331.0800 MST)]   Ken Hacker (930330) --

>Bill, I sense that you want me to say that I will attribute
>increased recall or retention to the learning program. Aha.
>BUT, I am not doing that. I am going to attribute whatever
>significant differences there are to what the subjects appear
>to need to learn [t]hat I am trying to teach them with the program.

>Essentially, the program is a self-paced tutorial where
>students learn how to use computer conferencing. I am trying
>to compare the program method (each student learns by him or
>herself with a pc) versus a lecture-by-the-expert method of training.

So if the subjects appear be able to satisfy what they need to learn better with the program method than with the lecture-by- expert method, it would be the differences between the methods (IV) that accounts for their increased ability to satisfy their needs (DV)?

You needn't be wary of a trap. It may well be true that the difference in method accounts for the greater ease of learning if such is observed, and that the greater ease may depend on the teaching method. The world is full of IV-DV relationships. The entire PCT model is built up out of IV-DV systems, in which the output of a function depends completely on one or more inputs. There is nothing inherently incorrect about looking for IV-DV relationships. We do it all the time in control-system modeling. For example, we show that the mean squared error signal in a control system is inversely dependent on the loop gain. We show that the measure of output is very closely dependent on the measure of the disturbance, and that the measure of the controlled variable is strongly dependent on the reference signal.

>You see, I am trying to find the best method of teaching-
>learning from the point of view that the student has needs and
>those needs will best be met by one method over the other (at
>least that's my hypothesis).

Here, however, we do have a problem, which I addressed by asking you if you have any data from a similarly-organized study that you've already done. If not, can you put your hands on someone else's data, so we can talk about the results? I listed the kind of information I would like to see. It would be better to continue the discussion with some actual results to look at.

What I would expect, just to avoid seeming to be setting another trap, is that your study will show that some of the people are able to satisfy their learning needs better with the programmed approach, some will do better with the expert-lecture approach, and some will do about equally well under either approach.

There is a methological problem here, which Phil Runkel addresses in some depth in his _Casting Nets and Testing Specimens_. Perhaps you have already addressed it. Studies that compare the effects of treatments by using two parallel studies suffer a sampling problem, in that the people in one study are not the same people in the other, so the significance of the results is weakened (often fatally) by the inability to say how much of the difference is due to a difference in the sample populations. One way to handle this is to put all the people through both treatments, randomizing the order.

But there is an even deeper problem, which I hope that Gary Cziko will also comment upon. It is that you seem (on the basis of what you've said so far) to be trying to decide which is the better method, with the implication that if one method produces signficantly better ease of learning for "the student," that is the method that ought to be used in the future. This sort of thing is often done by people interested in developing good teaching methods. Unfortunately, doing it this way puts the emphasis on the method rather than on the learning.

In a PCT-style approach, there would be no need for using one method with one group and another method with another. All the available people would be taught one thing by one method and another thing by another method, perhaps through a number of cycles and with randomization within the cycles. Then there would be an assessment of how well people did with the two teaching approaches. The people who did either very well or very poorly under either method would then be investigated further, to find out (if possible) why a given method worked better or worse for a given person. There might prove to be great differences in preference for one method or the other. There might prove to be hearing or vision problems, or problems with auditory versus visual comprehension, or differences in reading speed, or differences in attitudes toward using computers, or differences in typing skills. Some people may feel uncomfortable working all alone without interaction with other people, while others find that atmosphere facilitating.

With this approach, the emphasis is not on the method but on the learning, and on finding out what conditions of learning work best for each individual. Perhaps some of the investigations might appropriately use PCT concepts as a way of testing for the ability to control for various relevant perceptions, or as a way of deducing where people have set their preferences for the conditions under which they learn. But whether PCT is used or not, this general approach is organized around the individual, not around some measure of the average performance of a group of individuals. All studies like these give tests, and individual test scores are available. So the information is available that would lead to the next phase, and the next, in a study aimed at suiting the method to the individual.

Think about it. Even if you didn't do all the follow-up work, suppose you simply concluded that for THIS subgroup, the best method is evidently lecture-by-expert, and for THAT subgroup it is the program approach. You would clearly do better in future teaching by suiting the method to the student than by deciding which method is "best" and using it with ALL students, even those who did worse under the "best" method. For the students in the middle, of course, it wouldn't matter which method you used; for them, neither is better.

Also, if the study gives negative results, it could easily do so because half of the students did markedly better with lectures and the other half did markedly better on the computers, but the group as a whole didn't do significantly better or worse with either approach. By sticking strictly to the protocols of mass statistics, you would miss seeing the opportunity to enhance the learning of all the students (unless, as seems seldom to be done, you thought to look at the distributions).

Perhaps this is already your approach. If so, we have nothing to contend about.

Best, Bill P.

Date: Wed Mar 31, 1993 8:35 am PST
Subject: IV-DV

[FROM: Dennis Delprato (9303310] KEN HACKER [930330]:

>Bill, I sense that you want me to say that I will attribute increased
>recall or retention to the learning program. Aha. BUT, I am not doing
>that. I am going to attribute whatever significant differences there
>are to what the subjects appear to need to learn what I am trying to
>teach them with the program. Essentially, the program is a self-paced
>tutorial where students learn how to use computer conferencing. I am
>trying to compare the program method (each student learns by him or
>herself with a pc) versus a lecture-by-the-expert method of training.
You see, I am trying to find the best method of teaching-learning
>from the point of view that the student has needs and those needs will best
>be met by one method over the other (at least that's my hypothesis).

Profound. You show how actuarial research (necessary for addressing questions such as which service delivery system will be more effective over the long haul), although PROCEDURALLY following the classic IV-DV experimental model as well as hypothesis testing inferential statistics that behavioral scientists deftly integrated with the former, need not be INTERPRETED in the conventional way. As far as individuals, outcomes from actuarial research tell us only about control systems. You will find out whether or not one method is any more effective than the other at permitting the AVERAGE member of your population to adjust to particular circumstances according to particular criteria. Furthermore, if your experiment is flawed by confounding according to conventional definitions, you will not draw a reliable conclusion regarding the relationship between your manipulated variable and how it alone relates to adjustments of the average member of the population.

You will make no claims about learning anything about individual control systems (fundamental principles). The concern of the study at issue is actuarial.

Dennis Delprato psy_delprato@emunix.emich.edu

Date: Wed Mar 31, 1993 8:35 am PST
Subject: Movement

[FROM: Dennis Delprato (930331)]

>Martin Taylor 930330 16:15

>It tends to be ignored that we have other ways than movement of affecting
>our environment, and that these ways can also be important in controlling
>perception -- very, very important. Chemical output, which can be
>terribly important in mate selection, for example.   Other species
>perhaps make more use of this than we do, but chemical outputs are by
>no means negligible with ourselves. There may be others besides
>movement, but I can't think of one for now. We shouldn't ignore the
>possibility, though.

Point of clarification: My referent here is motor control of sensory feedback as perhaps most vociferously, if unheard, expounded by K. U. Smith. Thus, if we go and describe the details of chemoexcretions, we find motorsensory feedback control systems. I do not mean movement in only the sense of gross muscular activity. Taken this way, I find Powers's and Smith's as perhaps the best examples of the outcome of the historical behavioral movement. Motorsensory feedback control theory is behavioral in the extreme. But now (a) physical and biological behavior (events) are at the heart of all psychological events and (b) psychological behavior is inseparable from perceptual input. Of course, Smith and Powers are so far in advance of the behavioral movement that virtually no one finds their views even related to earlier versions of the behavioralization of psychology in one or another form of behaviorism.

Date: Wed Mar 31, 1993 9:16 am PST
Subject: Challenge 2, endgame

[From Bill Powers (930331.0900 MST)]   Greg Williams (930331) --

>>Perhaps there is a simple question I should have asked
>>originally, if I failed to do so: what kind of "understanding"
>>would be satisfactory to you?

>A kind which is not based solely on my introspective assessment
>that you have "correctly" identified the appropriate subset of
>my motivations....If I were really motivated to perceive X,
>then if you applied disturbances in an attempt to make me NOT
>perceive X, I would attempt to correct for those disturbances,
>and such attempted corrections would constitute evidence for my
>REAL motivations, independent of what I THINK my motivations are.

That is what I have been doing. However, since we are limited to verbal communication, there is no external counterpart of your controlled variables except as I can understand it by guessing at the meanings of your words, and the circumstances under which you make an effort to convey those meanings.

When I gave you the list of questions and the answers I predicted that you would and would not give, I was testing the hypothesis that you would find a way to show that all of my guesses were wrong. Under the hypothesis that the goal was to avoid being understood, or to show that the Test is difficult to use in the field, to have answered the questions in the same way you had answered other questions (as in my positive predictions) would have caused an error relative to the goal of showing that understanding is not possible by this method. You are smart enough to realize what I was getting at, but of course I had not said why I was making these predictions.

This does not prove my hypothesis correct, but it is an encouraging hint. You appear to be perceiving that the Test is very difficult to carry out in the field. I presume that you are motivated to perceive this X, as no definitive conclusion has yet been reached on this subject. On that hypothesis, I would predict that everything I to do alter that perception will be met by some move that will counteract my attempt, and support the idea that the Test is very difficult to use in the field.

One sign of this, as I interpret the proceedings, is that when I ask a very simple question like "Do you still intend to maintain this challenge," your answer has uniformly dodged the request to state your intention; after all, to do so correctly would be to show that the Test can succeed, and in a simple way at that. You say things like "Well, it may or may not have been my intention when the challenge was issued, and it makes sense to suppose that I still have that intention, but I have no awareness of my structure of motivations, so I can only guess as to whether the intention still exists, or is the same as it was, or will be the same the next time you ask."

By giving this sort of answer, instead of just thinking a moment and saying "Yes, the challenge is still on as far as I'm concerned," you exhibit a lot of carefully-thought-out reasoning leading to a response that shows that the Test is very difficult to apply in the field. All the extra complication in your answer, the invocation of things that might be, or are beyond your knowledge, or are uncertain, shows that a lot of effort is going on in order to preserve some particular perception. You don't spontaneously say such things about your own statements under ordinary circumstances, so Im not just sampling a sort of behavior that goes on all the time anyway, at least not on this subject. This pattern of answering doesn't happen when I ask you if you want something, casually, in other conversations. You don't expect me to react to your request about posting BURN and UNBURN by saying "Gee, I'm not in contact with my motivations, so I can't say whether I'll end up posting them or not, and even if I might (conceivably, hypothetically) have some such intention right now, I can't predict whether I will still have it ten minutes from now." I just say "OK." (By the way, I find that I already have all your materials, including the script, tucked away in a long-forgotten directory, so you don't need to mail the disk).

So it's evident to me that when I pose questions in this context, you're answering in an unusual way, and doing so only when I ask the sort of question I am asking. When I get too cocky, as by saying you're an easy case, you soon remind us all that it really does seem hard to apply the Test, judging from the progress we have made during this challenge. This is, of course, your judgment of the progress, which still meets your reference level for it.

I should point out that "facts" and "impressions" are very likely reference signals. When I turn the key in the ignition, the car starts. That's just a fact. If someone asks me how to start the car, I tell them to turn the key. But when the car DOESN'T start, it becomes immediately clear that this fact is a reference signal.

You may have the opinion that the Test is hard to apply in the field, and think that this is nothing more than a reasonable conclusion. But when something happens that tends to show that the Test can work in a rather simple and immediate way, that is an error, calling for immediate efforts to find a different interpretation. When something that has already been accepted as a fact begins to look nonfactual, a corrective effort is called for.

PCT claims that all people have reference signals and perceptions. My interpretations of it suggest that by various means they can become aware of them if they want to. But when one becomes aware of such things, they don't carry the label "this is a reference signal" or "this is a perception." They are just old familiar aspects of experience with no names as such. In seeing how PCT applies to oneself, the main hurdle to get over is to identify the terms of PCT with these nameless experiences that thread through everything. This is not a matter of finding some objective experimental proof; it's a matter of recognizing something that is already there and that can be observed quite easily -- but can't be made sense of so easily, without some organizing theory.

The point of my asking you about your subjective impressions has not been to prove that the Test is working. The Test works without knowing the subjective appearance of controlled variables or reference levels. It doesn't work infallibly, and without continued investigation it can yield only approximations to the real situation, but it does yield strong evidence relating to hypotheses, and can also yield strong evidence against them. We haven't carried this far enough for me to be utterly confident that I have identified any of your controlled variables, but a reasonably close approach has been made.

No, the point has been that in order to meet this challenge, I must get you to say that a correct identification has been made, near enough. And in order to do that, you must become aware of your own reference signals AS REFERENCE SIGNALS, and of your own controlled perceptions AS PERCEPTIONS. You must make the identification yourself. When you do, the outcome will be self- evident to you.

My eye is on your King.

Best, Bill P.

Date: Wed Mar 31, 1993 9:39 am PST
Subject: Martin's Thought Experiment

[From Rick Marken (930331.0800)] Martin Taylor (930330 11:20) --

>Let's consider a thought experiment to test this.

EXCELLENT THOUGHT EXPERIMENT! Now let's do it as a simulation, shall we?

>If I understand the
>claim, oft repeated, Rick means that no function that takes as input
>(a) the perceptual signal and (b) any other signal that is agreed to have no
>information about the disturbance can reconstruct the disturbance, but
>that nevertheless the disturbance is mirrored in the output.

I'll buy it.

>I'll leave the logical problem with the "mirroring" unaddressed, and
>assume that Rick accepts as correct what he says, that the output
>mirrors the disturbance.

What's the "logical" problem with the "mirroring"? You can look at the data from our tracking experiments and see that o = -d to within a few pixals throughout an experimental run. When both o and d are measured in screen units, the time traces of these two variables will be symmetrical about a line corresponding to the fixed screen position of the target. I call this characteristic of the graph "mirroring".

>In my thought experiment, I will take the ECS, and add a simple function
>that takes as its input the reference signal to the ECS (which I think we
>can agree has no information about the disturbance) and the perceptual
>signal, which Rick CAPITALIZES as having no information abou the disturbance.

Excellent! For simplicity, let's make the reference signal a constant when we simulate your model. But a variable r will work too -- just trying to keep it simple.

>Let us see whether a function can be constructed that takes these two
>inputs and produces a signal that matches the disturbance. If so, I
>would consider it conclusive evidence that information about the disturbance
>is to be found in the perceptual signal.

OK!

```
>             ------------> Signal X (which should match the disturbance)
>            |
>       mystery function M(r, p)
>        ^      ^
>        |      |
>        |      |        V  (reference signal R(t) into ECS)
>        |      |        |
>        |       <-------|
>        |               V
>        |---------->comparator------- error = P-R
>        |                           |
>    perceptual                   output
>     signal P(t)                     function O(error)
>        ^                           |
>        |                           V
>        |                      output signal
>        |             (accepted as mirroring the disturbance)
>
>-------------------------------------------------------------------
>
```

>If Signal X matches the disturbance, the perceptual signal must be the
>route from which the mystery function M(r, p) gets the information about
>the disturbance. Right?

Right!! I completely agree with your proposal as diagrammed above. I think a good first candidate for M(r,p) would be the function O(r,p), right? Ah, I see you think so too:

>Now let the function M be indentical to O(R-P). Signal X will then be the
>negative of the output signal, which is the disturbance.

It is at this point that experience will triumph over the "obvious" conclusions of your thought experiment. I think it's time to fire up the simulator; really!

>The only question
>here is whether O(error) is a function or a magical mystery tourgoodie.

Your magical mystery tour will really begin when you run the simulation!

>I prefer to think we are dealing with physical systems, and that O is a
>function. Therefore, information about the disturbance is in the perceptual
>signal, and moreover, it is there in extractable form.

>QED.

And a right excellent proof i'tis. Now try the simulation.

Best

Rick "There is no information about the disturbance in controlled perception"
Marken

Date: Wed Mar 31, 1993 10:22 am PST
Subject: Your controlled variables are showing

[From Rick Marken (930331.0900)]

From Ken Hacker [930330]--

>Apologies for unwarranted arrogance are just as bad as the arrogance
>originally identified. BTW, I was being exceptionally NICE by using
>the word arrogance. Other terms were not used when they could have been.

I never apologized for unwarrented arrogance because I never thought I was being arrogant -- warrented or not; that was your perception, remember? But why all the anger? Why the (surpressed) desire to call me names? Whatever I said must have really pushed on one or two controlled variables. I know that when this happens, it seems like one is being attacked personally, doesn't it? Actually, I don't think I have ever attacked you personally, have I? What I have "attacked" (it feels like an attack when someone disturbs a controlled variable) is the IV-DV approach to research in the behavioral sciences and the underlying justification (causal model) for its use. My guess is that you are trying to perceive the IV-DV methodology as a reasonable approach to achieving your research goals. One hint of this is in your reply to Bill Powers:

KEN HACKER [930330]:

>Bill, I sense that you want me to say that I will attribute increased
>recall or retention to the learning program. Aha. BUT, I am not doing
>that. I am going to attribute whatever significant differences there
>are to what the subjects appear to need to learn what I am trying to

>teach them with the program.

So, from your perspective,the IV-DV methodology you are using is fine, even when your goal is not to determine a causal relationship between the IV and DV. You will simply "attribute" the differences to needs inside the subjects. I suppose it would seem arrogant of me to ask you to describe your model of how the subject's need to learn is expected to interact with the learning program to determine the results you observe in the experiment?

Since we seem to be dealing with my criticism of IV-DV research on a purely verbal level (you seem to eschew mathematics and modelling, for example) then as long as you can generate sentences that let you perceive your controlled variables as under control, I can't be much more than an uninformative disturbance, and you should be able to maintain your perception of yourself as the winner of this "debate". So, congratulations.

>If you were more secure in what you are doing with PCT, perhaps you would
>not have to launch daily attacks on what other people with other
>perspectives are doing.

The only "attack" of this nature that I can recall is when we discussed the potentially severe, negative consequences that might result from basing individual behavior on statistical findings. This "attack" was not based particularly on PCT and I didn't really participate in it much. So I would really like to know what you consider a "daily attack on what other people and other perspectives are doing"? I have a feeling that what you consider an "attack" is what I would consider a "discussion". I really have no intention of attacking you, your perspective or your friends and their perspectives. I'm just trying to describe PCT. If you think I'm "attacking" please give me examples. Is Bill Powers attacking too, or am I the only attacker? Could you please give me an example of an "attack" from one of my posts. I will apologize, if I perceive it as an attack too; especially if it is a personal attack.

By the way, here are some examples of what I perceive as personal attacks:

>You apply what you know about psychology to all of social science. That
>is inane.

Using "wrong" instead of "inane" would have made this sound much less like a personal attack.

>There are many points of view in social and behavioral sciences
>and you continue the fallacy of monolithic sameness for them. It is
>not a big lie; it is a HUGE lie!

Again "wrong" might have sounded less personal than "lie".

>Dare I also suggest that the megalomania about Galileo is a distraction
>from what PCT is attempting?

I consider the implication that I am a megalomanic (who considers meself as great as Galileo) to be a personal attack. I'm just a sensative guy, I guess.

>Translation: Only Marken knows who to do research - c'mon!!
>Can we please drop this pretentious bullshit?

And I am pretentious too.

These are what I would call personal attacks. I don't mind them; I understand where they are coming from and I hope they suceeded in helping you control whatever it was you needed to control. But for discussions on CSG-L I prefer substantive discussions.

> What is the strength of PCT -- the limitations of
>non-PCT? Or the explanatory value of PCT within itself?

I believe that I have been trying to point out the strengths of PCT and the limitations of non-PCT models in behavioral science. You seem to think that I am just attacking people and perspectives. So help me out here -- just show me what you consider to be an "attack". Or is my asking for this an attack? or is my guessing about my asking about this being an attack, an attack? etc.

>I would hope the latter because I am convinced that much, not all,
>of the former is unmitigated bullshit.

If you think that much, but not all, of PCT is unmitigated bullshit (UB) then it would help everyone (but especially those of us who are promulgating this UB) if you would explain just what you consider to be the UB in PCT. It would also help if you would show why it is UB.

Best Rick

Date: Wed Mar 31, 1993 10:25 am PST
Subject: sociological questions

[from Kent McClelland (Wed 930331)] Re: Bruce Nevin (Wed 930331 08:37:59)

>To the point here, the study of socially standardized reference
>perceptions is not subject to the same experimental techniques as
>the study of behavioral outputs as byproducts of experimentally
>determined reference perceptions. The former has a "meta-"
>relation to the latter. Other sorts of questions must be
>addressed, questions about how people establish and change
>reference perceptions. . .

I think Bruce has put his finger on something here which I find useful as a sociologist in trying to understand the implications of PCT for my own field. Once you accept that behavior is the result of individuals' attempts to control their perceptions, the questions of greatest interest for sociologists seem to me to come down to the one that Bruce suggests (how different individuals develop highly similar control systems for perceiving the social environment, including language), as well as questions of the _distribution_ of such control systems across populations, and also the distribution of settings of control systems. (We may both have control systems for a given CEV in our shared environment, but if you and I are holding different reference values for our control systems, conflict is likely to ensue, as Bill Powers has pointed out many times). These sociological "meta-" questions can prompt further inquiries, e.g., how different empirical distributions of people actively trying to control similar perceptual variables in an environment might lead to predictable sequences of interaction.

For example, I've been attempting a PCT account of social power for several years, and my latest formulation goes like this: social power can be defined as occurring when two or more people share the "same" control systems and are actively controlling for the "same" reference values of those systems in a given environment. I would describe this phenomenon as "alignment" of control systems, and by my definition social power is synonymous with alignment. I put 'same' in quotes, because I see as an open question just how similar two control systems (and their settings) need to be before they can be regarded, for practical purposes, as aligned.

Some implications of this PCT definition of power:

power is proportional to the number of people involved in an alignment;

social power is not the property of any single individual, and our attributing of power to an individual is simply a way of indicating that the individual is acting as the spokesperson for an alignment;

power comes and goes as individuals change their allegiance to various alignments.

I have a number of other implications sketched out for a paper which I hope to have ready for the CSG meeting this summer. If things go well, I'll also attempt to produce some spreadsheet models or an experimental design for exploring the behavior of some simplified distributions of control system settings and alignments.

I hope this speaks to Ken Hacker's question on the net several days ago about how social scientists might try to apply PCT to their concerns. It might also have some relevance to the recent controversy between Bill Powers and Greg Williams on whether statistical methods will be necessary in the application of PCT to social scientific questions. I wouldn't want to make any categorical assertions, but I wouldn't be surprised if statistical tools are helpful in investigating the empirical distributions of control systems that I'm talking about. . . Best to all,

Kent

Kent McClelland Office: 515-269-3134
Assoc. Prof. of Sociology Home: 515-236-7002
Grinnell College Bitnet: mcclel@grin1
Grinnell, IA 50112-0810 Internet: mcclel@ac.grin.edu

Date: Wed Mar 31, 1993 10:44 am PST
Subject: Info about disturbance

[From Bill Powers (930331.1030 MST)] Rick Marken (930330.2100) --

>I don't think Allan is going to use information about o (the
>o generated in the real, closed loop run) at all. In fact, the
>way I conceive of Allan's study, you don't need to save o at

>all; all you need is d and p from the closed loop run.

As I understood it, Allen was going to assume that control was perfect, so o = -d with a constant reference signal. So the procedure depends on knowing o, which amounts to knowing d.

But let's wait to see what Allen says.

Martin Taylor

(930330.1615) --

>It tends to be ignored that we have other ways than movement of
>affecting our environment, and that these ways can also be
>important in controlling perception -- very, very important.
>Chemical output, which can be terribly important in mate
>selection, for example.

There are different kinds of "important." I think it's unlikely that there are control systems that vary pheromone output in order to attract a mate, even though it's true that emitting pheromones does attract mates and makes continuation of the species possible. The importance here is from the perspective of the scientist who realizes the role that pheromones play, and the essential effect on survival of the species. Some human females are attracted by the smell of sweat, and thus sweating contributes to the preservation of the jock species. But the control systems involved with perspiration control for temperature, not sex or survival.

It's just as much an error to attribute purpose to nonpurposive effects as it is to do the opposite.

(930330.1120) --

The diagram you gave, below, won't work:

```
             ----> Signal X (which should match the disturbance)
            |
       mystery function M(r, p)
        ^      ^
        |      |
        |      |        V  (reference signal R(t) into ECS)
        |      |        |
        |       <-------|
        |               V
        |---------->comparator------- error = P-R
        |                            |
    perceptual                     output
     signal P(t)                     function O(error)
        ^                            |
        |                            V
        |                      output signal
        |             (accepted as mirroring the disturbance)
```

If the reference signal is zero, the signal X won't mirror the disturbance. In general, it won't. What you need is

```
             ----> Signal X (which should match the disturbance)
            |
       mystery function M(O,P)          <-------------
        ^      ^
        |      |
        |       ---------------<-----------------   <-----
        |               V  (ref sig R(t) into ECS) |
        |               |                          |
        |               V                          |
        |---------->comparator------- error = P-R  |
        |                           |              |
    perceptual                   output            |
     signal P(t)             function O(error)     |
        ^                           |              |
        |                           V              |
        |                      output signal O -->-
        |              (accepted as mirroring the disturbance)
```

The output signal does NOT follow the reference signal, so must be sensed directly.

If we accept the perceptual signal as following the external controlled quantity qi, we have simply

p(t) = O + D, exactly

Thus, D = p(t) - O, exactly. If the signal X is to match D, the higher system must sense the output itself, which requires a sensor that the illustrated control system does not have.

In all our examples so far, the reference signal has been fixed at 0. According to your diagram, the mystery function would then receive the perceptual signal and a null signal, with only the perceptual signal variations then producing signal X. This would say that when the reference signal is zero, the perceptual signal is the same as the disturbance, which is not true. Check it out with Simcon.

Furthermore, the sensor should not really detect the output signal directly, but the EFFECT of the output signal as translated through the output transducer and through the environmental link between the physical output and the controlled quantity qi. But that's a side-issue.

I think that QE is not quite yet D.

In these developments, "information" is being bandied about loosely. From Garner, information is simply the negative log of the number of values that a variable can attain. If the disturbance can take on any value between 0 and R, and this range is divided into n intervals r, the information in the disturbance is just -blog(R/r). Let's assume a continuous disturbance with a bandwidth less than that of the control system.

The controlled variable, on the other hand, with a constant reference signal, has a range of R', which equals R/(1+G), G being the loop gain at the maximum frequency contained in the disturbance. The quantum of measurement (in a tracking experiment) is the same, r, the screen resolution. So the information in the input quantity is -blog(R'/r), which is equivalent to -blog(R/[(1+g)*r]).

The ratio of the information in the perceptual signal to the information in the disturbance is

```
                 blog(R/r) - blog(1+G)
                 --------------------
                       blog(R/r)
```

This ratio is self-evidently less than 1 with any loop gain of 0 or greater. With very large loop gains, the ratio can become very small: there is much less information in the perceptual signal than in the disturbance. In Ashby's terms, information from the disturbance has been "blocked" from reaching the controlled variable. In fact, the ratio can be negative: the control system itself can be putting information into the perceptual signal (oscillations due to the quantized nature of the controlled quantity).

This only treats the AMOUNT of information; it says nothing about the "contingent uncertainty," which depends on the correlation between the perceptual signal and the disturbance. The maximum information would flow if the perceptual signal were perfectly correlated with the disturbance. So the actual fraction of the information about the disturbance present in the perceptual signal must be even less than the ratio computed above. The greater the loop gain, the less the information about the disturbance that exists in the perceptual signal.

Clearly, if the perceptual signal can attain only n values while the disturbance can attain m values, and m > n, it is impossible to reconstruct the disturbance from the information contained in the perceptual signal. The larger m is in comparison with n, the poorer the reconstruction would be. When the range of the perceptual signal has been reduced to 1/50 or less of the range of the disturbance, practically nothing could be said about the disturbance from knowledge of the perceptual signal.

It seems to me that this argument, which hardly takes one step away from first principles, supports the case that Rick and I are making in a simple and direct way.

Best, Bill P.

Date: Wed Mar 31, 1993 11:56 am PST
Subject: Re: Info about disturbance

[Martin Taylor 930331 14:15] Bill Powers 930331.1030

```
>The diagram you gave, below, won't work:
>
>
>             ----> Signal X (which should match the disturbance)
>            |
>       mystery function M(r, p)
```

```
>        ^      ^
>        |      |
>        |      |        V  (reference signal R(t) into ECS)
>        |      |        |
>        |       <-------|
>        |               V
>        |---------->comparator------- error = P-R
>        |                            |
>    perceptual                    output
>     signal P(t)                      function O(error)
>        ^                            |
>        |                            V
>        |                       output signal
>        |             (accepted as mirroring the disturbance)
>
>
```

>If the reference signal is zero, the signal X won't mirror the
>disturbance. In general, it won't.

Have another look. There are two inputs to function M. But there is only one input to function O. The input to function O is (P-R). M is defined as equivalent to O except that it incorporates (R-P) as a first stage. It is functionally identical to O(R-P). If you want to be even more general about it, you can remove the requirement for odd symmetry in O, and define M as -O(P-R).

Why does the value R = 0 have any special quality? In what way will X not mirror O?

Notice that I never claimed O mirrors the disturbance. That's Rick's claim (I don't remember you making it without the necessary simplifying assumptions). The diagram and thought experiment is intended to show the inconsistency of simultaneously maintaining the two claims:

(1) The output mirrors the disturbance,
(2) There is no information about the disturbance in the perceptual signal.

>In all our examples so far, the reference signal has been fixed
>at 0. According to your diagram, the mystery function would then
>receive the perceptual signal and a null signal, with only the
>perceptual signal variations then producing signal X. This would
>say that when the reference signal is zero, the perceptual signal
>is the same as the disturbance, which is not true.

Neither is it claimed. Mystery function M is not a unity operator. It has every characteristic that the output function has, including any leaky integrators and nonlinearities. It has access, apart from the specific inputs described in the figure, only to that information available to the output function. If the output function contains or is affected by some kind of homunculus that can look to see whether gain should be changed, so does (is) the M function.

Function M simply takes the reference signal, which no-one has claimed to contain information about the disturbance, together with the perceptual signal, which some people have claimed to have no information about the disturbance, and reproduces a mirror image of the output signal, which some people have claimed reproduces the disturbance exactly.

There are three possibilities: (1) The reference signal contains information about the disturbance; (2) The perceptual signal contains information about the disturbance; (3) The output signal does not mirror the disturbance, and is uncorrelated with the disturbance. At least one of these three must be true. Rick's claim is that all are false. The Q that ED is that Rick's claim is false; accepting (1 and 3) as false requires (2) to be true.

QED.

>Furthermore, the sensor should not really detect the output
>signal directly, but the EFFECT of the output signal as
>translated through the output transducer and through the
>environmental link between the physical output and the controlled
>quantity qi. But that's a side-issue.

For the purpose of the argument, it's a side issue, which is why I so labelled it in my original posting; but it will loom large in the information-theoretic argument yet to come, once we get this roadblock out of the way.

ALL THE INFORMATION IN THE DISTURBANCE CAN BE RECOVERED FROM THE PERCEPTUAL SIGNAL (There--I can CAPITALIZE too).

>In these developments, "information" is being bandied about loosely.

For the purpose of arguing with Rick, I use Rick's definition. Besides, if we can recover the disturbance exactly out of the perceptual signal by a demonstrable mechanism, there seems little need to fluster about the abstractions of information theory, which are often misunderstood.

Your argument from information-theoretic principles is, so far as I can see, correct in demonstrating that the current value of the disturbance is not recoverable from the current value of the perceptual signal. But then, I know of nobody who thought it was. That's red herring that Rick keeps asking Allan and me to accept as a fish fresh for cooking. I won't, and I don't think Allan will, either. (Though I really like a good kipper, it isn't so nice after it has been dragged across the fox's trail).

I should actually compliment you on the IT analysis, because it saves me a bit of work later. I was intending to go through the points you raised, but maybe now I need only refer to your posting.

Martin

Date: Wed Mar 31, 1993 11:59 am PST
Subject: DME & HOMUNCULUS - RKC

[From Bob Clark (930331.1420 EST)] RE: Martin Taylor (930328 20:30)

Your Post reads:

>Count me among the non-acceptors. Your DME sounds very much like the old
>homunculus who sits behind the sensors and effectors, manipulating. How
>does he work? Does he have his own little hierarchy?

OBSERVER'S VIEWPOINT (Bob's)
From this viewpoint, this post is equivalent to two decisions: 1) "Count me among the non-acceptors" is equivalent to your (Martin's) having decided that the First Person Singular does not refer to a Decision Maker; and 2) the posting of your (Martin's) decision to the Net is equivalent to a second decision.

USER'S VIEWPOINT (Martin's) REQUESTED
Who, or where, is the "ME" included in your (Martin's) post and involved in creating it? Please explain your (Martin's) alternative(s), with or without using PCT. Remember, I (Bob) am assuming a situation where both 1) the established "nodes" or "choice points" are, for any reason, unable to provide a "decision" and, 2) "Intrinsic Error is neither present nor anticipated.

"HOMUNCULUS" -- THEORIST'S VIEWPOINT (Bob's)
My dictionary (Webster's Encyclopedic Unabridged Dictionary, 1989) gives: "homunculus, n. 1. a diminutive human; midget. 2. a fully-formed, miniature human body believed, according to some medical theories of the 16th and 17th centuries, to be contained in the spermatozoon."

I don't think you intend the term, "homunculus," to be taken literally per the definition above. More important, in my posts I have tried to restrict the capabilities of the DME to those without which it could not perform its defining function: making decisions. Can any of these be omitted? Should any others be added?

"HOW DOES THE DME WORK?" -- THEORIST'S VIEWPOINT (Bob's)
For the original suggestion, see Bob Clark (12052). Also the post to which you responded, Bob Clark (930327.1945 EST). Other posts have attempted to further clarify the idea.

Here is another attempt to describe the essential characteristics of a Decision Making Process -- a DME in operation.

It is assumed that:
A. no built-in automatic branch-point is available
B. no Intrinsic Error currently exists.
C. no Intrinsic Error is anticipated.

I have tried to limit this description to those items without which Decisions cannot be made. Thus the proposed items are:

1. Current Perceptions. The DME selects the signals to which it directs its attention. They are selected from among the incoming neural signals available. These signals are available for use as feedback signals if needed.

2. Current Objectives (reference levels, etc), if not already in operation, are selected from Recordings of past decisions, events, etc.

3. Past Perceptions -- Recordings (Memories). The DME finds Recordings both by named addresses and by similarities of content. They may result from simple "recognition" ("reminders"), or (more or less) extensive searches for relevant material.

4. The Recordings found are examined for relevance and possible application ("feasibility") to the Current Perceptions (Perceived Situation).

5. The Recordings are further examined, by imagination, for anticipated future effects as they relate to current, relevant reference levels ("objectives," "purposes," etc).

6. The entire hierarchy is available to serve as the Output Function for the DME. In ordinary situations, only limited, selected portions will be needed.

7. On the basis of the above examinations, etc, the DME selects and activates a Recording. The DME's selection can be arbitrary. The Recording selected can consist of revised and/or combined Recordings.

If this is a "Homunculus," so be it.

For Perceptual Control Theory of Behavior to be complete, it seems essential to me that "Decision Making Activities" are included somehow.

In addition, I think that these elements are consistent with most, if not all, of the ideas either stated or implied by Powers in BCP.

Finally, to repeat; Please let me know your (Martin's) procedures for making ordinary decisions, and what part(?) of you (Martin) does it.

Regards,   Bob Clark

Date: Wed Mar 31, 1993 12:36 pm PST
Subject: Sample data from tracking run

[From Bill Powers (930331.1300 MST)]

Allen, Martin, other interested parties.
FYI:

Here is a record of 2 seconds of a 60-sec experimental tracking run with a real subject (me) and a moderately hard disturbance. The scale is 1 unit = 1 pixel on the screen.

column 1 is the handle position
column 2 is the disturbance magnitude
column 3 is the cursor position

In this experiment, it happened that

cursor = handle - disturbance   (instead of plus)

o	d	qi	82	102	-20	66	122	-56
			81	102	-21	62	122	-60
95	93	2	80	108	-28	60	130	-70
90	93	-3	80	108	-28	59	130	-71
87	97	-10	76	115	-39	54	130	-76
86	97	-11	74	115	-41	52	139	-87
85	102	-17	73	115	-42	51	139	-88

50	148	-98	253	226	27
50	148	-98	252	226	26
50	148	-98	251	221	30
50	158	-108	251	221	30
50	158	-108	251	221	30
50	167	-117	245	214	31
51	167	-116	242	214	28
52	167	-115	241	206	35
52	177	-125	234	206	28
64	177	-113	230	206	24
70	186	-116	228	198	30
73	186	-113	227	198	29
74	186	-112	216	189	27
92	195	-103	211	189	22
101	195	-94	208	189	19
106	204	-98	207	180	27
108	204	-96	195	180	15
126	204	-78	189	170	19
135	211	-76	186	170	16
140	211	-71	175	170	5
155	218	-63	170	160	10
162	218	-56	167	160	7
166	218	-52	166	149	17
168	225	-57	157	149	8
179	225	-46	153	149	4
185	230	-45	151	138	13
188	230	-42	143	138	5
195	230	-35	139	128	11
198	235	-37	137	128	9
200	235	-35	136	128	8
201	239	-38	131	116	15
206	239	-33	129	116	13
208	239	-31	128	105	23
209	241	-32	127	105	22
210	241	-31	121	105	16
216	242	-26	118	93	25
219	242	-23	117	93	24
221	242	-21	111	81	30
228	243	-15	108	81	27
232	243	-11	107	81	26
234	242	-8	106	69	37
235	242	-7	105	69	36
237	242	-5	104	57	47
238	241	-3	104	57	47
238	241	-3	104	57	47
238	239	-1	102	45	57
242	239	3	101	45	56
244	239	5	101	33	68
245	235	10			
249	235	14			
251	231	20			
252	231	21			
253	231	22			

Best, Bill P.

Date: Wed Mar 31, 1993 12:49 pm PST
Subject: Re: Movement

[Martin Taylor 930331 15:30]   Dennis Delprato 930331

>Thus, if we go and describe the details of chemoexcretions,
>we find motorsensory feedback control systems. I do not mean movement
>in only the sense of gross muscular activity.

If that is so, then I retract in my differentiation of chemical from mechanical control, and (for now) accept that all actual control is by mechanical means.

Bill Powers also commented that he did not think we actually used chemical output as the output mechanism in the main hierarchy, but that chemical output was a kind of side effect (if I read him aright), much as the behaviour of the main hierarchy is a side effect of the reorganizing system. The result in each case is that intrinsic variables come to desired states (i.e. are under control) even though they themselves are not the percepts of any control system. I can accept this as a possibility.

But I can also accept the converse. How would you tell whether a chemical output was the effector mechanism of a control system (desirable member of opposite sex being in close proximity being the perceptual reference of some such control system). How do you test whether any particular muscle twitch is part of the effector system of any particular higher-level control system? All you can do, as far as I can see, is determine whether a putative CEV actually corresponds to a controlled perception, and the effector mechanisms that influence that putative CEV may well not be controlled. They may just be executed as outflow (especially if chemical!). Even if they are controlled, how does one test whether their control is related to the control of the putative CEV in question? The Test, as such, seems not to apply.

Martin

Date: Wed Mar 31, 1993 1:57 pm PST
Subject: Simcon simulation of Martin's setup

[From Bill Powers (930331.1430 MST)]   Martin Taylor (930331) --

Rick is right. Simulate your proposed setup with Simcon and see what happens. It is not what you say happens. The signal X does not reproduce the waveform of the disturbance. However, the variable q = p - o does reproduce it.

Here is a Simcon program to try:

```
title reconstructing disturbance from perceptual signal
#
time 10.0 0.05
t generator puls 2.0 6.0 10.0
y generator puls 4.0 8.0 -5.0
```

```
d summator t 1.0 y 1.0    # composite disturbance waveform
p summator o 1.0 d 1.0    # p = o + d
r const 5.0               # ref signal constant at 5.0
e summator r 1.0 p -1.0   # error signal
o integrator 0.0 e 1.0    # output signal
v summator r 1.0 p -1.0   # Mystery function
x integrator 0.0 v 1.0    # Mystery function, X
q summator p 1.0 o -1.0
group p e o
print d o x q             # show disturbance, output, X, q
plot
```

I leave it as an exercise for the student to explain why p - o is not the same as O(r - p). Hint: feedback has something to do with it.

Best,   Bill P.

Date: Wed Mar 31, 1993 3:22 pm PST
Subject: Re: DME & HOMUNCULUS - RKC

[Martin Taylor 930331 16:30]   Bob Clark 930331.1420

>USER'S VIEWPOINT (Martin's) REQUESTED
>Who, or where, is the "ME" included in your (Martin's) post and involved in
>creating it? Please explain your (Martin's) alternative(s), with or
>without using PCT.

Two questions and an assumption. I recognize the existence of consciousness in me, and extend you the courtesy of assuming it exists in you. I have no explanation of it, other than the simple presumption that its content must be based on signals in the hierarchy, and that it is not itself such a signal. Consciousness is a multidimensional experience. "ME" is an element of consciousness.

The assumption: that this is a User's viewpoint. What I mean by a "User's" viewpoint is that you are can take account of only the signals accessible at that point. The User's viewpoint of an ECS is not that of a person within whom the ECS operates. It is consideration of what is accessible at some point within the ECS, often the perceptual signal, but possibly one of the other signals. Your DME does not have a User's viewpoint of the action hierarchy. It has access to signals from all over the hierarchy.

>Remember, I (Bob) am assuming a situation where both 1)
>the established "nodes" or "choice points" are, for any reason, unable to
>provide a "decision" and, 2) "Intrinsic Error is neither present nor
>anticipated.

What is a "decision" WITHIN the control hierarchy? It must be something that happens at the program level or above (assuming Bill Powers' set of levels). Below the program level, there may well be multiple means to achieve any particular perceptual signal value, but the variation of means must be caused by differences in the reactivity of the world. The increase of difficulty (I sometimes say "impedance") of one lower-level control may mean that a higher-level perception is brought under a control by an entirely different set of actions. This is not "decision" as I understand it. It is a natural consequence of there being a non-linear system with more (in this situation) degrees of freedom for output than there are perceptual degrees of freedom being controlled at a high level.

Something nearer "decision" can occur within the hierarchy below the program level, when possible actions are played through imagination loops in various ECSs. I suspect this happens all the time, and is not switched. The effectiveness or otherwise of this imaginary control may affect the real gain of different ECSs, resulting in different real patterns of action when the imaginary control is actualized. Again, there is the appearance of decision without any actual decision.

At the program level, "decision" is intrinsic to the level. It is the nature of the program level to select among sequence reference levels. There, decisions have an explicit place within the hierarchy, and I would think that they would be accessible also at higher levels. But that's pretty high in the hierarchy.

Now back to your quote: If you are talking about a Powers type of hierarchy, you must be talking about the program level or above, because below this level there are no choice points. The PIFs do not permit them.

How does one "anticipate" intrinsic error? One can't even perceive it when it does occur, according to Powers. I don't think it is relevant to the issue of the DME.

>[HOMUNCULUS definition omitted]
>I don't think you intend the term, "homunculus," to be taken literally per
>the definition above.

No, of course not. One of the reasons that behaviourist psychology became popular in the early years of this century was that people saw that most of the 19th century psychological theories were recursive. To explain what a human did, they in effect passed the results of sensory processing to a "little man in the head" who decided which levers to pull and push to make the muscles work. All the issues of the psychology of the human were incorporated within the LMITH, and he was usually called "the homunculus."

>1. Current Perceptions. The DME selects the signals to which it directs
>its attention. They are selected from among the incoming neural signals
>available. These signals are available for use as feedback signals if
needed.

On what basis is this selection made? What is the perception that the DME is controlling by means of varying its choice of neural signals? "...for use as feedback signals..." in what control loop?

>6. The entire hierarchy is available to serve as the Output Function for
>the DME. In ordinary situations, only limited, selected portions will be
needed.

So the hierarchy is the environment on which the DME operates, exactly as does the Powers reorganizing system? Your seven characteristics certainly seem to indicate this. But how does the DME itself operate? Is it controlling anything? If so, what can it be controlling but its own perceptions? And if it is controlling its own perceptions, do not the same considerations apply to it as to the main hierarchy: it is a hierarchy of perceptual control systems, needing a sub-DME to make decisions on its behalf, such as what signals in the hierarchy to attend to?

>If this is a "Homunculus," so be it.

Well, it still sounds like one, in that it solves an acknowledged problem within the control hierarchy by replicating the problem at a new level. The recursion, as with the original psychological conception of the homunculus, is potentially infinite.

>For Perceptual Control Theory of Behavior to be complete, it seems
>essential to me that "Decision Making Activities" are included somehow.

Yes, but why must they be outside the control hierarchy? Isn't the program level adequate? Remember that in the Powers system, perceptual input functions may accept any neural signal as input, though in our diagrams and analyses we usually consider only the perceptual signals of the next lower level of control.

>Finally, to repeat; Please let me know your (Martin's) procedures for
>making ordinary decisions, and what part(?) of you (Martin) does it.

If I knew that, I would join the ranks of those making pronouncements about the truth of the world, and I might be rich, into the bargain.

Look, my problem with the DME as an entity isn't a matter of faith that everything can be solved within the main hierarchy (though I like to think that true, and is one reason I continue to think of local reorganization instead of postulating a separate reorganizing system). My problem with the DME is that it seems to do the same kind of job within the main hierarchy that the main hierarchy does in the outer world. That means that the DME must need its own DME, which needs its own DME, which... In other words, introducing the DME does not seem to solve the problem it addresses. If I misunderstand what the DME is supposed to be, then I'm quite happy to retract all I have said. But I have indeed read your postings, and refrained from comment for lack of time. I simply didn't want silence to be taken as acceptance when you made that an issue.

Martin

Date: Wed Mar 31, 1993 4:20 pm PST
Subject: Re: Martin's Thought Experiment

[Martin Taylor 930331 18:15]   Rick Marken 930331.0800

Rick accepts my thought experiment, and suggests doing a simulation, which he suggests will not turn out as I claim. Seems reasonable, though I am not clear how they could come out different.

>>I'll leave the logical problem with the "mirroring" unaddressed, and
>>assume that Rick accepts as correct what he says, that the output
>>mirrors the disturbance.
>
>What's the "logical" problem with the "mirroring"? You can look at
>the data from our tracking experiments and see that o = -d to within
>a few pixals throughout an experimental run.

The logical problem is that the output of the ECS, like the perceptual signal, is a neural current, whereas the disturbance is something in the world. The output being modelled by the function M(r,p) in my proposal is a neural current. The output one sees on the screen (Rick's "o") is related to that neural current, but may be affected also by the actions of other control systems, so there is no one-to-one relation between the output and the effect on the screen. There IS a correspondence between the output and the error, so this lack of a one-to-one relationship does not cause a problem for control (as it would if we were dealing with inverse kinematics). It does mean that the signal X in my model does not replicate the disturbance. It mirrors the output.

Martin

Date: Wed Mar 31, 1993 5:13 pm PST
Subject: Endgame?

From Greg Williams (930331 - 2) Bill Powers (930331.0900 MST)

>You are smart enough to realize what I was getting at, but of course I had
not
>said why I was making these predictions.

Thanks for the compliment (?), but my introspection tells me that I didn't realize what you were getting at.

>This does not prove my hypothesis correct, but it is an
>encouraging hint. You appear to be perceiving that the Test is
>very difficult to carry out in the field.

I'd call it an untested hypothesis -- I haven't seen the evidence yet to draw the perceived conclusion you seem to think I've already made.

>I presume that you are motivated to perceive this X, as no definitive
>conclusion has yet been reached on this subject.

That sounds more reasonable, and agrees with my current introspected motivations.

>On that hypothesis, I would predict that everything I to do alter that

>perception will be met by some move that will counteract my attempt,
>and support the idea that the Test is very difficult to use in the field.

This is sort of an aside, but if my motivation is REALLY as I introspect it and as you have hypothesized, why do you think that I must try to counteract EVERY attempt to apply the test?

>One sign of this, as I interpret the proceedings, is that when I
>ask a very simple question like "Do you still intend to maintain
>this challenge," your answer has uniformly dodged the request to
>state your intention; after all, to do so correctly would be to
>show that the Test can succeed, and in a simple way at that.

Well, let's see. I introspect currently (and honestly!) that I still maintain the challenge and that I intend to maintain it until I think you have met it or give up. I don't think I can be much more explicit about my introspected motivations. Your move.

>(By the way, I find that I
>already have all your materials, including the script, tucked
>away in a long-forgotten directory, so you don't need to mail the disk).

OK.

>So it's evident to me that when I pose questions in this context,
>you're answering in an unusual way, and doing so only when I ask
>the sort of question I am asking.

I promise to try to be as usual as I can be.

>PCT claims that all people have reference signals and
>perceptions. My interpretations of it suggest that by various
>means they can become aware of them if they want to.

What evidence do we have that our awarenesses of our motivations (which I agree we do have) are congruent with our actual "underlying" controlled variables at the level below the surface phenomena of behaviors? Here you are, still trying to get me to tell you what I introspect as the reasons for certain of my behaviors, when I've been telling you that the solution to the problem lies elsewhere, far from my introspections. What you need to do to convince me that you understand why I am behaving in a certain way is to show that when you apply the Test, repeatedly, I attempt to maintain control of certain perceptions. But it might not be an easy task to decide what perceptions I am attempting to maintain. I can foresee difficulty if the essentially the same outputs can be used to maintain VARIOUS reference levels.

>This is not a matter of finding some objective experimental proof; it's a
>matter of recognizing something that is already there and that can be
>observed quite easily -- but can't be made sense of so easily, without
>some organizing theory.

This I doubt, having just been reading through the literature on "waking suggestion," where subjects will swear up and down that they aren't aware of moving in certain ways (presumably to try to maintain certain "unconscious" perceptions?). If these subjects were asked "Why were you leaning forward?" they would answer "I wasn't moving forward."

As ever,   Greg

Date: Wed Mar 31, 1993 6:01 pm PST
Subject: Re: The Experiment

[Allan Randall (930331.1730 EST)]

This is a response to various posts from Rick Marken and Bill Powers.

Rick Marken (930329.2000) (930330.1130) (930330.1500)

I think you have understood most of the experiment correctly, but a few points need to be clarified:

>Let me see if I have this right. I assume that cutting the "perceptual
>line" means that the system is now operating open loop.

Yes. The system is now simply getting percepts and producing outputs. The output no longer has any effect on the input. The whole objective is to replicate the closed-loop result in an open loop, using only the percept.

>Also, what is the length of a perceptual input? Are you referring to the
>length of the vector of values that make up Pi?

Basically, yes. The length is the number of bits actually used to specify the percept in the computer simulation. This is not a single value of the percept at one moment in time, but a vector of values over a time period.

>...So if D had 1000 samples and I could recover all 1000 values by plugging
>in a Pi with 500 samples, then the entropy of D is 500? Is that right?

Yes, provided that 500 is the *minimum* required. That is why we need to start with the shortest and work up.

>Is this right? ... Is P coming in via the usual closed
>loop perceptual input or "for free" as an open loop input?

Open loop, for free. Your intuition here was exactly right. The procedure for H(D|P) is identical to that for H(D) except that we are given P for free. Both entropy calculations are done strictly open loop, since the algorithmic definition of entropy we are using is defined in open loop terms: input, program, language and output.

>I know you expect to need only one P0 value -- that any a P0 value of length
>0 can be added to P' and maintain the ability of P' to produce O=D.
>But you do have to try adding at least ONE P0 string (other than
>the zero length one) to show that this is true.

I think we have a notational confusion here. Pi, *not* P0, is my notation for an arbitrary value in the set {P0,P1,P2...}. There is only one P0, and it is of zero length by definition. It is the Pi values that can be of arbitrary length. Since the procedure is to look for the shortest successful string, there is no need to try *any* non-zero-length Pi's *if* P0 is successful. Otherwise, we try as many as needed to find a successful contender. The procedure is simple: try all Pi's in order of shortest to longest, and STOP as soon as one is successful. Also, I assume your P' is the same as my Pk - it represents the first successful Pi in an open loop search.

>>Do you disagree about the results that I have assumed we would get
>>from this experiment?
>Yes. But you'll see when you actually do the experiment.

As I said in a message to Bill, I think this experiment is almost redundant, since P' will turn out to be zero length (P0) for H(D|P), and this is the very first Pi we will try. This part of the experiment is essentially identical to Martin's proposal for a "mystery function." If I am right, the rest of the experiment will be unneccesary. If you are right, I will have to continue with the search. So I guess I'll just have to do the experiment.

Do you understand my assumption that we will not actually be looking for D in the open loop exponential search? Since you have already agreed that nearly 100% of the information about D is in the closed loop output, I have assumed it will be sufficient to look for replication of this closed-loop output in the open loop phase (without using the closed-loop, of course). I am trying to show your claim that the output has 100% of the information about D and the input has 0% of it to be inconsistent, so I think this assumption is valid for now.

>Rick ("There's no information about disturbances in controlled
>perceptions") Marken

Ashby's whole point was that perfect control was not possible, and that error control relies on the detection of minor imperfections to prevent major imperfections. This is why it's called "error control," after all. If perfect control were possible, then yes, there would be no information about the disturbance in the perception. This is the goal of the control system. But it will always fall short of this, since it controls via the detection of imperfections. As Ashby said, perfect control would perfectly block the very channel through which the control system gets its useful information. Thus, error control is inherently imperfect. Its very power lies in its inherent imperfection.

>...The length of the first candidate i vector that produces
>o values that match d perfectly is H(D) (There is a problem here;
>What if you don't get any perfect matches to d, Allan?.Nothing
>but an infinite loop gain system could produce outputs that
>match d perfectly anyway: how about doing this just until a
>candidate i vector produces o values that match the disturbance
>to the same degree as did the o values generated in the closed loop
>case.

This is basically what I did - I specified that, since we agree that the closed-loop output has near-100% of the information about the disturbance, then we are simply going to look for a replication of the closed-loop output, rather than the disturbance.

>Assuming you can find H(D) using this decidedly peculiar technique
>(why not just measure the variance of d?)

Because what I have described, as "peculiar" as it may seem, corresponds directly to the technical definitions you will find in the literature on algorithmic information theory. Variance does not. Information is defined in terms of probability, not variance.

>Allan is assuming that the output resulting from the original i vector
>along with the "null" (0 length) candidate i vector will produce an
>output that perfectly matches the disturbance (or, at least, matches
>the disturbance as well as the output did in the original run).
>
>Is this a correct description of the experiment Allan?

Yes, exactly.

Bill Powers (930330.2000 MST)

>Allan is using information about o in his method,
>as I vaguely understand it now, with your help. It shouldn't be
>surprising if he can also reconstruct d.

No, I am not using information about o to reconstruct d (or in this case the original o).

>By the way, I think this notation H(p|d) is not just an ordinary
>function, but represents some sort of probability calculation
>with base-2 logs and all that. Allen?

H(P|D) = -log prob(P|D) (see my original posting)

Rick Marken (930330.2100) responding to Bill Powers:

>>Allan is using information about o in his method,
>>as I vaguely understand it now, with your help. It shouldn't be
>>surprising if he can also reconstruct d.
>I don't think Allan is going to use information about o (the
>o generated in the real, closed loop run) at all. In fact, the
>way I conceive of Allan's study, you don't need to save o at
>all; all you need is d and p from the closed loop run.

Your description is essentially accurate, except that I suggested looking for the original closed-loop o, rather than an equal d-correlation with it, as you suggest. Either strategy would be okay as far as I am concerned. Given that we assume the original o has 100% of the information about d, it really doesn't matter terribly much which strategy we use.

>...are you with us, Allan?

I think so. As far as I can tell, your description is accurate.

Bill Powers (930331.1030 MST) responds:

>As I understood it, Allen was going to assume that control was
>perfect, so o = -d with a constant reference signal. So the
>procedure depends on knowing o, which amounts to knowing d.
>
>But let's wait to see what Allen says.

No, the procedure that replicates the closed-loop o in the open loop phase uses only the Pi value, not o. The closed-loop o value (or d if we choose Rick's version) is used only for comparison purposes to check the result. It is not used in any capacity at all to produce the open-loop result. I believe Rick has understood the proposal, and so I will now attempt to implement it. I probably will not post anything else on this subject until I have the results.

Allan Randall, randall@dciem.dciem.dnd.ca
NTT Systems, Inc.   Toronto, ON

Date: Wed Mar 31, 1993 8:27 pm PST
Subject: Misc; Challenge 2

[From Bill Powers (930331.1930)]   Martin Taylor (930331.1530) --

>How would you tell whether a chemical output was the effector
>mechanism of a control system (desirable member of opposite sex
>being in close proximity being the perceptual reference of some
>such control system).

By looking at the effect it has, and seeing whether interrupting that effect causes a counteractive change in the chemical output. For example: if perspiration might act as a pheremone that attracts potential mates or as a coolant, one could arrange for either of these effects to be disturbed, and see whether the result was more or less perspiration. Bringing the skin temperature back to the pre-perspiration level should increase the amount of perspiration if perspiration acts as a coolant. Keeping the jock groupies away should increase perspiration if perspiration is the output of a control system controlling for potential mates. The latter hypothesis, incidentally, seems to be bought into doubt by the sales levels of deodorants.

>Even if they are controlled, how does one test whether their
>control is related to the control of the putative CEV in
>question? The Test, as such, seems not to apply.

The Test focusses on the _effects_ of the action, not the action itself. If the _effect_ is disturbed, and is a controlled variable, the action should change appropriately.

Mark Olsen (930331) --

Application noted. Announcement will soon appear. We have 5 scholarships available this year, thanks in part to generous contributions from well-wishers. The final choice will be made June 1, so there's no rush.

Martin Taylor (930331.1815) --

>The logical problem is that the output of the ECS, like the
>perceptual signal, is a neural current, whereas the disturbance
>is something in the world.

This applies only to the model, not to the experiments. In the experiments, we measure the actual output (the signal produced by the control handle), not the neural currents or the muscle activities. The transformation from the model's error signal to its output is all collapsed into the output function we use in the model.

When a model that uses just a simple integrator in its output function matches behavior within certain limits, the result is to show how well a simple linear integrator represents all the processes between the neural comparator and the position of the handle. The model therefore makes a statement about nonlinearities, dynamics, effects of other control systems, and all that. The kind of fit we get says that these other considerations have only minor effects, if any.

Greg Williams (930331 - 2) --

>I'd call it an untested hypothesis -- I haven't seen the evidence yet
>to draw the perceived conclusion you seem to think I've already made.

Each time that I make a request for information, I'm inserting a disturbance. I make a prediction to myself about which requests will be complied with and which will not. Each time you reply to my request, you are either providing the information, showing that providing it does not constitute a disturbance, or resisting providing it, showing that you're taking steps to counteract the disturbance. This is a pure application of the Test.

The lowest-level evidence I have is simply that requests for information about your intentions have been -- until a day or two ago -- uniformly resisted. This tells me that supplying such information is a controlled variable, with a low reference level. It doesn't tell me why the variable is controlled, or why the reference level is set as it is. But it does indicate at least reasonable support of the hypothesis that you intend to keep subjective reports at a low level, even when requested. This became clear quite a while ago.

I did the usual variations, such as asking similar questions on different topics involving other perceptions, and encountered no resistance. I found that giving subjective reports was not absolutely forbidden; there were circumstances in which you freely told me what you thought. I also tried a disturbance that called for altering this reference level as a way of preserving what I was hypothesizing as a higher-level controlled variable, setting up a "fork". When I made those three predictions, I set it up so that if you continued to refuse to state what you wanted, you would have to admit that I had made a correct prediction (that you would not answer any of the questions Yes or No), whereas if you chose to give an answer that made my predictions wrong, you would have to answer yes or no to a question about your intentions. You answered just as I would have predicted under the hypothesis that you were controlling for failure of the Test: all my stated predictions that you would NOT answer yes or no were marked "wrong," and all my predictions that you would answer in an evasive way were also marked "wrong" because you had answered two of them with yes or no, and one with a point on a scale between "success" and "failure." So this disturbance showed that the lower-level reference signal was momentarily reset to a higher value, allowing a subjective report to be given (and perceived), in order to preserve a higher-level variable at its reference level: showing that the Test was failing.

>This is sort of an aside, but if my motivation is REALLY as I
>introspect it and as you have hypothesized, why do you think
>that I must try to counteract EVERY attempt to apply the test?

That would be too simple for Greg Williams. You don't counteract attempts to apply the Test; you simply think ahead and thwart them.

>Well, let's see. I introspect currently (and honestly!) that I
>still maintain the challenge and that I intend to maintain it
>until I think you have met it or give up. I don't think I can
>be much more explicit about my introspected motivations. Your move.

Thank you. If I were dealing with an innocent person, my move would be to say something like "Well, I don't know. This has been going on for a long time, and I don't really know if you're just keeping it up for my sake." My intent would be to communicate a mild inclination on my part to give up and stop participating in the challenge, without indicating that this is my actual desire. In other words, having heard a description of the stated reference condition, I would then proceed to test it by applying a disturbance that ought to change it if there's no counteracting response. From what you just said, I would expect some sort of encouragement to continue from you. During our conversations, I actually tried that sort of test a few times, to see if you really wanted to continue, and apparently you did because you resisted my suggestion -- actually, my taking up of your suggestion -- that we let me judge the outcome instead of you and make the outcome easy. When I explained why I felt you should continue to be the judge, you dropped the subject. However, when I suggested at one point that I wasn't going to tell you something (trying to see if you wanted to find out), you were too smart for me. You just said OK. I didn't believe you, but I got the message that you saw through my ploy.

Now, of course, I can simply ask in the shorthand terms with which we are both familiar. If I became reluctant to go on with the challenge now, would this feel like some kind of error to you? If, on the other hand, I started asking more and more personal questions, pressing to get you to investigate more and more subjective impressions, would that also feel like an error, but in the other direction? My hypothesis would be that you have decided raise your reference level for discussing subjective information, but not anywhere near as high as it could go. So attempts to change the perception in either direction would be resisted, as usual for a controlled variable.

>What evidence do we have that our awarenesses of our
>motivations (which I agree we do have) are congruent with our
>actual "underlying" controlled variables at the level below the
>surface phenomena of behaviors?

I take a deliberately naive view on this most pertinent question. I took it about 30 years ago. It is that everything of importance in the brain, _at the functional levels we experience_, is open to inspection if we happen to look. All we have to do is notice what is going on. This is complicated by the fact that the scope of attention is limited, and a lot that goes on is nonverbal. More than that: we have developed no way to fit many of these nonverbal experiences into our higher-level models of the world. It's like asking "What is the word for the way a forkful of spaghetti looks?" We know how it looks, but there's no word for it. Having no word for it, we can't incorporate that experience into our higher-level models of the world and reason about it in symbols as we do with other experiences that have names. Perhaps a better example would be "what is the word for the way a knot looks?" Most people, including myself, have no such word (it looks like a knot) -- but topologists have developed an extensive vocabulary about knot-experiences, and have a great deal to say about them that nobody else knows how to say. We all experience knots, but topologists can reason about them.

The PCT model provides some words that can be attached to experiences that are all familiar to us, but which either have never been labeled in a systematic way, or have been labeled too ambiguously to lead to any consistent way of reasoning about them. Think what the term "reference signal" does for us. It peels away all irrelevant side-issues that come attached to words like want and desire and longing and wish and intention and purpose, and points to a single central phenomenon: a specification for a perception. By doing so, it doesn't sweep the nuances of these other terms away; it merely shows us that they refer to something in addition to a specification for a perception. This makes it easier to understand what those added factors are, and see them and reason about them as phenomena in themselves. All the terminology of PCT works this way, giving us a handle on experiences that otherwise just happen, namelessly.

We speak a lot about generative models, as part of the formal methodology of PCT. But beneath that level, PCT is descriptive. What it describes is not the part of the world to which we assign the label "outside," but the totality of it, including the part that is "inside." The Test for the Controlled Variable is, in one context, a formal method that is part of the rational structure of PCT. In another context, however, it is a way of connecting terms such as "perception" and "reference signal" and "error" to otherwise unnamed experiences. When used as we're using it now, it is only in part a research tool. It's also a learning tool.

Is this terminology and the conceptual structure of which it is a part "correct?" I think that's probably the wrong question. The important question is, "Does it bring our experiences of ourselves and others into a common framework of understanding that is consistent both within itself and with experience?" Whatever one means by "correct," I think that if PCT gives an affirmative answer to the question, it will have earned its keep.

>What you need to do to convince me that you understand why I am
>behaving in a certain way is to show that when you apply the
>Test, repeatedly, I attempt to maintain control of certain
>perceptions. But it might not be an easy task to decide what
>perceptions I am attempting to maintain.

Given the conditions under which we're working, I think I have at gone at least part of the way to finding the kind of evidence you want. You may not agree with the words I use to express what I think I see being controlled, but once you see that the evidence does point to something being controlled, you can probably supply the correct meanings better than I could. The outside observer always operates under a handicap, in that any definition of a controlled variable is necessarily cast in terms of the observer's meanings. There's a limit to how much the external observer can learn.

>I can foresee difficulty if the essentially the same outputs
>can be used to maintain VARIOUS reference levels.

The outputs don't matter. All that matters is their relationship to the disturbance that's going on at the moment. The hypothesis always is balanced on the fulcrum of the controlled variable; the disturbance shows a relationship to any output only in terms of that controlled variable.

If the same output maintains a different controlled variable, this implies some change in other perceptions as well, and most probably indicates that the output is just one of several that contribute to the new variable. In other words, that the output is being seen at too low a level. If the reason is a change of organization, then the old variable will quickly prove not to be under control any longer by that output. Showing a lack of control is far easier than pinning down a controlled variable. But I think we will have enough to do for a while if we just try to indentify stable organizations.

>This I doubt, having just been reading through the literature
>on "waking suggestion," where subjects will swear up and down
>that they aren't aware of moving in certain ways (presumably to
>try to maintain certain "unconscious" perceptions?). If these
>subjects were asked "Why were you leaning forward?" they would
>answer "I wasn't moving forward."

Things like this do happen. They are fringe phenomena, a drop in the bucket. All our perceptions lie to us to some extent, judging by cross-checking them against other perceptions. But we are always cross-checking, always trimming up our perceptions to make them agree better with each other in terms of higher-level perceptions.

I think that the importance of the little weirdities of experience have been overblown, simply because the behavioral sciences have missed recognizing the vast number of regularities in which the weirdities are embedded. When you have trouble finding a defineable stimulus or a defineable response in the incessant behavior of an organism, unusual phenomena that happen in unusual circumstances might not be recognized as being so unusual.

Where do we stand?

Best to all, Bill P.

Date: Wed Mar 31, 1993 10:21 pm PST
Subject: Re: Info about disturbance

[From RIck Marken (930331.2100)]

Martin Taylor (930331 14:15) --

>ALL THE INFORMATION IN THE DISTURBANCE CAN BE RECOVERED FROM THE
>PERCEPTUAL SIGNAL (There--I can CAPITALIZE too).

That's your position, all right. And I agree that this is true as long as p = o + d and you have access to both p (the perceptual signal) AND o. But no information about the disturbance can be recovered from the perceptual signal all by itself (or, as you will see from the simulation, with the help of the output function).

>For the purpose of arguing with Rick, I use Rick's definition [of
>information]. Besides, if we can recover the disturbance exactly out
>of the perceptual signal by a demonstrable mechanism, there seems little
>need to fluster about the abstractions of information theory, which are
>often misunderstood.

This is a very helpful way to formulate the problem; it let's us get past definitions and into working simulations of control systems. I have developed a HyperCard stack that implements your ingenious method of testing whether or not one can recover the disturbance exactly out of the perceptual signal. Your method does not use knowledge of o to extract the information about d from p -- so I consider it quite a fair method. It does use information about the output function (that transforms error into the variable that affects the perception). But I think you imagine that it is the output function that extracts (recovers) the information from p. In fact it doesn't -- but that's what the simulations will show. You will find when you run the simulations that there is only one case in which you can recover the disturbance from the perceptual input (recover it from the output os what Bill calls the Mystery function. This occurs ONLY when R = 0 and P = O + D. Even in this case, however, when you introduce the realistic possibility that there will be neural and integration error added to the computation, the closed loop output, o, still approximates -D (with a very small RMS deviation); while the open loop output deviates from D by a huge amount.

The HyperCard stack lets you compute the match of the open loop output of the Mystery function and the closed loop output of the output function (both computed simultaneously, of course) under various conditions: fixed or variable reference signal (fixed at 0 or 50), linear or non- linear connection between o and p (ie. p = o + d is linear; p = o^3 + d is non-linear) and with noise (the same noise, just to make it fair) added or not added to the input to the output and Mystery function.

The results are quite interesting; but they do show that in one very specific condition (p = o + d, r = o, no error in the system) you are right -- the output of the Mystery function can be considered to be the extracted value of the disturbance. So this means that there IS information in the perception when the reference for that perception is 0, there is no error contributing to the signal going into the output function -- ie. no neural noise -- and the output of the system, o contributes linearly (and with a coeficient of 1) to the perecptual input variable. Given those conditions, you can fairly claim that in- formation about the disturbance can be extracted from the perceptual signal. But there is no such information available when any of these con- ditions don't prevail -- that is, when the reference is not a constant = 0, when there is neural noise or when o has non-linear (or linear with a coefficient not equal to 1) effects on p.

Would you like the stack?

Best   Rick