

Date: Tue Jun 01, 1993 8:46 am PST
 TO: Buhts (Ems)
 MBX: h2de9reb@orl41.orl.usace.army.mil
 Subject: PCT, AR 5-1

[From Dag Forssell (930601 09.45)]

Major Robert E. Buhts

Dear Bob:

Through Gary Cziko, I have learned a little about who you are and what your interests are. Just the other day, Gary said that you now are on Internet. I hope you will subscribe to CSGnet.

When you sent \$13 for our video tape, I had no idea, except that Gary had given you our address. You still owe us an evaluation of our video and written material! Please!!!

Recently, I have sent a (further refined edit of same) video and extended materials to "V.P. new technology" types in high tech companies. Will soon find out if that is a viable approach to spread the word about a new engineering science.

We are studying how to simplify the initial introduction to help people understand the advantages of this thinking faster.

Friday, Gary posted to CSGnet a commentary from yourself, which included the following:

>TQM is the new order of things in the Army [tell Dag to hit the
 >Military bases for traing opportunities]. We have a new management
 >regulation AR 5-1 that is only about one page long and actually
 >says to "do the right things, the right way, for the right
 >reasons, and to constantly strive for improvement." Now that is
 >change - hell that's revolutionary!

Questions:

Can you send copy of AR 5-1?
 How does one approach military bases? Whom to approach? How is
 such training funded? What budgets? What rates of pay? Who
 gets trained? At the top? Near the bottom? Top to bottom!!!!???

We will enjoy hearing from you!

Dag & Christine Forssell
 23903 Via Flamenco
 Valencia, Ca 91355-2808
 Phone (805) 254-1195 Fax (805) 254-7956
 Internet: 0004742580@MCIMAIL.COM
 MCI mail: Forssell, Dag or 474-2580

Date: Tue Jun 01, 1993 8:51 am PST
 Subject: Twain, truth and tempting titles

[from Joel Judd 930601]

Greg (030527) **DAG-THERE'S SOME STAT STUFF LATER ON-KEEP READING**

Re: Mark Twain's comment. I find his later writings generally insightful, but ultimately depressing (such as Letters From Earth). They're depressing because they come from someone who saw so much of life and earlier on wrote of its possibilities. I found your posting of his comments and Bill's post of the same day interesting in that they touch on the same theme. It seems that when people get up in years

(I won't put a number here at the risk of offending!) and reflect on life they either comment on how horrible human beings can be to each other and how glad they are life's ending, or how in spite of all the horrible things they have witnessed, life was a good experience and now they're ready for someone else. It seems to me that Twain falls into the former category. It also seems that those who debunk religion seem to end up in the former as well. But enough gross generalizations.

MT>When we believe in immortality we have a reason for it. Not a reason
 MT>founded upon information, or even plausibilities...

Sure we do. We have as much info about its existence as about its non-existence, especially since arguments against often take a default stand: there's no conclusive evidence FOR, so it must not exist. But I find in the Twain quote what I think is a common thread running through this type of viewpoint:

MT>(talking about "annihilation" or "pre-existence")There was a peace,
 MT>a serenity, AN ABSENCE OF RESPONSIBILITY (emphasis mine).

This sentiment is completely up-to-date. Ironically, the rest of the quote is also fairly descriptive of some characterizations of a pre-existence and reasons for having an earthly existence. Why do people want to abdicate responsibility? Why don't we always want to DO things, and take joy in the doing? Is it because we sometimes make MISTAKES, and hurt people (and ourselves), intentionally and unintentionally? For me, the purpose of life involves accepting the fact that we ARE, and becoming a better "ARE" through finding out how others became who they WERE (history), and actively helping those around us become the best they can be, too. This means taking in responsibility and, of necessity, making mistakes. It also means that if we have decided that this life is it, we may become too scared to do anything for fear of making a mistake, or preach a grand relativism for the same reason...(cont.)

Bill (930527)

Your post was a great insight on some of the feelings you dealt with in developing CT early on. (And perhaps why Glasser and others like him are where they are today). The dangers you mention are why I don't mind science and religion being (usually) separated.

I guess my concern for developing some sort of morality based on PCT is for the reason I have mentioned before. A model may be apolitical, amoral, etc., but as soon as it is sent off to behave along with others, then ethics and morality enter in. This is why behaviorism is so nasty from my point of view, not because there's a faulty underlying view of behavior, but because when applied to things like schools there are behaviorists who say think that they can inform classroom dynamics without getting involved in the "politics" of education. In fact, I believe that the behavioral theory one uses in understanding a situation like education can't help but get involved in the "politics." Therefore, I would prefer to have as accurate a theory as possible, and then be responsible enough to say something like "Even though I recognize the primacy of human autonomy as explained by theory X, I also recognize the need for humans to be exposed to certain principles, concepts, etc. in order to have a BASIS for optimal functioning as explained by theory X." In other words, left to itself, I don't think a human being will "be all it can be" in the span of life allotted. I don't think we should be reticent about making moral or ethical decisions (actually we make them all the time) with respect to other people, and particularly children. The question is HOW we make such decisions and what such decisions are based on.

Dag (930525?)

Re: I wasn't the one who posted something about busing, but I did mention a book a while back I'll post some excerpts from. Since the book is about twenty years old, the example are not very topical, but they're located around the page #s I'll give in case you look at the book:

Hopkins, Harry. 1973. The numbers game: The bland totalitarianism.
 Boston: Little, Brown and Co.

Chapter 5 enumerates problems with polling of the type Runkel refers to in Casting Nets.... For example, during the Edward Kennedy affair, the polls were busy asking people about different aspects of the events, trial, etc. Hopkins comments:

"We do not know what degree of attention or scruple each of the 1600 representative Americans pinned down by the Harris interviewers-- according to the sampling prescriptions--brought to the lengthy and exacting questionnaire that was put before them. Certainly, in the midst of their various preoccupations, it must have varied widely. Nor have we any means of knowing the nuances of interpretation each brought to the questions, nor the range of meaning implied in the answers" (p.103). With respect to the miniscule percentage one always finds in a poll with "no opinion":

(from a cited study)"Although tests of information invariably show at least 20% of the public totally uninformed (and usually the figure is much closer to 40%), the "no opinion" vote on any poll question seldom exceeds 15% and is often much lower..." To admit ignorance must go against the grain for most people" (p.108). There's a controlled variable to study empirically!!

Chapter 6 discusses the trend away from subjective "observation" in modern human affairs. A good quote from the heading from H.J. Eysenck, The Scientific Study of Personality:

"To the scientist, the unique individual is simply the point of interaction of a number of quantitative variables."

From the beginning of the chapter:

"In his introduction to Zeisel's well-known primer, Say It With Figures, a leading American sociometrician, Professor Paul Lazarsfeld, writes:

'Modern social life has become much too complicated to be perceived [sic!] by direct observation.' Of course some of us in occasional moments of what used to be called insight may still now and again feel that we perceive a thing or two, but as educated persons we know that, however authentic they may appear, such glimpses are merely impressionistic, tainted with the subjective, and even, very possibly, riddled with our value judgments... we realise that to penetrate social reality today requires a wide range of 'indicators,' mainly numerical, a sort of statistical radar screen monitoring our lives, installed, operated, and interpreted by the task forces of social science" (p. 122).

Regarding the preoccupation with measurement, Hopkins describes two main side effects:

"One class of side effects arises because what in the end is measured, as a result of the vast ingenuity and effort poured into the task, is often not quite what it is supposed to be, but something a bit different. A second source of confusion, compounding the first, is that factors which persist in defying the ingenuity of would-be measurers, which resist both the calipers and correlation, are quite likely to 'evaporate' and be effectively denied existence" (p. 125).

Problems with matching people and numbers:

"As already suggested, however skilfully questions seeking to define human attitudes are designed, the variables are so many and so volatile and the possibilities of misunderstanding so rich that the quantifying investigator is visibly impaled on the dilemma that words are hopelessly vague, and figures--into which they are somehow to be transposed--are impossibly precise. The choice seems to lie between a confession of failure and a confidence trick, both glossed over with the invaluable phrase: 'further research is needed'" (p. 139-40).

Chapter 8 on objectively viewing society:

"From this 'natural science' stance, looking down on it through the statistical lens, society inevitably appears as 'the given,' ideas appear 'caused' rather than causing; ideology is not 'ended,' it merely disappears from view. What WILL be seen will be some sort of functioning system--which statistical analysis and the multiple correlation mill can be trusted to clarify" (p. 155).

Chapter 9--influences of statistical manipulation:

"So statistical rationalism terminates in the policy of obscurantism. In contemporary Western society the economists and econometricians resemble the high priests of ancient Egypt, possessing the papyrus on which were inscribed the secrets of the rise and fall of the Nile. Just as the fellahin awaited the life-giving Nile Flood, so it is our lot to wait for 'the economy' to deflate, reflate, inflate, take-off, slide...Although in our case the economist-priests themselves do not seem immune from a certain bemusement before the semi-mystical automated processes they are believed to control. 'There may be a certain orbit speed that you have to get in order to make recovery self-sustaining,' mused Arthur Okun (Time, June 1971).

In such circumstances, the public--whose cooperation and initiatives are surely not expendable--is naturally reduced to the role of voyeur or statistic-gazer--or to that of

the established hypochondriac who has abandoned himself, body and soul, into the hands of the doctors" (p.181).

If people are interested, I can also mention Hopkins' treatment of IQ and other psychological variables at the end of the book.

Basically, one thought that came to mind while reading was that part of the effect of bemoaning the "complexity" of human intelligence, motivation, etc. is to subtly legitimize the quantification of such factors as a means of "controlling" or "dealing" with them.

Over-reliance on statistics contributes to difficulties in making decisions that affect real people (e.g., value judgments) because numbers don't reflect useful concepts of feeling, intuition, and other "subjective" variables.

Regards, Joel

Date: Tue Jun 01, 1993 9:18 am PST
Subject: Glasser; Disturbances about disturbances

FROM CHUCK TUCKER 930601

A NOTE ON EDUCATION POST

I was in error. The "education post" (as I call it) was not last Summer but in 1991. I have the date for the post as 29 July 1991 (which oddly enough is Stewart's birthday). I do not have a copy of the original post just my WS version of it.

GLASSER

I knew Glasser's work on "reality therapy" (words I don't use) before I knew of CT, PCT or HPCT so my angle on his work is perhaps different than Ed, Rick, Dag or Bill. I found his work useful along with that of Thomas Szasz, Erving Goffman, Thomas Scheff and, of course, everyone's favorite, George H. Mead when presenting my views on "mental health." His focus on responsibility, his dismissal of conventional psychiatric views and his attention to "behavior" in contrast to "attitude" were all quite useful to me (and still are). I have very little difficulty when discussing a work to tell myself and others (usually my students) to ignore certain statements or ideas of an author (has anyone yet picked up on my statement that I don't have my undergraduate students read Power's BCP because I have found that they require some "background" to begin to "understand" it [this is based on "testing" that I do]). I have found that some of what Glasser says about "therapy" is quite similar to what I read in Ed's books (not exactly the same but similar). Eventhough Glasser makes some serious mistakes as he writes about CT and he does not give adequate recognition to Bill for what he has taken from him if I stop reading everyone I know who makes these mistakes I would have to stop reading my own writing (which would not be allow me much more time in the day) AND almost everyone I find useful. Perhaps the case can be made that Glasser has stolen the intellectual property of Powers, I just don't know if that is so (it is serious if it is so) but that Glasser makes mistakes in presenting CT can be forgiven and overcome. Peckham has noted that the disciples never get the Master correct; in fact, it is rare that we can repeat exactly what we ourselves have said and written w/o just reproducing a statement from a post as we do on the net. Of course, if you think that Glasser's mistakes are so serious and/or you believe that reading him will present a distorted view of PCT, then just tell people (as I am sure some of you have) not to read his works.

DISTURBANCES ABOUT DISTURBANCES

I have found one set of statements that might be confusing to a person trying to understand PCT. Both statements come from the discussion of the L & L work. Here they are:

Rick (930525.1400) Reply to Satan's messenger

The "natural state" of the organism is seen to be one of controlling its input. The actions of the organism are CONTINUOUSLY countering disturbances to controlled variables. So the organism is ALWAYS precisely as active as are the disturbances to its controlled variables. (MY EMPHASIS)

The outputs of a control system (including those that become the goals of lower level systems) are CAUSED by disturbances to controlled variable. Output are NOT caused by goals, they are CAUSED by disturbances. (MY EMPHASIS ON THE WORD 'CAUSED')

And this is exactly what a control system DOES -- "in reality". It gets a goal specification (the reference input) and acts to make the perception match the reference (goal attainment) and, in the process, eliminates, NOT the disturbances themselves, but the effect of the disturbances on the perceptual variable. (MY EMPHASIS)

Bill (930526.0830 MDT) Response to Locke and Latham

Control theory shows that organisms are organized to control their own inputs, NOT their outputs (their behavior). Behaviorism has always maintained that the environment is ultimately in control of behavior. Control theory shows that behavior is PRODUCED by organisms to satisfy their own internal requirements, and the influences from the environment are REGULARLY REJECTED and OPPOSED IF THEY INTERFERE with meeting these internal requirements. (MY EMPHASIS)

I see a matter of focus when I compare the statements from Rick with the statements of Bill that could lead to some misunderstanding of PCT. Rick seems to place special emphasis (it almost appears to be causal in an S-R sense [sorry, Rick]) on disturbances while Bill's statement does not place an emphasis on any part of the LOOP but rather, notes explicitly, disturbances MAY be involved "if they (I would say "only if") interfere with meeting these internal requirements." This difficulty (if you think it is a difficulty and not just a nit-pick of mine) is one we all have when we write about PCT - we have to write in circular rather than linear statements with a language (at least the one I have limited knowledge of) which is basically linear. It is one of the reasons I believe that people have difficulty reading Powers (and Dewey, James and Mead): its expressions are circular as constructed by the writer(s) but appear to most readers (using linear transformations for perceptual signals) to be quite confusing ("Why doesn't he just say what he means and drop all that fancy language?"). This is also one of the problems with trying to state these ideas in "ordinary language" (James tried to "popularize pragmatism" and Pierce rejected it; Mead never did it and is very difficult for most to understand); it may be difficult to distinguish a "disturbance" from a "stimulus." If those notions can not be clearly distinguished then people will say (with justification) there is no difference that makes a difference between S-R and PCT. I don't believe this is the case but we must be careful in our statements so we do not leave others with that perception.

Regards, Chuck

Date: Tue Jun 01, 1993 1:18 pm PST
Subject: Looking under the hood

[From Rick Marken (930601.1300)]

So many posts, so little time. So I have to choose. And I choose CHUCK TUCKER's post (930601) because it provides a nice opportunity to show why it's the PCT model that's important and not so much the words used to describe it (though, of course, we have to try to choose words that evoke the right "meanings" about how the model actually works).

Chuck quotes the following from me (930525.1400)

> The outputs of a control system (including those that become the goals of >lower level systems) are CAUSED by disturbances to controlled variable. >Output are NOT caused by goals, they are CAUSED by disturbances. (MY >EMPHASIS ON THE WORD 'CAUSED')

and Bill (930526.0830 MDT)

> Control theory shows that behavior is PRODUCED by organisms to satisfy their >own internal requirements, and the influences from the environ- ment are >REGULARLY REJECTED and OPPOSED IF THEY INTERFERE with meeting these >internal requirements. (MY EMPHASIS)

Chuck says:

>Rick seems to place special emphasis (it almost appears
>to be causal in an S-R sense [sorry, Rick]) on disturbances

No apologies necessary. I was trying to say in words what is succinctly described in the "second law of PCT":

$$o = g^{-1}(r - d)$$

(the "first law" being $p = r$).

My comment above assumed that the goal (r) is a constant; if it's not, then goal variations contribute as much to output variations as do disturbance variations. However, when the goal is fixed (r is constant) then output variations are caused by disturbance variations exclusively, as I said. There is, indeed, a causal relationship between o and d but, of course, the causality does not run through the organism; rather, it runs through the controlled variable and depends on the feedback function connecting the output to that variable. The fact that the causal relationship between o and d runs through the environment and not the organism is captured by the fact that the function that relates o to d is the inverse of the feedback function (g^{-1}).

Bill said exactly the same thing using words that emphasized different aspects of the "second law". For example, saying that environmental effects are rejected or opposed by outputs simply refers to the negative sign in front of the d ; the negative relationship between o and d could certainly be described as rejection or opposition of d by o .

> Bill's statement does not place an emphasis on any part of the LOOP

Bill's statement (like mine) does focus on one point in the loop; the controlled variable (the knot, for rubber band freaks). We are describing the relationship between two variables that affect the controlled variable -- o and d (the two rubber bands).

> but rather, notes explicitly, disturbances MAY be involved "if they
> (I would say "only if") interfere with meeting these internal
> requirements."

Disturbances are ALWAYS involved. Bill said that disturbances will be opposed IF they interfere -- which is equivalent to saying that there will be no output unless $d \neq 0$; if there are no disturbances ($d=0$) then is no output necessary. This fact is explicit in the second law and implicit in my statement that disturbances cause output; when the cause (d) is zero, the output (o) is zero (or, at least, a constant).

> This difficulty (if you think it is a difficulty
> and not just a nit-pick of mine) is one we all have when we write about PCT

Yes. If your understanding of the PCT model is based only on the sentences that are used to describe it then there will always be arguments about who is describing PCT "correctly". I think there are better ways of describing PCT than others, but ultimately the evaluation of what constitutes a "better" explanation has to be based on an understanding of how the model actually works; the behavior of the working model is the best "description" of itself. It is also very important to be able to understand at least some of the basic math describing the control model. Math is a clear and conventionalized way of communicating which, although not perfect, is often a big improvement over sentences. I love language, especially when it is used well. But in PCT we should try not to lose sight of the fact that the language is being used to describe an actual, working model. This model is always available for inspection by anyone who wants to "check" their words. The model is available in the Demo and Little Man programs and in the spreadsheet hierarchy, to name a few relatively "user friendly" examples.

> we have to write in circular rather than linear statements with a language

Neither Bill nor I were using "circular" language to describe the behavior of a control system. The two laws of PCT can be described quite "linearly" -- $p = r$ is equivalent to saying "perceptual signal variations depend on reference signal variations" and $o = -d$ is equivalent to saying "output variations depend (inversely) on disturbance variations". Those statements sound pretty linear to me and they are correct descriptions of the relationship between some of the variables in a control loop.

> This is also one of the problems with trying to state these ideas
> in "ordinary language"

I think that the problem with trying to state PCT in ordinary language is the same as the problem of trying to describe ANYTHING in ANY language; the problem occurs when the one

doing the describing has little or no familiarity with the actual thing being described. If one's only experience with something is based only on other people's descriptions of it, then how well can one describe it? Suppose that someone knew about car engines only from what people told him about them (plus a few graphs and some demos of their behavior). How well could he then describe the operation of an engine to another person? Well enough so that someone who knows nearly nothing about engines could build one based on that description? I think not. A good description of the engine is important but it is even more important to take a good, long look under the hood -- to see what's actually going on under there.

Best Rick

Date: Tue Jun 01, 1993 3:07 pm PST
Subject: On Vacation 'till June 14

[from Gary Cziko 930601.2235 UTC]

Just a quick note to let people know that the CSGnet listowner will be leaving soon for Costa Rica to check out the volcanoes, monkeys, and beaches with my son. I will be back to sift through my hundreds of waiting e-mail messages on June 14.

--Gary

Date: Tue Jun 01, 1993 3:38 pm PST
Subject: spreadsheet modelling in Excel

[Avery Andrews 930602.0919 Eastern Oz Time]

I was fooling around with Excel, seeing if I could do a simple control system model, & found that it didn't seem to like the circular references. Does anyone know a way around this?

Avery

Date: Tue Jun 01, 1993 3:58 pm PST
Subject: Statistics

[From Dag Forssell (930601 16.00) Joel Judd 930601

>Re: I wasn't the one who posted something about busing, but I did
>mention a book a while back I'll post some excerpts from. Since
>the book is about twenty years old, the examples are not very
>topical, but they're located around the page #s I'll give in case
>you look at the book:
>
>Hopkins, Harry. 1973. The numbers game: The bland
> totalitarianism. Boston: Little, Brown and Co.

Thanks, Joel. I'll look it up and see if it fits as background information for me.

The one who did post something about busing was Kent McClelland. It came to me, so I sent him a query direct. He replied direct:

>This sounds like the so-called "Coleman Report" (James S. Coleman
>et al., Equality of Educational Opportunity, Washington, DC:
>USGPO, 1966), an exhaustive study of the achievement of black
>students done by a group of famous sociologists under government
>contract at the height of the Civil Rights movement. They failed
>to isolate any school-related variables that affected
>achievement other than presence of white students in the
>classroom, and the study did have a major influence on the
>development of the "forced busing" policy in the next few years.
>
>The study has been subjected to considerable criticism, starting
>with a volume edited by Frederick Mosteller and Daniel P.
>Moynihan (later Sen. Moynihan) called, On Equality of Educational

>Opportunity (NY: Vintage, 1972).

Thanks to both of you, Dag

Date: Tue Jun 01, 1993 4:18 pm PST
Subject: Responsibility, Clemens-style

From Greg Williams (930601) Joel Judd 930601

I appreciate your comments on the Mark Twain excerpt. If you have time, I'd still like to hear your answers to my questions about "truth" (earlier in the post which included the Clemens quote). Thanks.

>MT>When we believe in immortality we have a reason for it. Not a reason
>MT>founded upon information, or even plausibilities...

>Sure we do. We have as much info about its existence as about its
>non-existence, especially since arguments against often take a default
>stand: there's no conclusive evidence FOR, so it must not exist.

Would you agree that it IS conclusive that EITHER ALL OR ALL BUT ONE OF the theories of immortality (including the theory that it doesn't exist) MUST be wrong? Doesn't this follow logically from the facts that each of the theories are claimed by their believers to be incompatible with the others? Don't you also agree that many of those believers would claim to have more evidence EITHER for the existence OR for the non-existence of immortality? They certainly seem to ACT as if they had more evidence one way than the other! Do you think that they shouldn't act that way?

>This means taking in responsibility and, of necessity, making mistakes.

Taking responsibility is a fine thing, I think, if one is careful to avoid the sort of enthusiasm which can all too easily become fanaticism. This is the kind of "taking responsibility" which shouts at others that THEY are mistaken. One of Clemens' endearing traits, for me, was his apparently absolute inability to take ANY notion very seriously. I suppose that Rick Marken might call that the ULTIMATE in "taking responsibility"!

>It also means that if we have decided that this life is it, we may
>become too scared to do anything for fear of making a mistake, or preach
>a grand relativism for the same reason...(cont.)

I agree that fear of making a mistake can be tragically crippling, but I disagree that relativism in practice (deeds, not just words!) is necessarily (or even often) anywhere near as terrible in its consequences as has been absolutism. Of course, that's just my opinion.

As ever, Greg

Date: Tue Jun 01, 1993 4:54 pm PST
Subject: Re: spreadsheet modelling in Excel

[From Rick Marken (930601.1700)]

Avery Andrews (930602.0919 Eastern Oz Time)

>I was fooling around with Excel, seeing if I could do a simple control
>system model, & found that it didn't seem to like the circular references.
>Does anyone know a way around this?

You betcha. Click the "Iteration" box in the "Options" menu (under "Recalculate"). When "Iteration" is on, circular references are OK. I don't know where the "Iteration" setting is on the PC (which I presume is what you are using) but that's what you want.

Also, it is best to set the "Recalculation" to "Manual" so that it doesn't go off and try to recursively iterate every time you change an equation. I can send you a copy of my Excel hierarchy if you like. I sent one to Dag once and it ran OK on his PC (from my MAC). I forget how I sent it; I think it was as a SYLK file. Do you remember Dag??

Happy modelling. Rick

Date: Tue Jun 01, 1993 5:19 pm PST
Subject: Responsibility, Clemens-style

[From Rick Marken (930601.1800)] Greg Williams (930601)--

>Taking responsibility is a fine thing, I think, if one is careful to
>avoid the sort of enthusiasm which can all too easily become fanaticism.
>This is the kind of "taking responsibility" which shouts at others that
>THEY are mistaken. One of Clemens' endearing traits, for me, was his
>apparently absolute inability to take ANY notion very seriously. I
>suppose that Rick Marken might call that the ULTIMATE in "taking
>responsibility"!

Here, here!!

>I agree that fear of making a mistake can be tragically crippling, but I
>disagree that relativism in practice (deeds, not just words!) is
>necessarily (or even often) anywhere near as terrible in its consequences
>as has been absolutism. Of course, that's just my opinion.

Well said!! I don't know if anyone's noticed but PCT shows that relativism is not only a good way to go -- it's the ONLY way to go. You can watch relativism in action is you get a copy of my spreadsheet hierarchy -- an incredible matrix of sin and degradation. Maybe we can get it placed on the Index Librorum Prohibitorum; it worked for Galileo!

Best Rick

Date: Wed Jun 02, 1993 6:27 am PST
From: Major Robert E Buhts
EMS: INTERNET / MCI ID: 376-5414
MBX: h2de9reb@orl41.orl.usace.army.mil
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: PCT, AR 5-1

Dag, a copy of AR 5-1 is in the mail. Each military base has a training office usually associated with the headquarters. Fort Knox is currently involved in a big effort to train their civilian workforce in TQM. Army wide, the Commander of Training and Doctrine Command has proponency for all training. The address and POC is >> Commander, HQ TRADOC, ATTN: ATTG-CD (Mrs. Carberry), Fort Monroe, VA 23651-5000. This office was very interested in spreading TQM training throughout the Army and can probably give you more specific answers. The Office of Personnel Management (OPM) has negotiated contracts with most large consulting firms for TQM services. These set the maximum rates to be charged for ranges of TQM training and implementation services. The contract number is OPM-89-2884 (this may have been superceded). A POC (old) in DC is Vivian Bethea, Tel. 202-606-2240. I've included an old rate sheet from a large vendor along w/ the AR. Good Luck.

Bob Buhts

Date: Wed Jun 02, 1993 6:34 am PST
Subject: Social control systems; morals; explaining PCT

[From Bill Powers (930601.0600 MDT)] Bob Clark (930521.2145 EDT) --

>... people, at least adults, are familiar with many versions of
>hierarchical social organizations. These organizations consist
>of the structures of inputs, outputs and objectives required to
>form an effective system. And then HCPT comes along and
>proposes another application of the hierarchical concept!

I was trying to point out in my post on hierarchical social organizations that HPCT has some new things to say about social hierarchies. One factor that I left out was the reliance on direct physical force to make the social hierarchical system work. Inside an

individual, this is simply not a factor when the system is functioning normally; the nearest thing to one internal system forcing another to act is in a conflict situation, where two systems demand incompatible actions from lower systems. This normally leads to reorganization that eliminates the conflict.

In all social hierarchies of which I know, the use or credible threat of physical force is the main means of maintaining the structure. This is a fundamental design defect; it creates conflict automatically. It is very hard to think of any current social organization that does not rely on the use or threat of force -- from the legbreakers of organized crime to the threats of hellfire and damnation of certain organized religions. The "rule of law" means nothing without police packing handcuffs and guns.

 Thanks for the historical notes.

>You seem to be pointing out many of the reasons that social
 >systems cannot exist -- or at least cannot exist for long. And
 >you attack the "traditional command structure" as the cause of
 >more problems than it solves. Of course this happens -- but
 >the autocratic manager tends to get into difficulty one way or
 >another. Effective managers do not behave in that manner -- and never have!

You're invoking your own criterion of effectiveness, not the generally accepted one. I think that J. P. Morgan was considered an effective manager in his time, and is still so considered by some. If you have sufficient force behind you it doesn't matter how autocratic you are. Even the most effective manager, as Dag Forssell has pointed out, may find nothing wrong with reminding an employee that management giveth and taketh away jobs.

I agree that managers who rely on coercion tend to get into difficulties. Such difficulties, however, are rampant.

>Social control systems do exist. Some don't work very well, in
 >the sense that they don't "live" long. They routinely develop
 >conflicts of assorted kinds, both internal and external. They
 >are formed, develop, and disappear. I have personally
 >participated in the formation, development and operation of
 >several organizations of several types. I have written, and
 >helped to write, sets of by-laws etc.

Speaking literally, social CONTROL systems don't actually exist, and never have. All that actually exists are individuals trying to make their own worlds conform to their own preferences. What seem like social control systems are really people trying to maintain their own conceptions of a social structure. The clue that tells you that these are not really hierarchies of control like those in an individual is the reliance, in the final analysis, on coercion. By-laws mean nothing unless there is a penalty for failure to live by them, and penalties mean nothing unless an individual can be forced to pay them even unwillingly.

A social system is not necessarily a social control system. A social system is an agreed set of goals which each individual adopts because they further that individual's concept of a world worth living in. A social CONTROL system, on the other hand, necessarily employs physical force to make all members conform to a set of goals that only some members consider desirable. Physical force is necessary because there is no other way to control the behavior of another control system that refuses to do you would prefer. And the control that results is not social: it is not the emergent outcome of everyone's own control processes. It is control of some people by others, against their will. Social control systems boil down to individual control on the part of those who happen to agree on goals, and have the physical means to see that they are carried out by everyone.

Perhaps the new management concepts that are developing today are a sign that people are beginning to recognize the drawbacks of simple majority rule. I think this is equivalent to seeing the difference between a social system and a social control system.

 Joel Judd (930601) -- RE: immortality.

Considering the course that my life has taken, I would be hard for me to imagine wanting it to continue forever along the same lines. One person's life is a finite package, and there comes a time when one's own works are best handed on to others who are earlier in life's progression, so they can bring fresh ideas into them and correct one's mistakes. And I'm sure that they would look forward more to their own futures knowing that I will not hang

around forever being myself and ever more so, without respite. One's mentors should have the grace to die properly.

If my consciousness were to survive beyond death, I fail to see how it could contain much that I value now. It will certainly not participate in csg-1. It will not live in a nice house in Durango with a nice wife. It will not look forward to CSG meetings every year, and all the great people who attend. It won't sit around writing programs and getting big thrills when they work right.

And if all that is to end, who will it be that survives? Not me. Only some abstract consciousness without content. That has little interest for me now, and I can't imagine it suddenly becoming attractive.

I hope that when I die my life will have made a picture that is esthetically satisfying, however crude some of its stages along the way may have been. To think that all of it has only been a preparation for something else would be to devalue it.

On morality:

What is morality but a set of principles that hang together to form a satisfactory system concept? One of my more liberating experiences was taking a course in ethics from Paul Arthur Schilpp. This was the first time I realized that ethics could make sense -- that one could find reasons for moral and ethical behavior, reasons that go beyond the mere fear of punishment or hope for a reward (the only reasons that my religion had ever offered me). I was so taken by this new insight that I went to Schilpp and told him I had decided to become a philosopher. Schilpp looked at me for a couple of seconds, and then said "Wait until you're 40 years old, and the gonads have stopped sizzling."

PCT itself doesn't have any obvious moral system in it. By implication, however, it suggests that a properly-functioning human organism will have goals that are mutually consistent at all levels, so that internal conflict doesn't waste the system's capabilities through mutually-cancelling efforts. This means ALL goals, including those for how a society should be structured and how you treat other people, and they you. It means goals for right now, and goals that encompass imagined and future worlds. It includes goals that apply to yourself, and goals that apply to your place among other people. I think that when you try to put together an internal structure that is free of conflict under ALL circumstances, you end up with what most people would recognize as a moral system.

So basically I agree with you when you say

"Even though I recognize the primacy of human autonomy as explained by theory X, I also recognize the need for humans to be exposed to certain principles, concepts, etc. in order to have a BASIS for optimal functioning as explained by theory X."

That statement in itself is an example of goals concerning principles and system concepts. I find nothing inconsistent in proposing to others certain ideas about morality for their consideration. The same principles require giving consideration to their suggestions. Ultimately each person has to decide what makes sense.

If we simply extracted from all religions the moral principles and system concepts for human existence that they propose, and left out the stories used to make them seem more convincing, we would have a powerful set of concepts from which to choose. That's essentially what I've tried to do since my encounter with Schilpp and his course in ethics. The ideas are none the less powerful for being considered as suggestions by other human beings; they are far more effective in one's life for being adopted through one's own free choice and as an expression of one's own autonomy.

Chuck Tucker (930601) --

>This difficulty (if you think it is a difficulty and not
>just a nit-pick of mine) is one we all have when we write
>about PCT - we have to write in circular rather than linear
>statements with a language (at least the one I have limited
>knowledge of) which is basically linear. It is one of the
>reasons I believe that people have difficulty reading
>Powers (and Dewey, James and Mead): its expressions are
>circular as constructed by the writer(s) but appear to most
>readers (using linear transformations for perceptual
>signals) to be quite confusing ("Why doesn't he just say

>what he means and drop all that fancy language?").

It seems to me that the essence of the message of PCT -- the sine qua non -- is the circular way of thinking. Whenever anyone tries to make the message easier to understand by using a linear description (hard not to do that), the wrong message gets across.

Any idea that's worth learning is hard to understand at first (although the converse is not true). Most of the misrepresentations of PCT that we find in the literature arise because their authors try to make PCT easy to understand, which is only a short step from saying that PCT is just a technological way of saying something they've known all along. Which is hardly ever true.

Rick Marken (930601.1300) --

I think we have to try to deal with the causes of output in a more general manner:

>My comment above assumed that the goal (r) is a constant; if
>it's not, then goal variations contribute as much to output
>variations as do disturbance variations.

It gets somewhat awkward to assert that disturbances cause output, and then to have to backtrack and explain that this is true only when the organism itself isn't causing output, too.

It's probably easier to start right out by saying that actions result from the combined effects of the organism's changing goals and the changing environmental situation. If either one is fixed, the other explains variations in action.

It's also important, I think, to explain that achieving some FIXED goals requires CONTINUOUSLY CHANGING action -- like the goal of knitting a sock.

>I think that the problem with trying to state PCT in ordinary
>language is the same as the problem of trying to describe
>ANYTHING in ANY language; the problem occurs when the one doing
>the describing has little or no familiarity with the actual
>thing being described.

Astute comment. This is why demos are so important, if you can get people to sit still long enough to have the new phenomena pointed out in detail. How many times have we seen examples of control processes in which a pure stimulus-response phenomenon was described? The basic control relationships are so different from what behavioral scientists are used to that they would NEVER think of them spontaneously.

Best to all, Bill P.

Date: Wed Jun 02, 1993 8:02 am PST
Subject: Responsibility & Relativism

[from Joel Judd] Greg (930601)

You wanted to hear my answers to truth questions earlier in your post (930527):

Yes, I believe human knowledge is tentative. I do however like the distinction between the verbs KNOW and BELIEVE, where the former for me involves personal experience with the thing to be known and the latter more of a desire to know something.

>...you think that "truth" is a pragmatic measure of the usefulness
>of a particular belief for a person or persons

I would say that one measure of truth is its pragmatic usefulness.

>...then if a belief is considered by person A as being useless...
>then the belief can legitimately be called "false" by person A...

Knowing the nature of language, we can call anything whatever we want to. I prefer to believe that there are some principles which are "true" and exist in the same way as scientific "laws" do--"out there"-- and which we may "discover" if we wish, in the popperian sense of discovering gravitational principles, etc. A number of people on the net

are comfortable describing CT principles in numerical fashion, but most of the time the equations are not useful for me, so I can call them false? I don't think ignorance is a good criteria for determining the truthfulness of something.

>Would you agree that it IS conclusive that EITHER ALL OR ALL BUT ONE
>OF the theories of immortality...MUST be wrong.

Yes, although I don't think EVERY theory is necessarily COMPLETELY wrong. Some can be correct in certain ways. And people act as if they know one way or another, and yes, they should act that way (as long as they respect others' beliefs as well).

>One of Clemens' endearing traits was his apparent inability to take ANY
>notion very seriously.

That's why I feel depressed when I read his later work. I love cynicism as much as the next guy (ask Gary), but I find as time goes on it doesn't do much for the next person.

>...I disagree that relativism...is necessarily anywhere near as terrible
> in its consequences as has been absolutism.

Maybe not necessarily so, but I think that the kind of relativism described by Allan Bloom in The Closing of the American Mind is deadly, even if one doesn't agree with his solutions. Perhaps as bad as absolutism can be, it often polarizes people to act, whereas the apathy relativism can lead to is insidious and lulls people into thinking everything is OK. And so...

Rick (930601)

>I don't know if anyone's noticed but PCT shows that relativism is not
>only a good way to go--it's the ONLY way to go.

Ah, somebody said it. You want to unpack this a little? I don't know if I would want you on the local school board any more than the person who wants to burn Catcher in the Rye. You've mentioned having to deal with your children in the past. HOW do you deal with them, and on what BASIS? Do you think, for example, that a second grader can wade through the curriculum and derive relevant values for his or her life? How far does your relativism go in terms of individual development and societal permissiveness?

Regards, Joel

Date: Wed Jun 02, 1993 10:49 am PST
Subject: Stats again

[from Joel Judd 930602] Some further history twisting through stats:

After posting all the quotes from the book yesterday I went to the library and just happened across the following article:

Fallows, J. 1993. Vietnam: Low-class conclusions. Atlantic. 271(4), 38-44.

Apparently last fall an MIT professor and two graduate students published a new study concluding that the Vietnam War hit the High and Mighty as hard as the poor and working class (the latter being the generally accepted group most affected by the draft).

To do this the researchers took a famous random sample from records of the dead (around 1500 I think). They decided that the determining factor for class would be income. This would be derived solely from the deceased's address and the corresponding census income for that area. Bottom line? A significant finding that all income groups were hit equally hard.

This would sound stupid if it wasn't for the fact that Fallows points out how TIME and other magazines immediately jumped on the results and printed stories about how they had been wrong about the lower-class thesis, etc. obviously without carefully reading the study and examining its methods.

Regards, Joel

Date: Wed Jun 02, 1993 11:51 am PST

Subject: More truth

From Greg Williams (930602) Joel Judd [930602]

>I would say that one measure of truth is its pragmatic usefulness.

Ok. What are some other measures, for you?

>I prefer to believe that there are some principles which are
>"true" and exist in the same way as scientific "laws" do--"out there"--
>and which we may "discover" if we wish, in the popperian sense of
>discovering gravitational principles, etc.

Such as? I'd like to hear some examples of such principles, if you have "discovered" any.

If you want to save time, basically, in asking these questions, I'm attempting to discover whether you think any aspects of (any) revealed religion can ever be known in a non-hypothetical way.

>>Would you agree that it IS conclusive that EITHER ALL OR ALL BUT ONE
>>OF the theories of immortality...MUST be wrong.

>Yes, although I don't think EVERY theory is necessarily COMPLETELY wrong.
>Some can be correct in certain ways. And people act as if they know one
>way or another, and yes, they should act that way (as long as they respect
>others' beliefs as well).

But what if someone's theory includes the notion that the possibility of becoming immortal is fostered by killing infidels who believe "falsely"? What justifies addition of the clause about respecting others' beliefs, especially if it is contradicted by "revelation"?

>>One of Clemens' endearing traits was his apparent inability to take ANY
>>notion very seriously.

>That's why I feel despressed when I read his later work. I love cynicism
>as much as the next guy (ask Gary), but I find as time goes on it doesn't
>do much for the next person.

But Clemens wasn't JUST cynical. He was POSITIVE in his apparent idea that prodding their scared cows is good for (at least some) "next guys."

>Perhaps as bad as absolutism can be, it often polarizes people to act, whereas
>the apathy relativism can lead to is insiduous and lulls people into
>thinking everything is OK. And so...

But I don't see the relativists deeds playing out passively. Like the Marines fighting for the relativistic clauses of the U.S. Constitution.

As ever, Greg

Date: Wed Jun 02, 1993 12:15 pm PST
Subject: Relativism

[From Rick Marken (930602.1100)] Bill Powers (930601.0600 MDT) --

>It gets somewhat awkward to assert that disturbances cause
>output, and then to have to backtrack and explain that this is
>>true only when the organism itself isn't causing output, too.

Yeah. OK. It's just that when I read Locke and Latham they seemed to be talking about a goal as though it were a reference that was already "set". In that case, all of the variance in actions is determined by disturbances -- as in the rubber band demo with a fixed target for the knot.

>It's probably easier to start right out by saying that actions
>result from the combined effects of the organism's changing goals
>and the changing environmental situation. If either one is fixed,
>the other explains variations in action.

Or just say: $o = 1/g(r - d)$

I said:

>>I think that the problem with trying to state PCT in ordinary
>>language is the same as the problem of trying to describe
>>ANYTHING in ANY language; the problem occurs when the one doing
>>the describing has little or no familiarity with the actual
>>thing being described.

Bill replies:

> Astute comment. I agree.

Bill continues:

>This is why demos are so important, if you can get people to sit still
>long enough to have the new phenomena pointed out in detail.

Actually, I was thinking more of the model itself than the phenomenon. I think part of the problem we have communicating PCT to others is that many people know the model only as a set of sentences. These people, it seems to me, are then in a position similar to that of a person trying to communicate about the internal operations of an engine when then know about these operations only from verbal descriptions. These people might know all about what engines can do -- then know the phenomenon of power. But they don't know HOW the engine generates this power. The internal operations of the PCT model are open to inspection -- especially in your ARM model, DEMO program and my spreadsheet hierarchy. I think inspection of the dynamic behavior of these models and a lot of fiddling around to see how the model actually works is prerequisite to being able to talk coherently about PCT.

Joel Judd --

I said:

>I don't know if anyone's noticed but PCT shows that relativism is not
>only a good way to go--it's the ONLY way to go.

Joel asks:

>Ah, somebody said it. You want to unpack this a little?

Control systems AT ALL LEVELS must be able to vary their outputs in order to achieve goals that are themselves set by the variable outputs of higher level systems. The only ABSOLUTE goals in HPCT are the intrinsic references, which specify the survival requirements of the system itself. It would take a lot to convince me that "values" like "thou shalt not steal" are amongst this set of absolute intrinsic references.

>You've mentioned having to deal with your children in the past.
> HOW do you deal with them, and on what BASIS?

My children were born JUST before I learned about PCT. But I always dealt with my children as though they were autonomous purposeful systems -- and, sure enough, it turns out that they are. My main strategy for "dealing" with my children was the one suggested by O. Ivar Lovass -- behaviorist extraordinaire -- in a child development class I took: "lod 'em up wit luv" (he's Norwegian). That turns out to have been about the only thing that ol' Ivar was right about. My hope for my children was (and still is) that they be able to skillfully control their perceptual inputs. They are doing a great job. If I or my wife took responsibility for their wonderfulness we would be stupid, egotistical and wrong. The kids are making themselves into beautifully functioning little autonomous systems; all we did was keep the side effects of their controlling from inconveniencing us. It's really not that hard to do.

>Do you think, for example, that a second grader can wade through
>the curriculum and derive relevant values for his or her life?

I don't think a second grader is after "values"; maybe a 10th grader is and, yes, he or she could derive them just fine; who else could? My kids did it just fine; I did it just fine.

My wife did pretty well too but, then, she had all that help from the Catholic Church. People who think it's important to give kids (or other people) "guidance" or "help" are really just people who want to control other people. Think about it.

>How far does your relativism go in terms of individual development
>and societal permissiveness?

All the way.

Relativism: It's not just a good idea, it's the law.

Best Rick

Date: Thu Jun 03, 1993 5:00 am PST
Subject: short cuts

[from Joel Judd 930603] Greg (930602)

>If you want to save time...[can] any aspects of revealed religion ever
>be known in a non-hypothetical way?

If you mean the existence of God and other beings, or actual physical objects pertaining to a group's beliefs, yes. I have written statements of those, in many cases more than one person, who have seen and touched such things. Such people are, of course, a minority (like me). If you're talking about principles, well, by definition they're going to be hypothetical, right? Isn't the proof in the pudding, so to speak? You have to try them out, and your claim to "knowing" principle A is "true" is based on your experience with it so far in your life, and there may come a day when you decide you can no longer abide by it. To take an example from my church, there were a number of men (eleven, to be exact) to whom Joseph Smith showed the gold plates from which he claimed to be translating. Virtually all of these men became disaffected with Smith or the church or both sometime later, repudiating particular principles of the church. None of them ever denied having seen the plate (though some lived to be more than seventy) or their knowledge of the Book of Mormon.

But most religions require FAITH, and so their principles are going to be hypothetical, and some are going to say such principles are true, and others are going to say they're false. What would it take to make a principle or concept "true" anyway? Certainly not particular behaviors, because there are too many degrees of freedom in something like "Love thy Neighbor." So principles are simply subject to perceptual control like other aspects of our experience.

If I'm not getting at what you're looking for you might want to prompt me. Basically I've admitted that I haven't seen Christ personally and that religious principles are hypothetical, yet that one can have knowledge that a particular set of principles are "true" and will work for everyone, without making everyone believe such a thing by force.

Marken (930602)

Lovass? Did you make this up or is it pronounced differently than the way I'm reading it?

>I always dealt with my children as though they were autonomous
>purposeful systems--and sure enough, it turns out that they are.

We assume that they would be even if you HADN'T, right?

>My hope for my children was that they be able to skillfully control
>their perceptual inputs...they are making themselves into beautifully
>functioning little autonomous systems...

What's "beautiful" to you? Surely one could be skillfully controlling for "selling drugs" to make a few \$\$\$, or intimidating a classmate, or enjoying sexual intercourse, etc. all without conflict. Is that OK with you as long as they are functioning well PCT-wise?

>I don't think a second grader is after "values"

I don't either, but I sure as heck think the experiences during that time are going to influence the values developed later.

You are content with the situation in your family; do you think the majority of people (say in the US) are the same? Do you feel that development is a situation where "most" of the people, given a chance, will "turn out" alright? This was the question I asked a couple of days ago.

>People who think it's important to give kids "guidance" or "help"
>are really just people who want to control other people

Do you think most people would grow up to feel "responsible," "honest," "respectful," "kind," etc. if they are just left to their hierarchical reorganizations? From your comments I don't think that you do.

Regards, Joel

Date: Thu Jun 03, 1993 8:53 am PST
Subject: It ain't easy

[FROM: Dennis Delprato (930602)]

Couldn't resist passing on Oz's commentary. Qz is commenting on an American artform known as bluegrass music which has an e-mail network. I enjoyed the many correspondences between Qz's suggestions about the out-of-the-mainstream bluegrass music that is close to Qz's heart and the subject of CSG-L. As Oz says, he tries to get his friends interested "but it ain't easy." [BTW, I doubt that either BG music or PCT is experiencing zero population growth.]

Subj: New Blood

New Blood

> We have gotten along nicely outside the mainstream for quite a while
> now, and there's no reason to think we can't continue, but we do need
> new *true* fans-- our audience suffers from zero-population-growth.

The Good Mr Godbey speaks the truth! I love Bluegrass, and try to get my friends interested, but it ain't easy. Trad 'grassers won't allow electric basses and some festivals will only book certain types of performers and theres a subtle (sometimes not so subtle) undercurrent of "You don't understand it, so get out of the way."

My feeling is "there's room for us all!"

I applaud the radio folks that do their shows at little to no pay to try and promote the music. I applaude the festival promoters that try to book a mix of acts. I applaude the fan that takes the time to talk to newcomers about the music. I think everyone that expresses an interest at all deserves a big "HELLO! Welcome aboard! Let me show you around" attitude from someone that knows the music.

-Oz oz@pharlap.com

Date: Thu Jun 03, 1993 8:53 am PST
Subject: Raising Kids

[From Rick Marken (930603.0900)] Joel Judd (930603) --

>Lovass? Did you make this up or is it pronounced differently than the
>way I'm reading it?

Sorry. Misprint. It's O. Ivar Lovaas. I took child development from him at UCLA. I loved the class and was very impressed by his efforts to control autistic children. I am one of those rare birds who has become even more liberal as I get older; I still like the early Bob Dylan albums the best.

>What's "beautiful" to you? Surely one could be skillfully controlling
>for "selling drugs" to make a few \$\$\$, or intimidating a classmate, or
>enjoying sexual intercourse, etc. all without conflict. Is that OK with
>you as long as they are functioning well PCT-wise?

I think this is the typical assumption of the religiously minded (and others who think control is important). This the assumption that people will go "haywire" if left to their own devices. Another way of looking at it is that the religiously minded are afraid that kids will end up being happy controlling (with no errors) perceptions that are "wrong". The fact is, it is probably impossible to be a successful hierarchy of control systems (and success includes the ability to successfully interact with other control systems, of course) by controlling for some of the things mentioned above. Selling drugs is fine but the kid is likely to have constant errors resulting from having to deal with the law; intimidating is not nice for the intimidatee but the real problem is that if my kid controls for intimidating others he (or she) would probably end up having to deal with constant errors created by other control systems. Sexual intercourse is fine but then the kid would have to deal with the problems (assuming they cared) of pregnancy, venereal disease and reputation.

I think the religiously minded look at other control system (people) as though they were about to "go off" at any moment in the "wrong direction". That's why people need to have all that help and guidance (control) and the religious (and right wing political) types are all set to give it -- in the form of "values", "wisdom" and "morals". This is the most insidious form of control we've got -- and the real tragedy is that it is the main cause of the very problems that it so righteously claims to be the sole defense against.

>Do you feel that development is a
>situation where "most" of the people, given a chance, will "turn out" alright?

A big, life affirming, Molly Bloom "YES" on that one. (By the way, Molly is the Penelope of James Joyce's "Ulysses"; and "Ulysses" is one of the greatest -- if not the greatest -- novel written in the English language. Of course, efforts were promptly made by those who would control our morals to keep it from being distributed anywhere in the English speaking world.)

>Do you think most people would grow up to feel "responsible,"
>"honest," "respectful," "kind," etc. if they are just left to their
>hierarchical reorganizations?

I do.

>From your comments I don't think that you do.

I think the main cause of the problems you fear are the always well intentioned efforts of people to improve, help, guide -- ie. control -- other people. Your (not you personally -- generic you) efforts to control people are not guaranteed to screw them up -- some people might be thrilled to adopt the goals you want them to adopt -- but they are a step in the right (wrong?) direction.

By the way, you didn't ask me an obvious question; what would I do if my son came home from college and told me that he had become a Mormon? Since it is clear that I'm not a fan of organized (or disorganized) religion, shouldn't I try to guide my son back on the path of secular humanism?

Best Rick

Date: Thu Jun 03, 1993 12:40 pm PST
Subject: More on ECSG meeting

[From Bill Powers (930603.0800)]

This is a preliminary announcement of the first meeting of the European Control Systems Group (ECSG), to be held on 22-26 June, 1994, in Aberystwyth, Wales, UK. It is being organized by Marcos Rodrigues through his University in Aberystwyth. Marcos is writing to a list of potential European attendees. We want to get a more or less firm commitment from as many people as possible very soon, so this first meeting can be planned. At present, the following people have indicated their intention to attend (in no particular order):

Marcos Rodrigues and two associates (Youfu and Lee)
Tom Bourbon
Rick Marken
William and Mary Powers
Wolfgang Zocher

Because of the sponsorship, this meeting is likely to have a more technical flavor than the American CSG meetings have had. However, I hope that psychologists and other non-engineering types will attend and contribute, to help remind the engineers that they live in a real world and to uphold the interdisciplinary tradition of the CSG. I do recommend that those who choose to attend have sufficient acquaintance with PCT to understand what is being talked about in general, if not in technological detail (when the engineers start talking). Anybody who has subscribed to csg-1 for six months or more will have a sufficient grounding to enjoy the meeting.

Marcos will be expanding on this announcement with more details, organizational principles, and whatever else he thinks pertinent. For now, here are excerpts from two communications from Marcos, the first describing the venue and the second giving estimated costs and more details.

 FROM MARCOS:

Gregynog is an enormous house owned by the University to hold conferences and workshops. All facilities are available re. accommodation, meals, meeting rooms (several meeting rooms each one holding about 20-30 people) etc. I think the house offers the ideal setting for our meeting. I can describe it in detail later, but it has a wooden carved room about 400 years old (the four walls are carved in wood, it is really a very impressive work of art), some famous paintings hanging on the walls, etc.

The house is around one hour from Aberystwyth by bus, in the middle of the countryside with no neighbours or noise or disturbance whatsoever. It is the ideal place for a spiritual retirement of a few days... Ah, and because it is so isolated, the house has its own pub which, by tradition, is left at the meeting attendees' disposal (people serve themselves and pay the money in a box).

Aberystwyth is a Welsh name, as you would guess. Aber means the mouth of a river, and Ystwyth is the name of the river. Aber is pronounced as in "aberration" and Ystwyth as it is written. The stress goes on "Yst".

 FROM MARCOS A BIT LATER:

I now have all the details on Gregynog. I'm pleased that you liked the idea because it is really a nice place for our meeting. I'm as sure as I can that you are going to love the place. We've just sent out the booking forms for 22-26 June 1994.

Because our meeting is of an academic nature and it is not sponsored by any Society, they can offer better prices, that is, all their tariffs are VAT exempt (Value Added Tax, actually at 17.5%). All I have to do is to fill in some tax forms and open a special bank account against which all the incomes and expenses must be registered. The Taxman will inspect this account at a later date.

There is still another bonus, which is a special price for students. To be entitled to this further reduction, a student ID card or any other proof of student's status must be produced. The full breakdown for Gregynog is as follows (in pounds sterling):

Wed	Thu	Fri	Sat	Sun
(dinner)	(full board)	(full board)	(full board)	(bed & breakfast)
7.50	37.50	37.50	37.50	24.50

TOTAL 144.50

For students:

Wed	Thu	Fri	Sat	Sun
(dinner)	(full board)	(full board)	(full board)	(bed & breakfast)
6.50	26.75	26.75	26.75	15.50

TOTAL 102.25

The meals include: breakfast, mid-morning coffee and biscuits, lunch, mid-afternoon tea with cakes, and dinner. Children under 15 pay half the cost, under 5 pay 1/3 and babies come for free.

I'm planning to do two things outside the meeting: at the end of Friday afternoon, we could go sightseeing (it gets dark after 10 pm at this time of the year) and have a nice dinner in a nice restaurant. I've got the estimate for that: the coach will cost around 50-60 pounds, and the restaurant dinner 20 or less per head, depending on the number of people. If we do that, the dinner on Friday can be deducted from the Gregynog bill, at 7.50 (6.50 for students) and used to complete the 20 pounds dinner. The other cost is the proceedings; a nice one would cost 10 pounds.

Putting all together, the full cost for the conference would be (considering 4 pounds as the individual share for the coach):

Members of CSG: $144.50 - 7.50 + 4.00 + 20.00 + 10.00 = 171.00$
 Students: $102.25 - 6.50 + 4.00 + 20.00 + 10.00 = 129.75$

It might sound a lot of money, but I assure you the prices are VERY reasonable when compared to European prices. Accommodation and food are very expensive in Europe, as you will find out. Actually, I'm impressed that the above cost covers 4 nights. A few weeks ago I went to Edinburgh and paid 100 pounds just to sleep with breakfast! And that was a 2 or 3 star hotel, not the very best.

I don't expect these prices to vary from now to next year. Inflation in the UK is running very low, below 2% a year. The exchange rate at the moment is holding between 1.50 and 1.60 for several months now. This seems to be the underlying trend. When I first came to the UK 6 years ago, it was 1.50. Last week it was at 1.58 dollars per pound.

Your mentioning of the French funding support for international workshops on behavioural sciences reminded me of the European Commission. I have a very good friend at the Commission in Brussels (an ex-Aberystwyth student) and I'll ask her if there is any possibility of funding there. I know they fund human mobility across Europe for training purposes; they might have something for summer schools or workshops.

 SO -- please get in touch with Marcos very soon if you plan to attend. His e-mail address is

mar@aberystwyth.ac.uk

P.S. Please, also send to Marcos the names of any European contacts you know of who might be interested in the meeting. Marcos, please add to your list

Drs. Franz and Hettie Plooij (Ethology, infant behavior)
 c/o G.P.I. Ijsbaanpad 9 1076 CV Amsterdam, Netherlands

Looks like the meeting is in competent hands, doesn't it?

Best to all, Bill P.

Date: Thu Jun 03, 1993 3:11 pm PST
 Subject: What's good about morality?

[From Bill Powers (930603.1445 MDT)] Joel Judd (930603) --

Rick Marken appears to be arguing that if you just let children develop all by themselves, they would turn out to be successful moral beings, but I don't think that even Rick knows that. There's no way to test the proposition, short of dumping your children into a wilderness and coming back twenty years later to see if they're alive, and if so how they turned out. All parents teach their children about how to be (or how not to be).

However, most of what parents teach children, I am firmly convinced, has nothing to do with the words which which parents inundate their offspring. Children learn how to be by considering what their parents do far more than what they say. If you teach moral principles to children in a patient, kind, and open-handed way, the children may well learn to be patient, kind, and open-handed; if you teach the same moral principles in a stern, demanding, unforgiving way, the children may well learn to be stern, demanding, and unforgiving. That is, assuming that the children don't rebel or otherwise refuse to take

the parent as an example of a good way to be. The moral principles children are most impressed with are the ones that the parents live, not the ones they describe.

 >Surely one could be skillfully controlling for "selling drugs"
 >to make a few \$\$\$, or intimidating a classmate, or enjoying
 >sexual intercourse, etc. all without conflict. Is that OK with
 >you as long as they are functioning well PCT-wise?

Presumably, you're arguing that children need to be taught that selling drugs, intimidation, and promiscuity are wrong; you're suggesting that they might become organized to do such things in an unconflicted way, and so would not violate any principles of PCT even though they are doing something "wrong" in moral terms.

I think the crux of the present argument comes down to the justification one gives for moral principles. If there is something wrong with doing the above things, and many more such as murdering people, breaking your word, being covetous or jealous or envious or greedy, how do you explain to someone just what is wrong with doing these things?

The only answers I have seen coming out of religious teachings are that God has told us that such things are wrong. Such religious teachings seem to assume that if it were not for revelations handed down to man from God, nobody would have any basis for declaring any human behavior to be immoral. If God had not told mankind through Moses that it is a sin to murder, nobody would ever have figured out that murdering is not an acceptable social interaction.

Perhaps it is true that an ordinary person raised in an ordinary family and given an ordinary education does not learn any higher personal reasons for moral behavior. A person raised in such a vacuum might easily find moral principles difficult to understand, especially when they go against what one wishes to do in the here and now. Without any framework within which to understand why some principles work better than others, and to what end, a person might simply give up on the whole question, and adopt whatever is offered simply on the basis of threats of punishment or promises of rewards. You then decide not to murder people because if you do, you will be arrested and thrown in jail, or be executed. And if you happen to get away with it in this life, God will catch up with you when you die and you'll roast forever in horrible torment in Hell. Furthermore, if you do manage to keep all the commandments in practice and in your heart, God will reward you with everlasting peace and joy and your soul will spend eternity in an ecstasy of love.

That's all very well, but it still doesn't explain what is wrong with murdering people or breaking any of the other moral commandments, or what is good about keeping the commandments. To live up to the commandments simply because God told you that you had better if you know what's good for you is to give up on trying to make sense of them. You still don't really know why you shouldn't go around murdering people, stealing other people's spouses, etc.. You can find no reason in your own understanding for adopting any moral principles.

From a practical point of view, it's probably a good thing that people who have no inner basis for morality tend to accept the moral pronouncements of authoritative institutions that deal in such matters. A person who truly can't think of any good reason not to murder whomever he or she pleases ought to be sent immediately into some religious institution, and be convinced that murdering people will bring sure and awful retribution, the more awful for being delayed. Such a person ought to be told in the most convincing possible manner that there is a God who knows their every deed and thought, and who is keeping score for a final reckoning. The other people in the world would be fully justified in doing this simply to protect themselves against a psychopath.

On the other hand, I think there can be a basis for understanding morality. HPCT suggests a way. Moral principles are not simply invented at random; they are generalizations which, if maintained in practice, tend to create a social system in which every person respects the will of others and can expect to be respected in turn. When you think of morality as a means for maintaining a certain kind of system, an idea of a human world of which one is a part, it becomes clear that certain principles simply will not work to achieve such a system while others seem to work quite well.

Consider contracts. When a person makes a contract with another, the expectation has to be that the contract will be honored during its lifetime. Each person must realize that there must be a general acceptance of the principle of honoring contracts, for without it, no individual could rely on anyone else. When one person becomes known for breaking contracts, that person finds it impossible to enter into any new ones, and thus loses all the

advantages of being able to plan for the future and make bargains with others for one's own benefit. A principle that allows breaking a contract at a whim is simply not practical.

Or consider murder. If it is generally accepted that murder is an option open to any individual, then no individual is safe. Without a general agreement against murdering, you can never know whether the next person to be encountered won't take it in mind to win an argument or just express frustration with something private by pulling out a gun and blowing you away. Penalties for murder are always applied too late to help the victim. The only true safeguard against being vulnerable to murder is to shape a society in which everyone understands that this is not a practical method for solving problems, and why it isn't.

People without a clearly formed set of system concepts have no basis for choosing any particular morality. No basis, that is, except a belief that certain prescribed moral principles are enforced in some way beyond understanding, but also beyond escaping. There is nothing in such people to restrain them when a moral principle comes up against a practical immediate problem of comfort, health, or survival. A person with a clear concept of the kind of society he wants to be part of will be aware that violating the principle will violate something that is more far-reaching than the immediate problem. It makes the person into a member of a kind of society that is, above all, not to be encouraged or exemplified. Even a simple system concept of the kind of person one wants to be, without respect to what others want to be, can be a powerful influence on the outcome when principles clash with immediate needs. When one must alter principles, the existence of a clear system concept will make sure that the alteration does not create, even in principle, a way of being that is not viable for all.

System concepts are a higher level of perception and control than any others, including principles. A person who has a clear and consistent set of system concepts must necessarily submit to them: they are the person's own highest goals. To such a person, moral principles are not simply given as revelations from another world. They are the means by which the most important perceptions of all are maintained in a shape that is consistent, pleasing, and beautiful.

While we can't literally teach system concepts, we can describe, illustrate, and demonstrate them. We can rely on the capacity of a human brain, even a young one, to recognize consistency, elegance, beauty, workability. Even a child can see the difference between a playtime in which everyone squabbles over toys and always gets the favorite one snatched away, and another one in which each child can count on a turn with whichever toy is appealing. Parents and teachers who understand system concepts can show them at work, show how maintaining them results in a better life for even the least and weakest of the individuals.

Religious teachings, stripped of their explanatory frameworks, contain a great deal of wisdom and practical experience. Certain aspects of these teachings evoke in us a sense of powerful goodness, of rightness, as we recognize how much better the world would be if the teachings were followed by everyone. In my opinion, explaining these teachings by saying that they must be followed not because they make sense but because God in all his power and majesty commands that we follow them is to rob human beings of the chance to understand why they are so powerful and seem so good.

Perhaps there are people who are unable to comprehend experience at the level I call system concepts. I doubt that this is really true, but if it were true, there would be no recourse but to invent a God, a Heaven, and a Hell. Without some such constraint, people would pick principles without regard to the welfare of the system of which they are a part. They would choose principles that are to their own immediate advantage, and they would be unaware of what they were doing to the capacity of all people to get along together. The only way to enforce any sort of morality for the social good would be to substitute for understanding the idea that morality must be followed for the individual's good: to avoid personal eternal punishment, or to attain personal eternal reward. That is the sort of concept that a person without system concepts can understand.

But I do not believe that any significant number of people is incapable of grasping system concepts. I think that system concepts are simply not articulated, made clear, made real, taught. I think that children can learn to see why moral principles make sense in their own lives, when they reach the age where system concepts become formed (whenever that is).

Best, Bill P.

Date: Thu Jun 03, 1993 4:40 pm PST
Subject: Non-Tentative "Truth" (again)

From Greg Williams (930603) Joel Judd 930603

GW>If you want to save time...[can] any aspects of revealed religion ever
GW>be known in a non-hypothetical way?

>If you mean the existence of God and other beings, or actual physical
>objects pertaining to a group's beliefs, yes. I have written statements
>of those, in many cases more than one person, who have seen and touched
>such things.

Let's get radical (well, OK, PC-radical). Why aren't the seeing and touching also theoretical, and hence hypothetical? Earlier in this thread, I was somewhat surprised when you said that you think ALL knowledge is hypothetical. (I was worried about what your church leaders might think about a member who claimed that even "revelation" (that is, its interpretation) is tentative.) But now it appears that you think there is some class of knowing which ISN'T tentative. If so, I would be interested in hearing you decide whether an experience is theoretical (I think they ALL are) or not-theoretical.

>To take an
>example from my church, there were a number of men (eleven, to be exact)
>to whom Joseph Smith showed the gold plates from which he claimed to be
>translating. Virtually all of these men became disaffected with Smith or
>the church or both sometime later, repudiating particular principles of
>the church. None of them ever denied having seen the plate (though some
>lived to be more than seventy) or their knowledge of the Book of Mormon.

But isn't the claim that those plates were delivered to Joe Smith by a supernatural entity (Moroni, to whom they later were returned) THEORETICAL for EVERYONE, including folks who never saw any plates, folks who did see some plates, AND Joseph Smith, who said he was visited by Moroni? Maybe I'm not understanding you in this regard, but you sound as if you think that (in some cases, at least) believing you have had a vision equates to KNOWING ABSOLUTELY that you have had that vision (and perhaps more, like knowing absolutely that some notions associated with the vision but not directly experienced MUST be true. Couldn't the plate witnesses have seen plates contrived by Smith or some hoaxer out to fool Smith and his friends? I have read that some scholars have claimed that the BOOK OF MORMON was basically copied from a manuscript by a third party (no, not the ms. which the Mormons have shown to be different from the B.M., ANOTHER ms. by the same third party), so the possibility of an elaborate hoax doesn't seem entirely out of the question in this particular case.

>But most religions require FAITH, and so their principles are going to
>be hypothetical, and some are going to say such principles are true, and
>others are going to say they're false. What would it take to make a
>principle or concept "true" anyway?

Some people think that, say, the proposition "the angel Moroni delivered the plates to Joe Smith" would be true if it were the case that there were a supernatural agent Moroni which showed Joe Smith how to find the plates which he showed to some other folks. Others would have other requirements. I've been trying to learn what your requirements are.

>Basically I've admitted that I haven't seen Christ personally and
>that religious principles are hypothetical, yet that one can have
>knowledge that a particular set of principles are "true" and will work
>for everyone, without making everyone believe such a thing by force.

I'm just trying to understand how you can be sure of certain kinds of "truth." We seem to be making progress. Perhaps if we leave "truth" aside and just speak of "working" or not, then you might try to make the case that because people are sufficiently similar in certain ways, then some principles (in general enough ways, with appropriate caveats) should "work" for all humans. Would you?

As ever, Greg

Date: Thu Jun 03, 1993 6:25 pm PST
Subject: Re: What's good about morality?

[From Rick Marken (930603.1800)] Bill Powers (930603.1445 MDT)

>Rick Marken appears to be arguing that if you just let children
>develop all by themselves, they would turn out to be successful
>moral beings, but I don't think that even Rick knows that.

Actually, I was just trying to get out saying (not nearly so well) all the stuff that you said so nicely in your post. Of course, parents set an example and that's why my kids are so sweet and brilliant. I just think a lot of parents think they have to do what amounts to controlling their kids in order to get them to "come out right". If parents could just learn to forget the controlling (in the guise of guidance) I think they would generally be fine models for their kids.

>Religious teachings, stripped of their explanatory frameworks,
>contain a great deal of wisdom and practical experience.

I am underwhelmed by that "wisdom", frankly. I mean, who really needs to be told that it's wrong to eat pork, that you should treat your slaves well and that masturbating is a sin. This stuff is OBVIOUS.

Dennis Delprato (930602) --

Boy, am I a barometer of losing propositions. I love bluegrass music (one of my old (young at the time) girlfriends was a dynamite bluegrass fiddler; she still has a great band up in the bay area) and I love PCT.

Mary Powers (930603)--

>On free will. I may not have free will, but it sure feels like I
>do. I think free will is just that - a feeling. Specifically, how
>I feel when I am acting without conflict. When choices are
>obvious and easy, when I don't have to decide between this and
>that, when I don't have to force myself because I'd rather be
>doing something else. Maybe this looks as though I'm just being a
>tool of my intrinsic reference levels, but not really, since they
>are me.

This is really nicely put and it makes a hell of a lot of sense from a PCT perspective. Of course, free will is a perception -- the perception that get's to the right level when we are operating without conflict. I like to call it "grace".

Ed Ford (930529:2300) --

Sorry for the belated reply. And thanks for the transcript of the Reality Therapy talk. Those people are horribly confused, indeed. So confused that it is difficult to say just what they do and don't actually know about PCT. The whole, confused transcript is a prolonged argument for the importance of having a working PCT model in the room at all times.

But, then, they're the ones making the big bucks so who am I to say "nay".

Best Rick

Date: Fri Jun 04, 1993 7:36 am PST
Subject: A tentative answer (again)

[from Joel Judd 930604] (this will eventually get to Rick and Bill too)

Greg (930603)

>...aren't seeing and touching also theoretical...?

I can understand and accept that ALL our experience is derived from rather simple sensory input information, such that the "level" at which most adult human beings normally function deals with perception that have been well-processed, so to speak (e.g., language). When I tell you "Sit down" or "Check this book out" I don't expect, nor do I usually have, questions about whether the chair or the book are "real" or just figments of our

perceptions. In other words I have learned through experience with the sensations, configurations and all the rest that this thing I'm typing on is a "computer" and it serves to do a number of things I like to do.

So from a nervous system point of view, all things are theoretical, inasmuch as I can't guarantee that something which I know by my experience won't turn out different tomorrow.

Likewise, EVERYTHING in human experience can be reduced to electro-chemical processes; therefore, noone can know anything ABSOLUTELY, not the person who had the vision, or the people who believe someone had a vision. But I don't find it useful to spend a lot of time at this level; about the only thing practical it does is cast a skepticism over experience. Rather, I can say that I know something for sure because it works, as you say; it has always worked for me, for my parents, grandparents, brothers and sisters, many acquaintances, etc., and I have no reason (other than the one just mentioned) to believe that it won't work tomorrow for me or for others as well. Whereas before I could say this just because I could see others in action, now I think there's a basis for saying it because if everyone is similarly organized PCT-wise, then "it" will work for them as well.

One of the questions which has already come up several times then, is this (a particular systems concept and accompanying principles) the ONLY way to develop such principles? In terms of a religion, is being part of religion X the ONLY way to be a certain type of person? And the answer is no. There are many people, on this net, and others who have lived in places and times where religion X was not available or even in existence, who have grown up and lead (or led) exemplary lives. This leads to two related questions: Why participate in a religion if one can be a "good human being" without it, and, What good is religion in the first place? The first, and probably least convincing answer for netters, is that most religions are not just concepts, they're also institutions which administer rites or ordinances supposedly required by God (e.g., baptism). Just as one cannot go out and declare oneself a policeman, or a Ph.D. without passing certain requirements, one cannot declare oneself a Catholic, Buddhist, Mormon, etc. without complying with certain practices (of course one can SAY anything but let's go on). The second, and more relevant for me, is that a principled, responsible (fill-in your own adjective) religion will also help one (particularly parents) "hedge his bets" in growing up to make the most of life. That is, given an organization made up of individuals who have found success in applying certain principles to their lives and can talk about it, odds are that the same principles can work for me, too. And instead of just being a scattering of principles, they are principles united by a clearly understandable systems concept that provides justification for them.

Rick and Bill (930603) (Alternately, and actually beginning with the last half of the above paragraph)

Although religion started this whole thread (I think), I am particularly interested in the ideas discussed here with respect to aspects of society. So please don't think that I'm trying to hone my proselyting skills here.

RM>...an obvious question; what if my son came home from college and
RM>told me he had become a Mormon?

Let's cross that bridge when we come to it. One of the things that's fun about e-mail is that what one person expects the other doesn't see, so what was an obvious question to you wasn't to me. But I'm digressing.

I found Bill's post helpful and confirming in many respects.

BP>Children learn how to be by considering what their parents do far
BP>more than what they say.

Interesting, considering that we consider the behaviors INCIDENTAL to what's really important: controlled variables.

BP>[children might not] violate any principles of PCT even though they
BP>are doing something "wrong" in moral terms. I think the crux of the
BP>argument comes down to the justification one gives for moral principles.

I completely agree. And so if autonomy is fundamental (or life, liberty and the pursuit of happiness) then there's a basis for justifying certain moral values. But as you point out, there's (at least one) higher level of perception above this: systems concept. Perhaps the problem nowadays is that there's no good reason NOT to grab a gun and shoot, or steal a

car. What is being taught to others in terms of systems concepts? What is one to be a part of today? American? Family? Crip? "WASP"?

BP>Parents and teachers who understand systems concepts can show them
BP>at work, show how maintaining them results in a better life.

This is where I think I'm having trouble. One recent movement in education is "multicultural" education [BTW, where is Hugh Petrie? Is he available for comment?] where certain people are insisting on greater attention being given to more variety of cultural ideals, and, to a greater or lesser extent, on the worth of those ideals. One of the criticisms of this movement is that instead of being given a model against which to judge (in the U.S. usually an anglo-european one) children are being taught essentially that all cultures and ideals are equally valuable, with the result that they have almost no basis for their own experiences, in the sense that what has gone before was uniformly good and bad, and what really matters is only NOW and the gratification of one's own desires and needs.

Where's the development of a concept in this, then? Is it any wonder that those who have been unfaithful spouses, thieves, gamblers, drug users, sexually promiscuous and so on are idolized by youth and in the media?

SIDELIGHT: Much of the discussion about God often implies that He can do whatever He wants; it's us sinful mortals that have to be constrained by commandments. Would it be a surprise to suppose that God can't do whatever He wants, either? I believe that the "laws" which feel constrained to obey are eternal in the sense that we didn't invent them. Just as we can't violate the "laws" of physics without consequences, we can't violate moral laws without consequences. It's not that God is sitting around waiting to zap us when we disobey, the consequences follow from the behavior. We do it to ourselves (and others).

Regards, Joel

Date: Fri Jun 04, 1993 10:44 am PST
Subject: Higher levels

[From Rick Marken (930604.0900)]

I read an interesting thing in The LA Times this morning. In an editorial endorsing Woo for mayor, The Times said that he has shown the ability to "rise above his principles" in order to achieve common goals. The "rise above principles" quote was attributed to JFK who apparently used it to describe an attribute of successful leadership. Kennedy seems to have understood the point I have been trying to make about the relativity of perceptual goals; I understand JFK as saying that a person must be able to vary his or her references for principle perceptions in order to control higher level perceptions -- which we would call system concept perceptions (such as "the state of the society").

In order to be able to vary references for principles, one must be able to "go up a level" and see principles as adjustable means for achieving higher level perceptual goals -- system concepts. So "rising above principles" is the process of "going up a level" in PCT -- in this case going from principles to system concepts. Given that most people seem to find it nearly impossible to get from programs (rules) to principles, this seems like a pretty impressive observation.

Nice going, JFK.

Best Rick

Date: Fri Jun 04, 1993 2:00 pm PST
Subject: Laws of nature

[From Dag Forssell (930604 10.15)] Joel Judd 930604

>I believe that the "laws" which feel constrained to obey are
>eternal in the sense that we didn't invent them. Just as we can't
>violate the "laws" of physics without consequences, we can't
^
>violate moral laws without consequences. It's not that God is
^

>sitting around waiting to zap us when we disobey, the consequences
>follow from the behavior. We do it to ourselves (and others).

I am enjoying the thread on truth and morality very much.

Joel, It is a mistake to equate "laws of nature" and "moral laws."

A clarification on the difference between "laws" and "laws." No-one has *ever* violated any "law" of physics. There are no consequences of breaking them - you cannot do it. They have not been "legislated." "Law" is a misnomer. These are descriptions of phenomena which always apply. They *cannot* be broken. These are what some call first principles. As an example, the law of gravity always applies. Many people have not studied it [Newton's equations, physical experiments] and do not quite understand it, however. Therein lies the difference between empirical descriptions and generative theoretical modeling, which I have addressed in previous posts.

Joel, you may be familiar with the well known mormon Stephen R. Covey. His program: The seven basic habits of highly successful people is the best conventional leadership program I know of. He espouses several "principles," such as "Seek first to understand, then to be understood." While he speaks of his principles as "laws of nature" -- "based on scripture," not one of them are laws of nature. With any of his principles, you can do the opposite of what he recommends. They are recommended choices. The particular one quoted above even requires an agreement with another party.

Reading Webster's we find seven definitions of principle, including "rule of conduct" and "law of nature."

"Don't jump off the balcony" is a rule of conduct, a wise law perhaps taught by means of a grisly story of what happened to the fool who did. It is not the same as the law of gravity.

Some of the thread on morality, as I read it, says that it is better to teach about gravity than to tell stories about what happens to fools who are careless on balconies.

Then again, I have recently argued that we need more application stories a la Ed Ford's Freedom from Stress to draw attention to PCT -- truly a set of laws of nature.

Certainly, a story about the fool and the balcony can serve to draw attention to and justify the effort required to study the law of gravity.

The sad thing about the overwhelming majority of both religious and secular teaching in the behavioral sciences is that it stops with the story and never explains the underlying law of nature.

An enormously large quantity of stories are required to illustrate all facets of human interaction. Each person faces a very large task of integrating all the stories (full of inconsistencies and contradictions) into a coherent whole and ultimately, sense of self.

The beauty of PCT, as I see it, is that with a little study, anyone can learn the laws of nature, the underlying principles of the human condition, and as a consequence extract what can be learned from all the stories and personal experiences quickly and well. PCT allows a person to sort the wheat from the chaff better than any other teaching has ever been able to (religious or secular). Integrating conflicting teachings is no longer a problem, when you grasp the underlying causal relationships.

Best, Dag

Date: Fri Jun 04, 1993 2:54 pm PST
Subject: Unreality Therapy

[From Rick Marken (930604.1200)]

Here are some comments on the transcript part of the 1989 Reality Therapy Convention that Ed Ford sent to me.

On page 1 Glasser says " For them [meaning, apparently, us -- the control systems group] the needs are the disturbance".

This is incoherent. "Needs" , in PCT, are intrinsic reference signals; "wants" are reference signals in the perceptual control hierarchy. A disturbance is an environmental variable that influences the state of a perceptual variable. A disturbance is definitely NOT a need. It has effects on perceptions that may require actions that prevent the perceptions from straying from their "needed" value. I think we can tell the difference between a need and a disturbance; but, then, we're looking at models; Glasser's looking at sentences.

On page 2 Glasser says "The thermostat has an internal need too but we have put that in. By itself, it doesn't care what the temperature of the room is."

The thermostat never cares about the temperature of the room; it "cares" about a perceptual input signal which represents the temperature at the sensor. And it does this caring all on its own; the thermostat is an autonomous system; it cares about what happens to itself. It only perceives one thing that happens to itself (temperature at the sensor) but it cares about that one thing deeply -- and it continually acts to keep the perceptual signal that represents that thing (the temperature) at the reference level.

The reality therapy people seem to think that "needs" are very important. The "needs" are apparently a list of names that Glasser brought down on a tablet from Mt. Reality. Diane Gossen says (on p. 2) "There are many people within the control system group that believe in the needs". Apparently, the "needs" are something people have to believe in to be true reality therapists. What was apparently not "revealed" to Dr. Glasser (the guy has an MD, which probably explains a lot) was "the test for the controlled variable". Had Glasser paid more attention to learning from Powers and less to lining his own wallet he would have been delighted to find that there is a way to test assertions about what people need and/or want. Needs and wants are reference states of perceptual variables; an observer can test to determine whether these variables are controlled by applying disturbances and watching for systematic opposing action and/or lack of effect on the perceptual variable. But doing the test requires thought and work --which just wastes time that could be spent spewing BS and getting the suckers into the tent.

On p. 5 - 7 Glasser discusses his objections to the idea of levels of perception. He seems to object to the idea that we can control different perceptions -- and that we can do this in order to control other perceptions. He says: " If our goal is to get to the town, our behavior is to get there or we are going to die trying. That' s the strength of the control system. Systems are a totality, they don't dissect out". Glasser seems to be neglecting the fact that many perceptual variables, all subject to unpredictable influences from the environment, must be controlled as part of the process of controlling the perception that represents one's relationship to the town. While it would be possible to build a system that controlled only one particular variable (proximity to town) using automatically generated actions, such a system could only accomplish it's goal in environments where disturbances to everything but the final goal were non-existent.

Glasser's objections to the hierarchy are very puzzling; he seems to think that it is impossible to attend to one perceptual variable at a time (he doesn't really even seem to understand that perceptions ARE variables) even though he describes perceptions one at a time. On p. 6 Glasser says: "You can't just listen to the tuba when listening to a symphony orchestra". Glasser seems to think PCT says that when you become aware of one perceptual signal all the others go away. In fact, all the perceptual signals are always there and, if they are controlled, they are always under control. Glasser's complete lack of understanding of this point reflects his lack of interest in models; it takes a little effort to learn how several distinct perceptual signals can be controlled simultaneously; too much effort for Glasser apparently.

By abandoning the hierarchy or, at least, the basic notion behind the hierarchy (control systems controlling by varying the goals of other control systems) Glasser has abandoned the feature of PCT that is most relevant to clinical practice (which is what I thought Reality Therapists care about). If control is "wholistic" as Glasser describes it then what is a conflict? If control systems don't "dissect out" then why do we see people tied in knots trying to control incompatible perceptions or to control perceptions relative to incompatible goals? Seems like the control systems have been able to dissect themselves out well enough; why can't the good doctor do it?

What is Reality Therapy about, anyway?

Best Rick

Date: Fri Jun 04, 1993 3:19 pm PST
 Subject: Education/understanding

I have been dabbling in PCT over the last three years--particularly in how it might apply to higher level systems. I believe that PCT provides the best explanation for how humans function.

I have not been on the network much until this last month, but I do have a question that I hope some of you old-timers could help with.

As an educator, my concern is that my students would develop a sense of understanding regarding certain traditional (and non-traditional) educational concepts. How could one explain (or model) understanding (or lack of it) in terms of PCT? Is it plausible that humans have an intrinsic reference standard for understanding?

Tom Hancock
 Grand Canyon University (Summer Research Associate at Armstrong Lab, USAF)

Date: Fri Jun 04, 1993 6:45 pm PST
 Subject: Of "somethings" and "its"

From Greg Williams (930604) Joel Judd 930604

>... none can know anything
 >ABSOLUTELY, not the person who had the vision, or the people who
 >believe someone had a vision. But I don't find it useful to spend
 >a lot of time at this level; about the only thing practical it does
 >is cast a skepticism over experience.

I, myself, find such skepticism useful in suggesting that claims (of either individuals or, through the individuals making them up, institutions) -- words, if you will -- should take a back seat to other activities -- deeds. Doing thus and so "because God wants us to" or "because [authority figure] says we should" seems not only insulting to fully functioning adults, but also unwieldy in practice.

>Rather, I can say that I
 >know something for sure because it works, as you say; it has always
 >worked for me, for my parents, grandparents, brothers and sisters,
 >many acquaintances, etc., and I have no reason (other than the one
 >just mentioned) to believe that it won't work tomorrow for me or
 >for others as well.

I won't go back to quibbling about the "for sure," because I think this is basically the bottom line regardless of how it is labeled. Each of us looks around, sees the approaches of some others to life, and, to a degree, adapts (our interpretation of) certain of those approaches to our own living. Of course, we are ignorant about how well OTHER approaches which we DON'T see (very closely, on an ongoing basis) might work for us. And, too often, I believe, the result is a dismissal of those "distant" approaches. In a word, bigotry. (Moderating this tendency is claimed to be a primary goal of multicultural education.) The bigot just doesn't seem to understand that there are MANY folks with "alien" (to the bigot) "somethings" who believe that their "somethings" have "always worked" for them and their acquaintances, and who believe that it they probably continue to work.

>Whereas before I could say this just because I could see others in action,
 >now I think there's a basis for saying it because if everyone is similarly
 >organized PCT-wise, then "it" will work for them as well.

And if everyone is similarly organized PCT-wise, it appears that MANY, MANY "its" might work for ANY individual (I'll hedge here: to a degree; I suppose that because each person is NOT organized IDENTICALLY to others, some "its" might NOT work for some individuals, and some "its" might work better than other "its" for some individuals).

Some believe that ONE "it" is IT -- that it will work better (if not in the natural world, then in a supernatural world to come) than any other "it" for EVERYONE. I am wondering whether you have any sympathy for anyone making that sort of claim?

>The second, and more relevant for me, is that a
 >principled, responsible (fill-in your own adjective) religion will

>also help one (particularly parents) "hedge his bets" in growing up
 >to make the most of life. That is, given an organization made up of
 >individuals who have found success in applying certain principles
 >to their lives and can talk about it, odds are that the same
 >principles can work for me, too. And instead of just being a
 >scattering of principles, they are principles united by a clearly
 >understandable systems concept that provides justification for them.

Setting aside the question of just how "clearly understandable" are the systems concepts of some organized religions, I am wondering at this point whether you think that the odds you speak of are much reduced for all or some of those who shun such religions.

I know several secular humanist/agnostic/atheist types whom I think would satisfy most anyone's (but not God's, of course) desiderata for being admirable individuals in many ways. Do you suspect that they are statistical fluctuations, rare successful-in-spite-of-themselves freaks in a burgeoning crowd of despicable miscreants? Am I unjustified in supposing that (something along the lines of) some of their beliefs might "work" for me, too?

Unfortunately, there is also another side to this: the moral decrepitude rampant (not a new phenomenon) in many nominally principled and responsible organized religions. (I.e., lately Kentucky seems to be blessed with shepherds who like to molest lambs from their flocks. No Buddhist cases yet, but both Catholics and Protestants.) The plethora of organizations "made up of many individuals who have found success in applying certain principles to their lives" and ALSO harboring many individuals for whom the organizations' systems concepts apparently didn't work (in the sense of producing admirable principles, or at least admirable actions) doesn't give me much confidence that the organizations' systems concepts would work for me.

As ever, Greg

Date: Fri Jun 04, 1993 8:19 pm PST
 Subject: What Is Reality Therapy

from Ed Ford (930604:2100) Rick Marken (930604.1200)

>What is Reality Therapy about, anyway?

Rick, I enjoyed your review of Glasser's comments during Diane Gossen's presentation at the RT 1989 convention.

Reality Therapy was proposed by Glasser in his book, Reality Therapy, in 1965, long before he had heard of PCT. He actually explains it more clearly in later books, especially his chapter entitled Reality Therapy in his book, Identity Society. Basically it is a very effective counseling technique which proposed a sequence of ideas to helping clients. They are:

1. Involvement - which basically said that establishing a relationship with your client was critical to an effective counseling relationship.
2. Ask Client What Are You Doing? - which meant that the counselor should get the client to focus on their present actions.
3. Ask Client Is What You are Doing Helping You? - Here the client was being asked to focus on whether his/her actions were helping them with their problem. This was often viewed as the responsibility step since you were obviously asking clients to take responsibility for their actions. This step was also called, Ask For A Value Judgement.
4. Ask Client To Make A Plan - Once clients are willing to take that responsibility for their actions, then they should focus on making a new plan of action to help them achieve their goal.
5. Ask Clients To Make A Commitment (to the plan, that is) - which demonstrates their resolve.
6. Don't Accept Any Excuses - which involves not asking why because when you do, it leads to excuse giving which leads away from accepting responsibility.

7. Don't Punish - which means that this type of activity doesn't help clients take responsibility and obviously doesn't work.

When Glasser heard about PCT, although as you now realize he doesn't understand it, he did pick up the idea that we are internally driven and we set internal goals. So, he added the question as a second part to No. 1 above, What Do You Want? This addition made it much clearer and made his second step much easier to understand. He has since done away with the steps to reality therapy and has reduced it to just What Do You Want?, What Are You Doing?, and Is What You Are Doing Getting You What You Want?

He also states that everything we do is to satisfy internal needs. Those who are happy are satisfying those needs, those who aren't satisfying those needs aren't happy. He has pronounced those needs to be Love/Belonging, Power, Freedom, and Fun. I personally don't agree with his needs concept (See FFS Chapter 7).

Much of what he wrote in Stations Of The Mind was what Powers taught him. I just think he wrote those things but I really don't think he understood them. He once wrote me "leave perceptions alone, leave them to the theorists."

His counseling techniques became quite popular and useful to many of those in the counseling field, especially in corrections, schools and residential treatment centers. They are very effective tools and widely used and respected. Because of that respect, his mistaken ideas on PCT are also accepted as gospel and respected because that's all people have heard about. In my travels into school systems around the country, once they've heard what PCT is, I've had widespread acceptance. Obviously, many haven't learned it to the degree necessary for a really good understanding, but they've learned enough to perceive the "sense of it."

Will be off the net next week. Will be up in Green River, Wyoming, teaching the PCT Gospel According To WTP.

Best, Ed

Date: Sat Jun 05, 1993 11:23 am PST
Subject: BBS paper, version 4, incomplete

[From Bill Powers (930605.1300)]

The following untitled paper is draft 4 of a paper for BBS. It seems to be going better than the previous versions. As usual, I invite all interested competent parties to become coauthors and suggest changes from wording to organization to style to content. Also, this paper ends at the section where specific misinterpretations about PCT are to be discussed, with references and quotes. I will supply some examples that I have on hand, but request more from all you out there, the more recent the better, and better still from the pages of BBS. I have a tendency to misplace things, so don't worry about duplication. I would also appreciate suggestions about how to organize the last section. One way would be to pick major fields in which PCT has been mangled or criticized; another way would be to sort by the kind of misinterpretation. I'll take care of putting it all together.

If you suggest changes, please send them to me by e-mail, copying just the part to be changed and then adding your substitution. I won't guarantee to go along with every suggestion, but will accept anything well-considered.

I'm currently thinking of presenting the paper as authored by a rather long list of people, giving their specialties and affiliations. The point will be to leave an impression of the wide range of disciplines represented by PCTers. Alternatively I could appear as author along with an "Editorial board of the CSG", the listing being given in a note at the end. Any other suggestions also welcome.

Bill P.

Control theory has been knocking at the doors of the behavioral and life sciences for over 40 years. A few doors have opened; most have not. There is no single clear reason for the difficulties in assimilating this "new" idea into the mainstreams of science; on the surface, the reasons seem to shift with every different application. Perhaps one of the most illuminating objections to it came from Edwin Locke (199?), who proposed that because

control theory seems to be used to explain a wide variety of phenomena now explained in terms of different theories, it amounts to nothing more than a series of borrowings. The fact that control theory has been in existence much longer than most of the theories from which it supposedly borrows casts some doubt on Locke's interpretation, but Locke has put his finger on one major problem. It is not the details of control theory that cause the difficulties, but the fact that control theory offers explanations of a kind with which many behavioral scientists are unfamiliar. The conclusions that can be drawn from control theory often fit generalizations that others have drawn intuitively from their data. But control theory should not be evaluated solely at the level of conclusions; its real heart is at the next level down, the level of underlying regularities and relationships from which conclusions can be deduced and observations can be explained. It would be surprising indeed if a valid theory of the organization of behavior led to conclusions contrary to competent observations.

The proponents of control theory, or as it is now called to differentiate it from engineering, "perceptual control theory" (PCT), have only slowly come to realize what the major difficulty is. At first, of course, its developers and supporters spoke of it as a revolutionary concept, but they had no idea then of what "revolutionary" actually means. It has taken decades of struggling to introduce the ideas of control theory into the life sciences for PCT adherents to grasp the fact that PCT is not revolutionary in the sense of being a new product that will titillate the market for a while, but that it is revolutionary in the true Kuhnian sense. It is a departure from normal science. The extent of this departure becomes apparent only through interacting with supporters of normal science and gradually realizing the depth of the differences of interpretation.

The basic concepts of PCT meet with resistance in almost every field of the life sciences because they depart from a basic understanding of the nature of behavior that is common to all established lines of investigation (it will be described later). The effects of this departure take on different forms in different disciplines. In some, the methodology is attacked; in others, the conclusions; in still others, the premises; in many, the emphasis on a mathematical approach to psychological phenomenon; in yet others, the psychological interpretation of biological phenomena. PCT seems to get into every line of fire that exists between rival theories and rival subdisciplines. Where a discipline finds points of agreement with PCT, it claims that there is nothing new in it; where it sees departures, it interprets the meaning of PCT as placing it in some enemy camp and meriting the same (ready-made) criticisms. There sometimes seems to be no way for proponents of any mainstream discipline to see PCT simply in its own terms, as a single coherent approach to the organization of living systems.

What we propose to do in this paper, or attempt to do, is to present the basic principles of PCT and to point out some of the major misinterpretations of it. We will focus on misinterpretations not to embarrass anyone but simply to post warning signs showing how a wrong conception can be formed through incomplete understanding. While we will cite specific examples, the authors involved should not be blamed for inventing the misinterpretations; most of them have been circulating for two or three decades or more, and would be difficult to trace to their real sources. We hope that by showing through specific examples what PCT is not, we can make clearer what it is.

Let us begin, however, by examining the core concepts of PCT and outlining the nature of control as it is defined in PCT.

Control

The term "control" is used loosely in a multitude of different meanings. Outside of control engineering and the few branches of the life sciences where control theory is being knowledgeably applied, the meaning is interchangeable with the verbs "affect", "influence", "determine", or "cause." If one can account for all the variables on which some other variable depends, it is commonly said that the independent variables control the dependent one.

Rather than get into a verbal dispute over the meanings of words, we propose here to use the term "control" in a single narrow sense, and to use other words to refer to other usages. In the sense used in PCT, control is a process that arises from a relationship between a system and its environment. One of the main features of this relationship is the emergence of variables in the environment of the system whose state is determined unidirectionally by variables inside the acting system.

To say that much is not nearly enough to distinguish control from influence, determination, causation, and so forth. As stated, control could simply mean linear causation. To make

clearer what the relationship actually is, we must introduce some worst-case circumstances into the environment. Not every example of behavior will entail all of these worst-case conditions, but what matters is that such conditions are common, and control systems can work in spite of them where other types of systems fail. Living control system must work under all conditions, not just under specially-protected conditions.

Consider the old standby, the thermostat. The environmental variable that is controlled by a thermoregulator is, obviously, temperature.

In a home system, the means of control is a furnace that can circulate heat through a room. When the furnace spends most of its time turned on, the temperature in the room continually rises; when the furnace is on only a small fraction of the time, the room temperature falls.

At the same time that the furnace is acting to alter the room temperature, other factors also affect the temperature. Doors and windows open and close; the number of people giving off heat varies; a fireplace may be lit or an oven may be in use or the lamps in the room may be on or off. Heat is continually being lost through the walls, at varying rates depending on outside air temperature and heating by the sun.

The physics of the situation may get complicated, but it is not basically complex. The room temperature will come to some average value that depends on all the heat losses and all the sources of heat including the furnace. At equilibrium, the sum of all heat inputs will exactly equal the sum of all heat losses, in units of energy per time. When that is true, the temperature will be constant.

The primary point to see here is that the controlled variable, the temperature, is not simply the output of the control system. The furnace that is turned up and down by the thermoregulator is only one of many influences on the temperature. In general, the other influences are unpredictable, depending on factors that can't be anticipated in the design of the control system. There are innumerable events that can tend to make the temperature rise and fall, events as unknowable to the design engineer as the occurrence of a large birthday party or a spell of cold but sunny weather, with or without wind and passing clouds.

There is no practical way to design a home thermostat along the lines of traditional machines, by calculating and applying carefully-calibrated output effects. The only action the thermostat can take is to increase or decrease the output of the furnace, either continuously or by varying the on-off duty cycle. No one output command to the furnace can possibly result in a predictable room temperature. Not only are all the other sources and losses of heat unpredictable, but the furnace itself may produce a variable amount of heat as a result of a given command. The fuel BTU content can vary; the line voltage can vary and alter the blower speed; someone may push a piece of furniture or a rug over a heat vent.

One can, if freed from practical constraints, imagine an elaborate system equipped with sensors for everything: for the fuel assay, for the flame height, for the air speed in ducts, for the degree of openness of doors and windows, for the outside air temperature and solar input flux, for the number of people in the room and their degree of activity, and so on. From complete knowledge of all these factors, a very smart computer could deduce the exact commands to send to the furnace to make up the difference between any net heat gains and heat losses to produce just one particular room temperature. So "control" could be achieved by analyzing the environment and the furnace, and computing the commands required to produce a specified equilibrium temperature.

While this sort of temperature regulator is just barely conceivable, it has one great drawback as an explanation of how a real home thermostat works. When we look at the actual device, we find only one lone temperature sensor, located in the box on the wall. There are no sensors for all the heat sources and sinks, for the fuel energy content, for the duct air speed, for the outside air temperature, for solar heat, and so on. All those disturbing influences are unsensed, and their states are not taken into account at all.

Yet the temperature control system maintains the room temperature within a few degrees of any given set point despite all these variations, altering the average heat output from the furnace in exactly the way needed to counteract variations in all the other contributors to room temperature. Despite all the unpredictable and unsensed disturbing influences, including even changes in the efficiency of the furnace, the temperature of the room is maintained very close to the specified reference temperature.

This is the general case that unequivocally identifies a control system and differentiates it from any conventional kind of machine. A control system controls an outcome, not an output. It varies its output as required to compensate for the effects of all possible disturbing influences that can alter the outcome, the state of the controlled variable, even influences that have never occurred before and may never repeat.

Yet its design takes no account of disturbing influences at all. The sensor of a control system provides a direct measure of the state of the quantity to be controlled; the system acts by varying its own influence on that same quantity. As long as that loop is intact, so that the control system can affect the variable it is sensing over the required range, it does not matter in the slightest what disturbances are present, or why the controlled variable tends to change. As long as the sum of all disturbances does not exceed the capacity of the control system's output to oppose their effects, control will be maintained: that is, the controlled variable will be maintained near a specified reference state.

It is casually said that a home thermostat controls room temperature. If we were equipped with a hypersensitive thermometer, however, we would find that the degree of control depends on where in the room we measure the temperature. If we placed the probe over a heat register, it would show enormous temperature variations as the furnace went on and off. If we measured near the front door or a window, the variations would be appreciable as the door opened and closed, or as the wind outside varied in speed. By measuring all over the volume of the room, we could map the magnitude of temperature fluctuations. This map would show a reliable minimum fluctuation in just one location: at the sensor in the box on the wall. The very lowest temperature fluctuations would be found inside the temperature sensor itself: the bimetallic coil or the air capsule or the thermistor that is sensing the controlled variable.

So what does the temperature control system control? It controls the state of its own sensor. As we trace backward from that sensor toward the furnace, through the environment, we find that disturbances have a greater and greater effect, the effect being the largest at the output device, the furnace. Opening a window in the living room causes the sensor temperature to fall slightly; the air in the middle of the room gets even cooler; the air near the heat registers becomes tens of degrees hotter. The farther we get from the sensor, the greater the variations that are produced by disturbances of temperature. The greatest variations of all are found in the output device of the system, the furnace -- and of course in the command signals that are altering the output of the furnace.

So -- and this is the case for every true control system -- the control system does not control its output. It controls its input, the sensory indication that corresponds to the state of the controlled variable that we observe from outside it. The output of the control system, which we observe as its most obvious behavior, varies with every disturbance. The output of the system is just as unpredictable as the disturbances are; only the input can be predicted from knowing the set-point. To some extent, more distal variables on which the input state depends can be predicted, but the farther they are from the sensor, the more they will be affected by disturbances.

It's critical to understand that we are witnessing an apparent violation of cause and effect here, one that no traditional scientist would expect to find. We see a fluctuating output from the furnace. As we trace the effects of this output farther and farther into the environment, we find that the fluctuations become more predictable, until at one point downstream in this causal chain we find the most predictable state of all, that point being the location of the sensor.

It might seem that the fluctuations caused by the changing output simply die out or are averaged out as we proceed further into the environment. There is a simple way to show that this is not the right explanation. All we have to do is change the reference- setting of the thermostat, moving the set-point from, say, 68 degrees to 75 degrees.

We can predict with very high confidence that the temperature of the sensor will rise from a narrow range centered on 68 degrees to a narrow range centered on 75 degrees. As we move farther from the sensor, we find that the air temperature shows much larger fluctuations, and at the furnace we again find the largest fluctuations of all. Unless we can predict quantitatively every future disturbance that is going to occur inside and outside the room, we will be totally unable to predict the furnace output, or the room temperature at points far from the sensor. But we can predict very accurately the effect of changing the set-point on the input temperature. To make this prediction, we do not have to know anything about whatever disturbances may occur. All we need to know is the reference setting, and the fact that a control system is present and active.

This is the basic organization and mode of behavior of any control system of the type we consider in PCT studies. This sort of system is very different from the mechanisms traditionally imagined to be involved in behavior. There are aspects of control behavior that go against untutored common sense; to understand control behavior it is necessary to take a different approach toward analyzing it, different from those traditionally employed in the life sciences.

The methodology of PCT

PCT is not just a theory that organisms are organized as control systems. It is a model that has properties, and those properties can be compared with measurable properties of real behavior to determine whether or not a control system is present. Each time the hypothesis of control is applied to behavior, the facts can easily disprove it, if certain tests are failed.

The basic test is nothing more than an attempt to find the relationships just described above. The core of the test is a search for a controlled variable. A controlled variable is some measure of the environment near the supposed control system. This variable must satisfy a number of constraints:

1. It must be affected by an output of the behaving system, and also by independent disturbing variables. A test for control requires the presence or application of known disturbances.
2. The effects of the output must remain nearly equal and opposite to the effects of disturbing variables, so that the putative controlled variable is maintained near some specific value.
3. If the output is prevented from affecting the controlled variable, that variable must change according to the effect of disturbances on it.
4. If the control system is prevented from sensing the state of the controlled variable, the effects of unpredictable disturbances on the variable must no longer be canceled by the variations in output of the behaving system. When these four conditions are met, the putative controlled variable is accepted as being actually under control by the behaving system. All the defining relationships of control have been observed. Note that they are all public observations, entailing no model of the insides of the behaving system.

There are complications that can make this determination equivocal. The main one is that reference settings in living control systems are not fixed, but can vary widely. To make the clearest determinations, it is necessary to persuade subjects to maintain a constant reference value for some time (or to vary the reference setting in a predictable pattern), long enough to perform the test. Fortunately, this is not difficult to do through instruction and demonstration. From quantitative measurements during such tests, the parameters of control can be measured, and can be shown to be extraordinarily stable over long periods of time, as long as five years, even when new patterns of disturbances are employed (Bourbon, ---).

Once a control system has been identified and measured under these somewhat artificial conditions, it can be used as an element in more extensive models, multi-leveled models in which higher control systems act by varying the reference settings in lower systems already measured. In principle, this method can lead to empirically-based models containing any number of levels of control, with multiple systems acting in parallel at each level. In practice, followers of PCT are barely at the beginning of such investigations; the best current models are single-system one-level models. However, such models can be proposed for any level of organization; it is not necessary to work strictly from the bottom up, or to know all the hierarchical arrangements that are possibly present. The signature of a control system is defined by the four steps above; there is nothing that limits applying these steps to any one type of control process.

The control-system model

The picture of control processes sketched in above is not a model; it is a definition applied to specific observable circumstances. The relationships involved are clear and quantitative, identifiable by empirical procedures. But to understand these control processes, to link them with other knowledge about the nervous system and subjective experience, it is necessary to propose an organization internal to the organism that could

account for what we observe from outside. That proposed organization is called PCT, or in its hierarchical embodiment, HPCT: hierarchical perceptual control theory.

For reasons of parsimony, we usually employ the simplest canonical model possible, and express the parameters of control in terms of this model. Many equivalent arrangements in the nervous system would be possible; without PCT-informed exploration of the nervous system, however, there is little point in trying to propose a one-and-only organization that will account for the observations. The canonical model treats the inner organization in a simple way that can be used as a starting point to suggest avenues of research. In the simple kinds of experiments that have been done by PCT researchers, it is capable of representing some behaviors with excellent accuracy -- correlations between modeled and real behavior consistently above 0.99, and errors of point-by-point prediction of dynamic behavior in the range of 3 to 5 percent. The research bibliography at the end will give an entry into such PCT-related literature as exists, and a few leads into other lines of quantitative work using control theory.

However, the point of this article is not to elaborate on the control model itself, but to focus on the phenomenon of control that this model is intended to explain. The phenomenon we saw outlined in the discussion of the thermostat is ubiquitous in the behavior of organisms, but it is all but unknown to life scientists. The literature of the behavioral sciences is full of misapprehensions about how control works and what the capabilities of control systems are. There are what can only be called superstitions against acknowledging the major features of control -- some biologists, for example, seem to have a severe allergy to the concept of a set-point, being seemingly unaware that set-points are quite adjustable and by no means imply a static state of affairs even when they are constant. The literature of the "softer" sciences of behavior is full of bad guesses about the properties of control systems -- we will get into that. Control theory in general is often treated as if it had just been invented a year or two ago, and as if anyone's guess about the properties of control systems is as good as anyone else's. None of that is true. Control theory was developed as a formal system over 50 years ago; the analysis of systems of this type is a mature discipline. There is very little guesswork left about the properties of control systems. Applying control theory to the behavior of organisms is largely a matter of identifying the key relationships in behavior, and applying well-known techniques to modeling them.

We shall turn now to the relationship between PCT and some of the traditional approaches to behavior.

Engineering models and PCT

A major source of misunderstanding of perceptual control theory comes from the literature of control engineering itself. Control engineering results in devices meant, in many cases, for human use. The terminology reflects the fact that human beings tell artificial control systems what to do, and are interested primarily in the visible outcome of control. As a result, when a layman sees engineering descriptions of control systems, it is all too easy to get the correspondences with the organization of living control systems wrong.

In engineering parlance, it is customary to speak of the reference setting of a control system as "the input" to the system; in the thermostat, the knob that adjusts the temperature setting would be called the input. A careful examination of any higher organism, however, will reveal a total lack of such adjusting knobs on their surfaces. In every case where a control subsystem has been traced out in neuromuscular circuits, the reference input exists in the form of a signal descending from higher structures in the brain. It is not accessible from outside the organism.

To compound the confusion, the engineering custom is to refer to the controlled variable in the environment as "the output" of the control system. Again, this is understandable in the practical products of control engineering, which provide to a customer a means of stabilizing something in the environment and allowing for its adjustment by means of the reference input. The "output" of the thermostat, from the customer's point of view, is the temperature that is maintained at whatever set-point the customer dials in. The actual output device of the temperature controlling system, where large physical effects are created and cross the boundary between the electrical circuits and the environment, is the furnace; by engineering custom, this device is treated simply as part of the control system and its immediate effects disappear into the defining equations. The definition of "output" is moved downstream in the output cause-effect chain to the position of the variable that is actually under control.

And finally, to round out this rich source of misapprehensions, control engineers commonly omit the actual sensing devices from their diagrams altogether. The output is commonly shown as a line exiting the control system, with a sensory takeoff shown merely as another line branching off the output line and reentering the control system. The so-called output is treated as entirely equivalent to the sensory signal representing it.

This engineering type of control-system diagram is found very early in the literature of living control systems; it appears, for example, as Fig. 4, p. 132, in Norbert Wiener's (1948) *Cybernetics: Control and communication in the animal and the machine*. It is found throughout the literature of engineering psychology dating from the same approximate time, and has been reproduced innumerable times in publications in the life sciences since then.

The omission of explicit sensors from engineering diagrams has left the way open for behavioral scientists to assume that the engineering "input" to a control system, actually the reference input, arises from an organism's sensory receptors. It has also led to overlooking the role of the actual sensory receptors, which are located not where the reference input is shown, but where the "output" becomes a line re-entering the control system. In many cases, it has been assumed that the reentrant line in engineering diagrams exists inside the organism, not even crossing the environment-organism boundary. A typical consequence of this misreading of engineering diagrams can be found in almost any current article on control theory that is not informed by PCT: for example, Fig. 1 below, from -----.

If one ignores the reentrant pathway from the output, the arrangement of Fig. 1 is simply that of a cause-effect device. An input from the environment enters at the left; the resulting sensory effects pass through the system and an output exits to the right. Thus the control system becomes assimilated to a perfectly conventional picture of behavior: an input-output, stimulus-organism-response, independent-dependent variable, antecedent-consequent conception of the organization of behavior -- with just a little hook at the end to indicate the re-entrant effect of the output.

In PCT, we do a topological rearrangement of Fig. 1 to produce Fig. 2. There is no difference at all in the basic organization, but by relabeling the input, output, and reference, we make explicit what is tacit in the usual engineering diagram and show the reference input in a way that can't be confused with a sensory input. By orienting the diagram vertically, we show the reference input arriving from above, implying correctly that it originates in higher systems in the organism and not in the environment.

Furthermore, we separate the actual effector output from its consequences in the observable environment, showing that there are environmental functions inserted between the immediate physical output effects and the controlled variable at the input. All these effects of the environmental linkage are properly taken into account in engineering analyses; they are simply lumped into the system's effector. But doing that has hidden them from the view of behavioral scientists trying to apply control theory to the behavior of organisms.

And finally, every PCT diagram of behavior contains at least one explicit disturbing variable with a linkage giving it an effect on the controlled variable. This continually reminds us that independent disturbing factors exist, and that they act directly on the controlled variable, tending to alter it. You will notice no such representation of disturbances in Fig. 1. This also reminds us of the primary reason that control rather than just response is required: the outputs of an organism are not the only influences on the controlled variable.

Something was clearly lost in the translation from engineering diagrams to applications of control theory in the behavioral sciences. PCT grew out of a line of analysis in which the lost details were restored and made explicit. However, by the time PCT had become a coherent theory of behavior (Powers, Clark, & McFarland 1960), the mistakes based on the initial misinterpretations had proliferated and attained the status of an authoritative analysis. This greatly steepened the slope up which PCT had to progress.

PCT and behavioral theories: general observations

There have been two major types of theory used to interpret the behavior of organisms. One can be called input-output theory, which contains all theories that trace the causes of behavior back to environmental effects on the organism. The other is not greatly different, but it focuses primarily on internal processes, some of which may arise spontaneously (oscillators, pattern generators, and symbol-processing programs, for example) and which create behavior by sending commands to the output apparatus.

Input-output theory sees behavior as resulting from the convergence of many influences, with behavior being the resultant in the manner of vector addition. Many theories of behavior in this class see the final result as a balance between opposing forces and tendencies. This is, by and large, the biological view of behavior, which influenced psychology enough to create the school called behaviorism. The logic behind this sort of model is straightforward: regular consequences follow lawfully from regular causes. In principle, according to this approach, if one observes the surroundings of an organism carefully and over a long enough time, correlating environmental events with the actions of organisms, one will eventually discover the regularities in the input-output function and will thenceforth be able to predict how organisms will behave under similar circumstances. The vast bulk of behavioral research is organized around this principle.

Theories that focus on internal processes are now known generically as cognitive science, although there are many aspects of cognitive science that resemble input-output theory. Neural networks such as multilayer perceptrons with backward propagation do attempt to reproduce internal learning processes in a brain, but once the networks are functioning they are thought of as converting sensory inputs into behavioral outputs: responses, for example, differentially appropriate to specific patterns present at the sensory inputs.

Cognitive science has made the concepts of goals and purposes more or less respectable again, after a hiatus of a century or so. It has encouraged a large amount of productive research into functions of the brain and nervous system. It has put the person back into the human organism.

But cognitive science shares a misconception with input-output theory. The main problem occurs on the output side of the organism-environment relationship. It is a result, oddly enough, of an intuitive recognition of the difference between output effects created by an effector and the controlled consequences of those output effects in a disturbance-prone environment. This problem can be seen in Fig. 2 (but not Fig. 1), or inferred from our earlier example of the home thermostat.

The problem is this: there is, generally speaking, no one effector output that an organism can produce that will have any reliable consequences in the environment. The reason is to be found in the representative disturbance shown in Fig. 2, and in the unpredictable effects of heat sources and sinks and of changes in the furnace characteristics in the thermostat example. For any given command delivered to the output device, the controlled variable might be brought to any state whatsoever, depending on all the other influences that are acting at the same time.

So it is not possible under conventional interpretations for a stimulus to create a reliable behavior, or for a cognitive plan or pattern generator to create a reliable effect on the environment. Of course special circumstances can be arranged, such that the output actuator has constant properties over some short period of time, and such that no external disturbances are allowed to occur. But such circumstances merely conceal the main problem that confronts every organism operating under normal conditions. There is simply no way to predict the effect of sending a specific command to the muscles, regardless of the source of that command. There is no universally reliable way to predict the effect of any action.

This fact, as mentioned, has been intuitively recognized. We do not describe most behaviors by describing the action produced by an organism. Instead, we describe a consistent outcome of variable actions. When we see a person extending an arm outstretched, we do not say that this person is tensing muscles, but that the person is extending the arm horizontally. In fact, the action of the system is producing an equivalent force acting straight upward, at right angles to the direction in which the arm is extended. There is no force in the direction of extension. We characterize the behavior not in terms of the output that is bringing it about, but in terms of a consequence of adding the actual outputs to whatever other forces are also acting. We do not need to know whether other forces are tending to pull the arm downward or lift it upward or push it sideways, nor do we need to understand that the output forces are counteracting those other forces. We do not need to know that the muscles are gradually fatiguing and losing their sensitivity to neural signals. We look only at the outcome -- and because it is a controlled outcome, we can confidently expect it to be repeatable under a variety of circumstances.

Almost any behavior picked at random will prove to have this character: what is named is not the action being created by the organism, but a stable consequence of that action, an outcome of the action. From one instance of a behavior to the next, the action might change radically, but the "behavior" remains the same. The behavior could not remain the same if the action repeated, because in general the environment's properties change and disturbances act in different directions and by different amounts. The actions of an

organism must be continually changing in order to produce a constant outcome in the real world.

Both cognitive theory and input-output theory tacitly assume that regular consequences -- behaviors -- follow from regular commands that produce actions. It would be illogical to propose that a stimulus situation could produce some typical behavior if the link from the nervous system to the final observed pattern were not repeatable; it would be equally illogical to suppose that a cognitive computer could compute an output act having a given result if that result were subject, at the same time and in unpredictable ways, to external interference. Both of these models, as currently conceived, must assume a regular connection between act and consequence. But that connection does not in fact exist.

The only reason that these two major kinds of theories seem to work is that the behaviors they try to predict are not actually outputs created by the nervous system, but controlled variables resulting from variable outputs being added to variable disturbances in just the right way.

In a more sophisticated kind of thermostat, there can be a "cognitive" system that contains information about the time of day and which can set the reference temperature to a comfortable level during the day and reduce it to an energy-conserving level at night when everyone is in bed. This higher-level system appears to compute the desired temperature as a function of time, and then to emit commands that cause the temperature to change as required.

But these commands never get to the furnace. They enter the temperature control system (similar to the previous example) as a reference signal, a signal that performs the same function as the knob on the cheaper thermostat but without any moving parts. If the temperature is to drop from 70 degrees to 60 degrees, the reference signal simply changes at the appropriate time from the equivalent of 70 to the equivalent of 60 units. It specifies the level at which the (now-electronic) temperature signal from the sensor is to be maintained. It says nothing at all about the state of the furnace output. The furnace may already be off because of a temperature disturbance that has raised the input temperature above 70 degrees. In that case it will simply stay off until the temperature has dropped to 60 degrees. On the other hand, someone may have left the front door open, so the furnace is spending a large part of the time turned on. In that case, the proportion of on-time will decrease until the temperature has dropped to 60 degrees, then increase again to prevent any further drop (anticipation mechanisms, by the way, see too it that the furnace begins cycling on and off before the final temperature is reached). This decrease and subsequent increase of on-time happens without any explicit instructions from the higher-level system. It is brought about by the automatic action of the temperature control system at the lower level. The reference signal drops to indicate 60 degrees and remains there; it does not have to vary to create the required variations in the furnace output.

So the reference signal from the higher system does not command output at all. It specifies the required level of input, the signal from the temperature sensor. The output is automatically varied by the lower system to achieve the specified input level.

This concept, expanded, utterly changes the meaning of a cognitive decision to act. When one reasons that the proper action is to buy 100 shares of IBM at no more than \$75 dollars per share, this decision is not converted into an action that will result in the specified purchase. It becomes a specification for a lower-level perception. The lower-level control system produces all the variations in action that are required (assuming that such systems have been learned), such as looking up "Stockbrokers" in the Yellow Pages, dialing the number, uttering strings of vocalizations, and thus launching series of consequences in the environment. The control system receiving this specification from the cognitive level in question manipulates all the lower-level details (involving, potentially, many levels of control similarly related) until there is a perception heard as "You now own 100 Shares of IBM purchased at 74-1/4 dollars per share," which matches the reference specification.

Carrying out this decision at different times might well entail completely different actions on each occasion; the cognitive decision does not specify the actions, but only their outcome. If the input ultimately received does not match the specification -- you hear that you have bought 200 shares at \$85 per share -- an immediate series of error-correcting actions will ensue, without further instruction from the cognitive decision-maker. The reference condition remains "buy 100 shares of IBM at no more than \$75 dollars per share." It will remain that way until the highest cognitive control system is satisfied; then it will be terminated.

The control-system model has somewhat different implications for the input-output type of model. In most studies of behavior of this type, the actual inputs to the senses of the

behaving system are not known. Instead, aspects of the environment are manipulated and ensuing changes in behavior are observed. It is simply assumed that whatever changed in the environment must have affected the inputs of the organism in some way, the proof being that there was a response to the change. From the standpoint of simple empiricism, it is sufficient to note that the environmental change is provably related to the behavioral change.

It should already be apparent that the regular link to behavior that is assumed is not consistent with the facts. In addition, however, there is no longer reason to assume a priori that an organism simply sensed the environmental change and reacted directly to it. PCT offers an alternative suggestion: that the environmental change amounted to a disturbance of one or more variables under continuous control by the organism. If that is the case (and there are methods for finding out if it is the case, or not), then the apparent stimulus is really just a disturbance of some other aspect of the environment that the organism is sensing and controlling. The action of the organism is then not simply a response, but a way of preventing the controlled variable or variables from being materially affected by the disturbance. The organism, in fact, may not have sensed the disturbance at all.

The four-step process described above can be used to test various hypotheses about potential controlled variables. If the test is failed, the input-output hypothesis remains tenable. But if a controlled variable is found that passes all parts of the test, the input-output interpretation has been proven false and a different explanation of the behavior must be adopted.

It is the contention of those who have pursued PCT in some depth that there is ALWAYS a controlled variable: that the input-output interpretation is never correct, for any organism. That contention goes considerably beyond what can be justified, but it is actually a reasonable extrapolation. After several decades of investigations, there has been no instance of an apparent input-output relationship that has not revealed a controlled variable. This does not prove that there are no correct input-output interpretations, but it does reduce one's expectations that many valid ones will be found.

Misinterpretations and misapprehensions

Date: Sat Jun 05, 1993 11:28 am PST
Subject: On Bill Glasser

[from Mary Powers 9306.05] more on Bill Glasser:

"I also appreciate you and what you've done. I'm still speaking more than ever on control theory and I never fail to give your name and location as the person who got me started and certainly the person who really took the ideas from nothing to a very powerful place. I say I've expanded and clarified them and we may not agree on everything I did, but I think you'll find that you taught me well and most of what I'm doing is psychologically sound."
Glasser to WTP 1/21/93

So this is Glasser's position, right to Bill P.'s face. He may acknowledge Bill, but he maintains that control theory is difficult and obscure, and he has done this great big favor of "expanding and clarifying", so his audience can just relax and soak up the improved version from him.

Mary P.

Addendum from Bill P.

The above was in reply to a letter I sent in condolence for the death of Glasser's wife, Naomi. I'm not particularly interested in making an enemy of Bill Glasser. He may eventually come around; I've even considered inviting him to a meeting of the CSG, which would be a distinct change from the sorts of meetings he now attends. I don't generally like to give up on anyone as a lost cause. An educated Bill Glasser would be of far more help to our movement than an ignorant and alienated one. He is a smart person with a great sense of the practical. Let's not write him off.

Date: Mon Jun 07, 1993 12:02 pm PST
Subject: Its and laws, cont.

[from Joel Judd 930607]

Dag (930604)

Re: laws. I'll go along with the mislabelling comment about "laws of nature", inasmuch as what we have "discovered" is, in fact, an incontrovertible law. Maybe it would be clearer for me to say that jumping off a balcony would be a behavior inconsistent with the law of gravity? But your point is well-taken.

I'm familiar with Covey but have not taken the time to go back and read church members like him since becoming interested in PCT. Someone who I do enjoy for his experience and manner of expression is Neil A. Maxwell.

>...with a little study, anyone can learn the laws of nature, the
>underlying principles of the human condition...

This is what I'm still wondering about--the APPLICATION of this knowledge. I agree with you about sorting the "wheat from the chaff" but only once one has a BASIS for doing so (which PCT provides in theory). I don't think simply knowing ABOUT PCT provides this basis.

Greg (930604)

>...claims...should take a back seat to...deeds

No problem: "But be ye doers of the word, and not hearers only, deceiving your own selves. For if any be a hearer of the word, and not a doer, he is like unto a man beholding his natural face in a glass: for he beholdeth himself, and goeth his way, and straightway forgetteth what manner of man he was...Even so faith, if it hath not works, is dead, being alone. Yea, a man may say, Thou hast faith, and I have works: shew me thy faith without thy works, and I will show thee my faith by my works. (James 1:22-24; 2:17,18)

>Doing thus and so "because God wants us to"...seems...insulting to
>fully-functioning adults...

Well, here we face one of the fundamental problems with being "learned" and being "religious": reconciling the two with humility. If we insist on everything that we believe or might believe conforming to what we know, then we are unlikely to ever have the kinds of experiences spoken of by "religious types." Most every religion requires humility of the type that asks that one to believe that there exists Someone who knows more. It's not a blind, unthinking obedience, but it does require submission, among other things. This can come about either by force or by individual choice. If we believe that because we have come upon a powerful theory of human behavior which can explain why so many organized social institutions wreak havoc among individuals we should avoid such institutions, or certain of their principles, and so on, then we are unlikely to be willing or able to submit to such an institution.

>Of course, we are ignorant about how well OTHER approaches which we
>DON'T see might work for us.

Sure. But I don't think one should confuse open-mindedness with a persistent skepticism (maybe not the best word?). I also don't think it's necessary to spend one's entire life trying everything out. I can be ignorant without being a bigot; I can believe in my systems concept (do I have a choice?) and in its usefulness for my children, without forcing it upon them and without knowing every possible alternative for raising children.

Which brings up the developmental thread again--is there a good reason for treating, say, pre- and post-pubescent control systems differently in a systematic fashion? Do I really do a child more good by exposing him or her to a number of "ways of being" in somewhat ambiguous terms instead of saying "here's a reasonable concept for being based on the following principles" which he or she may later decide is incomplete or inappropriate?

Again, it appears that REAL teaching becomes more and more of a risk-taking adventure. You can't force everyone to behave the same as the teacher, but neither can you present all experience to students in order for them to sort out the useful from the useless. Instead, you sort of throw out knowledge and skills a little at a time, and deal with the (unpredictable) returns as students attempt to deal with, or avoid, what is presented. But as with parents, you should be willing to say "that's a 'good' or 'bad' way to do it; these are possible consequences, are these what you want? All education can't just be based on whether or not the student solves the problem or acquires the skill, can it?"

Re: good people on their own and bad people inside organizations. I don't know what else to say about this. I believe it's possible for an organization to have a systems concept such that "it" would work for everyone. I know there can exist people outside of such an organization who, one way or another have incorporated some or all of the principles of the organization into their lives. I also know that such an organization may have members for whom the concept doesn't seem to work, who give it a bad name, and can even hurt or kill others as a result of their self-destructive behaviors. Either way, one is not going to be convinced of the rightness or wrongness of said institution.

Regards, Joel

Date: Mon Jun 07, 1993 1:29 pm PST
Subject: Understanding

[From Rick Marken (930607.1100)] Tom Hancock --

> How could one explain (or model) understanding
>(or lack of it) in terms of PCT? Is it plausible that humans have
>an intrinsic reference standard for understanding?

I have to make some assumptions about what you mean by "understand" . In one sense a simple control model understands how to control the controlled variable; the understanding is built into the comparator and output functions that make it possible for the system to make its perception match its reference signal.

There is also a more "cognitive" sense of "understand". In this case, understand refers, not to the ability to control a variable, but to the ability to explain how it is controlled. So I might "understand" (first sense) throwing a ball, because I can throw it, but I might not "understand" throwing it (second sense) because I can't explain how it is controlled (the neural processes, the muscle forces generated, the physics of the flight of the ball, etc). I think this second sense of "understand" corresponds to control of higher level perceptual variables -- like the program and principle level. So both aspects of understanding are part of PCT; understand (1) is being able to control; understand (2) is knowing how this control happens. In both cases, understanding involves control. In the first case it's obvious; understanding is evidenced by the fact that variables can be controlled (you can throw the ball). In the second case understanding is evidenced by the ability to control the program and/or principle perceptions that satisfy the references for a "correct" analysis of throwing a ball.

Does this make any sense, Tom?

Best Rick

Date: Mon Jun 07, 1993 3:11 pm PST
Subject: Re: Gap between higher and lower levels

[Martin Taylor 930607 18:00]

I'm trying to get a handle on what has been going on in CSG-L and elsewhere since I went away. I promised myself not to post until next Wednesday (16th) since I have another major presentation (largely on PCT) to prepare for next Monday, and I won't be back till the 16th.

But I can't let this pass...
=====
(Bill Powers 930511.0800 MDT)

According to the current HPCT model, when you have isolated a particular one-dimensional variable as a variable controlled by some person, the implication is that you could put an electrode in that person's head in the right place and measure the magnitude of the perception that represents the state of that controlled variable. To encompass a complex perception you would have to find the degrees of freedom involved in it and through the Test discover the specific state of each dimension that is desired; then you would need an electrode for each dimension. If, from all these dimensions of the perception, the person is perceiving some unitary thing, then we would expect that at a higher level there would be a single perceptual signal that is a function of all the measured ones. An electrode at that level in the right place would measure the magnitude of that single perceptual signal.

This is quite unlike the "distributed" concept of perception that others have offered (Martin Taylor, for example). So I suppose that some day we will find out which model is right.

=====

If that's what you think, you must have quite misunderstood my side of the "distributed control" discussion. I always accepted the first paragraph of the above quote. The only point of possible argument was that there might NOT be a only one perceptual signal that took into it all of the effects of the dimensions inherent in the lower level perceptual signals. In general, I would assume that you would also not claim that in any one perceptual control level there is only one ECS that provides reference signals to a particular subset of ECSs at a lower level. For the most part, any perceptual signal at one level contributes to many at a higher level, and the reference signal at a low level is derived from the outputs of many higher level ECSs. So one must consider either the actions of a single ECS as "seen" from within that ECS, or the actions of several, as "seen" from outside, by a Theoretician, Engineer, or what have you. The "distributed perception" is in this outsider's view (but, as we have seen in recent Science articles, it can be seen by neurophysiologist outsiders quite directly).

I see no difference of models (in this, anyway). Do you?

Martin

PS. I still promise myself not to reply until June 16 at the earliest, but I probably won't be able to resist temptation.

Date: Tue Jun 08, 1993 7:09 am PST
 From: tbourbon
 EMS: INTERNET / MCI ID: 376-5414
 MBX: tbourbon@heart.med.uth.tmc.edu

TO: * Dag Forssell / MCI ID: 474-2580
 Subject: RE: PCT acceptance

In Message Fri, 28 May 1993 22:25:00 GMT,
 Dag Forssell < writes:

>Chuck and Tom: Can I count on your permission to quote you in this fashion?

Certainly!

> Do you have requirements on how you are identified?
 > Tom is a professor of psychology and researches neurology.

Your need to identify my position for your readers has finally forced me to think of how I define my professional identity. For 25 years I taught psychology, starting as assistant professor and eventually for 18 years as full professor. All of that occurred at a small state university.

What am I now? My official title is Research Instructor in Neurosurgery. (Not as "impressive" as Full Professor, but I am in a major medical school so I suppose there is a good trade off!) My work is closer to that in experimental human neurophysiology than in psychology. I am paid to conduct research on humans in a clinical setting.

Let me know if you can use that information to fashion a description for your literature.

Too bad I can't just say I am a perceptual control theorist and leave it at that.

I received your tape. For the past year, during our period of many moves, our daughter has retained custody of our VCR. I will try to locate one here and view the tape.

Best regards, Tom Bourbon
 Department of Neurosurgery
 University of Texas Houston Medical School Phone: 713-792-5760
 6431 Fannin, Suite 7.138 Fax: 713-794-5084
 Houston, TX 77030 USA tbourbon@heart.med.uth.tmc.edu

Date: Tue Jun 08, 1993 9:11 am PST
 Subject: draft 4

[from Joel Judd] Bill,

Excuse me for using the net but the server didn't want to deliver this to your address.

I like the emphasis placed on environment and a living control system's interaction with it. Every opportunity to make this clear should be taken. Just above the "PCT and Behavioral Theories..." section why not add a clause like the following to the end of the next to last paragraph:

"...disturbances in Fig. 1. This also reminds us of the primary reason that control rather than just response is required: the outputs of an organism are not the only influences on the controlled variable [NEW], because the organism is in constant interaction with its environment."

Re: the last section, I like the idea of mentioning things by field; in fact, I thought about sorting criticisms/misinterpretations by journal, but assume that this would be "unprofessional" or counterproductive(?).

Joel

Date: Tue Jun 08, 1993 9:35 am PST
 Subject: BBS Paper, Who am I?

[From Rick Marken (930608.0800)]

Bill Powers -- I'll try to get comments on the BBS draft to you this weekend (if the computer is up). My real quick reaction: I would cut the discussion of "what PCT is" to 1 paragraph. I would then go through the litany of misconceptions about PCT. This should include misconceptions about what PCT is about (not overt behavior but control), how it works and how it's done. The discussion should include explicit PCT perspectives on a number of "hot" topics like language (controlled perception of "meaning"), symbol grounding (Harnad's favorite -- symbols [perceptual signal values] are grounded by the perceptual function), chaotic attractors (reference states of controlled variables), etc. It should present explicit examples of the kind of data PCT is NOT interested in explaining (perhaps the correlations offered by Bandura, and some significance test results) and some examples of the kind of data PCT DOES explain (little man arm trajectories, Bourbon two-person cooperation data) and show why the latter data is valuable and , most important, WHERE YOU CAN GO WITH IT!! We must explicitly explain to these people what the PCT research program might look like. After all, the potential readers of this article only know how to make a living using the conventional methodology; it may be crap scientifically but it puts bread on the table. The article will be successful if there is a section describing one or two directions for PCT research based on the research AND MODELLING that has already been done.

Joel Judd (930607) --

> Do I really do a child more good by exposing
 >him or her to a number of "ways of being" in somewhat ambiguous terms
 >instead of saying "here's a reasonable concept for being based on the
 >following principles" which he or she may later decide is incomplete
 >or inappropriate?

There is a phenomem that is so reliable that I think it is worthy of consideration as a real fact for PCT; the phenomenon is the extraordinarily high correlation between the professed beliefs of parents and their children. I would guess that the correlation is on the order of .999 (world wide). There may be local regions where this correlation is substantially lower (for example, in regions of the western world where liberal education still survives). But world wide I think that my estimate of the correlation between parent and child beliefs is probably correct.

Now, Joel, you seem to think that you have the "right" set of religious beliefs for yourself (and, possibly for your children); but the data suggest (based on a non-PCT, causal model) that your beliefs are very likely to be the result (somehow) of your parents' beliefs. If you had been raised by muslim parents, you would VERY probably have had muslim beliefs and looked at the beliefs you now have as those of a misguided infidel. Do believe

you would have ended up with Mormon beliefs if you family were Muslim and living in rural Iran?

PCT would say that the high correlation between parent and child beliefs is not a causal one; it is a side effect of efforts by the child to control intrinsic variables and, given the environmental constraints (including the parents beliefs) , the result was systems that happened to control for some variables at levels similar to those of the parents. If this is true, then you should be able to at least IMAGINE yourself controlling your religious beliefs at a different level. So, give it a try, Joel. Try to imagine yourself controlling for something other than being Mormon -- something REAL different (I don't know what that is; how about Hindu or Athiest).

I imagine that you might find it difficult and unpleasant to do such imagining; after all, you really ARE a Mormon, right? That's because the setting for that reference comes from above the level of control at which you can become conscious (that's my hypothesis, anyway). If you could become conscious at the level above the one receiving the reference for Mormon you would be able to see Mormon as a choice rather than a requirement and you would see how, in fact, you could have "chosen" to become a Buddhist or an Atheist if that were what was required in the environment in which you were raised (to satisfy your intrinsic needs for love, affection and respect from your parent, perhaps).

There is nothing "wrong" with adopting any particular system concept, from the PCT perspective. What PCT can help you see (and, possibly experience) is that these system concepts (like all perceptions) are variables that can be maintained at different reference levels (obviously, look at all those religions out there). Your religion isn't "caused" from outside (by your parents or by god); it is caused by neural signals in your own brain which have selected that level of the system concept perception as the one that's "right" for you -- in your circumstances. PCT doesn't say you should select level X of reference signal Y; but it does suggest that you can see the selection of level X as something YOU DO in order to satisfy OTHER GOALS that are also PART OF YOU. And, if you can get your consciousness up to the level where you can see the setting of the system concept reference as an option, you can the CHANGE these references, if you need to, in order to solve internal conflicts that might be the result of the particular system concept references you've selected. It seems like the Mormon reference works for you, Joel, and that's great. But if you ever do develop some persistent internal conflicts, having the ability see the reference for Mormon as an option rather than a requirement could be quite therapeutic.

Best Rick

Date: Tue Jun 08, 1993 11:03 am PST
Subject: upward and outward

[from Joel Judd 930608] Rick (930608)

I've been trying to steer my posts away from a religious base and onto a more developmental-educational one. I should have made clear that the part of my last post you cite at the beginning was referring to current educational practices (granted there exist church-sponsored public schools but I wasn't referring to them). I was trying to formulate a question about development in general enough terms for anyone to answer. Several times lately the concept of a "well-functioning" control system has been mentioned; I want to determine how such a system comes to be. It has been accepted that certain controlled variables are more desirable than others; if so, how do we increase the likelihood that a developing system settles on such variables, for example.

BTW, I would guess that the correlation between PROFESSED beliefs of parents and children would be quite low; the correlation would be high between something more subtle, deeds or whatever. My wife and I (n=2) are quite contrasting in parental beliefs. I pretty much fall into your characteration. My wife became LDS, got married, had children has stayed married, etc., none of which her divorced parents did (except her mom how is a member of the church as well).

>...you would be able to see Mormon as a choice rather than a requirement...

I do see it as a choice, as well as a requirement. The difference is in how one deals with such a systems concept in terms of OTHERS. So I have been trying to educationalize the thread not only because I think we have exhausted the intellectual possibilities of religious examination, but because I don't want to be perceived as continually bringing religion up to insist on my systems concept as being "right." I think there are some

serious societal issues that could use some PCT "light," education being the one with which I am most familiar, and with which I am most interested.

Date: Tue Jun 08, 1993 12:46 pm PST
Subject: Inheritance of Acquired Beliefs

[From Rick Marken (930608.1200)] Joel Judd (930608) --

>It has been accepted that certain controlled variables
>are more desirable than others;

It has not been accepted by me. What makes one controlled variable more "desireable" than another is the context of other variables being controlled by the system. There is no perceptual variable that is "absolutely" more desirable than others; and there is no level of a perceptual variable that is absolutely the most desirable level -- at least, not in terms of the current version of the HPCT model. I'm sure this is getting to be like a broken record but you should really study the details of my spreadsheet hierarchy to see that I am not saying this in order to sell liberal propoganda or satanic relativitic morality. I'm takling about how the HPCT model really, quantitavely WORKS.

When I say "Relativism: it's not just a good idea, it's the law" I mean it. If you come to the CSG meeting this summer I'll show (with the spreadsheet model) exactly what I mean when I say that what constitutes a "desireable" controlled variable or a "desirable" reference setting for such a variable DEPENDS on the context of environmental disturbances and other variables that the system happens to be controlling.

>if so, how do we increase the likelihood
>that a developing system settles on such variables, for example.

That is NOT THE GOAL. The goal is to have systems that can CONTROL successfully and this means teaching them how to VARY perceptual goals in order to control other perceptual goals. We can't know IN ADVANCE what variables the developing system should "settle on". Deciding what variables the system SHOULD control (because YOU know they are the best) is precisely what is done in education NOW. If you really stick to the idea of increasing the liklihood that the developing system "settles on" such variable you are just trying to control the system -- and the results are as likely to be a successful control system as a barely functioning neurotic (that is, the result of education will be the same as they are now).

>BTW, I would guess that the correlation between PROFESSED beliefs of
>parents and children would be quite low;

Aren't most jewish children jewish, hindu children hindu, shinto children shinti, moslem children moslem, etc etc? I know there are exceptions. Most of the people I know (or care to know, really) are exceptions. But I'd guess that out of, say, 500,000,000 Hindus in India, probably AT LEAST 499,000,000 had Hindu parents. Out of 1,000,000,000 catholics in the world AT LEAST 999,000,000 had catholic parents. I bet if you calculate the Speaman rho on this you'll get a measure of association on the order of .999 or higher. If your guess (of a low correlation between PROFESSED beliefs of parents and children) were right, there would be all sorts of religions running around india, in about equal proportion; europe and america would have nearly as many jews, hindus, moslems and shinto as it does christians. I'd be happy to get the real data from the encyclopedia on this but, aside from the few people I know (and even several million is "few" in the contex of world population) who don't profess the beliefs of their parents, most people obviously do.

Best Rick

Date: Wed Jun 09, 1993 8:46 am PST
Subject: Re: Reply to Satan's messenger

[Martin Taylor 930609 11:10] (Rick Marken 930525.1400)

>>The natural state of the organism is seen to be one of
>>motionlessness or rest.

>

>This is factually false. The "natural stste" of the organism
>is seen to be one of controlling its input. The actions of the

>organism are continuously countering disturbances to controlled
>variables. So the organism is always precisely as active as are
>the disturbances to its controlled variables

It's an easy mistake to make. I did the same (mis)analysis of the consequences of HPCT early in my acquaintance with CSG-L. But your answer can go even further than you did, because even in the absence of disturbance the "natural" state will not be one of rest.

What are the real top-level references? I think they have to be the "intrinsic" variables, whether (as Bill P says) they are controlled only through reorganization of the main hierarchy, or (as I prefer to think) they are actually the references for top-level ECSs in the main hierarchy. Either way, the following argument works.

If the organism is at rest (unless it is a very particular kind of plant), it will not be acquiring energy from the environment. Since its chemical processes require energy to proceed (as well as other resources acquired from the environment), they will alter in some way, over time, and perhaps stop. Anyway, the levels of intrinsic chemical variables will alter, presumably acting as a disturbance to either the reorganizing system (Bill P) or top-level ECSs (me). If the disturbance is in the reorganizing system, the shape of the hierarchy changes, and for sure there will be ECSs whose perceptual signals now do not match their references, so the organism will "spring into action". If the disturbance is in a top-level ECS, that ECS will have an error signal, and thus output, which alters reference signals all the way down the hierarchy, so the organism will act.

Either way, the "natural" state of the organism is to continue to act so as (at minimum) to acquire energy from the environment. The alternative is death, which is, I suppose, the most "natural" state for any lump of organic matter, in the end.

=====

To follow this up, I have recently found that a very good way to get people to see the necessity of PCT is to make the argument from evolution. Any organism that tends to act in a way that has no detectable consequences for it is wasting resources, and an organism that is like it but refrains from those actions will have an evolutionary advantage. More descendants of the latter will tend to be found at a later date. If the organism tends to act so that the consequences of action are detectable but neutral, the same argument applies. If they are detectable but detrimental, the selection against it will be stronger. The organisms that will be seen are those whose actions most probably will have consequences that are detectable and beneficial. In other words, those that evolve negative feedback perceptual control systems are those that will remain after a few (thousand) generations. And the ones that will be preferred are those whose perceptual control systems use actions with the least wasteful side-effects (neutral or unobservable consequences).

This argument seems to hit home with people I have recently tried it on, especially the more senior psychologists. I haven't yet read the BBS draft, but I think such an argument should be in it.

Martin

(Now I will withdraw again until tempted or until June 16, whichever comes first).

Date: Wed Jun 09, 1993 11:26 am PST
Subject: Bill P, Tom B on words

[From: Bruce Nevin (Wed 930609 07:11:43 EDT)]

(Bill (930528.1930) Tom (930528) Bill (930528.0930) Tom (930529.0310)

I believe I see one confusion of mine that may have contributed to misunderstandings. I was thinking category detector recognizes an exemplar of the category by a match to its reference. I see that it does by a match to its input function. This is why I thought the reference for category recognition had to be derived from memory, and had to be stable rather than driven by higher-level error. Variations in the reference signal I now suppose correspond to how good a match is sufficient at the input function for the ECS to recognize that input as an exemplar of the category. (I have some difficulty distinguishing this functionally from variations in gain, but I will let that be for now. Maybe someone can clarify the distinction?)

The issue is still how people recognize repetition vs. contrast in their language. This is learned (they do it for words of their own language, not for the different phonemes and words of other languages). And two people communicating with one another in a given language will agree that two utterances are repetitions or that they are different words.

My expectation is that a linguist working to catch on to an unfamiliar language is analogous to a child working to catch on to the language of its parents. Let's consider what that analogy would look like.

Harris's test for contrast/repetition is a form of the Test for a Controlled Variable. Ask a native speaker to control a perception of "repetition" or "same word(s)". Produce two pronunciations (perhaps recorded) that are alike, except that in the second some portion is replaced by a similar portion of a different utterance (perhaps by splicing segments of recordings). The native speaker assents that they are repetitions, or dissents, saying that they are different words.

By such substitutions, we are able to create a linguistically relevant segmentation of utterances. The segments are relevant because they represent the contrasts between utterances and locate the points of contrast within utterances.

Analysis of the segments in terms of sound features or articulatory features and in terms of their distribution relative to one another in utterances can enable us to devise alternative ways of representing the contrasts by phonemic elements (successive segments, simultaneous components, combinations of these). There are always alternative ways of representing the contrasts. Some ways are more advantageous for writing, some are more advantageous for understanding and describing relations among words and parts of words in utterances, some are more advantageous for understanding relations between similar dialects and languages and for studying the changes of languages through time, and so on. But because each such alternative system is a representation of the contrasts between words in the language, they are interchangeable in respect to our observational primitives, the contrasts. In each case, substituting one element for another yields the representation of a different word. Conversely, any two words that native speakers perceive as different are represented differently from each other, no matter which system of representation is used.

These elements--call them phonemes, or distinctive features, or what you will--are derivative from the contrasts between utterances. Bill, you suggested (930528.1930 MDT) that contrast is a non-phenomenon. Easier to say that the phoneme is a non-phenomenon. One phonemic system may look quite different from another, but each is a representation of the contrasts and by that criterion they are interconvertable.

Just so, people speaking "the same" language may have developed different ways of representing the contrasts to themselves. These differences may have (minor) diverse consequences in their acoustical outputs, in the kinds of misinterpretations they make of others' pronunciations, in their accommodation to dialect differences, with consequences for language change over time, and so on. Some people pronounce American English /r/ with retroflex tongue tip, others with the tongue tip lowered and the blade raised behind the alveolar ridge, with lip rounding. I associate this (probably wrongly) with lisping. How could this come about? Let's look at the other side of the analogy, language learning by a child.

Children learn to control words first as events without controlling the phonemic distinctions that adults do. As they develop the ability to factor words into their points of contrast with one another, their active vocabulary growth suddenly explodes. There are some constraints as to which perceptions are easier to control or less subject to disturbance, but prima facie there is nothing (beyond the premature claims of Universal Grammar) that says one child's representation of the points of contrast should be the same as another's or the same as an adult's, so long as the word distinctions are preserved, and so long as the means used to distinguish different words have recognizably similar acoustic consequences from one individual to another.

You (Bill) have drawn a parallel to Land's theory of color perception. This theory proposes that distinctions in color categories are perceived as relative deviations from an average (in the ratios of long and short wavelengths) over the entire visual field. The colors are defined relative to the average, rather than in directly contrast to one another. This is possible because the differences between colors are defined in terms of one ratio. The colors are not being factored into characteristics that cross-classify them in intersecting ways, e.g.

	Red	Orange	Yellow	Green	Blue	Violet	...
property a	+	+	+	-	-	-	...
property b	+	-	+	+	-	-	...
property c	+	+	-	-	+	-	...
property d	+	+	-	+	-	+	...

Such a cross-classification is not only possible but natural for speech sounds, however. We can factor out perceived features of speech associated with the segments identified by the Test. We can then use these features as symbols to represent the phonemic distinctions in words. And these perceptual parameters do have the effect of cross-classifying the segments that they identify. A gesture of velum-lowering, and the noise characteristic of nasalization, is associated with m, n, and ng in English. (It is also associated with nearby segments when one of these three nasal stops occurs.)

Perhaps we do this because we must produce phonemically distinct pronunciations of words in a way that is recognizable to our fellows. This requirement doesn't arise for color perception. People don't have to produce a biological display that will be perceived as this color or that by others. Even if they did, the number of color categories is orders of magnitude fewer than the number of words.

The orthodox view in linguistics today is that only a small set of these perceptual parameters is available, universally. At first (in the 1940s and 1950s), it was assumed that they must be acoustic parameters, because of the social character of language. But it turns out that there are no reliable invariants in the acoustic signal. Some of the material I posted about Liberman and Mattingly's reworking of the "motor theory" of speech shows this. So the universal alphabet of distinctive features is now defined mostly in articulatory terms. But it is also difficult to identify invariant features in the articulatory data. Therefore, L&M argue that the relevant parameters are intended articulatory gestures. This is of course just the way we like to think of things in PCT.

But it seems doubtful that any closed set of intended gestures is universal. The TA for my undergraduate intro course in linguistics had worked with missionaries in Brazil. He told me of a Brazilian language in which bilabial closure (p) contrasted with a speech gesture in which the tongue was entirely protruded with the lips closed over the blade of the tongue and then released. The Ubykh language of the Caucasus has (or had: it became extinct last year) 81 consonants, and some African languages have more. It is true that there are parallels across languages, suggestive of universals. But it seems to me that these would follow from various constraints that are universal, rather than requiring that there be biologically innate feature-recognizers, some of which atrophy if they are not used in the language(s) that the child learns. The acoustic properties of the vocal tract are such that within the regions of the vocal tract associated with p t k there can be considerable error in the location of the occlusion with little or no discernable acoustic effect (similarly but less so with ch and back-velar [uvular] q). The fact that differences between gestures are used to maximize contrasts between words represents a constraint of another sort, in that the space of choice for gestures is defined by categorial polarization rather than continua. The appearance of universals is probably a byproduct of control under constraints of this sort, rather than a reflection of biologically innate feature recognizers. However, my opinion here (and this is all it is) could be wrong, and indeed there could be a combination of constraints with some feature-recognizers, e.g. for voice-onset time.

I am not clear where phonemes or phonemic distinctions fall in the perceptual hierarchy. Phoneme perception is generally taken to be categorial. Perhaps this is incorrect, and they are configuration perceptions that just happen to be multi-modal: sound feature and intended gesture linked in the perceptual input function of a single configuration detector. I have argued that, developmentally, feature/gesture detectors arise out of control of word perceptions. Once the control of features/gestures has developed, absence of perceptual input for a given feature/gesture element is compensated. One obvious mechanism would be for word-detectors to fill in the blanks by calling for the imagination connection in feature/gesture detectors that didn't get actual input. Something like this seems to be necessary, given the degeneracy of the sensory input, i.e. the unreliability of cues in the acoustic signal or in a trace of articulatory movements alluded to above. The appearance of controlling syllables may be a byproduct of control of gestures. However, it sure looks like control of sequences of gesture-categories like consonant and vowel and subcategories of consonants, going along in parallel.

As children start learning language, the words are few, and so the means to differentiate them can be much sparser than is required for a larger vocabulary. Toddler: "Mi-mi." (pointing to cup). Visitor: "`Me' ... Oh, that's your cup, is it?" Toddler: "Mi-mi."

(pointing to cup). Mother (passing through the room): "Miriam's cup. `Mi-mi' is her `Miriam'." Visitor: "Oh. `Miriam's cup.' Yes, that's your cup, Miriam." (This was an actual encounter about a week ago, when I wrote the example down.) Then as they acquire the contrasts they might not get the segments in the right order for some time. Where for us a word like "stand" or "snow" begins with s plus a consonant, my daughter Katrina as a toddler (about 25-36 months) pronounced these words with the s at the end of the syllable: "t/dance", "now-s", "newsman". The word "stand" contrasted with the pronunciation of "dance" for her, just as surely as for you and me--the unaspirated t or voiceless d that characteristically follows s in our speech was the initial sound for her. The perceptions that she controlled as means for controlling word contrast were apparently the same as those adults control, merely deployed in different sequence. (Observations like this are missing from the literature on child language acquisition because researchers are looking for stages of acquisition of universals predicted by their theory.) "Wrong" word shapes may persist lots longer. Katrina is now almost 6 and "except" for her is /Ovsept/ (O the open o of bomb contrasting with the a of balm and the o of bow), apparently analyzed as "of" plus unique word "sept". But the means to distinguish words recognizably to others have been in place since sometime in her third year (after 2nd birthday).

What we would hear as nonsense syllables "bava" a speaker of the Southern Paiute language hears just as unequivocally as "baba". There is no contrast between words whose only difference in pronunciation is that which we can represent by b and v. (Really, it's the bilabial fricative of Spanish caballo, but that difference makes no difference to us English speakers, we hear it as v.) If a word that begins with b is fricative of Spanish caballo, but that difference makes no difference to us English speakers, we hear it as v.) If a word that begins with b is combined with a prefix that ends in a vowel, the b is normally pronounced v between vowels, and in other words where b might occur between vowels we normally hear only v. With special emphasis, isolating the syllables, the Paiute speaker may in fact say ba-ba. The difference makes no difference between words for him.

When you object to the proposal that toddlers are able to achieve social agreements, you seem to be limiting the notion to verbally negotiated agreements. This is a recurrent problem. Either I need a different term for agreements nonverbally negotiated and effected, or you need to expand the possibilities for the word "agreement".

=====

Here's another pass over some of the same issues. If there is a contradiction with the above, the above is more recent.

You propose that the appearance of word-contrast, syllable-contrast, etc. is a byproduct of phoneme recognition. This seems to entail the further proposal that the process in children that results in phoneme recognition in adults is not dependent upon the perception of contrast. The second proposal concerns the process by which a child comes to have phoneme recognizers attuned to the adult language. The first proposal assumes that phoneme recognizers exist, attuned to the phonemes of the given language. (The number of phonemes, roughly 10<n<100, depends upon the language.)

Suppose that the second proposal is true. Then it follows that one should be able to determine the contrasts between words in a language by examining the acoustic record of utterances in the language, or by examining a trace of articulatory movements used in producing them.

Generations of linguists worked very hard and with great skill at the problem of representing unambiguously the appearance of word contrast in various languages. The aim is a set of graphical cues from which anyone could produce variable repetitions of words as a native speaker does, rather than imitations of the original pronunciations of which the graphical cues are a representation. (Of course, as in any science, "anyone" means anyone with appropriate training and practice. Training includes acquiring the ability to step outside the sound-categories of one's own language, the ability to interpret the graphical cues and accompanying descriptions of sounds and articulatory movements appropriately, etc.) One reason they worked very hard at this was that they assumed that a child learning the language had only the phonetic data to go on, equivalent to the linguist's acoustic record and articulatory trace. Another reason was that reliance on perceptions of the meanings of utterances is notoriously treacherous. The meaning-assignments of words differ between languages, and the meanings to which words might be assigned differ between cultures. It is perhaps difficult for a North American to appreciate this, we have grown so accustomed to all the world eagerly translating into English and mapping their perceptions into American categories because it is politically and economically advantageous for them to do so.

But it has turned out not to be possible to determine the phonemes of a language without reference to native speakers' word-perceptions. Harris showed that this can be done without knowing the meaning-associations in advance, or even where the word boundaries are, because all you have to ask the native speaker is whether two utterances are repetitions or not. By substituting a perceived feature within one utterance A for a similar feature within another utterance B, and determining whether the two renditions of B are perceived as repetitions or not, one can locate the contrasts between utterances at points within the utterances, and identify them with features of sound and articulation--the phonemic distinctions.

Exactly this, the child may do also, experimenting with different substitutions so as to factor word events into sequences of sound/gesture perceptions (categories? configurations?) that are the same across all words. Initially, word events can only be sequences of sounds/articulations, without categorization (imitation). The development of categories to represent the distinctions appears to be by successive approximation. Perhaps someone here knows more about the development of category perceptions.

So far, this concerns the process of learning. What about after the phonemes have been established as contrasting sound/gesture perceptions?

| (Bill Powers)

V=====

If we say that a phoneme-signal is a function of some set of lower-level variables, doesn't that say all that is necessary?
^=====

The rub here is that the acoustic signal and visual cues about the articulatory gestures are not sufficient. One must imagine by what gestures one might produce like sounds. "Like" here means not contrasting. The sound/gesture perception of phonemes is a representation of contrasts.

=====

This has been hanging around on my PC for too long. I'm not going to get a collected chunk of time to try to get more coherent. Let's try if this makes better sense to you.

Bruce Nevin bn@bbn.com

Date: Wed Jun 09, 1993 4:05 pm PST
Subject: Re: Off to see the wizard

[Martin Taylor 930619 18:30] (Rick Marken 930527.0830)

>I have met MANY scientific psychologists who were OPEN enough to
>have read BCP. As you can see from Locke's article, there are also
>many people who cite Powers as a reference. It's not the willingness
>to READ a book that I consider OPEN. It is the willingness to try to
>UNDERSTAND what is being communicated that is obviously missing
>from those of the Locke/Latham ilk.

I don't think you are entirely fair. I hope you will grant that I have been TRYING to understand over the last 2+ years of reading a megabyte per month of PCT talk, and of trying to develop and test out my own ideas. Maybe I don't understand yet fully (I don't feel I do), and I know I say things that make you think I don't understand PCT "properly", but I try.

A failure to understand does not necessarily imply a failure to try to understand. I think most of the comments that so annoy you (and me, too, nowadays) come from a belief that they DO understand, having tried. We think that they don't, but we do not control either their perceptions or their reference signals, so we cannot make them have the error signal that would correspond to their recognizing that they don't understand. They work from a perfectly satisfactory position: "I have read this work, and see what it implies. This is what it is, and here are its errors." What else can they do, if that's the way they perceive the world and their understanding of it?

If you were one of them, would you see a problem with this position? We learn only when we understand that we do not understand. When we think we do understand, we stop. Isn't that a rather familiar pattern, to a PCT-er? I sometimes think it would be nice if you understood information theory. Your comments in that area remind me a lot of comments like those of L&L on control theory. You don't see that, and neither do they see their "problem," because, like you, they have no problem.

Martin

Date: Thu Jun 10, 1993 11:21 am PST
Subject: Re: Off to see the wizard

[From Rick Marken (930610.0800)]

I said:

>It's not the willingness
>to READ a book that I consider OPEN. It is the willingness to try to
>UNDERSTAND what is being communicated that is obviously missing
>from those of the Locke/Latham ilk.

Martin Taylor (930619 18:30) replies:

>I don't think you are entirely fair. I hope you will grant that I have
>been TRYING to understand over the last 2+ years of reading a megabyte
>per month of PCT talk, and of trying to develop and test out my own ideas.

You are absolutely right. I was not being fair -- but I was not thinking of you (or anyone else on this net or who attends CSG meetings) when I made the comment. I am sure it is possible to get a good understanding of PCT with just a willingness to do so. But why not join with others (as you have, for example) who have been working for years doing the work that will help one achieve that understanding? And if one has disagreements about the model (as you do, for example -- and me too occasionally) why not hash them out with the people who are working on that model (like you do). I think it's pretty clear that science benefits from being a social enterprise and you are definitely a part of that enterprise.

The people I was thinking of when I made the comments above are people like Carver, Scheier, and several others (including Locke) who use the PCT model (or specifically object to it, as does Locke) as the basis for their research but don't bother consulting with those of us (mainly Bill P, of course, but there are many people in the Control Systems Group) who have been working on and thinking about what we now call PCT for YEARS. I don't understand why Carver or Scheier have never bothered to attend a CSG meeting or participated in net discussions. Of course, anybody can use the PCT model who wants to and they can do whatever they want with it, I suppose, including use a distorted version of it for personal financial gain. I am not interested in having those who use the PCT model hew to some PCT orthodoxy in their "preachings". I don't want to make them attend meetings or kneel down before my spreadsheet as I press the recalculate button. I'm just curious about why these people seem to avoid interaction with the person who developed the model (William T. Powers) and those of us who are working with him to extend it. It seems a bit like going off and promulgating relativity theory while avoiding any contact with ol' Albert or any of his associates doing relativity research. This might actually have worked in physics because most physicists understand the mathematics of the model. But with PCT we've got people going around talking authoritatively about how PCT works -- and these people don't know zilch about how the model works or even what it is designed to explain (purposeful behavior).

Anyway, let me repeat that I was most emphatically NOT referring to you in my comments above. In fact, you are precisely the OPPOSITE of the type I was thinking of -- not because you agree with everything Bill or Tom or I or others say about the PCT model but because you are REALLY working with the PCT phenomena (demos), the PCT math, the PCT models and the PCT people. You are honestly trying to understand and DO PCT science; the people I was thinking of are already convinced that they KNOW PCT so they don't need to talk with those of us who down here under the hood trying to figure out how it works.

Best Rick

Date: Thu Jun 10, 1993 12:25 pm PST

Subject: PCT research

From Richard Thurman (930610.1300) Rick Marken-

I'm curious about the 'control of sequence' vs 'control of configuration' study you were about to do last March. Anything to it yet?

Richard Thurman

Date: Fri Jun 11, 1993 5:38 am PST
Subject: Re: BBS4

From Tom Bourbon (930611.0750) Bill Powers (930605.1300)

>If you suggest changes, please send them to me by e-mail, copying
>just the part to be changed and then adding your substitution. I
>won't guarantee to go along with every suggestion, but will accept
>anything well-considered.

[* Bill's original appears without quotes, brackets or other
identifiers, while my comments are inside the *]

Control theory has been knocking at the doors of the behavioral and life sciences for over 40 years. A few doors have opened; most have not. There is no single clear reason for the difficulties in assimilating this "new" idea into the mainstreams of science; on the surface, the reasons seem to shift with every different application. Perhaps one of the most illuminating objections to it came from Edwin Locke (199?), who proposed that because control theory seems to be used to explain a wide variety of phenomena now explained in terms of different theories, it amounts to nothing more than a series of borrowings. The fact that control theory has been in existence much longer than most of the theories from which it supposedly borrows casts some doubt on Locke's interpretation, but Locke has put his finger on one major problem. It is not the details of control theory that cause the difficulties, but the fact that control theory offers explanations of a kind with which many behavioral scientists are unfamiliar. The conclusions that can be drawn from control theory often fit generalizations that others have drawn intuitively from their data. But control theory should not be evaluated solely at the level of conclusions; its real heart is at the next level down, the level of underlying regularities and relationships from which conclusions can be deduced and observations can be explained. It would be surprising indeed if a valid theory of the organization of behavior led to conclusions contrary to competent observations.

[* Bill, on this point, Locke is RIGHT. I have posted about this on csg-1 a few times, but no one else picked up the thread. Locke focused, understandably, on CT or CST as it is presented in the bastardized, but popular, literature. There, he sees C & S, Hyland, and others, all going out of their ways to show that CT- CST is "just like everything else" and that it just offers another "perspective-viewpoint-framework- etc." They say that and show that. It is the key to their success at publishing. It is what psychologists are accustomed to and what they will accept. All of those people borrow from and cozy up to contemporary "theories" and in that regard Locke is correct. None of the popularizers, or their detractors like Locke and Bandura, have an inkling of what hard-science PCT is about or of what modeling means, but the detractors are right when they say bastardized CT-CST merely borrows and restates. That being the case, I am afraid the carefully stated points in your draft will go right past both sides of the debate over CT-CST, but that is no reason to remain quiet. *]

..

Consider the old standby, the thermostat. The environmental variable that is controlled by a thermoregulator is, obviously, temperature.

[* Your discussion of the thermostat-furnace-room-universe system is the clearest yet. Along the way, you finesse the gambit in which people trot out the idea of "complex control systems" that monitor, and have a "plan" for, *every* relevant variable that affects the output. It is hard to imagine a rational person who would miss the point of your remarks about the fact of control and its essential features. My only reservation concerning this section is that it is *too well-reasoned*. It is lucid and thorough, which means it is

necessarily deliberate and long. If remarks by our previous reviewers and editors, including professor Harnad, are any clue, work like that does not go down easily in modern psychology. This time might be different, given that the article will contain data. The editor ought not be able to repeat his claim that PCT (CST, in previous submissions) is "just another way of talking about things we already know."

It would break the nice flow of your presentation, but maybe somewhere up front there should be one of those brief statements that, "this is what I am going to tell you -- control, PCT methods, major misconceptions about PCT, examples of how PCT methods refute the misconceptions, etc. *]

..

The methodology of PCT

..

The control-system model

..

Engineering models and PCT

..

PCT and behavioral theories: general observations

[* Your logic and presentation in those sections is excellent. Once again, that is a source of my doubt about the manuscript ever going for the first round of review at BBS. I have read (more recently, glanced at) BBS from the first issue. I do not recall one target article that explained a theory in detail and presented clear, unambiguous examples, using as much space as necessary to assure that readers understood. I doubt the deficiency continued for all these years simply because no one from among the host of authors ever thought of doing things the right way. *]

Misinterpretations and misapprehensions

[* Rick made some good suggestions about how to present the agenda for behavioral research, as PCTers see it. The following are some of my favorite misunderstandings, misrepresentations, and misapprehensions (the MISes) along with some of the demonstrations and figures I used to counter them while I was still teaching. Whenever possible, I used modeling to make my case. Several of the refutations depend on work by Bill, Rick and Greg.

MIS A. "Commands (which I use here to also mean feedforward, plans, programs, heredity, genes and similar directives) determine behavior, therefore there is no need for a negative feedback model." The cognitive model in "Worlds" refutes this possibility. ARM demonstrates there is no need for commands. Rick's blindmen also address this. One of the significant advantages of using modeling to evaluate various popular "theories" and "models" of behavior is that it is easy to see that nearly all of them are variations on a single simple theme:

Behavior = a function of Antecedent(s).

Modeling quickly reveals the fallacy of such models, but the idea of linear causality for behavior is strong, as is seen in the next item.

MIS B. A corollary to A is, "Since commands determine behavior, if "sensory feedback" is eliminated, then behavior will remain as it was before the elimination, therefore, there is no need for a negative feedback model." Again, the cognitive model in "Worlds" refutes this possibility, as do Rick's blindmen. So, too, do many accurate observations and interpretations of studies with deafferented animals. (References are available.) And so does a simple tracking task, when half way through the run, the experimenter blanks the screen, or covers it with a piece of cardboard. Only the zealots hold to this oldie.

MIS C. "Negative feedback models cannot represent living things. Those models go to equilibrium then stop behaving." A simple demonstration of a PCT model for tracking with a disturbance present refutes this claim, as does a model with a changing or variable reference signal.

MIS D. There are many claims that positive feedback must be added as a feature of any model of behavior. Most of those claims seem to spring from one of two bases: (a) a belief that negative feedback produces static equilibrium and that positive feedback is necessary to create changes in behavior (this is similar to the argument in C); or (b) negative feedback is somehow less aesthetically pleasing than positive feedback, or is uglier.

I have no ready reply to (b), but for (a), the demonstrations used in C refute the claim. Also, it is easy to demonstrate what happens when a negative feedback model goes positive. I run a routine tracking task and the accompanying reconstruction of the run by a simple PCT model in which handle position is computed by:

```
New Handle := Old Handle - (k * Error).
```

I let the model's results appear on the screen, using the same time base as was used for the person. Of course, the model reconstructs the run to the expected .99 level and the model's run looks like one by a person. Then I start a second reconstruction by the model. Everything looks fine, for the first few seconds, then the algebraic sign of the term for changing the position of the handle goes from negative to positive:

```
New Handle := Old Handle + (k * Error).
```

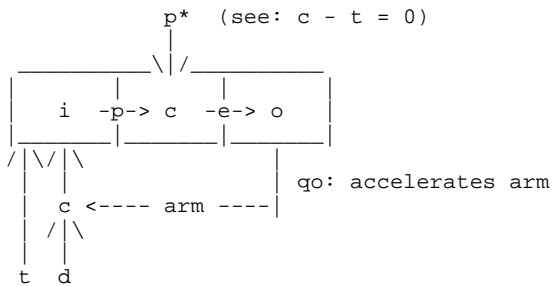
At that moment, the model's behavior "blows up," with the cursor and a simulation of the model's handle position going off of the screen in a fraction of a second. After that, an occasional blip appears on screen when the handle or cursor happen to move into range. Everything looks chaotic, in the popular sense of the word -- random, unpredictable and out of control. So much for positive feedback.

But there is more. Careful inspection of the blips after positive feedback is introduced reveals that the model's behavior might actually be chaotic in the technical sense of the word -- seemingly without pattern, but actually patterned. If that is the case, then an unexpected (by me) result of using the PCT model to examine the effects of positive feedback is a realization that chaotic behavior might result from positive feedback in a system designed to control by negative feedback. (Now *there* is a trendy topic if I ever saw one!) I have not done the necessary plots, but I suspect this simple demonstration can generate data that duplicate some of the classic (pretty) pictures in the (very popular) literature on chaos.

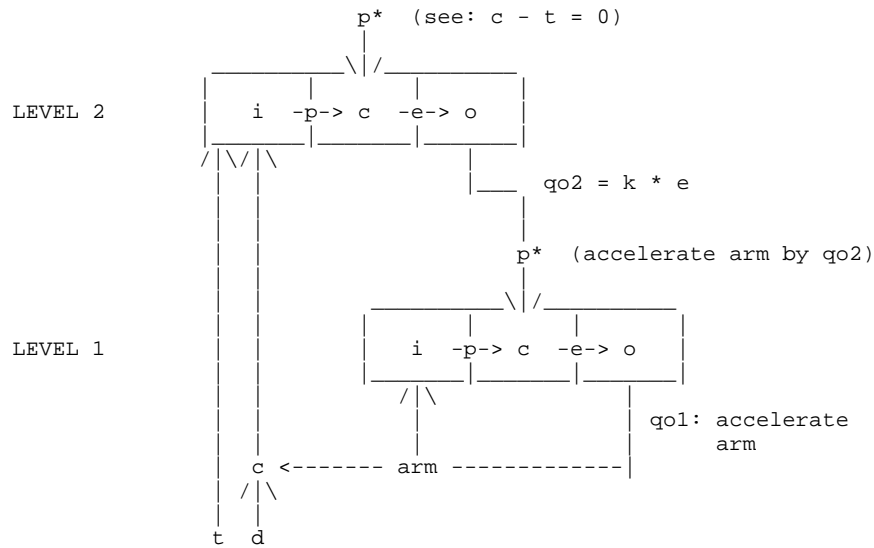
I used this demonstration in classes for about six years. It always seemed to get the point across.

MIS D. "Social systems (sometimes the topic is social hierarchies) are really PCT systems, with a different person playing the role of each function in a PCT loop." Rick's spreadsheet model, ARM, GATHER, and my interactive tasks all offer evidence against this claim. With my classes, I also found it useful to trot out a series of diagrams resembling the ones below. I let students work through them in sequence. Whenever possible, I moved to the next one only when they understood the present one.

+++++
CASE 1: simplest possible PCT model for a person keeping a disturbed cursor even with a target.

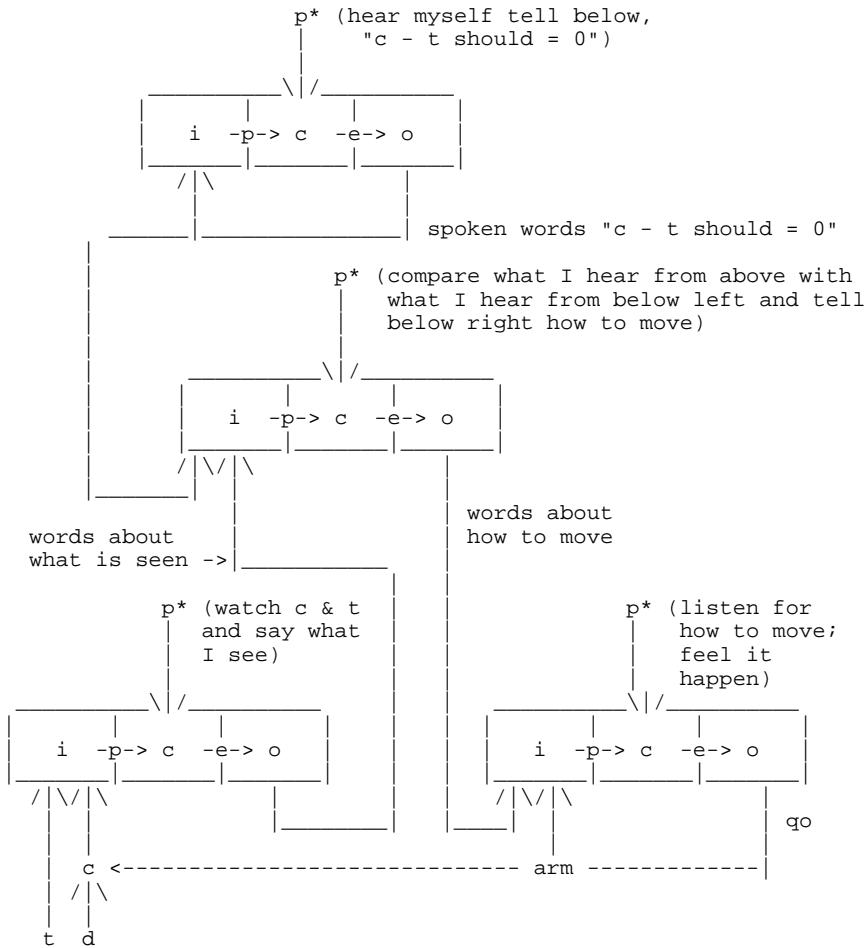


+++++
CASE 2: Minimal two-level hierarchical PCT model for the same pseson modeled in CASE 1. Output from the top level (visual relationship) IS the reference signal for the bottom level (arm movement).



+++++

CASE 3: Approximation of minimal connections between independent PCT models that would satisfy the popular belief that social hierarchies are really PCT systems, with a different person playing the role of each function in a PCT loop, which in this example is the minimal PCT model shown in CASE 1. A separate system represents each of p*, c, i, and o. Output from each system must be turned into words spoken to the appropriate other person; unlike the situation in CASE 2, no person has direct access to the reference signal of any other person. (A similar diagram for CASE 2 would be even more complex, with each function in each model in CASE 2 represented by a different PCT model, each of which should be hierarchical.)



Discussion of these figures can go in a number of directions. I typically allowed the direction to emerge naturally out of the spontaneous remarks and questions that always accompanied a presentation of the diagrams. Most people in my upper-level undergraduate classes got it; so did many in my freshman classes. Of course, for BBS, or elsewhere, we would pick the direction we prefer. (Do you think BBS would come up to at least the level of my former freshmen?)

Oh, I forgot. *]
 Until later, Tom Bourbon

Date: Fri Jun 11, 1993 6:50 am PST
 Subject: FLYWHEEL - RKC
 [From Bob Clark (930611.1015 EDT)] Bill Powers (930527.1100 MDT)

I am intrigued by your reference to the "flywheel effect."
 I don't remember using this term -- today my first reaction is that I would call it "momentum." However, that is a linear concept referring to a sequence of events where an object continues to move in the same direction. And I am reminded of your recent use of the term, "impetus" as related to Newton's Laws. Taking a second look, I suggest that "angular momentum" may be a better analogy for your purposes, since the interactions within such a group are likely to include a good bit of "recirculation." This does resemble the action of a "flywheel" as it uses the momentum (and energy) from one explosion within a cylinder to produce the effects needed to prepare for the next explosion.

This now seems to me a bit more apt than my initial interpretation, above.

This "effect" may apply to two situations: first, a "corrective" effect such as you note, tending to keep members of the group on the same path, with fewer straying from the "revealed word," and, second, making it more difficult for members to "understand" (ie, "fit into their existing intellectual framework") new proposals to modify, extend and apply the developing ideas.

I am intrigued, again, Bill, that you end, addressing Joel Judd, with:

>Thank God for the flywheel of the CSG, and the scientific world in
>general. One does need something to thank for a narrow escape from
>self-delusion. It is, you know, very tempting. I sometimes wonder
>whether Christ actually succeeded in resisting it, during his 40
>days, or whether it was only his followers who succumbed.

That is consistent with your Subject, "Revelations," but I wonder how you would express this sequence and combination of concepts in HPCT terms.

Regards, Bob Clark

PS -- Personal Note

You remark:

>It would be interesting to know how Bob Clark saw this young man in
>the throes of revelation; I'm sure there were times when it was very
>irritating.

Bill, I never thought of you as a "junior," "a young man," in "throes" or not. I found your ideas very stimulating and interesting -- and I hope you found some of my suggestions appropriate. I always felt that we were co-equals in every way in our struggles to work out these concepts and, later, present them to others. After all, I also thought I was "throeing" revelations. Of course, from an administrative standpoint ("viewpoints," also entered the picture back then) there were differences of several kinds.

These were some of the most fascinating and enjoyable times of my life. It is always enjoyable to be working in close cooperation with other people and feeling a unity of experience and comprehension. We had our differences and arguments, but managed to find ways to resolve our conflicts without even approaching violence.

My best, again, to you, Bill, and the Net, RKC

Date: Fri Jun 11, 1993 6:57 am PST
Subject: Re: Off to see the wizard

[Martin Taylor 930611 10:30] (Rick Marken 930610.0800)

>I said:

>

>>It's not the willingness

>>to READ a book that I consider OPEN. It is the willingness to try to

>>UNDERSTAND what is being communicated that is obviously missing

>>from those of the Locke/Latham ilk.

>

>Martin Taylor (930619 18:30) replies:

>

>>I don't think you are entirely fair. ...

>

>... I was not being fair -- but I was not thinking

>of you (or anyone else on this net or who attends CSG meetings) when

>I made the comment. ...

>...

>

>The people I was thinking of when I made the comments above are

>people like Carver, Scheier, and several others (including Locke) who

>use the PCT model (or specifically object to it, as does Locke) as

>the basis for their research but don't bother consulting with those

>of us (mainly Bill P, of course, but there are many people in the Control
>Systems Group) who have been working on and thinking about what
>we now call PCT for YEARS.

Thanks for the vote of confidence, but I really was thinking of THOSE GUYS when I made my comment about "what else can they do." I hadn't taken your comments personally. I referred to myself because I figured that it would help in making the distinction between "trying to understand" and "understanding."

If the "bad guys" in question think they do understand PCT, they have no error signal that would lead them to correct a perception of non-understanding. From their point of view, thinking they understand, why would they not present criticisms? The fact that from your point of view they don't understand is your problem, not theirs.

Granted, they might benefit from attending CSG meetings and participating on the net, but is the desire that they do so any different from the desire of the Mormon missionary to convert the lost souls who clearly don't understand the truth? They presumably think that the discussions on the net and at the meeting will be as unprofitable as the Mormon would presumably think a session of bathing in the Ganges to be. Why bother?

Martin

Date: Fri Jun 11, 1993 7:47 am PST
Subject: Re: PCT Research

[From RICK Marken (930611.0800)] Richard Thurman (930610.1300) --

>I'm curious about the 'control of sequence' vs 'control of
>configuration' study you were about to do last March. Anything
>to it yet?

I haven't worked on it at all. My current priority is trying to write a book about PCT; something along the lines of "Perceptual Park" with huge, cloned control systems running around terrifying visiting behaviorists and cognitive scientists. Actually, I am trying to write a "lay" introduction to PCT and this has been the focus of my (meager) brain power for the last month or so.

I wish you would come to the CSG meeting in July. I'll bring the program that is sort of the prototype of the kind of experiment I had in mind. I think we could make much better progress if we could sit down together and brainstorm about possible ways to do the experiment -- and possibly test some ideas right there using HyperCard or Basic or Pascal. I really would like it if you (and Tim Hancock?) could take the initiative on this research; you guys are a lot smarter than I am anyway and I will probably be preoccupied with writing (as my moonlighting activity) for the next six months or so.

If you (or Tom H.) can't make the meeting, I'm certainly willing to try to brainstorm on the net about possible research projects. Some of the "scholarship" on the configuration/sequence control stuff already exists in my "Hierarchical control of perception" paper; did you get a copy?

Best Rick

Date: Fri Jun 11, 1993 8:47 am PST
Subject: Re: Off to see the wizard

[From Rick Marken (930611.0900)] Martin Taylor (930611 10:30) --

>Granted, they might benefit from attending CSG meetings and participating
>on the net, but is the desire that they do so any different from the
>desire of the Mormon missionary to convert the lost souls who clearly
>don't understand the truth?

Yes; it is different, I think.

The people we are talking about (Carver, Scheier, et al) don't need to be "missionized"; they already believe in the Gospel According to PCT and claim to be PCTers (or, at least, Powersian Control Theorists; hey, another meaning of PCT). While I do desire that they

come to CSG meetings and participate on the net, I desire it, not because I want them to join the fold (they are ostensibly already in the fold); I just want to engage in an exchange of ideas with colleagues. Since these people are already converts, I am puzzled by their reluctance to come to "church". I don't want to force them to do it; I just don't understand the reluctance. It is really kind of a strange thing. The analogy seems to me to be that Carver and Scheier et al are like self-converted Mormons who are enthusiastic about being Mormon's but refuse to go to the Mormon temple. In fact, they continue to schmooze at their old church -- Our Lady of Lineal Causality. It's just puzzling to me, that's all.

>They presumably think that the discussions
>on the net and at the meeting will be as unprofitable as the Mormon would
>presumably think a session of bathing in the Ganges to be. Why bother?

I agree. So, again, it is puzzling. They should know that the CSG meeting is filled with fellow Mormons (er... PCTers) and that no Ganges bathing is done at all -- yet, indeed, they treat their fellows (us) as though we were the infidels. It's just puzzling. If they don't really like PCT then why don't they just use a different model? Why do they use the Powers model (at least, the terminology -- reference signals, perceptual signals, error, output, disturbance, etc) and treat the CSG (dedicated to testing Powers' model) as though it were anathema. It's just wierd -- understandable, perhaps, but wierd.

Best Rick

Date: Fri Jun 11, 1993 8:51 am PST
Subject: Joint paper; July's meeting

[From Bill Powers (930611.0915 MDT)]

Hi, folks. Western Colorado has been off the net for a few days (something blew up in Grand Junction), and Fort Lewis College picked the same time frame to change some hardware, so I've been off for a while. I've been working on the joint paper -- many thanks to all for excellent suggestions. But there will be more work to do. I combined JP2 and JP3, joint paper version 2 and version 3, (which latter I confusingly referred to as Draft 4 -- let's call it version 3), writing in most of the suggestions and adding new material, to produce a total of about 30 pages. When I read the pages yesterday they had suddenly become a long, pendantic, and excruciatingly boring paper that will certainly put any external audience to sleep. Somebody has evidently been tampering with my computer during the night. It just will not do.

I'm not going to be very responsive for a while, as I try to find the right introduction and tone, and get some notion of organization into the paper. Dennis Delprato suggested some history, and I think that may be a good way to start, to show that PCT didn't come out of nowhere. But the main emphasis should be on misunderstandings and misrepresentations, which really are the crux of the article. My problem with them is that I get so furious reading the stupid things that are done to PCT in the literature that when I write about them I become super-dispassionate and write like a long yawn.

I think it would be helpful if various of you would take on the job of pinning down specific citations, with quotes and page references, and laying out exactly what is wrong with them. I'll put in a general note that in many cases not EVERYTHING a person says is wrong, but that we're discussing just the major problems. And of course I'll edit and add my own ideas. So you just need to pick out the things that have been said in print that would be misleading to a reader who remembered nothing else, and think of a succinct way to say exactly what is wrong. Schmidt, for example, makes some really wrong statements about the effects of delays, and then later on hedges by citing some contrary evidence without really saying that the first statements were wrong. A reader looking for material with which to shoot down PCT will remember only the wrongest statements and forget the hedges. So we need only to deal with the worst statements.

One thing that would be very helpful would be to come up with a list of the major players in this game. When we analyze errors, they should be errors made by influential people, not simply others who might cite them. Avery Andrews, you seem to have some material on this. Tom Bourbon, Rick Marken -- who else? Some time ago, Greg Williams, you indicated that you might be willing to track down chains of references to see where the bad ideas came from -- this could be really useful.

It occurs to me that we may be writing a book, not an article for BBS. That would certainly expand the possibilities. But let's go on as before for a while. I'll post the results of my attempt at a new intro soon.

Best, Bill P.

=====
P.S. We now have 16 people registered for the CSG meeting. It may be necessary to cancel the banquet, as we must pay for 40 dinners regardless of the actual numbers. We may have to settle for a smaller room than otherwise. It's really too bad that people like Tom Bourbon, Wayne Hershberger, Kent McClelland, Clark McPhail, Dan Miller, Hugh Petrie, Dick Robertson, and Greg Williams will, judging from the lack of registrations from them, not be attending. Or does PCT really mean Postpone Commitments Totally?

Date: Fri Jun 11, 1993 10:27 am PST
From: CHARLES W. TUCKER
Subject: Look OK to me - thanks

Dear Dag,

I started SS this week so have not had time to answer your question. Your statements look fine to me; I don't think they are "sour grapes." Thanks for the file.

Regards, Chuck

Date: Fri Jun 11, 1993 1:12 pm PST
Subject: wierdness, flywheel

[From: Bruce Nevin (Fri 930611 15:43:51 EDT)] Rick Marken (930611.0900)

> It's just puzzling. If they don't really like PCT then why don't
> they just use a different model? Why do they use the Powers model (at
> least, the terminology -- reference signals, perceptual signals, error,
> output, disturbance, etc) and treat the CSG (dedicated to testing Powers'
> model) as though it were anathema. It's just wierd -- understandable,
> perhaps, but wierd.

Functionally, it is as though they are enacting a particularly devious form of resistance: cooptation.

They win acceptance and maybe even some endorsement by "mainstream" folks by "normalizing" a potentially troublesome alternative paradigm, reducing it to a repackaging in engineer-o-babble of what the "serious researchers in the field" really do.

Reminds me of the manufacturing of consensus in the field of linguistics.

This predicts that Carver&Scheier et al. will resist your blowing their cover, something that you say you are reluctant to do because you don't want to alienate them. It predicts that they will avoid circumstances where they might be compelled to pay attention to discrepancies between PCT and their rendition of CT. Makes it less puzzling. Parasitic relationship.

Too cynical?

=====
Flywheel, spiritual discernment:

Greg, you might be interested in Beyond Majority Rule by Michael J. Sheeran (I think that's the correct name), a Jesuit who studied the Philadelphia Yearly Meeting of the Religious Society of Friends (Quakers) and Quaker history and organization in general. The relation of individual "leadings" to the corporate process of seeking unity (not consensus) is a most fascinating one, and he (Sheeran) believes that George Fox and the Quakers did succeed in re-establishing practices in spiritual discernment by these means that were prevalent in the primitive Christian church, and that the Jesuits needed very much to re-institute for themselves. I believe he goes into the history of James Naylor and other early Friends in a most insightful way. He does lay out very clearly the process of Quaker meeting for worship for the conduct of business, meetings for clearness, etc. The book is published I think by Philadelphia Yearly Meeting, and is or was available from the Friends'

Book Store (or maybe Friends General Conference bookstore) on Race St. in Philadelphia. I suppose BIP would list it. I'll look in my copy at home if you'r interested and can't find it.

Off to see my family ... have a great weekend, everyone!

Bruce Nevin bn@bbn.com

Date: Fri Jun 11, 1993 3:56 pm PST
Subject: PEOPLE & BIKES - RKC

[From Bob Clark (930611.1700 EDT)] Bill Powers (930529.0800 MDT)

Regarding "ordinary people," you point out:

>But people's mixing PCT concepts with others often shows that they
>are unaware of the conflicts involved.

True -- but perhaps we could (should?) undertake the resolution of such conflicts. As you know, recognition of the existence of conflicts can lead to ("motivate" ?) increased interest in ("attention to" ?) suggested alternatives. Perhaps calling the attention of S-R people to their own every-day operations might get increased willingness to seriously examine PCT.

BIKES -- AGAIN!

Observing your descriptions of "keeping the bike upright," and the system controlling "the angle of the handlebars," it seems to me you are examining the detailed operations of some of the skills required for successful bike riding. I have no particular quarrel with your descriptions -- but that is not what I was talking about.

You remark:

>Output variables must be adjusted to keep the bike upright, but the
>control system concerned with uprightness does not perceive (or pay
>attention to) them. My preference is to stress the way control
>systems directly observe and control the outcome, not the means by
>which the outcome is achieved.

The control system with which I am concerned is the one that decided to ride the bike "somewhere." Its attention is directed, from time to time, to each of many different aspects of the sequences of events involved in progressing toward that objective. Among these is, occasionally, the "uprightness" of the bike. Almost all the time, the "uprightness" systems operate on "automatic," but an occasional stone, bump, or other "disturbance" calls for readjustment of the "uprightness" (&/or lower) reference levels. Likewise for the other muscle skills needed for successful bike riding.

Further:

>I don't like to characterize the operation of a system like this in
>cognitive terms, because that obscures the immediacy and
>automaticity of the operation of the control loop.

To what extent the operation of the control loop is "immediate and automatic" depends on which loop is considered, and on what time scale. On a time scale appropriate for "riding the bike 'somewhere,'" the bike-riding skills are, indeed, pretty much immediate and automatic. And the rider's attention can "wander" among other interests and concerns. (Left-over problems, the pleasant sunshine, anticipation of the planned visit, etc.)

For the "keeping the bike upright" loop, a much shorter time scale is more appropriate. It is not clear to me why you focus on a single bike-riding skill (selected from several), and don't see that I am trying to place these skills within a larger context based on a longer time scale.

It is among these higher level systems that I find the process of extrapolation from memory leading to selection and adjustment of lower level skills.

If you choose to call this process "reasoning," I guess I don't object. Usually I think of "reasoning" as involving more than decision making by simple correlation described here. If this is "cognitive," ["cognition, n. 1. the act or process of knowing; perception. cognitive, adj."], I'm not concerned.

Later you note:

>My own preference is to fit the description of the control process
>to the level of abstraction, and the degree of awareness that is
>probably involved.

That is pretty much what I thought I was doing, but I had in mind a higher "level of abstraction" than the systems making up the rider's bike-riding skills.

Continuing:

>I see no "events" in keeping a bicycle upright, and surely to
>analyze a control system as a sequence of events is incorrect!

The "events" I had in mind were those changing perceptions as the bike moves along the street. Surely these are "events?" And, correspondingly, the handlebars and other items are adjusted (compensating for minor disturbances) without requiring the rider's attention. But larger scale events, approaching a patch of gravel (a possible skid and spill), or traffic requiring coordination in time and position are also significant, and may require conscious evaluation.

Referring to my remarks about "associated memories being brought to mind, selected and applied as needed," you ask:

>And while this is occurring, what is the error signal doing? What
>output is occurring? etc

The higher order system is operating on a slower time scale. The lower order skill systems to which you are referring above, continue to operate on their faster time scale. They operate in terms of the continuing values of their reference levels, error signals etc. After the higher order system determines what adjustments, if any, are applied, the lower order reference levels are adjusted and they continue to operate accordingly.

A further remark:

>I just feel that you're trying to pack into a single control process
>operations that are better modelled as many control systems
>operating in parallel and many levels of control,

That is close to what I'm trying to do: I'm trying to use the rider's viewpoint in my description of the ongoing sequence of events (namely, the bike travelling down the street, the rider perceiving his surroundings, thinking his thoughts, remembering his memories, using his skills, etc etc). Most of the time he is not much concerned with the details of the bike's operation. (It's a bit late in the discussion, perhaps, but I am also assuming the rider is not a beginner.)

"Interpersonal Relationships" (also from your 930529.0800 post) calls for a separate post.

Regards, Bob Clark

Date: Fri Jun 11, 1993 4:01 pm PST
Subject: Information Theory (again)

[Allan Randall (930611.1700)]

Hi everyone,

My apologies to anyone who has been waiting for a response from me on the information theory thread - I've been taking advantage of Martin's absence by stealing his computer, so I've been busy with other things. I hope there are still some who are interested in keeping this discussion alive, even if at a slower pace than before. I certainly don't consider the issue settled. Excuse the length of this article. I'm catching up, so this is more like three articles rolled into one.

First, a note to Rick Marken:

Rick, I am still willing to attempt a disturbance-reconstruction demonstration. However, I am waiting to receive from you a simple control system for which such a demonstration would be convincing. You have already stated that the simple examples from the Primer that we were using are insufficient (due to lack of nonlinearities). I am not going to do a demo, only to be told it isn't nonlinear enough. I need an example of the simplest control system for which a demonstration would satisfy you (in the form of a computer program of some sort, ideally, but a description of the system will suffice if you're pressed for time).

There may still be some misunderstanding about this reconstruction issue. In one of your postings (930408.2100) you said:

```
>In the quoted section above it sounds VERY MUCH
>like Allan is saying that, if you CANNOT reconstruct D from P given
>the assumed conditions (r, O() and C), then one could "assert that
>there is no information about the disturbance in the percept".
>Since we have shown that, indeed, you cannot reconstruct the disturbance
>from the percept under these conditions then I think it is fair to assert
>that THERE IS NO INFORMATION ABOUT THE DISTURBANCE IN THE PERCEPT.
```

When did you show this? You showed that it was not possible to perform the reconstruction? I have no memory of such a demonstration from you. Please refresh my memory. I remember you showing a lack of correlation, but do not remember any proof that this implies non-reconstructability. Remember, once again, that correlation is not information. Period.

And to clarify, yes, I was saying that there would be no information present if it was not possible to reconstruct the disturbance with P and the stated legal variables WITH FEWER ADDITIONAL BITS than would be required without P (note that's FEWER additional bits, not NO additional bits, which is what you were asking for in your challenge for Martin and myself to do the reconstruction blind).

Rick, in spite of our misunderstandings, I do think that we have some kind of an understanding. You have agreed to the list of rules I gave for the disturbance reconstruction - about what is "legal" and "illegal" to use for free in the reconstruction. If you are right, I should be unable to write a reconstruction program with P as an input parameter that is shorter than the one without P for input.

All I need is an example of a control system that will satisfy you, and we can settle the issue, right? I'm serious about this. If we are reasonable people, we should be able to settle this matter, since we have agreed on the terms. In spite of your comments about IT people believing whatever they want no matter what, I am perfectly willing to admit I am wrong - and since we have a doable experiment for which we have already agreed on the terms, it would be a shame not to finish this.

Bill Powers (930418.1500 MDT)

Bill, I have a real problem with this whole $\log(D/r)$ thing. I'm not sure where you got this definition, but I do not think it can be right. The way I see you using it, it sounds more like a measure of channel capacity than information. This is unfortunate, because it has lead you to conclude all kinds of erroneous things about information theory that might actually be perfectly correct if $\log(D/r)$ actually *was* a proper definition of information. For example, you say:

```
>In order for two signals of different amplitudes to have the same
>information content, they must be exact scale replicas. And this
>scaling must include not only the amplitude range D, but the size
>of the minimum step in amplitude, r. To preserve the information
>content,  $\log(D/r)$ , BOTH D AND r MUST BE SCALED BY THE SAME
>AMOUNT. This may provide the missing piece in the puzzle.
```

Unless I'm completely misunderstanding how you're using D and r, this is simply not the case, so I think you are probably using an incorrect definition of information. I think you need to chuck this whole $\log(D/r)$ business and start over. If two signals have wildly different scalings of amplitude, but are otherwise identical, then I can write a very short program to convert one to the other, no matter how long the signals are (they differ, after all, by only a simple scaling). The size of this program, by definition according to algorithmic information theory, is a measure of information loss due to the transformation, and is very low. So scaling of amplitude alone is an information preserving transformation,

contrary to what you claim. This is not the case for $\log(D/r)$, so I do not see how this can be a measure of information.

To clarify this issue, so I understand exactly what your understanding is, tell me how to compute $\log(D/r)$ for the following sequence S:

.0 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .2 .7 .7 .7 .7 .7 .7 .7

I am guessing (it is only a guess) that you would say that D (amplitude range) is $.7 - .0 + .1 = .8$, and r (the resolution) would be $.1$. So:

$$\log_2(D/r) = \log_2(.8/.1) = \log_2 8 = 3 \text{ bits/sample.}$$

$$H(S) = 3 \text{ bits/sample} * 25 \text{ samples} = 75 \text{ bits.}$$

So would you say there are 75 (binary) bits of information in this sequence?

In what follows, I will describe how to actually compute the information measure H for this rather simple example, with a minimum amount of theory. Hopefully, this will put the whole definitional problem to rest once and for all. I know I've tried to explain all this before, but I don't think I managed to get the message across. So I'm trying yet again with a different strategy.

The problem with computing H(S) as given above is that there is really no such thing as H(S). Remember that *something* is always assumed. So we really compute H(S) with respect to some language, or model, or probability distribution, L: $H(S|L)$.

We'll consider this from both the algorithmic and the more traditional Shannonesque probabilities. First, the algorithmic:

L will be considered as a computer language. Assume language L1, whose programs are defined by pairs of numbers (n x). n is 5 bits (1 - 32) telling how many copies of x to generate. x is 3 bits (.0 - .7). This is 8 bits per (n,x) pair. Thus, a minimal program for S in language L would be:

(1 .0) (16 .5) (1 .2) (7 .7)

$$H(S | L1) = 4 \text{ pairs} * 8 \text{ bits} = 32 \text{ bits.}$$

This is a far cry from 75 bits. Note, however, that 75 bits could be correct given some other language, say L0, whose programs are defined by a sequence of 3 bit numbers (.0 - .7). The minimal L0 program is identical to the original description of S above:

$$H(S | L0) = 75 \text{ bits.}$$

So when you talk about D and r, you are giving information according to one particular language (again, I'm guessing this is how you're intending it). It is also a language that gives a result very similar to channel capacity - not a particularly useful definition of information.

Looking at it from the more traditional probabilistic perspective instead of the algorithmic one, we would look at the source of S, and measure the probability P(S) of the source producing this particular S. Say, for simplicity, that there are only two possible sequences that the source could ever produce:

.0 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .2 .7 .7 .7 .7 .7 .7 .7
.0 .0

The first one, S, occurs 1 time in 8. The second one, S0, occurs all other times (or at least, this is our model of what happens). This gives

$$P(S) = 1/8$$

$$P(S0) = 7/8$$

Given this rather simple probability distribution, call it D, the information measure would be:

$$H(S|D) = -\log_2 1/8 = 3 \text{ bits}$$

So L0, L and D all give quite different information measures for this sequence. So what good is a measure that is so arbitrary? Well, it just means you have to decide what to take as given. For an ECS, the hierarchy between it and the world is a given (it is provided to the ECS "for free"). Only features of the world (CEVs) that can be encoded by the hierarchy in few enough bits to be controlled by the output channel will be useful ways to model the environment.

If an outside observer uses a different language to describe the environment, he may define what he thinks is being controlled (what he imagines as the CEV) in his own terms. But these "CEVs" might require huge bandwidths to control, and this observer will decide that error control is not practical and be disenchanted with control theory. The way for him to really understand what is happening is to drop his preconceived model of the world and do a test for the controlled variable, so he can figure out what model the organism is using to encode the world (which might be different than his own). The organism may be encoding the world into CEVs that actually require much lower bandwidths to control than those the observer was initially imagining. This is what Ashby's Law is all about - the real meat of it. Note that the language being used for encoding is pretty trivial for a single ECS with no hierarchy. Information theory is not particularly useful until you have a hierarchy.

Note that we can talk about CEVs from the point of view of the ECS alone. Talking about disturbance, however, requires a minimal external viewpoint (from the organism's perspective, there is no reason to even put this "disturbance" thing into our HPCT diagrams). However, most of us probably feel that this minimal external viewpoint is very helpful in understanding what is happening, as long as we understand what we're doing. This is why I have said in the past that talking about disturbance requires a mixing of viewpoints, which is not in itself wrong, but must be recognized. "Pure" PCT would not even be able to phrase a question about disturbance, let alone talk about whether there is information about it in the percept. The moment you put an arrow in your diagram labelled "disturbance," however, you are taking a minimal, and understandable, external view. This is why I insist on the "free" use of a programming language such as C in my reconstruction attempt, but do not allow the disturbance itself, or any other specific external models of the environment.

Bill Powers (930416.0800 MDT)

>...Why not reserve the much more
>complex and opinion-weighted and assumption-sensitive arguments
>of uncertainty theory for situations in which the variables are
>in fact unpredictable and irregular on the scale of interest?

You have to understand, as I've explained above, that a probability distribution just means establishing a model with which to encode the data. You are interpreting the term "probability" in too narrow a sense. Its really just a way of saying that the system does not treat all signals as equal (as it would if it just used the channel capacity as its model). The signal does not need to be literally nondeterministic for the ideas of probability theory to apply. All you need is to be able to declare some signals as more probable than others, which is simply a model of how the world works.

>Reducing IT to practice
>appears almost too difficult to do at all; just look at Allen's
>proposed proof of something about entropy, which requires solving
>an NP-hard (one might even say NP-impossible) problem on the way.

This is a terribly quick dismissal. Just because the theory involves an NP-complete problem does not mean it is useless. Sure, it would be great if we could always compute the *real* information measure H. But we can't. The theory is still valid, however, and approximations will often suffice.

>As a practical approach to explaining behavior, I am not impressed with IT.
>I am not one of those who loves complexity for its own sake.

And I am not one of those who loves simplicity for its own sake. I would prefer things to be as simple as possible, but not more so, as Einstein would say. You assume because something is hard for you to understand, given your own background and prior knowledge, that it is therefore wrong or misguided - that it is being made complex for the pure love of complexity. Judge not, lest ye be judged. Complexity is in the eye of the beholder (eg: "this continuous treatment of everything is so complicated - why do we need it? What does it add to the true discrete nature of what's *really* going on? Why can't we just use a discrete analysis - everything would be so much simpler.")

Bill Powers (930420.1900)

>If the output of the system cancels nearly all of the information
>coming into the control system from the disturbance, then how can
>there be enough information left in the perceptual signal from
>which to produce an output containing nearly as much information
>as the disturbance?

There really is no paradox. I know it *sounds* paradoxical, though, and I have been struggling these past few weeks to reword this. I think that there is a problem with Ashby's terminology. I will try to rephrase it.

The control system should not be said to "block" information. I no longer like this terminology. If the disturbance has been reduced to only a few bits through successful control, then an output containing "nearly as much information as the disturbance" will also contain only a few bits. The information about the disturbance has not so much been "blocked" as it has been "compressed" or "reduced." I know this sounds the same as "block," but keep in mind that I can greatly reduce the bits required to describe something without information loss! Now let's rephrase your question, and I think the paradox will disappear:

"If the output of the system cancels nearly all of the information coming into the control system from the disturbance, then how can there be enough information left in the perceptual signal from which to produce an output containing nearly as much information as the disturbance?"

|
v

"If the disturbance as represented in the percept contains very few bits, how can there be enough bits in this percept to describe the disturbance?"

|
v

"If something contains very few bits, how can it be described with very few bits?"

No paradox here - more like a tautology.

To say the channel is "blocked" implies that the disturbance actually still requires the higher number of bits dictated by the external viewpoint. But while the control system *has* reduced the information IN the disturbance, it has *not* blocked information FROM the disturbance at all! It doesn't need to block when it can reduce, or compress. In fact, it could have nearly complete information about the disturbance, its just that there is little information about the disturbance to be had, due to good control!

I believe I have used very poor wording to describe this in the past - I hope this is better.

>...If most of the information from the disturbance is missing from the
>perceptual signal, how can perfect control even be approached?

I think it is wrong to say that the information is "missing," since we are taking the viewpoint of the control system. The information has been "made to be low," "reduced," or "compressed," but not "blocked," or "missing." It is only missing from the viewpoint of the external observer who ignores the hierarchy.

Bill Powers (930422.0030 MDT)

>What this implies is that talk about the information content of
>physical variables is empty. For all practical purposes, their
>objective information content is infinite. How much information
>we can obtain from them depends strictly on the dynamic range of
>the sensor.

This is true - objective physical information content is a meaningless notion. The information *is*, as you say, dependent on the sensors that collect the data (and also on the control hierarchy that processes this data).

In the limit of the truly continuous signal, the information *capacity* is infinite. But a finite system cannot model a continuous signal to an infinite degree of precision. The

number of bits in the signal, from the point of view of a finite observer, will always be finite. And it will be different for different observers.

I hope I have explained this stuff better this time around. I have spent some time trying to rephrase these issues over the past few weeks. I hope we can finally understand each other at some point.

Allan Randall

Date: Fri Jun 11, 1993 4:18 pm PST
Subject: INTRPRSNL RLTNS - RKC

[From Bob Clark (930611.1745 EDT)] Bill Powers (930529.0800 MDT)

Regarding my suggested assignment of "Control of Interpersonal Relationships" to Sixth Order, you remark:

>I don't see control of interpersonal relationships as a basic level
>of control. We control many kinds of relationships, some of them
>being interpersonal.

After several examples of various "relationships," you add:

>Relationship control is a very broad category, one that's easy to
>take for granted as a given of the environment, so I definitely
>think it has to be recognized as indicating a type of perception.

Bill, I agree with what you say about "Relationships" as a generally useful term. Indeed, it might be considered as an alternative to "Configurations" at that lower order.

But I am amazed that you seem to over-look the phrasing of "Interpersonal Relationship," as including both words to form a single, restricted, category. To me, it is similar to your using "System Concept" to characterize that category. The same comments you make about my "word pair" can be applied directly to your "word pair."

As I have indicated various places, I find this restricted concept very useful in distinguishing that group of perceptual variables from those that do not involve people. It also provide a framework within which further distinctions and perceptions can be identified.

In terms of viewpoints, I find it useful when using the Observer's Viewpoint, when using the Engineer's viewpoint, and when using the User's (is "Driver's" or "Operator's" or what, better?) viewpoint.

It also seems to involve rather slower time scales than those needed for Lower Order Systems.

Enough on this for now -- Regards, Bob Clark

Date: Fri Jun 11, 1993 9:17 pm PST
Subject: Information Theory (again)

[From Rick Marken (930611.2200)] Allan Randall (930611.1700)]

>Rick, I am still willing to attempt a disturbance-reconstruction
>demonstration. However, I am waiting to receive from you a simple
>control system for which such a demonstration would be convincing.

You can use any kind of control system you like; or any method you like whatsoever -- it doesn't even have to be a control system. The paragraph of mine that you quoted spelled out the rules of the game:

> reconstruct D from P given the assumed conditions (r, O() and C)

I will give you a sequence of numbers, P, that was an actual time sequence of perceptual inputs to a control system. I will also give you a number, r, that is the fixed value of the reference signal in that control system. I will also give you the output function, O(),

that converted error values (the sequence of r-P values) into outputs. I forget what C was but I presume it was the actual time sequence of values of the controlled variable (the cursor in a tracking task). I think I actually gave this data to you (and I gave Gary Cziko the actual values of the sequence of D values) but I never got the reconstructed values of D. I will happily post these data again if you like. There are absolutely no restrictions on the means you use to reconstruct D from the given values of P and the information about r, O() and C (well, it's no fair to use extortion or threats).

I also said:

>Since we have shown that, indeed, you cannot reconstruct the disturbance
>from the percept under these conditions then I think it is fair to assert
>that THERE IS NO INFORMATION ABOUT THE DISTURBANCE IN THE PERCEPT.

Allan asks:

>When did you show this? You showed that it was not possible to perform
>the reconstruction?

Bill Powers and I both posted data showing that there can be a very weak relationship between output and disturbance when there is a varying environmental non-linearity between the output and the controlled variable. Since you (and the control system) have no information about this non-linearity neither you (nor the control system) can reconstruct D from P.

>And to clarify, yes, I was saying that there would be no information
>present if it was not possible to reconstruct the disturbance with P and
>the stated legal variables WITH FEWER ADDITIONAL BITS than would be
>required without P

Without P there is no way at all to reconstruct D except by chance, right? What is this fewer additional bits thing? I'll give you a string of numbers which I swear is P. I'll also give you r, O() and C (I really forget why I agreed to giving C if C is values of the controlled variable; it can hardly be of much help since in my examples C will be a string of numbers that is exactly equal to the string of numbers that is P). You give me back a string of numbers that is D. You should be able to do this because, you claim, there is something (that you call "information") in P that let's the control system know what D is. That, according to you, is how the control system knows how to generate outputs (or effects on the controlled variable) that counter D. I am saying that it doesn't work that way; the control system knows nothing about D and can know nothing about D. The control system generates outputs (or effects on the controlled variable) that counter D because it continuously produces whatever outputs are necessary to keep r-P = 0. When there are non-linearities in the environment, some of this "necessary output" is going to make up for this non-linearity; some is going to counter the disturbance and some is going to move P towards the possibly varying value of the reference signal (but the reference signal is constant in this test).

>All I need is an example of a control system that will satisfy you, and
>we can settle the issue, right?

It's absolutely and completely up to you.

All I'm proposing is giving you a list of, say, 50 P values, eg.

2 4 6 7 8 9 8 7 6 4 3 4 5 6 8 9 10 12 etc...

and tell you $r = 0$, $O := O + k(r-P)$ and C is the same as the P numbers above. If there is information about the disturbance in the series of numbers above, you should be able to return a series of numbers that comes very close to the actual disturbance present at the time the P numbers were obtained. That is, you should be able to send me a series of 50 D values, eg.

3 5 7 8 9 6 5 3 2 2 3 3 4 5 7 8 9 etc...

And I should then be able to look at the file that contains the actual D numbers (or Gary Cziko could then put them on the net) and they will be :

3 5 7 8 9 6 5 3 2 2 3 3 4 5 7 8 9 etc...

And we will see that you have, indeed, been able to reconstruct D from P and we will have to concede that there is information about D in P and that our notions about how a control system works are wrong. But that would be OK because we will have learned how to get information about D out of P (that is, if you'll tell us how you did it).

Here we go again. Wheeeeeee.

Best Rick


Date: Sat Jun 12, 1993 8:17 am PST
Subject: Reconstructing disturbance

[From Bill Powers (930612.0930 MDT)]

Rick Marken (930611) and Allan Randall (930611) --

Before you two guys start going around again on reconstructing the disturbance, I think you'd better make sure you're talking about the same thing. It's this old business of what you mean by "the disturbance."

Consider the rubber band experiment, that magical instrument that seems to explain everything.



When we say "disturbance" in equations, we mean the STATE OF THE DISTURBING VARIABLE: that is, the measure of the position of the experimenter's end of the rubber bands above, with respect to any arbitrary zero point (here, a position at the extreme right). The action is a similar measure relative to the same zero point. We do NOT mean the amount by which the knot (X) is displaced from the target position. The position of the knot is a dependent variable, depending on the positions of the two ends of the rubber bands.

The subject is maintaining a specific relationship between the knot and the dot. The relationship shown above might be it, or might represent an error. There is no way to tell without knowing what the reference relationship is.

Allan, what the control system perceives is just this:



It does not perceive the experimenter's rubber band connected to the X, or the rubber band connecting the control system's own output to the X.

The claim that Rick and I are making is that simply from perceiving X o, the control system (or anyone inside it) could not possibly reconstruct the POSITION OF THE EXPERIMENTER'S END OF THE RUBBER BANDS. Even knowing that the desired perception is X o, this would not be possible.

The reason is that the control system receives no information about the characteristics of either rubber band, or about the state of its own output variable, here a hand position, or about the form of its own output function or input function. The experimenter's rubber band could be thick and the subject's thin, or vice versa. The input function could be linear or logarithmic. The output function could be a leaky integrator with any time constant up to infinity (a pure integrator). None of that information -- in the sense of knowledge -- is present in the perceptual signal.

If you're going to reconstruct the disturbance on the basis of the perceptual signal, you have to do it the same way the control system would have to do it: using only the available information that is represented by signals in the system. This is why I felt that Martin was cheating: the form of the output function is not represented anywhere by a signal. No function is so represented. MARTIN knew what that function was, but THE CONTROL SYSTEM did not.

Both Rick and I have send you some strings of numbers representing the behavior of the perceptual signal in an actual operating control-system simulation. This puts you in the

position of the control system, knowing only and exactly what the control system knows. Our challenge was for you to tell us how the disturbing variable was behaving during the simulation. I fully expected you to come back saying that of course it was impossible -- you didn't have enough information on which to base a reconstruction (either that, or that you would dazzle us by actually coming with with a correct reconstruction by some means we couldn't conceive, thereby proving that information theory could do something that control theory can't).

But you've ignored those strings of numbers, as far as I remember. You certainly haven't sent back any string of numbers representing the behavior of the disturbance that got to ME. If you claim that the strings of numbers are not sufficient for reconstructing the behavior of the disturbing variable, then we have made our point: there is insufficient information in the perceptual signal to allow reconstructing the disturbance. The only way you can disprove our point is to come up with the actual reconstruction, which I see no way of doing.

Best, Bill P.

Date: Sun Jun 13, 1993 6:14 pm PST
Subject: Re: SOCIAL SYSTEMS - RKC

[From Bob Clark (930613.2145 EDT)] Bill Powers (930601.0600 MDT)

In your discussion of social organizations, social control systems, individual control systems and HPCT, you emphasize

>"the reliance on direct physical force to make the social
>hierarchical system work."

You proceed to note:

>"Inside an individual, this is simply not a factor when the system
>is functioning normally; the nearest thing to one internal system
>forcing another to act is in a conflict situation, where two systems
>demand incompatible actions from lower systems. This normally leads
>to reorganization that eliminates the conflict."

True, of course, for a "normally functioning system." When "functioning" is abnormal we have assorted problems, some minor, some major. Sometimes conflict is avoided, sometimes it is ignored, sometimes it is "repressed" or "suppressed" (I don't like these terms myself, but they are used). These can be observed in various forms of mental illness, and other self-defeating behavior.

You seem distressed with the situation:

>In all social hierarchies of which I know, the use or credible threat
>of physical force is the main means of maintaining the structure.
>This is a fundamental design defect; it creates conflict automatically.

Yes, Bill, of course. I don't like these effects either. But neither you nor I created these social hierarchies. They are created by people cooperating in an attempt to achieve common goals. Whatever else? Of course their control systems are often poorly devised, resulting in assorted unanticipated conflicts. We see this far too often in governmental attempts to control phenomena that are intrinsically uncontrollable, and/or, using erroneous and incomplete theories of behavior as guidelines for design and operation of these hierarchical control systems.

"Effectiveness."

Referring to my statement:

>>-- but the autocratic manager tends to get into difficulty one way
>>or another. Effective managers do not behave in that manner -- and
>>never have.!"

You comment:

>You're invoking your own criterion of effectiveness, not the

>generally accepted one.

Dictionary: "effectiveness, n. derived from effective, adj. 1. adequate to accomplish a purpose; producing the intended or expected result."

Evaluation of effectiveness depends on the "purpose." For a short term purpose many poor management styles may be sufficient. For a long term purpose, autocratic and other management styles are found less effective in that they create unnecessary and destructive conflicts within the organization. I've seen examples of various management styles, in several different hierarchical social organizations.

Another remark:

>A social system is not necessarily a social control system. A
>social system is an agreed set of goals which each individual adopts
>because they further that individual's concept of a world worth
>living in. A social CONTROL system, on the other hand, necessarily
>employs physical force to make all members conform to a set of goals
>that only some members consider desirable.

I like your distinction between "social systems" and "social control systems." Indeed it is my impression that social control systems commonly develop from social systems. After a "club" (minimal organization) is formed, a question may be raised about expulsion of members who "misbehave" in some sense. Questions of definition of "misbehavior", expulsion procedure, appeal, reinstatement etc are raised. The entire subject may be rejected as unsuitable for one reason or another. If any such restrictive suggestions are accepted and implemented, the "social system" has been transformed into a "social control system."

These agreements (to cooperate) are formed regularly. And they are transformed regularly. It is predictable. Individual control systems seek ways to improve (simplify, accelerate, reduce force/energy/power needed, etc) their control of their environments -- including the human parts thereof.

Elsewhere you remark:

>Speaking literally, social CONTROL systems don't actually exist, and
>never have. All that actually exists are individuals trying to make
>their own worlds conform to their own preferences. What seem like
>social control systems are really people trying to maintain their
>own conceptions of a social structure.

This seems rather odd to me. Do you mean they "don't exist" BECAUSE of their ultimate "reliance on coercion?" For the one "being coerced" it surely seems pretty real -- no matter how unjustified, unwise, or whatever.

Regarding the "reality" of social control systems, I think it is pertinent to ask "where" they are. Their physical existence is sometimes very tenuous, often being expressed in the form of marks on pieces of paper -- or equivalent. The only important place where they exist is in the minds of the affected individuals, especially the participating individuals. In "minds" I include all learned systems, primarily, but not limited to, the verbal systems. As long as the systems in the minds of the participating individuals are consistent, they tend to work pretty well in reaching and maintaining their common goals. However there is much room for misunderstanding and conflict of many kinds.

"Real" or not, I think we can learn a great deal about how people think about and control their environments by studying their Social Control Systems. And, if we don't "like" their conflicts, we should study "METHODS OF CONFLICT RESOLUTION." I've been doing a bit of that in the City I'm working with.

Regards, Bob Clark

Date: Mon Jun 14, 1993 3:28 am PST
Subject: no witnesses

[From: Bruce Nevin (Mon 930614 06:57:04 EDT)]

In Closed Loop 3.2:49 Bob Clark describes transferring ball-throwing skill from his right hand to his left by analytically replicating the familiar sequence. He observes that he had to have no witnesses to do this successfully.

I think that this is because whenever we are in company of others we imagine how our actions appear to them. This becomes a distraction when we are on unfamiliar ground.

This imagining how we appear to others is the basis on which individuals may control perceptions of conformity (or not) to social norms. It is the basis for development of norms in the individual and in the evolution of social groups.

The gather program shows that some "social behavior" at least requires no perceptions of social norms. This proposal predicts that other "social behavior" does, and it suggests something about the origin and character of perceptions of social norms. An example of origination might be the means by which the individual control systems in the gather program come to have the same setting of a reference perception of proximity to other individuals.

Bruce bn@bbn.com

Date: Mon Jun 14, 1993 5:18 am PST
Subject: interpersonal relationships

[From: Bruce Nevin (Mon 930614 08:36:04 EDT)]

(Bill Powers (930529.0800 MDT)) -- Bob Clark (930611.1745 EDT)

If my supposition posted earlier this morning is correct, then interpersonal relationships are of a different order than other relationships, such as the relationship between a fingertip and an object to which one is pointing. In the latter case, one need not imagine how one's behavioral outputs appear to another. This is also shown by the longer delay in correcting disturbances to interpersonal relationships, as opposed to tracking a moving object with a fingertip. They have little in common other than the word "relationship," perhaps.

(Allan Randall (930611.1700) Rick (930611.2200) Bill (930612.0930 MDT)

To add my \$.02 again, information becomes apparent in a time-series of values. The control system in question doesn't perceive a time series, it perceives only the instantaneous value at time t. There is no information in the single, instantaneous value of p at time t for any given time t.

A higher-level control system CS' could perceive a time series in the inputs to control system CS; in this it is just like any other observer from a viewpoint outside CS. Information theory requires a point of view from outside the control system; it cannot be applied to the control system itself from its own point of view.

Bruce bn@bbn.com

Date: Mon Jun 14, 1993 6:22 am PST
Subject: Social control systems; H(S); evolution

[From Bill Powers (930613.2200 MDT)] Bob Clark (930613.2145 EDT) --

>When "functioning" is abnormal we have assorted problems, some
>minor, some major. Sometimes conflict is avoided, sometimes
>it is ignored, sometimes it is "repressed" or "suppressed" (I
>don't like these terms myself, but they are used).

I don't like such terms, either. They sound like explanations but actually explain nothing: they define a cause in terms of its effects.

>You seem distressed with the situation:

>>In all social hierarchies of which I know, the use or credible
>>threat of physical force is the main means of maintaining the
>>structure. This is a fundamental design defect; it creates

>>conflict automatically.

>Yes, Bill, of course. I don't like these effects either. But
>neither you nor I created these social hierarchies.

What does not liking them, or who created them, have to do with anything? A design defect is, in this case, a feature of the social organization that continually threatens to destroy it, and almost inevitably will. Not what you call a steady-state solution.

>They are created by people cooperating in an attempt to achieve
>common goals. Whatever else?

I guess that as an engineer, I am pained when people cooperate in an attempt to achieve common goals using means that will almost certainly, in the end, frustrate the attempt.

>>You're invoking your own criterion of effectiveness, not the
>>generally accepted one.

>Dictionary: "effectiveness, n. derived from effective, adj.
>1. adequate to accomplish a purpose; producing the intended or
>expected result."

Yes, that's the meaning I assumed. What most people mean by "effective" is doing whatever it takes to get the job done right now. Hitler was greatly admired by many Americans in the 1930s as an effective leader. Perhaps the distinction I'm trying to make is between short-term and long-term effectiveness, as you suggest. I don't see many signs of long-range effectiveness as a goal in government or organizations: more like, how much money can we make by next week?

>>Speaking literally, social CONTROL systems don't actually
>>exist, and never have.

>This seems rather odd to me. Do you mean they "don't exist"
>BECAUSE of their ultimate "reliance on coercion?" For the one
>"being coerced" it surely seems pretty real -- no matter how
>unjustified, unwise, or whatever.

See Tom Bourbon's diagram of how a "social control system" has to be implemented. I simply mean that the only control systems in a social organization are the individuals in it. If coercion is applied, it is applied by somebody to somebody. People often try to use the concept of social rules and social demands as a way of avoiding responsibility for the things they do with their own arms, legs, and voices. They like to pretend that they are governed by something outside themselves; in that way they can do what they really want to do, but when problems or objections arise they can pass the buck. "I'm only following orders."

>Regarding the "reality" of social control systems, I think it
>is pertinent to ask "where" they are. Their physical existence
>is sometimes very tenuous, often being expressed in the form of
>marks on pieces of paper -- or equivalent.

I am thinking very literally about this. Marks on pieces of paper can do nothing at all but be marks on pieces of paper. They are not control systems: they have no perceptions, no comparators, no outputs, no goals. They can be used by human beings in the course of human control behaviors, but by themselves they can do nothing. I think we will make the most progress in understanding social systems if we deal with them the same way we model individual behavior: with models that we proposed to be literally, physically, correct.

>The only important place where they exist is in the minds of
>the affected individuals, especially the participating individuals.

Yes, exactly.

>"Real" or not, I think we can learn a great deal about how
>people think about and control their environments by studying
>their Social Control Systems.

That's the way social systems have traditionally been studied: in terms of metaphors. As long as we have a new theory of human nature that is based on a model intended to be

literally true, why not try that same philosophy with social systems? If we keep insisting that each individual is responsible for all that individual's actions, maybe some day people will start to believe it.

Of course when you say "their" social control systems, you probably mean each person's conception of a social control system -- so my proposal agrees with yours.

Allen Randall (930611.1700) --

>Bill, I have a real problem with this whole log(D/r) thing. I'm
>not sure where you got this definition, but I do not think it can be right.

I got it from Martin Taylor. D is the range of the disturbing variable d, and r is the resolution with which the disturbance is measured (or perceived open-loop). This allows for D/r different values of d, with an information content of log(D/r). I'm just following orders. I don't agree with them, but I'm just trying to do what they seem to say.

>If two signals have wildly different scalings of amplitude, but
>are otherwise identical, then I can write a very short program
>to convert one to the other, no matter how long the signals are
>(they differ, after all, by only a simple scaling).

Suppose you have the relationship $X(t) = Y(t)/10$. In your computer program, Y(t) will be some input waveform, and X(t) a similar output waveform with one tenth the amplitude. No matter how you write the program, when you compute X from Y you will get a waveform with less relative resolution than there was in Y. If you multiply X by 10, you will not get Y back; you will get a coarser representation of the original waveform, with about 1/10 of the original resolution. Whenever you scale down a variable having a finite resolution, significant digits are lost forever; you can't get them back unless you're measuring the scaled-down variable with smaller resolution units than you're using to measure the full-sized variable. That's all I was trying to say.

>To clarify this issue, so I understand exactly what your understanding is,
>tell me how to compute log(D/r) for the following sequence S:

> .0 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .2 .7 .7 .7 .7 .7 .7 .7

Before I could do that, you would have to tell me the size of r. There isn't any inherent "resolution" in an arbitrary sequence of numbers like the one above. What they mean depends on the physical situation they are taken to represent. If this is a series of measurements of some physical variable, then we have to talk about the measuring device's resolution. If the least significant unit of a measurement is 0.001, then D/r for the above series is 800. If the least unit is 0.01, it's 80. The upper limit on r would be 0.1, because if it were any greater you couldn't distinguish both odd and even numbers of exact tenths. But the fact that the series happens to contain only exact tenths of a unit doesn't imply that a measurement is only accurate to 0.1. The above could be an unlikely-seeming sequence to obtain with a device that can measure to 0.001, but probability-wise it's no more unlikely than the sequence .001 .533 .533 etc. A specific waveform described by the sequence with three-digit resolution would also be described, at lower resolution, by numbers accurate to 0.1.

You say

>You assume because something is hard for you to understand,
>given your own background and prior knowledge, that it is
>therefore wrong or misguided - that it is being made complex
>for the pure love of complexity. Judge not, lest ye be judged.

Fair enough, but I think there's something more than this going on.

My problem is that I can't find any link between the manipulations you talk about and any PHENOMENON. You seem to be setting up arbitrary hypothetical examples and then applying mysterious calculations to them for no reason I can figure out (except that it can be done).

I think you express my problem very well when you say

>The problem with computing H(S) as given above is that there is
>really no such thing as H(S). Remember that *something* is

>always assumed. So we really compute $H(S)$ with respect to some
>language, or model, or probability distribution, L : $H(S|L)$.

This seems to explain something. When I've asked for instructions about how to compute information-related quantities for a given set of experimental data, the only answer I've received so far is "Well, that depends on what you assume." There seems to be no rationale for making any particular assumption that would lead to a definite prediction. It seems that no matter how you conceive of the situation to define a probability, there's always another equally plausible way to conceive it, leading to a completely different value of probability, $H(S)$, or whatever. How do you decide what is the most plausible way to set it up for a control system model of a specific example of behavior? So far I'm drawing a complete blank on that. And, apparently, so are you.

In a control-system model, we make assumptions about the form of the model, and within those assumptions we calculate the values of the parameters that fit the model to data from experimental runs. So we have a criterion for comparing different assumptions: how well the best-fit model having the assumed organization predicts the data. If we assume a simple linear system with an integrator in the output, we find a certain best degree of fit. If we add an input delay to the model, we get a measurably better fit, for every subject. So this tells us that the model with the delay is better than the model without it. We already know that many other models are ruled out because they don't behave in any way resembling the real behavior -- a positive feedback model, or an on-off model, for example.

What is it that you can use to decide which assumptions about the situation yield the most realistic or predictive measures of $H(S)$? What's your criterion for deciding that one set of assumptions is better than another? If the value of $H(S)$ that you calculate depends on your assumptions, then you have to have a way to say that one number is a more plausible representation of the real situation than another, so you can decide that one set of assumptions is better than another. If you can't say that, all you have is an all-purpose calculation that can be applied to any data set and that can yield any number you like today. You're not constrained to the real universe; you can't tell whether your results apply in this universe or only in an imaginary one.

In your example of computing $H(S)$, you set up an arbitrary situation in which one sequence of numbers occurs 1/8 of the time, and a different sequence the rest of the time. Naturally, you calculate that there are 3 bits of information. It doesn't matter what the events are, whether sequences or just occurrences of a single event. So if you're given the rule, you can calculate $H(S)$. After all, the way you set up the example completely defines the probabilities, so it is easy to apply the calculations.

But how do you get from the physical situation to the correct definition of probabilities? If you consider only examples in which it's already been decided what the probabilities are, it doesn't seem much of a feat to then proceed to calculate a measure that's a function of the probabilities. If you happen to get the rule wrong, you'll still be able to calculate $H(S)$. Maybe the actual probability isn't 1/8, but 1/7, or even $t/8$, where t is the elapsed time. Or maybe it's $\text{rand}(t)/8$. The possibilities are infinite; you need some way to determine the probabilities (if that's even the right conception of the physical situation) before you can do any calculations of $H(S)$. Of course even if the distribution isn't actually probabilistic but follows some perfectly regular rule, you can use the measures to form a data set and then compute $H(S)$. But how will you know that this is an inappropriate treatment of the data? Your calculations won't tell you that, will they?

You and Martin don't seem to care much for actual experimentation with real systems. This purely mathematical approach, however, is full of pitfalls. In a mathematical approach you're setting up premises (axioms) and then following out their implications. But without experimentation, you can't tell when you select axioms having no counterpart in, or even contradicting, experiential reality. A conclusion that is provably true in mathematical terms may not apply to any real systems, because the underlying axioms may be false in this universe (even though you can accept them as ground rules for the game). If you're doing the mathematics simply because you enjoy that sort of intellectual pursuit, fine by me, have fun. But if you want to apply the mathematics to any real systems, you have to be selective about the axioms and make sure they aren't false to fact. The only way to do that is to do experiments to see what constraints nature imposes. You can't reason them out.

Bruce Nevin (Wed 930609 07:11:43 EDT) --

>Harris's test for contrast/repetition is a form of the Test for
>a Controlled Variable. Ask a native speaker to control a
>perception of "repetition" or "same word(s)". Produce two

>pronunciations (perhaps recorded) that are alike, except that
 >in the second some portion is replaced by a similar portion of
 >a different utterance (perhaps by splicing segments of
 >recordings). The native speaker assents that they are
 >repetitions, or dissents, saying that they are different words.

Here's my problem. If I present you with a picture of a grape and a picture of an elephant, you can distinguish between them; the perceptual input that allows you to perceive the grape does not respond to the elephant, and vice versa. So I have established that you have two perceptual functions, one for each picture (sort of). Am I then justified in saying that you are perceiving something called "contrast" between the grape and the elephant? Or is the notion of contrast an interpretation by the observer, me?

>The segments are relevant because they represent the contrasts
 >between utterances and locate the points of contrast within utterances.

But isn't it really that the listener's perceptual functions make a distinction, rather than that there are objective contrasts in the sentences?

>But because each such alternative system is a representation of
 >the contrasts between words in the language, they are
 >interchangeable in respect to our observational primitives, the
 >contrasts. In each case, substituting one element for another
 >yields the representation of a different word. Conversely, any
 >two words that native speakers perceive as different are
 >represented differently from each other, no matter which system
 >of representation is used.

When you say "any two words", how does anyone (including the experimenter) know that they are distinct words? By applying the same kinds of input functions that the subject applies. The method you describe is ingenious; it even seems objective. But it still is defining "segments" in terms of human perceptual functions, not the other way around. What grates on my tender sensibilities is speaking of the contrast as if it existed in the utterance, or pairs of utterances. Maybe this can't be helped. But it still reminds me of studying stimuli to see what they have in common, instead of studying perceptual systems to see how they create variables out of inputs. >Bill, you suggested (930528.1930 MDT) that contrast is a non- >phenomenon. Easier to say that the phoneme is a non- >phenomenon.

Not that it isn't a phenomenon; just that it isn't necessarily what the subject perceives when stating that one perception is not another one, or that two perceptions are actually only one, repeated. If I have two experiences and they actually involve only one perception, I classify this situation as "same." If there is more than one perception, I classify it as "different." But before I can perceive which category of situation it is, I must know already whether one or two perceptions were involved: I must know that both experiences came from this perceiver, or that one came from this and the other from that. If this sort of discrimination hasn't already been made, perceptually, there is no way to decide on the category "same" or "different."

It's on this basis that I maintain that "contrast" isn't an explanatory term, but only a descriptive one. We don't perceive that two things are the same or different because they ARE either the same or different. We can only make that judgment after the discrimination has been made at lower levels. If we perceive two things via two input functions, we conclude that they are different; if both perceptions come from the same input function we call them the same.

It seems to me, therefore, that Harris' study does reveal something about the "tuning" of input functions, but that the use of the term "contrast" objectifies something that is really a judgment after the fact. Also, as in the case of phonemes, the actual degrees of freedom of the segments remain unknown. All we end up knowing is that when segments from different words are combined, a person either perceives something new or does not. What makes the difference remains undetected. The hypothesis that the differential responses are due to something called "contrast" isn't very informative.

 My citation of Land's experiment wasn't meant to be taken very literally. I was simply trying to find an example that might support the concept of perception through contrasts, by proposing a mechanism for normalizing a set of perceptual signals to make differences between them more important than their absolute values. After normalization, one perception is in effect being judged in relation to others, because the other perceptions have influenced the average. Recognizing a melody played in different keys would be a simple

example. The relationship of one note to another is vastly different in two repetitions of the melody (different frequency ratios) unless you somehow factor out the change of key. Just an idea.

 >But it turns out that there are no reliable invariants in the
 >acoustic signal.

That is an instrumentation and theory-based conclusion. Obviously, there ARE invariants, because we derive subjective invariants of perception from the acoustic signal. It's just that these invariants aren't derived from the signal in any way that anybody had guessed so far. If you use the wrong input function, you won't detect the invariants. We have to conclude that research so far has been constructing invariants by a method that doesn't yield the same invariants that real human input functions yield. Maybe the sound spectrograph isn't organized the way auditory input functions are organized. From what I read, it seems that self-organizing neural nets CAN find the invariants that are phonemes. They probably don't look anything like what we would guess by using an instrument like a sound spectrograph.

 >When you object to the proposal that toddlers are able to
 >achieve social agreements, you seem to be limiting the notion
 >to verbally negotiated agreements.

At some early age, infants are unable to perceive sequences such as patty-cake. Perhaps they don't perceive or control any kind of sequence as a sequence, intentionally. However, an observing adult who takes sequentiality as a given property of the world can easily see sequences, even controlled sequences, in the movements the infant habitually makes while exploring the world. The adult says "But look, first the baby drops the toy, then it reaches for it, then it cries. That's a perfectly clear and repeatable sequence!" The only trouble is that the baby knows nothing of that sequence or any other. That's just how things happen to work out, usually.

I think it's important to search assiduously for the LOWEST level at which to interpret behavior. If you can explain what you see in terms of controlling relationships, don't invoke control of categories or anything higher. The only reason to use a higher-level explanation is that something is left unaccounted for without it.

To seek and achieve an agreement requires being able to perceive an agreement as such, and to act to create agreement if none exists. Whether the agreement is stated verbally is irrelevant; this is an abstract kind of perception which, I think, requires higher-level systems than a toddler has yet developed. I can see toddlers controlling relationships, such that when you say something the toddler says it back, or says something else that completes the relationship, and I can see an adult interpreting this to mean that the toddler perceives this as an agreement to be sought and carried out. The adult sees a real agreement there, but it's only metaphorical: the toddler is behaving _as if_ seeking an agreement.

To find out whether toddlers actually control for something called an agreement would require some pretty ingenious experiments. I don't think they've been done yet.

 Second pass:

>You propose that the appearance of word-contrast, syllable-
 >contrast, etc. is a byproduct of phoneme recognition.

That's not my proposal, although I believe that what you say is true. I do think that words are recognized as functions of sets of phonemes and phoneme transitions -- although not any functions that anybody has been able to define yet. But that wasn't my point.

I'm saying that what is perceived is a syllable or a word, not a contrast. We might try to explain the fact that an input function responds specifically to a given word or a given syllable by looking for a contrast with other words or syllables, but in PCT terms that would be, or could be, a mistake. Two input functions respond differently to inputs because they compute different functions of the inputs. When one input combination occurs, one function produces a larger perceptual signal than any other (ideally). The functions do not respond by reporting "I perceive a contrast." If that were so, they would all respond the same way. Each input function responds or doesn't respond, or responds more or less.

Also, when you speak of contrasts, which contrast do you mean? Every possible discriminable segment differs from all other discriminable segments, simultaneously. If you have 3000 discriminable syllables or words in the working set, you have 4,498,500 dyadic contrasts in the set. Does each perceptual function have to search through 2999 contrasts to decide what syllable or word is being heard? Is a person using a working vocabulary of 3000 words actually working with 4 million contrasts?

The problem with perception of anything in terms of contrasts is really a logical one: this would mean that you couldn't perceive anything unless something contrasting were simultaneously present. You couldn't hear "dog" unless someone were simultaneously saying "dig" or "dug" or "Dag." The fact that you can find contrasts between such words doesn't mean that any one of them is recognized because of contrasts. The contrasting words are not present at the same time. How can you perceive a contrast between one perception?

I think it quite likely that there are phoneme-recognizers tuned to respond to the same phonemes that a person recognizes. I also think that words are probably detected by input functions that receive many phoneme signals and compute functions of them that are word-perceptions. The negative results from linguistic searches for invariants simply show that the wrong invariants have been tested.

>Suppose that the second proposal is true. Then it follows that
>one should be able to determine the contrasts between words in
>a language by examining the acoustic record of utterances in
>the language, or by examining a trace of articulatory movements
>used in producing them.

Neither of these conclusions is right unless you add a rider: unless we use the right function of the acoustic record obtained from the right instrument, and unless we examine the right function of measures of articulatory movements. You have to put these measures through the right perceptual functions before you can get any indication of "contrasts." Using the wrong perceptual function (or just the raw data itself) will just yield mush. There's no reason to think that the weightings given to such measures, or even the unweighted measures themselves, will prove relevant to the problem. Why are sound spectrographs used? Not because there's any reason to think they represent sound the same way the ear does; they probably don't. They're used primarily because somebody invented them and nobody invented anything else. When it was found that you can't obtain phoneme invariants from those instruments, somebody should have concluded that they were the wrong instruments, or that their output was being wrongly transformed, instead of concluding that such invariants don't exist.

>Generations of linguists worked very hard and with great skill
>at the problem of representing unambiguously the appearance of
>word contrast in various languages. The aim is a set of
>graphical cues from which anyone could produce variable
>repetitions of words as a native speaker does, rather than
>imitations of the original pronunciations of which the
>graphical cues are a representation.

Generations of psychologists have worked very hard and with great skill (or at least persistence) at the problem of showing what stimuli give rise to what responses. So what? If you're using the wrong model, all the results mean something other than what you think they mean, if they mean anything at all. Anyway, who gets to judge how skillful the linguists were? Other linguists?

I'm not disputing the data. I'm just suggesting that the term "word contrast" is probably the wrong way to think about them -- like looking at all the valleys and not seeing the hills that create them (or vice versa).. The data might well contain evidence about human perceptual functions. It would be nice to see PCT being applied to them, instead of conventional concepts.

I think that if you really want to do PCT and not just use PCT translations of existing linguistic concepts, you have to start looking for alternative explanations of the data. I keep trying, but you're the one with the data.

Martin Taylor (930609 11:10) --

Nice observation about control systems conserving resources.

There's a related evolutionary reason for control being the organization of choice. Natural selection takes place through effects of the environment on organisms. Any organism capable of controlling some of these effects would be less likely to be selected out than an organism which had no control over them. Ergo, there is a strong evolutionary bias in favor of creating control systems.

Well, in a very sketchy and somewhat sleepy way that more or less catches me up with some of the posts I've been ignoring, so I can get back to Version 4. As my boss Allan Hynek once said to Dick Aikens and me when we were in the middle of explaining some project to him in his office, "Keeping talking, I'll be right back."

Best to all, Bill P.

Date: Mon Jun 14, 1993 6:24 am PST
Subject: Re: SOCIAL SYSTEMS - RKC

From Tom Bourbon (930614.0840) Bob (930613.2145 EDT)] Bill (930601.0600 MDT)

Bob:

>Elsewhere you remark:

Bill:

>>Speaking literally, social CONTROL systems don't actually exist, and
>>never have. All that actually exists are individuals trying to make
>>their own worlds conform to their own preferences. What seem like
>>social control systems are really people trying to maintain their
>>own conceptions of a social structure.

Bob:

>This seems rather odd to me. Do you mean they "don't exist" BECAUSE
>of their ultimate "reliance on coercion?" For the one "being coerced"
>it surely seems pretty real -- no matter how unjustified, unwise, or
>whatever.

>

>Regarding the "reality" of social control systems, I think it is
>pertinent to ask "where" they are. Their physical existence is
>sometimes very tenuous, often being expressed in the form of marks on
>pieces of paper -- or equivalent. The only important place where
>they exist is in the minds of the affected individuals, especially
>the participating individuals. In "minds" I include all learned
>systems, primarily, but not limited to, the verbal systems. As long
>as the systems in the minds of the participating individuals are
>consistent, they tend to work pretty well in reaching and maintaining
>their common goals. However there is much room for misunderstanding
>and conflict of many kinds.

Me (now):

Bob, I think Bill is saying they do not exist because they do not exist, not because of their reliance on force. "They" (entities called social control systems -- social systems that control) are actually individuals, each of whom has goals or intentions that bring them into common action around a set of shared variables in their environments. From their various view points, among those shared variables might be some other individuals (outcasts, aliens, downtrodden, criminals, or just plain them-not-us), to whom the actions of the many who act in concert might easily be taken as proof that an autonomous "social system" is in the neighborhood. But I think those kinds of social systems are no more or less real than the particles, forces, masses, and other beasties described by physicists. They are all part of a model of how perceptions come to be. For some reason, or reasons, people who recognize that taxonomies in biology, like periodic tables of the elements in physics and chemistry, are parts of models intended to make sense out of perceptions, balk at the idea that the administrations of medical schools, or gangs on the street, are also parts of models intended to make sense out of perceptions. Their physical existence is tenuous to the point of non-existence. As you say, "The only important place where they exist is in the minds of the affected individuals, especially the participating individuals."

Until later, Tom Bourbon

Date: Mon Jun 14, 1993 6:33 am PST
 Subject: Levels' point of view

[From Bill Powers (930614.0800 MDT)] Bruce Nevin (930614.0836) --

>If my supposition posted earlier this morning is correct, then
 >interpersonal relationships are of a different order than other
 >relationships, such as the relationship between a fingertip and
 >an object to which one is pointing.

Good. We need some creative thinking about "the" levels. What would you propose makes the difference? For example, what sort of perceptual capability is required to distinguish between another person and a non-person?

>Information theory requires a point of view from outside the
 >control system; it cannot be applied to the control system
 >itself from its own point of view.

Hooray. Someone else once referred to this problem as that of "smuggling intelligence into the system." I've been trying to make this point, but you said it much better.

Best, Bill P.

Date: Mon Jun 14, 1993 7:43 am PST
 Subject: Re: interpersonal relationships

From Tom Bourbon (930614.0920) Bruce (930614 08:36:04) Bill (930529.0800 MDT))
 >(Bob Clark (930611.1745 EDT)) --

>
 >If my supposition posted earlier this morning is correct, then
 >interpersonal relationships are of a different order than other
 >relationships, such as the relationship between a fingertip and an object
 >to which one is pointing. ...

If by "are of a different order" you mean "often require higher levels in an HPCT hierarchy," no problem. That is *probably* true, but *whether* it is depends on what the parties are controlling. If they control the relationship of several lines on a computer screen (several cattle in a field, several aircraft in formation, charted cash flows from several ventures), the differences might be small. Even in more "complicated" relationships, eventually relationship as relationship must be perceived and controlled relative to a reference value. At that point, there is no difference.

>... In the latter case, one need not imagine how
 >one's behavioral outputs appear to another. ...

Fine. But a decision to imagine how your actions might look to another need imply nothing radically different about control. All it means is that the actor will control (his or her own) perceptions of his or her actions, relative to an imagined reference. Isn't that the way it *always* goes? The perhaps important fact that the actor now has a reference for the impossible task of controlling how another person sees his or her actions does not change the control process.

>... This is also shown by the
 >longer delay in correcting disturbances to interpersonal relationships,
 >as opposed to tracking a moving object with a fingertip. ...

The longer time course need imply nothing other than the control of a higher order perception, perhaps "the progress of my attempt to win the affections of X," rather than "cursor - target." But at the level of perceptual relationship (proximity of cursor to target; proximity of me to X), no difference.

>... They have little in common other than the word "relationship," perhaps.

I'm not so sure about that.

At the risk of producing another crashing silence on the net, I want to re-post the concluding sections of one of my earlier replies to you. The subject was language and PCT fundamentals. In your post, you had drawn what seemed to be (to you) an important

distinction between a person who IMITATES what another has said and a person who REPEATS what another has said. You asked how "you" (probably generic, but I chose to take it literally) model that distinction.

Absent any clarifying context, I took the question to be one of how a person controls perceptions of ones own actions to match present perceptions of what another is doing, as contrasted with controlling perceptions of ones own actions to match remembered actions of another. I cast my reply as an attempt to model imitation and repetition in a tracking task. (What else? I care that others perceive me as nothing more than a stick wiggler. How am I doing?)

That previous reply, with the same examples, would serve as my reply to your present post. It does not matter, at this level of control, whether I am imitating or repeating simply because I want to and do not care whether you see me do it, or because I have another agenda, such as impressing you or offending you. An attempt to model control relative to a reference for "being noticed" or for "being seen as doing what someone else does or does not want to see me doing" may imply references above, and perhaps in parallel with, controlling the relative positions of C and T, but it need not imply a whole new breed of cat.

Now, the old reply.

=====

Subject: Re: Language and PCT fundamentals From Tom Bourbon (930529.0310)

Bruce says:

How do you model "repeating X" rather than "imitating X"? Sounds like X is a category. How do you model "person B repeating the X that person A produced" as distinct from B imitating A's behavioral outputs? Is it not the case that both A and B must have a prior agreement as to what behavioral outputs constitute an X? That is, what behavioral outputs advertise the intention to produce an X?

+++++

Tom (now)(then):

At last, earth beneath my feet! If, for "X," you allow me to substitute "movement of a cursor up and down on the screen in a triangular waveform," I do it like this.

A. For Imitating,

There must be something to imitate. Let that which is to be imitated be the momentary position of a target, the position of which is driven by a computer program step and which, as a function of time, is the triangular waveform shown below. Now:

$$\begin{array}{c} \diagup \quad \diagdown \quad \diagup \quad \diagdown \\ / \quad \backslash \quad / \quad \backslash \end{array}, \text{ and so on.}$$

Cursor := Handle + Disturbance (here, let Disturbance = zero);

Handle := Handle - k * ("perceived position of cursor re target" minus
"intended position of cursor re target")

This is a standard, undisturbed pursuit tracking task, which PCT models with great accuracy. The participant can make the cursor accurately track the target, which is to say, the positions of the cursor imitate those of the target (incidentally, so, too, does the waveform of the participant's hand movements -- the person's actions).

B. For Repeating:

For Repeating, there must be something to imitate. Let that be the person's remembered positions of the cursor as a function of time during the Repeating task. (The target is not on the screen.) Now the model step for the person changes, because the source of the reference signal is different, but the cursor is still determined as it was before.

Cursor := Handle + Disturbance (again, let Disturbance = zero);

Handle := Handle - k * ("perceived position of cursor at time x" minus
"remembered position of cursor at time x")

where the reference position at any moment is the appropriate time-indexed amplitude on the person's remembered (imagined, intended) waveform. The person can do this easily.

For either Repeating or Imitating, the same actions will occur. So, too, will the same positions of the cursor as a function of time.

Now let us imagine a new Imitating condition, with a new participant (A), in which yet another person (B), not a program step, determines the position of the target at each moment. Let the Disturbance remain zero. Assume the target-controlling person creates the triangular path shown above. (I stay with that because I am not good at drawing random paths in ASCII.) In the Imitating task, which must necessarily occur first, the result would be exactly the same as before. Given that the cursor remained close to the position of the target throughout the run, we could say -- .

What *could* we say? Could we say that the participant imitated the actions of the target-controller? That would be true only if the dynamics of the devices by which the two people controlled their respective marks on the screen were identical.

(The person *might* make his or her handle movements match those of the target-controller. More on that later.)

Assume identical control devices. If a random disturbance had acted on the target, so that in order for the target-controlling person to create the triangular path of the target the person was required to move the hand to oppose the disturbance, then the target-controlling person and the participant would have created the same patterns of movement for the cursor, by disparate actions of their hands (their manual articulators), and I can say: If the triangular pattern is "X", then it is not the case that both B (the target-controlling person) and A (the participant) must have a prior agreement as to what behavioral outputs constitute an X. Neither is it the case that any behavioral outputs advertise the intention to produce an X.

It is still possible that the participant had been making her or his handle movements match memories of those by the target-controlling person during the Imitating task. That would still be a perceptual tracking task, but one that might come a little closer to the contention of the theory that says when people converse, they duplicate ideas or hypotheses about one another's actions. We can test for whether the participant is controlling actions, or a perceived consequence of the actions.

Let the participant once again use remembered triangular waveform of the cursor during the undisturbed Imitation task as the time-indexed reference signal for cursor position. (As the reference signal, we could also use memories of the handle positions of the target-controlling person, which were identical to the positions the target, which were nearly identical to the positions of the cursor) Let the Disturbance be non-zero, and a random function of time. Now the participant produces the same triangular waveform for the cursor, but by actions that never occurred before -- the participant's actions must negate the effects of the disturbance *and* produce the remembered waveform. The model step that duplicates this result is identical -- absolutely identical -- to the one used before for the Repeating task.

The actions vary and oppose the disturbance; consequently, the perceptual consequence of the actions remain as intended by the participant. The participant imitates and repeats perceptions, not actions. The PCT model does the same. That is how I would model imitating and repeating.

=====

Tom (in the real now):
And it would not matter whether I did either of those (imitate or repeat) just for the heck of it, or to satisfy an intention to be seen in a particular way by you.

Until later, Tom Bourbon

Date: Mon Jun 14, 1993 9:52 am PST
Subject: UNDERSTANDING/EDUCATION

{From Tom Hancock (930614) To Rick Marken,

I appreciate your response (930607.1100) to my query (930604) of "How does one explain (or model) understanding (or lack of it) in terms of PCT?"

Yes, your post makes some sense to me. In a PCT view, we could say that understanding is when the perception is made to match its reference signal. Also, in a cognitive sense we could say that there is understanding when one knows how control happens and can explain how a variable is controlled.

I am hoping you can help me perceive this with more of a sense of understanding. That is, when I consider students in my classes, who are perceiving something when I speak or demonstrate, what could be happening when they have 100% confidence of understanding ("I get it perfectly") versus when they have 50% ("I sort of understand") or 0% (I am totally confused). They are perceiving movements and speech sounds, prior perceptions are being activated in association with the speech sounds and body language, and they are functioning in imagination loops I would guess. I suppose they are exerting some effort to understand. Are they controlling for a perception of an "integration" of prior perceptions into a distinct perception (a partially "cognitive" view)? As such what would be an example of a reference, and at what perceptual level?

In a general sense, would you say that degrees of understanding is simply indicative of unsuccessful attempts to bring perceptions into line with references, and as such would the degree of understanding be inversely related to the magnitude of some persisting error signals?

Rick, could you help out here and articulate a more specific HPCT example of what a model demonstrating degrees of understanding would look like with an individual who is not involved in a physical procedural task such as throwing a baseball or threading a needle, nor is he explaining anything, but is just sitting in a chair, listening and "trying to understand" the teacher.

Do you have a good sense of understanding my concerns?

My concerns are ultimately empirical, as well as applied, but at this point my model is fuzzy. I value your perceptions. Thanks for your words!

Tom Hancock
Grand Canyon University (Summer Research Associate at Armstrong Lab, USAF)

Date: Mon Jun 14, 1993 11:38 am PST
Subject: Social Control Systems

[From Rick Marken (930614.1200)]

In the discussion of social control systems, I heartily align myself with the positions expressed by Tom Bourbon and Bill Powers (why am I not going to have a heart attack and die from the surprise!) as opposed to those of Bruce Nevin and Bob Clark (who seem to believe that social control systems either exist or are useful metaphores). In particular, I like these comments from Tom and Bill:

Tom Bourbon (930614.0840)

>Bob, I think Bill is saying they [social control systems] do not exist
>because they do not exist, not because of their reliance on force.
> "They" (entities called social control systems -- social systems
> that control) are actually individuals, each of whom has goals or
> intentions that bring them into common action around a set of
> shared variables in their environments.

Bill Powers (930613.2200 MDT)] --

> People often try
>to use the concept of social rules and social demands as a way of
>avoiding responsibility for the things they do with their own
>arms, legs, and voices. They like to pretend that they are
>governed by something outside themselves; in that way they can do
>what they really want to do, but when problems or objections
>arise they can pass the buck. "I'm only following orders."

I would like to contribute a personal anecdote that confirms some of Bill's opinions about social control systems.

When I was moving back to LA from the mid west I interviewed for a job with a law firm. It was never clear what the job was to be but they wanted a psychologist for some reason. One of the questions during the interview was something like "do you believe that companies or the individuals in those companies are responsible for what the company does". I had not thought much about the issue currently being discussed on this net (social control systems) but it seemed pretty obvious to me that companies ARE the individuals in them. So it's the individuals, not the company, that is responsible. I recall explaining this fact to the interviewers at some length (it seemed like an important issue to them). I, of course, didn't get the job. I now realize that this was probably a law firm that was busy trying to get individuals (like the president of Union Carbide) off the hook for the disastrous consequences of his or her actions (or lack thereof) by saying that it was really "the company" that did it. While I'm not a big fan of punishment, I do think that people try to get away with doing some pretty miserable things by claiming that it was really the "social system" (of which they are just a thoughtless cog) that "did it"; they were "just following orders", as Bill said.

Social systems are real perceptions -- they are an important and interesting byproduct of many people simultaneously controlling their own perceptions. But social system's (and what they do) are just a side effect of the operation of individual control systems. These side effects can be quite useful to many of the individuals involved in creating the side effects of their mutual efforts-- for example, the bridge that get's built, the opera that get's performed, etc. But the explanation of the social phenomenon exists in each of the individual control systems and their (mutual) relationship to the variables they are controlling.

Tom Bourbon (930614.0920) --

Great discussion of imitating vs repeating!

Best Rick

Date: Mon Jun 14, 1993 3:19 pm PST
Subject: Understanding

[From Rick Marken (930614.1400)] Tom Hancock (930614)--

> In a PCT view, we could say
>that understanding is when the perception is made to match its
>reference signal.

My first definition of "understanding" was that it is "knowledge" embodied in the comparator and output functions that makes it possible for the control system to match its perceptual to its reference signal. What you describe above already has a name, I think -- it's called "control".

>when I consider students in my classes,
>who are perceiving something when I speak or demonstrate, what
>could be happening when they have 100% confidence of understanding
>("I get it perfectly") versus when they have 50% ("I sort of
>understand") or 0% (I am totally confused).

There are many possible answers to this question, many of which are more questions. I presume you have asked the students to rate their confidence that they have "understood" you. What does each of the students think he or she is rating? Each probably has a different idea of "understanding". Maybe they are trying to give a number that is an estimate of what their score on a test of the material might be? Maybe they are trying to impress you-- or win your sympathy. Or maybe they have no clue what you are talking about except that you want them to say a number; and they say it.

I think the problem may be that "understanding" is just a word; what is the phenomenon to which it refers, for you?

> Are they controlling for a
>perception of an "integration" of prior perceptions into a
>distinct perception (a partially "cognitive" view)?

Maybe. That would be a tough one to test until it were made more concrete. How can I observe the degree of integration of prior perceptions into a distinct perception? If this is what they are controlling, it is likely to be something that you can perceive as well. It should also be a perception that can vary and, if it is controlled, is maintained in a reference state against disturbance. I've never had a perception that I would describe as "integration of prior perceptions into a distinct perception". Maybe there's a better way to describe it.

>In a general sense, would you say that degrees of understanding is
>simply indicative of unsuccessful attempts to bring perceptions
>into line with references, and as such would the degree of
>understanding be inversely related to the magnitude of some
>persisting error signals?

No. I think what you are describing is just "poor control", not lack of understanding.

>Rick, could you help out here and articulate a more specific HPCT
>example of what a model demonstrating degrees of understanding
>would look like with an individual who is not involved in a
>physical procedural task such as throwing a baseball or threading
>a needle, nor is he explaining anything, but is just sitting in a
>chair, listening and "trying to understand" the teacher.

I think you have to get a better handle on the phenomenon that you are trying to study. My guess is that you are interested in what happens when a person gets some verbal communications about how to do something (such as build a radio or solve algebra equations but it could also include giving the right answers on a test). A person "understands" to the degree that they can create, in their imagination, an image of themselves succeeding at the thing that they are supposed to do after the lecture. So, in your class, I presume most of the students are trying to produce lots of correct answers on your test. As you talk they are trying to memorize what you say, maybe determine what you might ask on the test, how you might ask it, whether their replies would satisfy you (so they ask questions or volunteer information as a test what you might consider to be "right" answers). So they are controlling, in imagination, for how well they can do on one of your tests. This is just a bunch of guesses, but some might be testable (using The Test).

Maybe you could try explaining once again what you are trying to understand. What have you observed that is interesting to you and that has motivated your research? What have you seen that you want to explain? If you have seen "understanding" then please explain in concrete terms what it is that you've seen.

I think you are interested in a very interesting phenomenon. I just think that a big part of the problem of studying it in PCT terms is identifying the perceptual variables that are involved (both for you and for the people you are studying) in this phenomenon that you call "understanding".

Best Rick

Date: Mon Jun 14, 1993 6:12 pm PST
Subject: Repeating/imitating

[From Bill Powers (930614.1900 MDT)] Tom Bourbon, Bruce Nevin --

Just a comment on the repeating-imitating issue (Bruce, correct me if I misunderstand what you meant).

I think that "repeating" means paraphrasing: conveying the same meaning but in different words. Imitating means reproducing what one heard. You can imitate without understanding, but not repeat without understanding.

Imitation means using the same means to the same end, reproducing not only the higher-level goal, but the lower-level subgoals used to achieve the goal. Repeating means repeating the higher-level end-result by any means available. If I understand Bruce correctly, right now I am repeating what he meant to convey, but not imitating what he said.

Best, Bill P.

Date: Tue Jun 15, 1993 1:15 am PST
 Subject: Re: social control

[Hans Blom, 930615] (Rick Marken (930614.1200))

This is impossible to resist. The discussion of social control systems reminds me of some of the statements that my biology and physiology colleagues often try on me. Let me try the following paraphrases on you, where the 'individual' is a cell and 'social' refers to an organized collection of cells, i.e. a (human) body:

Tom Bourbon (930614.0840)

> "They" (entities called body control systems -- body systems that
 > control) are actually individual cells, each of which has goals
 > or intentions that bring them into common action around a set of
 > shared variables in their environments.

This sounds all right, doesn't it?

Bill Powers (930613.2200 MDT)] --

>Body cells often try
 >to use the concept of body rules and body demands as a way of
 >avoiding responsibility for the things they do with their own
 >actuator mechanisms. They like to pretend that they are
 >governed by something outside themselves; in that way they can do
 >what they really want to do, but when problems or objections
 >arise they can pass the buck. "I'm only following orders."

But this sounds funny, isn't it? Responsibility isn't a term we use for cells. Could it be that notions like responsibility are orthogonal to our field of study? Or is a cancer cell worthy of our reproach?

>The body's systems are real perceptions -- they are an important and
 >interesting byproduct of many cells simultaneously controlling
 >their own perceptions. But the body's system's (and what they do) are
 >just a side effect of the operation of individual control systems.
 >These side effects can be quite useful to many of the individuals
 >involved in creating the side effects of their mutual efforts-- for
 >example, the food that is obtained, the move to a better feeding place,
 >etc. But the explanation of the body phenomena exist in each of
 >the individual control systems and their (mutual) relationship to the
 >variables they are controlling.

Let's try not to forget that when we study PEOPLE, we focus on THEM as our natural units. We do not regard their components as units of discourse, nor the multi-unit assemblies that they might form in order to make life easier for each. In my opinion, there is no reason NOT to talk about social control systems; 'higher level' concepts often emerge when many units interact; in statistical mechanics you can talk about concepts like temperature and pressure, concepts which do not apply to single molecules.

Greetings, Hans Blom

Date: Tue Jun 15, 1993 1:48 am PST
 Subject: Re: social control

[From Oded Maler 931615] Hans Blom, 930615

* >Body cells often try
 * >to use the concept of body rules and body demands as a way of
 * >avoiding responsibility for the things they do with their own
 * >actuator mechanisms. They like to pretend that they are
 * >governed by something outside themselves; in that way they can do
 * >what they really want to do, but when problems or objections
 * >arise they can pass the buck. "I'm only following orders."

Etc.

This is exactly what I was going to say, but it is useless. It turns out that there is something in those cell assemblies (aka BP and RM) which resists the idea that their level of organization is not the ultimate one. While they can look *down* and see how their own selves emerge from autonomous cells, who control some biochemical individual perceptions, they cannot look *up* and see the emergent social, historical entities, for which "autonomous" human beings are nothing but sensors, actuators and comparators.

--Oded

Date: Tue Jun 15, 1993 3:53 am PST
Subject: repetition != paraphrase

[From: Bruce Nevin (Tue 930615 07:22:35 EDT)]

(Bill Powers (930614.1900 MDT)) --

> I think that "repeating" means paraphrasing: conveying the same
> meaning but in different words. Imitating means reproducing what
> one heard. You can imitate without understanding, but not repeat
> without understanding.

> Imitation means using the same means to the same end, reproducing
> not only the higher-level goal, but the lower-level subgoals used
> to achieve the goal. Repeating means repeating the higher-level
> end-result by any means available. If I understand Bruce
> correctly, right now I am repeating what he meant to convey, but
> not imitating what he said.

Thanks for making your paraphrase explicit. It is indeed a misunderstanding.

A repetition is the same words in your pronunciation of them. It is not a paraphrase (same meanings in your words and sentences).

This is concealed by our use of written words to communicate with one another here. A limited analogy can be made even so. Your screen is amber and black (say), whereas mine is a black on white window on a workstation. Or suppose we are examining hard copy. Mine is printed on a QMS laser printer, their version of "courier" font, 11 point, whereas yours is on some other printer in some other point size. This analogy is limited because it assumes the identical segmentation into a linear sequence of letters, whereas my segmentation of words into sequential and concurrent phonemic elements may differ in number and in kind from your segmentation. Nonetheless, I perceive your pronunciation of words as repetitions of words that I have said, or that I might say, or that I have heard said. They are not imitations of my pronunciations, nor are my pronunciations imitations of words I have heard spoken (though in the very initial stages of learning, before phonemes, they were).

When you paraphrase what you understand me to have said, you use words that I perceive as repetitions of words that I can say. Even though what you think of as a paraphrase is in fact not (does not repeat the intended meanings using different words and sentences), it does nonetheless repeat words that I know and can also repeat, each word, in my own pronunciation (or handwriting).

Meaning has nothing to do with it. Phonemic repetition is below the level of meaning, though not so far below meaning as phonetic imitation is. I can repeat the word "beet" or "beat" after you without knowing or needing to know which you have said (although I will assume one), and when I do this I am repeating in my own phonemes, not imitating. I can repeat the word "cow" when I hear my brother say it. Or I can imitate his pronunciation, in which case it sounds and feels to me more like "kayo" (K.O.). He has lived since 1953 or 1954 (since about age 9) in Florida and Georgia.

Meaning is essential to initially learning to control phonemes as means for controlling contrast between words (as representations of word-contrasts). You have to be told (directly or implicitly) what is repetition and what is contrast (different words), and that is the irreducible minimum of meaning: a difference that makes a difference.

I have been nibbling away at a reply to earlier posts from you and Tom. Perhaps by miracle this will obviate the need for that reply. If this brief clarification makes a difference for you, I hope you will look at your last posts in this thread and see if you would say anything differently now.

Bruce bn@bbn.com

Date: Tue Jun 15, 1993 8:56 am PST
 Subject: Scholarships; Social Control; Repetition & paraphrasing

[From Bill Powers (930615.0900 MDT)]

Note on the meeting. All 5 available student scholarships have been awarded; notifications going out today or tomorrow. If you did NOT send in a registration applying for a scholarship, you do NOT get one: you will have to pay the \$145 (double occupancy) and, like those who got the student scholarships, your \$5 membership. A number of people said some time ago that they wanted scholarships, and were assured they were available, but some of them never returned any of several applications sent to them. They're out of luck. The people who actually applied got them. We are a not-for-profit organization: when the money's gone, it's gone.

Many thanks to Martin Taylor for a generous contribution to the scholarship fund.

 Hans Blom (930615) and Oded Maler (930615) --

Hans said

>Responsibility isn't a term we use for cells. Could it be that
 >notions like responsibility are orthogonal to our field of
 >study? Or is a cancer cell worthy of our reproach?

And Oded agreed.

You're both missing the point. The cells in a body (for example, liver cells) receive reference signals from superordinate systems like the pituitary that specify for those cells the level of their controlled variables that they are to maintain (blood glucose concentration). Liver cells are responsible for maintaining the specified level of input, but they have no say about what level will be specified. Similarly, reference signals for the spinal control systems tell those systems what level of muscle tension to maintain; those system are responsible for maintaining that level (they are the immediate cause and provide the immediate means), but they have no choice but to try to produce whatever level of sensed tension is specified.

The whole human hierarchy is, according to HPCT, organized this way: systems at each level of organization receive reference signals directly from superordinate systems (except for the top level where a different reference-setting mechanism is obviously required). A subsystem doesn't have to be persuaded to go along with the goal it is given; it simply does so because that's what it's organized to do. It is a subordinate system by the very nature of its relation to higher systems.

At the top level, reference signals are selected by reorganizing processes or something similar; the criterion for reorganization relates to internal states of the organism that are specified by inherited reference signals. So the origin of the highest level reference signals is basically experimental, with the criteria for selection from the results being an expression of evolutionary history. The line of command does not pass outside the organism at all. There is no way for reference signals to be received from outside.

This means that however the environment (including other people) interacts with a whole organism, it cannot do so by adjusting reference signals inside the individual. It can only interact through physical effects, either directly on the body or through the lowest levels of perception, the sensory organs. The internal principle of organization of the individual, according to HPCT, does not extend to relationships among individuals; no individual has the capacity to set reference signals for another one. All direct interactions between people happen at the lowest levels of organization.

This is not to say that societies don't have emergent properties; the properties would be those typical of an interacting set of independent control systems, as in the Gatherings simulation. There can be conflicts, oscillations, positive feedback, and all sorts of other interactions. But there are no superordinate systems in a society to act on a person as the pituitary acts on the liver -- to coordinate reference settings in different individuals and adjust them so that some higher-level social outcome will be maintained against

disturbances. All there are are people, and all they can do is control their own perceptions relative to their own internally-generated reference signals.

This is what the HPCT model implies. If you want to imagine that there are superordinate intelligences setting all of our highest reference signals, that is up to you. I would have to see some convincing evidence of their existence before I took them seriously, but you're welcome to try. If you want to junk the HPCT model and propose a better one, feel free. Maybe you have a better idea. But the above is what comes out of the HPCT model as far as I can tell, and until something better comes along I don't see any reason to abandon it.

Bruce Nevin (930615.0722 EDT) --

>A repetition is the same words in your pronunciation of them.
>It is not a paraphrase (same meanings in your words and sentences).

>When you paraphrase what you understand me to have said, you use words that I perceive as repetitions of words that I can say. Even though what you think of as a paraphrase is in fact not (does not repeat the intended meanings using different words and sentences), it does nonetheless repeat words that I know and can also repeat, each word, in my own pronunciation (or handwriting).

>Meaning has nothing to do with it. Phonemic repetition is below the level of meaning, though not so far below meaning as phonetic imitation is. I can repeat the word "beet" or "beat" after you without knowing or needing to know which you have said (although I will assume one), and when I do this I am repeating in my own phonemes, not imitating.

All this, to me, seems very similar to paraphrasing -- but perhaps that's because of how I think of paraphrasing. I see it as controlling one perception by means of varying others. I can control the perception of a word by varying my perceptions of phonemes; there is a range of phonemes that will lead to perception of the same word, so in that sense a Northern pronunciation is a paraphrase of a Southern one: they both have the same meaning, the word, and I have a choice of different ways of saying what I perceive as the same word. At a higher level, we can control the perception of a mental image by manipulating the words used to evoke it; many sets of words can evoke the same image, so are considered paraphrases.

It seems to me that there are lots of phenomena treated as separate in linguistics that reduce to the same phenomenon in HPCT. Why not take advantage of the theory to simplify the conceptual structure?

You're still talking about "contrasts." Are you seriously proposing that contrast is the variable that is perceived and controlled?

Best to all, Bill P.

Date: Tue Jun 15, 1993 9:04 am PST
Subject: July Meetings - Procrastinators 'R Us

[Dan Miller (930615.1200)] Bill Powers,

You've got my number. I have postponed this commitment too long, but not totally. My conference fees are in the works. The check will be in the mail ASAP. Barbara and I will be arriving in Durango on Tuesday, July 27th, hungry and thirsty after a day in some Anasazi ruins.

Also, I will spend a couple of days in Santa Fe talking to the complexity wonks. They have it all wrong. Their social and economic systems ideas are more scattered than those of structural sociologists. We will see if there is an elective affinity between them and PCT.

Hans Blom and Oded Maler,

Your analogies from cells to biological systems (to individual and social systems) is very interesting. It was enjoyable to read, but I'm not altogether convinced. I would prefer to discuss these problems as social interactions or social processes. Interacting control systems and collective control systems can engage in some interesting, important, and essential social acts. But I do not think that social organizations like the Control Systems Group act like control systems. But I am willing to be convinced otherwise. Discussions at this year's conference should be lively.

Later, Dan Miller millerd@dayton.bitnet

Date: Tue Jun 15, 1993 9:14 am PST
Subject: Cells and People

[From Rick Marken (930615.0900)] Hans Blom (930615) and Oded Maler (931615)

Let me see if I understand what your analogy where you compare individual cells to individual people. I think you are pointing out that both cells and people are control systems -- which I accept as true. The behavior of collections of cells (such as the cells that make up an organ, like the liver) is regulated by systems external to the cells. So, for example, the collective behavior of the cells in the liver is regulated via hormonal and neural control systems in the brain. This shows (I think you are saying) that the collective behavior of humans is regulated via control systems (Oded's "emergent social, historical entities" and others' "social control systems") that are outside of the collective. Is this what you guys are saying?

Assuming that it is what you are saying, let me try to explain where I have problems with the analogy. I agree that the behavior of collections of cells is regulated by other systems -- the collectives are part of a higher level control system. But this higher level control system must be using properties of these cells that the individual cells themselves are not controlling. For example, the individual cells may be controlling the differential concentration of electrolytes across the cell wall -- doing this by varying the permeability of the cell wall. This means that the concentration of electrolytes outside of each and every cell is part of a variable controlled by each cell -- possibly relative to a slightly different reference specification. The fact that electrolyte concentration is controlled by the individual cells means that a higher level control system cannot "use" the collection of cells to control a variable that strongly depends on electrolyte concentration. The higher level control system can really only effectively use uncontrolled variable aspects of cell behavior as part of its means of controlling whatever variable(s) it controls. For example, it can use the cell's actions that control electrolytes (the permeability changes) as part of its means of control because the higher level system can make unresisted variations in this variable by disturbing the variable being controlled by the cells -- cell permeability is an uncontrolled action of the cell.

Suppose that changes in permeability required to control electrolyte concentration cause changes in the size of the cell. Then the higher level system can use changes in the size of the cell to control whatever variable it is controlling -- like net blood flow through the collective of cells (the liver). So the higher level system can vary the concentration of electrolytes outside of the cell collective in order to influence the net size of the cells in the collective and influence blood flow through the liver -- which is the variable controlled by the higher level system.

Things could work this way with collectives of people; there could be a higher level control system, outside of the collective of people, that uses side effects of the control exerted by the individuals in the collective as a means of controlling whatever variables are being controlled by the higher level system. I'm perfectly willing to believe that there are such higher level social control systems. I just don't see any evidence of them. I see no evidence of side effects of collective action being under control. The results of collective action that I have seen controlled seem to be controlled by the individuals themselves. No outside control system is using the side effects of people controlling other variables to create a skyscraper or a symphony. In each case each individual is controlling a variable specifically in order to produce the intended collective result.

If you really think that there are higher level social control systems that are using side effects of human control to produce the results they (not the individual people) intend, then I would really like to see some evidence of it.

Oded claims:

>there is something in those cell assemblies (aka BP and RM) which resists
>the idea that their level of organization is not the ultimate one.

The reason we "resist" is because we have seen no evidence. If you have some real strong evidence of control by a higher level control system then please present it. The evidence for the existence of higher level control would be a variable that is clearly maintained against disturbance by effects produced by a collection of individuals who have absolutely no interest in the state of that variable; it is not controlled by the individuals, it is maintained AGAINST DISTURBANCE by some other control system.

Happy hunting Rick

Date: Tue Jun 15, 1993 10:43 am PST

[From Richard Thurman (930611.1430)] Rick Marken (930611.0800) --

>>I'm curious about the 'control of sequence' vs 'control of
>>configuration' study you were about to do last March. Anything to it yet?

>I haven't worked on it at all. My current priority is trying to
>write a book about PCT; something along the lines of "Perceptual
>Park" with huge, cloned control systems running around terrifying visit

Hey! What a great idea! You could beat everybody to the evolutionary chase. Why mess around with writing roach control algorithms (a la Beer bugs) when you can cut right to flesh eating carnivores!

>I wish you would come to the CSG meeting in July. I'll bring the
>program that is sort of the prototype of the kind of experiment
>I had in mind. I think we could make much better progress if we
>could sit down together and brainstorm about possible ways to
>do the experiment -- and possibly test some ideas right there
>using HyperCard of Basic or Pascal.

Unfortunately I have other commitments and I will not be able to attend the CSG meeting in July. I agree that the efficiency of any research efforts would increase dramatically if we could spend some good solid time working together in a setting such as the CSG meeting. However, at this point, we will just have to work it out via the net.

>If you (or Tom H.) can't make the meeting, I'm certainly willing to try
>to brainstorm on the net about possible research projects. Some of
>the "scholarship" on the configuration/sequence control stuff already
>exists in my "Hierarchical control of perception" paper; did you get a copy?

I have a copy of a paper called "The Hierarchical Behavior of Perception" which you put on the net last October. It has a very helpful section called "Levels of Perception" in which 'Perceptual Speed Limits' are discussed. I will try to get more familiar with it over the next few days.

As far as brainstorming over the network goes... I think what I need is a few more ideas on how move from researching lower level skills to doing higher levels. Both Tom H. and I (oops... maybe I should just speak for myself) are steeped in the cog. sci. tradition and may not always see how to apply 'tracking task' studies to more 'cognitive' tasks. Also I have to admit a weakness in moving from a more traditional group-means-statistical approach to a modeling based approach. Am I the only one, or are others on the net struggling with a change in research paradigms?

On another note -- I just cant help myself from butting into someone else's thread.

Rick Marken (930611.0900) replying to Martin Taylor (930611 10:30) --

>The people we are talking about (Carver, Scheier, et al) don't need
>to be "missionized"; they already believe in the Gospel According to
>PCT and claim to be PCTers (or, at least, Powersian Control Theorists;
>hey, another meaning of PCT). While I do desire that they come to CSG
>meetings and participate on the net, I desire it, not because I want them

>to join the fold (they are ostensibly already in the fold); I just want to
 >engage in an exchange of ideas with colleagues. Since these people are
 >already converts, I am puzzled by their reluctance to come to "church".
 >I don't want to force them to do it; I just don't understand the reluctance.
 >It is really kind of a strange thing. The analogy seems to me to be that
 >Carver and Scheier et al are like self-converted Mormons who are
 >enthusiastic about being Mormon's but refuse to go to the Mormon temple.
 >In fact, they continue to schmooze at their old church -- Our Lady of
 >Lineal Causality. It's just puzzling to me, that's all.

Mormons used to (around the turn of the century) have a term for such individuals. They referred to them as Jack Mormons. I believe the implication was that just as you couldn't get a jackass to come around to your way of doing things, some Mormons just wouldn't budge from some preconceived notions. They would just sit there braying and holding up traffic ... so to speak.

Does this mean we need to start referring to the Carver and Scheier crowd as Jack PCTers?

Martin Taylor (930611 10:30) --

>>They presumably think that the discussions
 >>on the net and at the meeting will be as unprofitable as the Mormon would
 >>presumably think a session of bathing in the Ganges to be. Why bother?

Rick Marken (930611.0900)

>I agree. So, again, it is puzzling. They should know that the CSG meeting
 >is filled with fellow Mormons (er... PCTers) and that no Ganges bathing is
 >done at all -- yet, indeed, they treat their fellows (us) as though we were
 >the infidels. It's just puzzling. If they don't really like PCT then why don't
 >they just use a different model?

I know that the gist of this thread is winding in another direction, but I would just like to say as convert to (and former missionary for) the Mormon Church I have immensely enjoyed the PCT and religion discussion. I think Joel Judd, Greg Williams, Bill Powers, and Rick Marken (did I miss anybody) did a great job of keeping the tread on a PCT relevant level. It helped me see a little more clearly just how PCT can be integrated with (as well as assess) ones personal belief system.

PCT and CSG-L -- ain't they great! Where else could the topic of scholarly conversation range all the way from 'perceptual theme parks' to 'non-Ganges- river-baithing Mormons' to 'perceptual speed limits!'

Richard Thurman Air Force Armstrong Lab BLDG. 558 Williams AFB AZ. 85240-6457

(602) 988-6561 Internet: Thurman@192.207.189.65

Date: Tue Jun 15, 1993 10:50 am PST
 Subject: Re: Repeating/imitating

From Tom Bourbon (930615.1235)

My access to the net went south yesterday morning and it just came back.

>[From Bill Powers (930614.1900 MDT)]

>

>Tom Bourbon, Bruce Nevin --

>

>Just a comment on the repeating-imitating issue (Bruce, correct
 >me if I misunderstand what you meant).

>

>I think that "repeating" means paraphrasing: conveying the same
 >meaning but in different words. Imitating means reproducing what
 >one heard. You can imitate without understanding, but not repeat
 >without understanding.

>

>Imitation means using the same means to the same end, reproducing
 >not only the higher-level goal, but the lower-level subgoals used
 >to achieve the goal. Repeating means repeating the higher-level
 >end-result by any means available. If I understand Bruce

>correctly, right now I am repeating what he meant to convey, but
>not imitating what he said.

Bill, I thought of those meanings, also. As I mentioned yesterday, I elected to apply the more literal meaning -- mainly because I could more easily model that one! I was hoping Bruce would be the one to "set me straight." In that case, my reply would have been that I know of no working PCT model that produces paraphrasing, but that the *productions*, or *actions*, would occur as controlled perceptions relative to reference signals in a time series like the triangular wave in my example of repeating. The reference for the paraphrasing would play the role of the remembered, or imagined, reference waveform in my example.

I was also thinking that if, by "repeating" Bruce meant "paraphrasing," and I replied as though he meant "repeating," he might vary his outputs toward the end of seeing my replies coming more in line with what he intended, thereby providing an example of one of my other points in my reply to him the week before last.

Until later, Tom Bourbon

Date: Tue Jun 15, 1993 11:53 am PST
Subject: Re: Cells and People

From Tom Bourbon (930615.1402)

Holy smokes! Let your server go down for a day and look what happens on this net! As a point to rejoin this discussion, I will use Rick's post.

>[From Rick Marken (930615.0900)] Who replied to:

>

>Hans Blom (930615) and Oded Maler (931615) --

>

>Let me see if I understand what your analogy where you compare
>individual cells to individual people. I think you are pointing out
>that both cells and people are control systems -- which I accept
>as true. The behavior of collections of cells (such as the cells that
>make up an organ, like the liver) is regulated by systems external
>to the cells. So, for example, the collective behavior of the cells
>in the liver is regulated via hormonal and neural control systems
>in the brain. This shows (I think you are saying) that the collective
>behavior of humans is regulated via control systems (Oded's "emergent
>social, historical entities" and others' "social control systems") that
>are outside of the collective. Is this what you guys are saying?

Is that what you guys are saying? Bill Powers (930615.0900) also posted along the same line as Rick. I think he might also ask you this question.

>If you really think that there are higher level social control systems
>that are using side effects of human control to produce the results
>they (not the individual people) intend, then I would really like to see
>some evidence of it.

As would I.

>Oded claims:

>

>>there is something in those cell assemblies (aka BP and RM)
>>which resists the idea that their level of organization is not the
>>ultimate one.

Will anyone ever catch on to the fact that I also "resist" those claims?

>The reason we "resist" is because we have seen no evidence. If
>you have some real strong evidence of control by a higher level
>control system then please present it. The evidence for the
>existence of higher level control would be a variable that is clearly
>maintained against disturbance by effects produced by a collection
>of individuals who have absolutely no interest in the state of that
>variable; it is not controlled by the individuals, it is maintained

>AGAINST DISTURBANCE by some other control system.

Exactly!

Hans and Oded, in lieu of working models, why don't you cast your ideas in the form of a few simple PCT diagrams that map out the interactions you think occur among cells in a person, and among people in a group, so that we can *see* why the two of you think they are the same? As an example of what I mean, here is a re-post of part of Tom Bourbon (930611.0750) in which I offered my idea of a comparison between people as control systems and social groups as "control systems."

=====

Subject: Re: BBS4 From Tom Bourbon (930611.0750)

=====

Hans and Oded, if each of you were to cast your ideas into similar diagrams, I might more easily see what you mean when you equate people in social groups with cells in a person. When I use PCT to model individuals in groups, I think of an organization like that shown in CASE 3. When I think of where a cell fits in the body, it is more like one element in the output box at the lower level of CASE 2. How do you envisage them?

Until later, Tom Bourbon

Date: Tue Jun 15, 1993 11:56 am PST
Subject: Moving Mountains; Wegener & Powers

[from Gary Cziko 930615.2000 UTC]

Just back from Costa Rica where my son and I saw in the mountains, valleys and active volcanoes lots of evidence of restless, dynamic earth.

That's the lead-in. Now the quote from

Hallam A. (1975, February). Alfred Wegener and the hypothesis of continental drift. Scientific American.

"Perhaps the long travail of Wegener's hypothesis can be best explained as a consequence of inertia. A geologist at the 1928 symposium of the American Association of Petroleum Geologists is reported to have said: "If we are to believe Wegener's hypothesis, we must forget everything that has been learned in the past 70 years and start all over again." It should also be remembered that to the geologists of the time Wegener was an outsider; they must have regarded him as an amateur. Today, of course, we can see that his position was an advantage because he had no stake in preserving the conventional viewpoint. Moreover, we can see that he was not an amateur after all but an interdisciplinary investigator of talent and vision who surely qualifies for a niche in the panteon of great scientists."

Do I just have a wild imagination, or can others see some parallel here involving Wegener and Powers, continental drift and PCT.--Gary

Date: Tue Jun 15, 1993 12:34 pm PST
Subject: Cells, People, PCT Reseach

[From Rick Marken (930615.1300)] Bill Powers (930615.0900 MDT)--

> The cells in a body (for example, >liver cells) receive reference signals from superordinate systems

Excellent discussion of the cell analogy. I just want to point out that my discussion was based on 1) a somewhat mythical physiology that was developed to 2) give Hans and Oded the benefit of the doubt; that is, my model of the cells was based on accepting the analogy of cells to people. There is no outside access to human reference signals so I assumed no outside access to cell reference signals. Given this assumption, my description of how collectives of such autonomous cells could be "used" by a higher order control system is correct. But your example (where the cells receive reference signals from the higher level system) is surely a more accurate description of how collectives of cells (we both picked

the liver -- what could we be thinking of?) are used as part of the hierarchical control process.

Richard Thurman (930611.1430)]

>I think what I need is a few more ideas on how move from researching
>lower level skills to doing higher levels.

I think this is a LOT easier than I used to think it was. Let's start with some examples from the current literature of what you might consider to be studies of higher level skills. How about the Tversky type decision making stuff? Or some problem solving experiments. If conventional methodology can address "higher level" skills than PCT methodology (which is really only subtly different) should be able to handle it too.

>Also I have to admit a weakness in moving from a more traditional
>group-means-statistical approach to a modeling based approach. Am I the
>only one, or are others on the net struggling with a change in research
>paradigms?

I had (have?) the same problem. The problem is really that we have learned that "research" means controlling for a particular result (a big difference between means of groups of numbers, correlations between measures of variables, etc). And we have been taught how to control for those results. Among the most powerful means we have of controlling for these perceptions of "good research results" is statistics. Once you've learned to play with numbers using statistics it's hard to stop -- especially if you get good at it (and I got to be VERY good at it) -- because you are IN CONTROL.

Learning to do PCT research requires some reorganization -- which means losing control for a while with no guarantee that you will get it back. My reorganization process has led me to focus on "stick wiggling" research; I can control the results of that pretty well now. But I would like to explore the possibility of doing QUANTITATIVE research on higher levels of control; I have done it a little (when I was a professor and had students who would do the research for me).

One thing I can tell you about research on higher level variables; it will be an ITERATIVE PROCESS. The first thing to give up when you do PCT research is the idea that you can sit in your arm chair and dream up the great experimentum cruxis. When you actually get down to the business of doing experiments, things will almost surely NOT come out as you expected. This has to be treated as a problem with one's own conception of the experiment (the controlled variable, what disturbs it, what the result of disturbance should be, etc) -- not a job for statistical cleansing. Keep fiddling with things until you get HIGHLY RELIABLE results. Then you will know that you have actually discovered a FACT about human behavior.

But try not to get frustrated when you start exploring PCT research. At least you will be getting a first hand look at the process of reorganization; the reorganization of your own concept of what it means to do good research.

Best Rick
Date: Tue Jun 15, 1993 5:22 pm PST
Subject: Myths, Ed Ford

[From Dag Forssell (930615 18.10)]

Subjects: 1) REQUEST FOR HELP WITH MYTHS TYPE 2
2) REPORT ON ED FORD SEMINAR
3) ED FORD BOOK REVIEW

REQUEST FOR HELP WITH MYTHS TYPE 2:

I have followed with great interest the definition and gathering of myths and misunderstandings about PCT.

I have also considered how these myths and the counter-arguments can be used to make a case for PCT. Bill's PCT joint paper is an outgrowth of this thread.

In the last two days, I have reviewed the posts with the intent of composing a two page promotional piece. I suddenly realize that the myths we deal with are of a particular kind: Truths about Control Theory which are misunderstood. These I now label MYTHS TYPE 1.

What I began to visualize turn out to be of another kind: MYTHS TYPE 2: _Commonly accepted truths which are demonstrably false._ A clear articulation of these may be much more useful to promote PCT. It would be much more upsetting and demanding of attention to show that what people think they know is false, than to say that something they never heard of, have gotten along fine without, and don't care about is right.

The one I noticed in our listings and discussions of myths is:

```
>Tom Bourbon (930611.0750)
>MIS A. "Commands (which I use here to also mean feedforward,
>plans, programs, heredity, genes and similar directives) determine
>behavior, ....
```

Rick (by phone) suggested: "Reinforcement selects behavior."

I now visualize a chart of MYTHS TYPE 2 as follows:

Generally accepted understanding of human behavior:

1) Thinking and behavior (Labeling unnecessary; depends on layout
and correspondence to list on second page)

- a) Psychological (cognitive) science says:
"(Internal) Commands determine behavior"
- b) Common leadership understanding:
"People do as you tell them" (Not a good translation, eh)?
"People do what they want to do" (better)?

2) Repeat performance

- a) Psychological science says:
Reinforcement selects behavior
- b) Common management practice:
Rewards make people continue to perform

Perhaps a page of 8-12 statements arranged as columns a) and b) will allow the reader to identify and agree with the portrayal of contemporary wisdom, and see the influence of psychological theories on "common sense understanding" and "management practice".

(Relating to these):

Our parenting, educational, managerial and leadership practices are based on these understandings. They all apply in our private lives, in business and in government. When we have problems in our families, in our schools, in business and in society, we solve these problems by applying our understandings to the best of our ability.

(Then some statement to the effect that):

Each of these accepted "truths" listed above is DEMONSTRABLY FALSE. It is very easy to demonstrate that they are false. All that is required is to recognize and understand the phenomenon of control. Control as a phenomenon has been clearly understood in the engineering sciences since 1927, but is not well known and has been misunderstood by many (under the label of cybernetics). Scientists (in the behavioral sciences) have been victims of overly simplified explanations of control. Many behavioral scientists claim that control cannot explain human activity and they are right. The way they understand and describe it, it cannot work and cannot explain.

In spite of historical misunderstanding, a new science, explaining how living organisms work, has been developed based on sound engineering principles and biology. Developed into a detailed model by William T. Powers, Perceptual Control Theory (PCT) offers a clear understanding of what actually is going on when people do what they appear to do, by themselves or interacting socially.

PCT gives clear insight into the phenomena listed above (other side) but with a very different interpretation.

Revised understanding of human behavior:

1) Thinking and behavior

- c) Perceptual Control Theory says:
 "Behavior satisfies internal wants despite (and compensates for) external disturbances.

2) Repeat performance

- c) Perceptual Control Theory says:
 People will always do what they can to satisfy internal wants. Reinforcement as such is an illusion. Rewards actually serve to reduce behavior if they help the person satisfy the internal want.

With a clear science of behavior that lends itself to testing and validation, the path is clear to unprecedented progress in parenting, educational, managerial and leadership practices.

It is very exciting to participate in the development of a true science of life, based on engineering principles. It does require clear thinking, because both the explanations and the conclusions are different from what is widely understood (but incomplete or false) today.

To understand the situation we find ourselves in today, imagine that you live in the 17th century and study the science of chemistry. Contemporary science is based on practical know how, developed by trial and error over 2,000 years, resulting in capable metalworking, plating and more ordinary chemistry of many kinds. Mixed with the practical recipes passed down from past generations of scientists is some philosophy, astrology and mysticism. Yes, alchemy works, but the scientists who know what to do have no clear understanding of the underlying processes. They cannot know in detail why and how the chemistry works.

You live in the middle of a society accepting and dependent on alchemy, where the scientists *know* what they know, are proud of it, respected, and the authorities on their specialty. They write the textbooks used in chemistry school, and referee and edit the scientific journals. It is very difficult to imagine a different chemistry with different ground rules, different explanations and different results.

Fast forward to the late twentieth century! The science of chemistry is now based on clear engineering principles, including recognition of the basic elements of the periodic table, (beginning with oxygen, identified in 1774). Chemists can predict results and design new compounds even before they mix chemicals, because they have a carefully tested and validated theory that explains what goes on as the elements interact.

In the same late twentieth century, the contemporary sciences of psychology are based on practical know how, developed by trial and error (including statistical studies) over centuries, resulting in a variety of therapies and management practices. Mixed with the practical suggestions passed down from past generations of scientists is some philosophy, descriptions and symptomatic explanations. Yes, psychology works, but the scientists who know what to do have no clear understanding of the underlying processes. They cannot know in detail why and how human behavior works.

You live in the middle of a society accepting and dependent on our contemporary behavioral sciences, where the scientists *know* what they know, are proud of it, respected, and the authorities on their specialty. They write the textbooks used in schools of psychology and sociology, and referee and edit the scientific journals. It is very difficult to imagine a different behavioral science with different ground rules, different explanations and different results.

Fast forward to the late twenty-first century! The science of psychology is now based on clear engineering principles, including recognition and an *accurate* explanation of the phenomenon of control. Parents, spouses, managers, leaders, workers - everybody is better able to develop satisfying, productive lives, because people have a carefully tested and validated theory that explains what goes on. PCT is so basic to human growth and development, and so easy to understand, that it is taught in elementary school.

It is a curious byproduct of the PCT revolution that historians studying the (confusing and ineffective) 19th and 20th century approaches to psychological therapies and management practices can clearly understand why all the various therapies and leadership approaches had some effectiveness. (All had some satisfied customers and none could be ruled out as invalid).

Very simply, because (as we now understand) all people *are* control systems and *reorganize* when control is ineffective, people would naturally seek more effective ways to control their perceptions and spontaneously reorganize regardless of whether the therapist or manager focused on and discussed their dreams, actions, emotions or desires, told them what to do or let them work things out for themselves. It was generally understood at the time, of course, that some therapy had to go on for very long time periods to give results.

Another way to express the last paragraph:

Any therapy works...
Any management style works...
... because people *are* control systems.

Any mistakes our parents / leaders / counselors / teachers make are compensated by the fact that we *are* control systems.

Whatever beliefs we hold about human nature makes little difference because we *are* control systems.

Any counseling, and management approach - of whatever flavor - works because the person considers wants in the course of any conversation and re-organizes naturally.

How about it? I hope for some suggestions. Please post what comes to mind as MYTH TYPE 2:

- a) Psychological (cognitive, behavioral) science says:
- b) Common leadership understanding:
- c) Perceptual Control Theory says:

I look forward to comments on my rambling musings above.

REPORT ON ED FORD SEMINAR

Ed presented his four day program on "effective parenting" or "teaching responsibility thinking" based on PCT to a group of teachers, a parent, a social worker and several principals last week in Green River, Wyoming.

Christine and I were privileged to attend. We were particularly interested in learning from Ed some ways to simplify our presentation to the industrial audience and improve our presentation in any other way.

We were not disappointed. In particular, I saw how Ed collects feedback at the end of each day and *responds* the next morning to any questions. On the morning of the second day, this took 1 1/2 hour, much less on subsequent days. This involves the audience and shapes the program.

Ed also has a good way to illustrate the control process by outlining a scenario of wants, actions, controlled variable, disturbance and perception with a simple story of what a student wants and does, how the teacher interferes, etc. Quite humorous.

The teachers gave Ed rave reviews. I could see that Ed's background as a teacher, counselor and parent has given him a rich understanding of the school environment. Ed earned the teacher's respect for his competence and understanding of their situation. (At one point he gave them suggestions on how to use the school district lawyer in ways they had not thought of). He also flows well with *their* concerns, not his own, with his very casual, laid back and respectful approach.

ED FORD BOOK REVIEW

When I first visited Ed (a full day, 3 1/2 years ago), I bought every book and tape in several copies. Ed told me very clearly that Freedom from Stress was the latest, most

comprehensive and most accurate. I read it and have learned much from it. I have recommended it and distributed it ever since. I never got around to read Ed's earlier works, though a friend and neighbor of mine did. He told me that he could discern Ed's growth in the sequence of books.

In *Green River*, I found that Ed had sent *_Freedom from Stress_* (1989) AND *_Love Guaranteed_* (1987) as required reading. I read *_Love Guaranteed_* on the plane back. This book is an introduction to Control Theory just like *_Freedom from Stress_*. I think it has a lot of merit, and will start recommending it. The foreword by Bill Powers is excellent and any technical flaws are of no consequence. It is much less voluminous (estimated 40% of *_Freedom_*) and comprehensive. It does not cover all aspects of PCT (such as reorganization) but focuses on a straightforward explanation of basic PCT (with 10 levels of perception) and a single, clear application: Building relationships through quality time. Thus it covers the most important relationship application clearly and simply.

Those on the CSGnet who profess to be interested in applications of PCT owe it to themselves and PCT to read Ed's books! I am glad I finally discovered Ed's earlier, simpler and more straightforward effort.

Thanks, Ed! Best to all, Dag

Date: Wed Jun 16, 1993 1:20 am PST
Subject: social / cell control

[Hans Blom 930616] (Rick Marken (930615.0900))

> I think you are pointing out
>that both cells and people are control systems -- which I accept
>as true.

This is not exactly what I'm trying to say. What I mean is that both cells and people can be MODELLED as control systems, and in quite comparable ways. What cells and people ARE will, in my humble opinion, remain beyond my understanding (i.e. modelling powers) forever.

> The behavior of collections of cells (such as the cells that
>make up an organ, like the liver) is regulated by systems external
>to the cells.

That is not what I perceive. I see that cells live in and depend upon an environment (mainly of other cells), and that they, in turn, shape that environment for the other cells in the collection. My perception is more that of a collection of interacting autonomous agents, very much like in the CROWDS or GATHERINGS programs. If the cell does not agree well with its environment, it may become sick, and the products of its sickness in turn may deteriorate the other cells' environment. What all individuals share is a common environment that works two ways: it influences and is influenced by each individual. Through this common environment, cells must necessarily influence each other. I therefore do not agree with your statement

> I agree that the behavior of collections
>of cells is regulated by other systems -- the collectives are part of
>a higher level control system.

There is no higher level control system, there is just a common environment. I agree with you when you see no evidence of higher level control systems. So why postulate them?

> I'm perfectly willing to believe that there
>are such higher level social control systems. I just don't see any
>evidence of them. I see no evidence of side effects of collective
>action being under control. The results of collective action that I
>have seen controlled seem to be controlled by the individuals
>themselves.

Collectives build common environments and if they cooperate (have common goals) they can, indeed, create a skyscraper or a symphony, something that no single individual would be able to do. Thank you: a skyscraper or a symphony is a nice example of an 'emergent' result of cooperation of individuals!

(Bill Powers (930615.0900 MDT))

> The cells in a body (for example,
>liver cells) receive reference signals from superordinate systems
>like the pituitary that specify for those cells the level of
>their controlled variables that they are to maintain (blood
>glucose concentration).

Why do you say that they receive REFERENCE levels? In my opinion, cells receive perceptions. Their reference levels are 'built in'. The analogy with people is, in this respect, complete. I'm very curious about your reasoning on this point.

> Liver cells are responsible for
>maintaining the specified level of input, but they have no say
>about what level will be specified.

Who tells the liver cells that they are responsible for other cells or for some grand collective? I think that you are anthropomorphizing. Let me assure you, liver cells LIKE to do what they do! They are not slaves that are responsible but have no say. Neither would they want to be, say, brain cells :-)

Greetings, Hans Blom

Date: Wed Jun 16, 1993 1:37 am PST
Subject: Re: Cells and People

[From Oded Maler (930616)]

Sometimes I forget that the meaning of Control is very precise among the hard-core participants, and this might cause misunderstanding. So I should separate my "claim" into two parts:

- 1) There "are" *systems* at a higher levels of organization for which humans, their artifacts, etc., play the role of cells or molecules inside a human.
- 2) Those systems are control systems.

I was criticizing your refusal to accept (1). I was not claiming explicitly that these systems are built exactly as control systems in the sense of an individual person.

What is clear to me is that there are collective perceptions for such entities, there are some collective reference, and there is some mechanism that connects them thru actions. It is very hard for us to imagine what are the perceptual variables for such systems because they may extend in time (thousands of years) and space, and may be realized by many generations of their mortal components. The whole thing operates on a different scale and the relation between a collective signal and its individual realization is maybe like the relation between a perceptual signal within an individual to the identity of ions that flow thru the relevant synapses.

I don't know whether these systems achieve their "goals" by by influencing their environment or by constantly creating new perceptual variables and updating old ones. Maybe the time-scale for "normal" functioning is yeares, and for reorganization it is decades. All this is very speculative, I know. I don't intend to give any "model", and if you insist I stick only to (1) and withdraw (2) altogether.

--Oded

p.s.

I think there is some confusion between two different notions of levels. The simple notion is the relation between higher and lower signals. Although a signal for, say, a sensation, and a configuration may have some different features concerning their update rate, etc., they are creatures of the same sort. On the other hand a biochemical signal that "tells" a cell to synthesize this or that protein at a given moment in the life-cycle is (a mon avis) qualitatively different from the perceptual signal that this cell might be implementing in case he or she happens to be a neuron.

Date: Wed Jun 16, 1993 6:30 am PST

Subject: Jack PCTers

[from Joel Joel 930616]

Ah but controlled variables must be satisfied, right?

>...Jack Mormons. I believe the implication was just as you can't
>get a jackass to come around to your way of doing things...

My understanding of this term has been different, but perhaps more intriguing from a PCT point of view. Jack Mormons referred to those who wished to receive the perceived "benefits" of church membership, but without complying with the prerequisites (i.e. baptism, tithing, etc.). Sometimes they were also called "dry-land Mormons" for their aversion to water.

But this seems to have many parallels in social organizations, doesn't it? How many times have we, or others we see, want to claim the advantages or benefits which accrue from belonging to a particular group, but without fulfilling particular requirements? I'm certainly guilty of this.

Regards, Joel Judd

Date: Wed Jun 16, 1993 9:01 am PST
Subject: Re: Cells and people

[From Bill Powers (930616.0900 MDT)]

Still just hitting the high spots while writing the Paper.

Hans Blom (930616) --

>Why do you say that they [cells] receive REFERENCE levels? In
>my opinion, cells receive perceptions. Their reference levels
>are 'built in'. The analogy with people is, in this respect,
>complete. I'm very curious about your reasoning on this point.

I think it will be productive to consider that some inputs to cells are sensory inputs, and some are reference inputs. In chemical systems the situation is quite different from what it is in neural systems. Neural signals can be isolated so they go only where they are supposed to go and are insulated from each other. In chemical systems, all the signals are dumped into the same soup pot and the pathways have to be sorted out through their chemical origins and selective binding processes at the destinations. In the nervous system it's fairly easy to sort out the reference signals from the perceptual signals because they come from geometrically different places. In biochemical control systems we have to try to identify the roles of chemical signals in some way other than circuit-tracing.

There is more and more work on biochemical feedback systems. It's pretty clear that negative feedback is present at all organizational levels from organs to detailed processes inside cells, even involving DNA and pritein synthesis. A couple of years ago I came across a book (_Systems analysis of enzyme systems_, I believe) by a couple of Japanese, one a biochemist and the other a control engineer, in which many biochemical feedback systems were modeled. Unfortunately, the authors were most interested in how such systems reach a final operating condition starting with wildly out-of-equilibrium conditions, with the model behavior being shown only until the stable state had been reached. But enough of the stable condition was shown in some cases to reveal a perfectly good control system, reference signal and all. The authors actually described one of the signals as determining the level of the controlled concentration to which the system converged, but didn't follow up the implications.

The comparator of the best control system was an allosteric enzyme. This enzyme was "switched" (actually, driven with high amplification) into the active state by one chemical substance that acted to represent the controlled variable (the gross output of the catalyzed chemical reaction), and was driven into the inactive state by another substance that depended on other concentrations outside the closed loop. The net activity of the enzyme thus depended very sensitively on the difference of concentrations between the independent chemical substance and the other substance whose concentration varied with the controlled concentration (i.e., a sensory signal). The enzyme catalyzed the reaction that

converted a pool of substrate molecules into the output concentration that was sensed and controlled: that was the output function.

By varying the concentration of the independent substance, which is the reference signal, one could cause the closed-loop system to produce the same variations in the controlled concentration, quite independently of drains on the chemical product and variations in the substrate. The authors didn't think of trying that with their model, but it was obvious that it would work, from the way their model snapped into a stable state once the initial gyrations came close enough. If they had then varied the reference concentration, they would have seen that the controlled concentration tracked it very nicely, with no wild variations in the other concentrations.

It's sort of funny that they missed this; the reason is the same one for which Ashby missed the real significance of control in behavior. If you think of a control process as one by which a system gradually finds its way toward a state of zero error, then you get interested in the trajectories. But if you assume that the systems have all been in operation for some time and have reached their stable states, from then on the only "trajectories" that occur result from adjustments in reference signals, with the controlled variables always closely tracking them. This is how I think of the hierarchy: all errors are maintained at essentially zero all of the time, with behavior always being under active control on the time scale appropriate to the level of control.

The same might well apply to biochemical control systems. If these biochemical control systems work as they seem to work, the rules connecting various concentrations become very different from those of the basic chemical reactions involved. The concentration of one high-energy substance can be made to depend on the concentrations of low-energy chemical signals in a way that ignores the obvious chemical rules and imposes new ones. You can begin to see an organization in the biochemical systems that is simply invisible if you concentrate too closely on the individual chemical reactions.

I put out some feelers to see if we could recruit some biochemists familiar with this kind of modeling, so we could collaborate on investigating these control systems further from the PCT point of view. But there haven't been any biochemists who have gone far enough with PCT to see the relevance of the ideas and to be willing to look at their work from an outsider's unorthodox point of view. Same story as in psychology, basically.

Actually, there are lots of control systems at higher levels of organization of the organ systems. The pituitary seems to contain a whole bunch of chemical comparators, the outputs being error signals going to various organs and the outputs of those organs being sensed (as negative feedback) by the pituitary. Reference signals enter the pituitary through the neurohypophysis, the signals originating in the brainstem. Every organ that produces a chemical output is shut down by excesses of its own output (negative feedback), and every organ receives both neural and chemical reference signals from higher systems -- i.e., signals that tend to increase the output of the organ. The control-system organization is obvious.

I think there's a lot more to the organization of the body than just a lot of cells living their independent lives in the midst of a lot of other cells. That happens, too, of course, but it's far from the whole story.

>Who tells the liver cells that they are responsible for other
>cells or for some grand collective? I think that you are anthropomorphizing.

Then you don't understand what I'm saying.

Oded Maler (930616b) --

>What is clear to me is that there are collective perceptions
>for such entities, there are some collective reference, and
>there is some mechanism that connects them thru actions.

If I'm anthropomorphizing, then you are mysticizing. Where do these mysterious entities live? Are they floating about invisibly in the atmosphere? By what mechanism do they have any influence on an individual?

I think that these apparent collective properties simply emerge from the interaction of individual control systems, as measures like entropy and pressure emerge from the interaction of individual molecules. I think they live in a conceptual space, which exists inside your head. You're creating an allegory or a metaphor, not a literal model of how

things work. You are reading something into the world that actually comes from your own imagination. I say that there is absolutely no evidence for collective perceptions, collective reference signals, collective error signals, or collective mechanisms for action. I go along with Tom Bourbon and Rick Marken: what's your evidence?

Best, Bill P.

Date: Wed Jun 16, 1993 9:04 am PST
Subject: Re: social-cells-people

From Tom Bourbon (930616.1020)

Replying to Hans Bloom [Hans Blom 930616] and Oded Maler [From Oded Maler (930616)].

Both of you are pressing the idea that cells in bodies are like people in social systems. That may be true in all cases, in some cases, or in no cases. It would be easier for some of us (me, in particular) to imagine what you are saying were you to present your claims in the form of PCT models, even if they are only preliminary drawings, as I suggested yesterday. Drawings like that are a first step to nearly all of my attempts at modeling; they are my first try at "checking out" my ideas, intuitions, and inspirations. The majority of my brilliant ideas about modeling have died at that stage -- ideas that somehow "felt" right simply could not be turned into workable models. On the other hand, things I originally thought I could not model, like some of my interactive tasks, turned out to be easy, once I had thought my way through drawing the relationships.

I know that both of you are familiar with this first step in modeling. Could the two of you go at least that far toward clarifying your thoughts? At least one of us out here would have a better idea what you mean.

As an example, it may be true that there are Social Reference Signals (Great Big P*s), developed and handed down over thousands of years, or even over the brief life span of a fashion fad. If so, where do I draw P* in a diagram that shows, in a simple case, two people who are affected by P*? And where do I draw the arrows from P* to the simple loop representing each of the people? If I can't imagine that preliminary step in modeling, there is no chance at all that I can go further and produce a formal model and simulations.

As another example, I can easily imagine at least two ways aggregates of cells in a body might act in concert. They also happen to be two ways aggregates of people might act in concert. It is possible to produce diagrams of very simple examples of the two cases and those diagrams can be turned into extremely simple models that can be run in simulations.

CASE A. Each entity (cell, person, model, device, etc) has its own ref. sig. (p*) and controls its perception (p) specified in p*. In the process, each entity also affects at least one variable controlled by at least one of the other entities in its neighborhood. (This is like the relationships in CROWD-GATHER and in my articles on interference among control systems.) In this case, through its actions, each entity controls its perceptions and interferes with variables affected by its neighbors, which also control their own (disturbed) perceptions. It is easy enough to diagram two or three entities in such an interaction.

How might sources other than the entities or the variables they control affect this relationship? There are at least two possibilities. First, a "locally universal" disturbance might affect the variables controlled by all of the entities (smoke suddenly fills the room, the experimenter jostles the petri dish, toxins flood the region, the earth shakes beneath our feet, the programmer set the program so that at t120 the positions of all stationary people in CROWD-GATHER are randomly reassigned ...). Easy enough to diagram and (in simple cases) to model. Second, if the proper connections exist, a "remote" source might "reset" the reference signals inside each of the entities. Easy enough to diagram and model, assuming the proper connections. (Diagramming begins to force the issue of the plausability of connections we assume.)

CASE B. Each entity (cell, person, model, device, etc.) at level 0 has its reference signal set by error signals from entities at level 1. The p* at level 1 specifies a perception that does not exist at level 0. Similarly, the error signal from level 1 is the p* for a perception that does not exist at level 1. Diagrams? Models?

Unlike CASE A, when I try to draw the diagrams for this case, I have some problems. If the entities at level 0 are, say, muscle cells, everything is fine; if they are people, I encounter big trouble.

Hans and Oded, could you give us at least a few simple diagrams of the relationships you are asking us to consider? If they look manageable, someone might try to produce models and simulations. If the ideas passed that step, then it would certainly be true that people in social groups and cells in bodies can be modeled by the same PCT systems. If we don't at least try to reach that step, we will merely exchange opinions on the subject. That might be enjoyable for all of us, but it will do nothing to resolve the important question.

Until later, Tom Bourbon

Date: Wed Jun 16, 1993 9:23 am PST
Subject: Re: Cells and people

[From Oded Maler (930616.II) Bill Powers (930616.0900 MDT)

- * If I'm anthropomorphizing, then you are mysticizing. Where do
- * these mysterious entities live? Are they floating about invisibly
- * in the atmosphere? By what mechanism do they have any influence
- * on an individual?

I mean entities like the US, the scientific community, the belief in a single god/theory, the white race, communism, etc.

- * I think that these apparent collective properties simply emerge
- * from the interaction of individual control systems, as measures
- * like entropy and pressure emerge from the interaction of
- * individual molecules.

- *I think they live in a conceptual space,
- * which exists inside your head. You're creating an allegory or a
- * metaphor, not a literal model of how things work. You are reading
- * something into the world that actually comes from your own
- * imagination. I say that there is absolutely no evidence for
- * collective perceptions, collective reference signals, collective
- * error signals, or collective mechanisms for action. I go along
- * with Tom Bourbon and Rick Marken: what's your evidence?

In fact, what you suggest is sociological reductionism: "every social-historical phenomena can be explained in terms of individual humans". As if you could analyse the information-processing PCT using the underlying bio-chemical or physical terms. A perceptual signal in your theory is a huge abstraction step -- no molecule in the body has any "evidence" concerning the existence of perceptual signals, not to mention whole persons/hierarchies.

Best regards Oded

Date: Wed Jun 16, 1993 9:24 am PST
Subject: PCT and Plate Tectonics, Social Control

[From Rick Marken (930616.0800)] Gary Cziko (930615.2000 UTC) --

>Do I just have a wild imagination, or can others see some parallel here
>involving Wegener and Powers, continental drift and PCT.--Gary

The parallel is astounding! Nice find. It sustains one's hope that we might see the life sciences might come to their senses within our lifetime; but, probably not.

Hans Blom (930616) --

> My perception is more that of a collection of interacting autonomous
>agents, very much like in the CROWDS or GATHERINGS programs.

> I therefore do not agree with your statement [that cells are used as

>part of a higher level control process]

So you are saying that there is NO higher level control systems that use cells are part of their means of controlling the variables that they control? So the fact that a combination of neural and muscle cells maintains a perceptual variable (a cell firing rate) at a fixed reference level in the face of disturbance is not evidence that there is a control organization above the level of the individual cells involved in this process? So control is a side effect of the interaction of autonomous control systems? I'm with Tom Bourbon on this -- could you please make a diagram of what you are talking about. That might help clear things up a bit.

>I agree with you when you see no evidence of higher level
>control systems. So why postulate them?

I said there is no evidence of control systems that are "higher" than the level of organismic control systems (like people). There is TONS of evidence for control systems that are "higher" than the level of cells. The nervous system is exhibit A. Here is a collection of cells (neurons) that is organized to control the firing rates of afferent neurons. The neurons themselves are control systems -- controlling variables that matter to themselves as cells (like nutrient intake rates, electrolyte balances, etc), but that do NOT matter to the operation of the control organization of which they are a component (such as the rate at which spikes are generated as a result of dendritic stimulation).

>Thank you: a skyscraper or a symphony is a nice example of an 'emergent'
>result of cooperation of individuals!

Not really. Those are controlled results of cooperative efforts (like the three line pattern in Tom's cooperation experiment). If the efforts of the collective were producing a cacophony rather than a symphony, the members of the collective would perceive the error and do what they could to correct it -- a process which would be implicitly taking into account the corrective actions being made by everyone else in the collective. Sometimes it helps to have coordinators (like conductors) who try to evoke in all members of the collective a common vision of the intended result. Although each individual is trying to produce some component result on his own, the overall result (the symphony) is unquestionably a controlled variable (check out a rehearsal and the effort that goes into correcting disturbances to the final symphonic result). Tom Bourbon's posts describing his research on cooperative control describe the model of this phenomenon quite clearly; it may be stick wiggling -- but the principles (once you get them) are VERY deep.

The emergent behaviors that are seen in the CROWD program are interesting precisely because they are NOT controlled (intended) results of the interaction of autonomous agents. The neat "ring" for example, that is formed around a "speaker" is a side effect -- not something the individuals are trying to do. You could test this by applying disturbances to the ring. For example, you could add a couple people who "break up" the ring pattern; the other people do nothing to compensate for this disturbance and maintain the ring; the ring configuration emerges "accidentally" as a side effect of the autonomous control exerted by the individuals in the crowd; the ring is, demonstrably, NOT controlled. The symphony, however, demonstrably IS controlled; try to disturb the result by inserting people who play the wrong notes or bang on the walls. Everybody will take steps to correct for these disturbances.

Oded Maler (930616)--

>Sometimes I forget that the meaning of Control is very precise
>among the hard-core participants

Is the precise meaning a problem? Would you, like humpty dumpty, want "control" to mean just what you want it to mean each time you use it, no more and no less?

>1) There "are" *systems* at a higher levels of organization for
>which humans, their artifacts, etc., play the role of cells or
>molecules inside a human.

>I was criticizing your refusal to accept (1).

I know.

>I was not claiming explicitly that these systems are built exactly as
>control systems in the sense of an individual person.

Nor was I, though they must contain the basic components of a control system if they control. All I was saying was that there is no evidence of the kind of controlled variables that you propose; no controlled variables, no control system (of any sort).

>What is clear to me is that there are collective perceptions for such
>entities, there are some collective reference, and there is some
>mechanism that connects them thru actions.

What you are describing is something like the symphony or the skyscraper. Actually, I bet you are talking about controlled collective perceptions like religion or culture. Controlling a religion is like controlling a symphony; people have practiced the same rituals and said the same prayers that make up the religion for centuries. Unlike the symphony, they have often done this "without a score". The control of these religious or cultural activities does require collective action; everyone who is controlling for the result must play their part -- just as each musician in the orchestra must play his or her part (as s/he understands it). But it is the individuals who control for the result -- nothing outside of the individuals is controlling for the occurrence of the symphony or the religion. This is evidenced by the fact that when the individuals change their mind about what they want to control (or when they die out) then the controlled result disappears. Thus, there are no more Zoroastrians (as far as I know) just because no groups control for that anymore; there was no control system around that kept Zoroastrianism "played". When people stop wanting to hear a symphony (live), it won't happen anymore; and no "god of the symphony" (cousin of "phantom of the opera") will come down and whip some available orchestra into a rendition of Beethoven's 9th (correcting for the disturbance created by its not having been played in some time). Similarly, when there are no more people around who want to play "mormon" or "jewish" the results that we recognize as the practice of these religions will not happen any more; at least, that is my prediction. I think many religions actually believe that their "god" is controlling for the practice of their religion and will maintain his (the god's) perception of it against disturbance. I am sure that I cannot "disabuse" those who believe this of that belief; the "existence" of the religious practice now (and possibly over a long period of time) is ipso facto evidence that the practice is under control by a diety. I suppose that's why people like to say "my religion is older than yours". It's their version of the test for the controlled religion. It just looks to me like the controlling is being done by the people themselves. And long term existence alone is not a guarantee of control; I think there is evidence that the ancient egyptian religion lasted a good 5000+ yrs but nobody's controlling for the worship of Isis anymore.

The idea that people are part of the means used by larger control systems to achieve their ends is not only demonstrably false (unless you can show, as I said, a controlled variable that is a result of human actions but is of absolutely NO INTEREST to -- not controlled by-- the humans themselves) but it is also dangerous (as Bill Powers pointed out). It lets people think that they are not really responsible for the controlled results that they produce; it lets people think that the results of their actions are really part of the means used by some "uber" control system to achieve its ends. This is what is happening all over the world today -- people thinking that their violent, despicable treatment of other human beings is really just part of the means used by a higher level control system (called "nationalism", "god's will", "ethnic purity", "manifest

destiny") to achieve its ends. What a crock.

Best Rick

Date: Wed Jun 16, 1993 10:08 am PST
Subject: Look Ma, I'm a Sociological Reductionist

[From Rick Marken (930616.1030)] Bill Powers (930616.0900 MDT)--

> I go along with Tom Bourbon and Rick Marken: what's your evidence?

Oded Maler (930616.II) --

>In fact, what you suggest is sociological reductionism: "every
>social-historical phenomena can be explained in terms of individual humans".

That's "evidence" Oded; not "non-sequiter".

But you are right about one thing; given your definition of it, I am definitiely a "sociological reductionist" (at least until you provide EVIDENCE of a "social-historical" variable that is under control, but not by the people who make up the society - history of which it is a function).

Best Rick

Date: Wed Jun 16, 1993 12:08 pm PST
Subject: Re: Cells and people: emergent phenomena vs control

[From Bill Powers (930616.1230 MDT)] Oded Maler (930616 II) --

>I mean entities like the US, the scientific community, the
>belief in a single god/theory, the white race, communism, etc.

So do I. These are ideas in people's heads. They have no objective existence. To make them exist, you have to get people to AGREE that they exist -- i.e., individually believe it. As soon as everyone stops agreeing that they exist, they disappear. That makes them quite different things from what we mean by "the planet Mars" and "peanut butter."

>In fact, what you suggest is sociological reductionism: "every
>social-historical phenomena can be explained in terms of
>individual humans". As if you could analyse the information-
>processing PCT using the underlying bio-chemical or physical
>terms. A perceptual signal in your theory is a huge abstraction
>step -- no molecule in the body has any "evidence" concerning
>the existence of perceptual signals, not to mention whole
>persons/hierarchies.

I didn't say that every social-historical phenomenon can be explained in terms of individual humans. I said that every social-historical phenomenon of CONTROL is so explained. Many phenomena are not control phenomena: I mentioned pressure, which is a phenomenon arising out of the behavior of collections of molecules. The molecules, through their properties and interactions, produce a stable phenomenon called pressure. The pressure doesn't cause the molecules to move as they do; it's the motion of the molecules that gives rise to the phenomenon of pressure. The molecules are not caused to move because something in nature wants pressure to appear (unless you count human intentions). But pressure is a perfectly real phenomenon: anyone who knows how can observe it.

Conflict is a social phenomenon that arises when two independent control systems with incompatible reference levels interact. Following, avoiding, leading, and gathering are also emergent phenomena. These phenomena are completely explained by the way control systems with certain reference levels interact. They are not in themselves control phenomena, because they are not established by any higher system that perceives them and adjusts the reference signals in the participants to make sure the result appears as desired. A leader may want to lead, but success depends on the existence of followers who want to follow. Nobody is creating both the leaders and the followers in order to produce the social phenomenon of following or leading. That social phenomenon results from the fact that some people want to follow and some people want to lead; when they get together, they each get what they want. There's no superordinate control system that sets one person's reference signals to "follow" and another's to "lead" in order to create the phenomenon of following/leading.

In the internal organization of an individual, there ARE superordinate systems that establish reference levels for subordinate systems, precisely in the way assuring that their combined actions result in the perception that the higher system wants to experience. The higher systems are not just emergent phenomena that can be seen in the interactions among lower systems. They are physically distinct systems which, by sensing and acting, impose coordinations on collections of subordinate systems through directly adjusting their reference signals (and sometimes their properties). This relationship simply does not exist between organisms.

>A perceptual signal in your theory is a huge abstraction step
>-- no molecule in the body has any "evidence" concerning the
>existence of perceptual signals, not to mention whole persons/hierarchies.

Not so. A perceptual signal in my theory is in principle, and often in practice, an observable signal in a physical pathway. It is defined as a perceptual signal because of

the way it is generated. The subsystem itself doesn't know about perceptual signals, of course; that knowledge ABOUT the subsystem must always be formulated from an external point of view. But that's what a theory is, as various contributors to CSG-L have been pointing out: an external view of a behaving system (but residing in the brain of each viewer).

I don't think that physical reductionism is any more profitable than sociological-historical reductionism.

When we talk about evidence, we mean to ask for descriptions of perceivable situations like those on which we normally base our conclusions. If there are social CONTROL systems, we should be able to see evidence that they exist, evidence that refutes the claim that apparent social CONTROL phenomena are simply emergent consequences of interactions among individual control systems.

Best, Bill P.

Date: Wed Jun 16, 1993 2:23 pm PST
Subject: Re: Levels' point of view

[Martin Taylor 930616 17:15]
(Bill Powers 930614.0800 to Bruce Nevin 930614.0836)

>>Information theory requires a point of view from outside the
>>control system; it cannot be applied to the control system
>>itself from its own point of view.

>
>Hooray. Someone else once referred to this problem as that of
>"smuggling intelligence into the system." I've been trying to
>make this point, but you said it much better.

And I've been trying, with apparently no success at all, to make the opposite point. One CANNOT make a consistent view of information at a place in a control system by taking a viewpoint outside the control system. It HAS to be applied to the control system FROM ITS OWN POINT OF VIEW.

All the confusion about information seems to come from the idea that information has to do with a transmission channel, both of whose ends can be seen by an outside observer. Information can be meaningfully discussed only from the viewpoint of the one making the observation. Shannon himself made this point. It does not preclude one from determining such things as the capacity of an information channel, but it ensures that one can see such concepts as limiting factors rather than as the central notions of information theory. I sent Bill P a copy of the section of Shannon's book about continuous information a couple of months ago, in which he makes the derivation. It really should not be so confusing as it apparently is.

Yesterday, I had completed a long response to Bill and Rick's misreading of Allan Randall's Saturday posting, when we had a power outage, and I lost it all. Every time we get into this information question, it seems we have to ever deeper into basics, to try to get the fundamental notion across. I'll try again, soon, to regenerate that posting. No time (or inclination) now.

Martin

Date: Wed Jun 16, 1993 3:14 pm PST
Subject: Controlled and Emergent

[From Rick Marken (930616.1500)] Bill Powers (930616.1230 MDT) to Oded

>I didn't say that every social-historical phenomenon can be
> explained in terms of individual humans. I said that every
>social-historical phenomenon of CONTROL is so explained.

What I think is being missed by the "social control" advocates is the distinction between controlled and uncontrolled results of collective action. This is precisely analogous to the difference between controlled and uncontrolled results of individual action (which I'm sure everyone who has been reading CSG-L for more than a year now understands perfectly -- right?). In the discussion of social control I think the word "emergent" has been used

somewhat ambiguously to refer to both controlled and uncontrolled results of collective action; I would prefer to use "emergent" to refer only to uncontrolled results of social actions; the controlled results already have a name -- "controlled variables". In individual behavior, we call emergent results "irrelevant side effects".

I think the difference between controlled and emergent (my meaning) results of collective action is best illustrated by Tom Bourbon's cooperation experiment. The configuration of the three lines on the screen is unquestionably a controlled result of the collective action of two people; indeed, if these people did not control cooperatively, the three line configuration would not have occurred. There are many different emergent results of Tom's experiment. One is the angular position of the two subjects relative to each other; because the subjects were looking at the same screen (and because of the location of the chairs) lines through the center of each person converge at the screen. Add a line connecting the two subjects and you have a triangle pattern; this triangle is an emergent result of the collective actions of the two subjects -- but it is not controlled, a fact that can be tested by disturbing the triangle (by changing the orientation of the chairs, for example); there is no resistance to disturbances of the triangle. There are many other possible emergent results of Tom's experiment; some could be made quite eye-catching (like when you make a single subject write his or her name unknowingly as s/he pulls on the rubber band to compensate for your disturbances to the knot in the rubber band demo).

The "ring" formed in the CROWD demo is an emergent (uncontrolled) result of collective action; the distance between control systems can be considered a controlled result of collective action in CROWD because the resulting distances depend on the controlling done by at least two control systems -- ie. a collective.

If this helps clear things up at all, then the "social control" issue seems to me to boil down to this; are the emergent (uncontrolled) results of collective action actually under control by a social control system that is outside of the collective? If a result is a controlled result of collective action then we already understand how it occurs without the necessity of invoking "social control systems". We don't need to add a social control system to CROWD to make the rings appear; we don't need to add a social control system to Tom's cooperation models to control the configuration of lines. But if we found that the emergent (uncontrolled) results of collective action were actually controlled -- and we could NOT attribute this control to the individuals in the collective -- THEN we would have evidence of a result of collective behavior that is controlled by something OUTSIDE of the collective. If, for example, the triangular orientation of the subject's in Tom's experiment WERE controlled -- and there was no evidence of this control being exerted by the subject's (for example, it was controlled even though the subject's were unable to perceive the result (triangle) and were deprived of the ability to have any effect on the result (by being curarized, perhaps)) -- then we would have evidence that another control system was controlling the triangle.

If I discovered such a controlled emergent variable, by the way, the first thing I would do is head down to my local house of god and apologize profusely.

Best Rick

Date: Wed Jun 16, 1993 3:16 pm PST
Subject: Conflict and the Judge
[Martin Taylor 930616 18:30]

In PCT, conflict is defined as a condition in which two ECSs both use some environmental variable in their perceptual input function (PIF) in such a way that they cannot simultaneously satisfy their reference signals. It is normally assumed that a consequence of conflict is one of two results: (1) One ECS is sufficiently stronger than the other in its access to output power that it can overwhelm the other. The winner satisfies its reference while the other ineffectually maintains as high an output signal as it can; (2) Neither ECS can attain a perceptual signal equal to the reference, and both escalate their outputs to their maximum levels as the integration of error continues.

What is the effect of a Judge? It seems not only that the conflict is (usually) resolved, in that neither party maintains output in opposition to the other (though one may continue to control the relevant perception, the other does not). Has the loser changed the relevant reference signal by virtue of the Judge's decision? Why does the conflict stop? Is the perceived threat of social force that would be applied if the loser continued the conflict in some way responsible for changing the conflicted reference? I don't see that as likely, given that the same result could be obtained if the parties agreed on accepting

the word of a mutual friend who has no backup social threat system. In some way, the mere fact of the judgment seems to alter the reference signal for some perception, provided that the judgment is "accepted," a term that seems to be a simple description of the result, rather than an explanation.

I can see several possibilities here, but none that is obviously a natural fallout from a simple-minded application of HPCT.

Bill? Greg? Oded? Rick? Anyone?

Martin

Date: Wed Jun 16, 1993 3:22 pm PST
Subject: Reconstructing the disturbance

[From Rick Marken (930616.1530)] Martin Taylor (930616 17:15) --

>Yesterday, I had completed a long response to Bill and Rick's misreading
>of Allan Randall's Saturday posting, when we had a power outage, and
>I lost it all. Every time we get into this information question, it
>seems we have to ever deeper into basics, to try to get the fundamental
>notion across. I'll try again, soon, to regenerate that posting. No
>time (or inclination) now.

Hey! I've got an idea!

Why don't you guys just forget about information and information theory for a second and just answer these three questions:

1) Can you, given a sequence of perceptual values, p , that were the input to a control system, reconstruct the sequence of disturbance values, d , that were present at the time? Yes or no?

2) Can you do it if I also give you the sequence of reference values present at the time? Yes or no?

3) Can you do it if I also give you the output function that was in effect at the time? Yes or no?

(The last two are freebees; you would have information that the control system itself doesn't have. I agree with Bill that if you take this information you are cheating -- but I'm just an easy going kinda guy).

Best Rick

Date: Wed Jun 16, 1993 3:54 pm PST
Subject: Re: Social control systems

[Martin Taylor 930616 19:30]

I think this discussion of social control might be helped by an insightful quote from an old posting by Bruce Nevin...

Bruce Nevin (Thu 920709 09:13:52)

>Coming at this from a different direction: the analogy is often made
>between cooperation of cells in an organism and the social cooperation
>of animals, including people. If this analogy were valid, it might work
>like this:

>

> 1. Each cell in an ECS can control itself using intracellular control
> mechanisms such as ion exchange across membranes.

>

> 2. A cell cannot control another cell with these mechanisms.

>

> 3. Yet cooperating cells can together constitute a control mechanism
> of a higher order.

>

> 4. This higher order of control is invisible to and does not itself

> affect the constituent cells.
 >
 > 5. However, chronic error in the higher-order control system is
 > perceived by the cells in the form of environmental factors that
 > are distressing to the cells.
 >
 > 6. In consequence of such distress, the cells may alter their
 > relations with one another, while still controlling for inter-
 > cellular cooperation. This constitutes reorganization at the
 > higher order of control.
 >
 > 7. One may substitute "person" for "cell" in propositions 1-6 (naming
 > appropriate mechanisms in proposition 1).
 >
 >Two questions:
 >
 >Is (2) true of cells? Or can one cell truly control another (in
 >the same intra-cellular terms in which a cell controls itself)?
 >
 >Does the specialization of cells for cooperating functions have a
 >parallel in human differences of temperament, talent, etc., as well as
 >in educative specialization for social function?

I don't know whether Bruce would still own to this, but I like it.

Martin

Date: Wed Jun 16, 1993 3:57 pm PST
 Subject: Re: Reorganization

[Martin Taylor 930616 18:40]

Going back deep into history to reopen a discussion never completed...

After our power failure of yesterday that inhibited (and lost) my information theory posting, I have started going back through postings I saved for later consideration. Here's the earliest.

(Bill Powers 920704.0800) to me RE: Reorganization
 >>I still like the idea that one is inserting levels,
 >
 >I don't, particularly. There are reasons for which I like it, but more for
 >which I don't. How do you open up the connections from a higher to a lower
 >system to insert a complete control system with all its connections to and
 >from both the higher and the lower systems? This idea seems to me to entail
 >enormous difficulties, whereas building from the bottom up eliminates those
 >particular problems completely.
 >...
 >This is, in fact, how I think higher levels of control come about.
 >situations arise in which existing control systems can't correct intrinsic
 >error well enough any more, even though they're all keeping their own
 >errors small most of the time. They're not controlling the right ASPECT of
 >the environment. So a new level of control begins to form, setting the
 >formerly fixed or randomly varying reference signals in a systematic way to
 >control new perceptual variables that are functions of the old ones. This
 >adding of levels continues as long as there is a need and as long as there
 >are available neural components of the right kind still unorganized. You
 >arrive at the top level when you run out of new layers of neurons that
 >permit new types of control.

After nearly a year, I still like the insertion of new ECSs into the hierarchy in arbitrary places. In particular, I like it between a fixed "top level" consisting of the control of intrinsic (largely chemical) variables and a growing hierarchy of controlled perceptual variables. What I realize is that since Bill wrote that posting, I have more or less absorbed into my understanding of PCT that the problem Bill raised does not seem to exist. We are, in fact, intending to incorporate this method of learning into our syntax control system (though not in its initial incarnation.) I think (and hope) that the scheme fulfils what Bill says in the second quoted paragraph above.

If I understand correctly, Bill (I address this to you specifically), level-jumping connections of perceptual signals pose no problems.

Perceptual signals can go from anywhere to anywhere, provided that they are useful (possibly with the caveat that they must go upward in the hierarchy). On the other hand, output signals cannot usefully jump levels, even though they may not be prohibited from doing so. This restriction occurs because a level-jumping output would be affecting the same variables (probably the reference signals of low-level ECSs) that would be affected by intervening ECSs, and control by means of those intervening ECSs would be more effective than affecting the low-level references directly.

I think I have represented Bill P's position as best I remember it. In any case, I will take the foregoing as a reasonable position.

Now let us consider inserting a new ECS between two cleanly separated levels, M and N ($N = M+1$). The new ECS does not have a level designation. I imagine this new ECS as receiving its reference signal from the outputs of some of the level N ECSs, and its sensory signals from some of the level M perceptual signals as well as (possibly) level M-1 perceptual signals.

The new ECS has to be considered as being at level Q, which is M+1 and N-1. In other words, it disturbs the labelling of levels M and N, which are no longer consecutive. Now, $N = M+2$, at least in those parts of the hierarchy to which the new ECS is connected. At this point, we have level jumping signals in both directions between levels M and N, and for some of those down-going signals, reference levels at M are set both from N and from Q.

What happens? Reorganization, in the sense of both smooth modification of perceptual (and reference) weights, and in the sense of reconnection within the hierarchy. At least that is what will happen if the new ECS sets up conflict somewhere, such that a sustained error is introduced into part of the hierarchy. If it doesn't, then it is controlling something orthogonal to whatever was being controlled before, and the whole hierarchy has become more capable.

Before considering the probable kind of reorganization, reconsider the reason for having levels in the first place. It is the non-linearity of the overall control system, which includes the part that goes through the world. If the world's responses to outputs were linear, a one-level system of orthogonal ECSs would be perfectly adequate. But it is not, and so we need higher-level ECSs to shift reference levels according to the current state of the environment. And I speculate that we probably need more of them at any level than we have effectors at the interface to the environment.

Now return to our new ECS at level Q, which is now assumed to induce some conflict in the hierarchy (the contrary having been considered above). The existence of this new ECS could be valuable to the hierarchy only if it allowed better control of the level N perceptual signals. It can do this if it represents a perceptual signal that coordinates the actions of level M ECSs in a way that corresponds to some nonlinearity of the world--a new CEV that is usefully controlled.

One of the accepted ways an ECS "learns" is by smooth modification of its Perceptual Input Function in response to its own failure to control. Bill has said that he has successfully simulated this as one aspect of reorganization. If, with its (by assumption) randomly selected input and output connections, it does not attain control, it will be likely to change its PIF until it finds something that it can control. Given that its reference signals are set from level N, that something will be useful to level N ECSs. (Presumably it will also be changing its output connection signs and/or weights at the same time, to achieve control, affecting how it interacts with level M).

If the new ECS cannot achieve control, the question remains of whether ECSs in the hierarchy die. Neurons, or at least neuronal connections, seem to, especially early in life, as I understand it. I see no reason why useless ECSs (or at least useless connections) should not also die. If the new ECS at level Q attains good control, Bill has argued that the level-jumping output connections from N down to M would become less useful, and might die. If many new ECSs came into being at level Q, then most of the level jumping connections would be reduced in usefulness and would probably die. Eventually, level Q might become a proper level in the hierarchy, rather than a single anomalous ECS dropped in an awkward place.

Now more practically, how would a new ECS at level Q be constructed? What we propose in our model is to use a kind of genetic algorithm, in which the connections of the new level Q ECS to level N are copied from subsets of the connections of several level M ECSs, and its connections to level M are made to those same ECSs from which it borrowed the level N connections. We may change this proposal, but that's the present idea.

(Bill)

>There's some developmental evidence that, if interpreted in a certain way,
>indicates that you may be right. For now, however, I don't see how you CAN
>be right.

Does this help? Martin

Date: Wed Jun 16, 1993 6:35 pm PST
Subject: social control

[Hans Blom, 930617]

Let me try to make clear what I mean when I talk about social control emerging from individual control systems. I will use a constructive (modelling) approach that anyone can easily test.

Take a number of control systems. In this example I will use 'robots' that move about. Give each of them some control laws: don't get too close to others, don't get too far away from others, but otherwise you can move freely (use random numbers to create a random walk for each individual, for instance). The details of the laws don't matter much: under the most general conditions a flock, herd or some such will be formed. Do NOT introduce non-local control laws; each individual is supposed to use only its own sensors and actuators to establish its position.

Each individual is a separate and autonomous control system, controlling only for its own goals. But the flock will prove to be a 'higher level' control system, although no higher level rules are programmed in in any way. Try to use The Test, in whatever way you want, to disturb 'flockness'. Pick up an individual and drop it again some way off. If the other individuals can still perceive the lost one, the flock will reform. Limit movements by introducing obstacles or corridors: the flock may temporarily separate while moving around an obstacle, but it will reform again.

You can call the flock a social control system. No single individual has a notion of what a flock is: it just controls its own movements. Yet a (moving but stable) flock emerges.

Many more examples of emergent behavior can be found in: From animals to animats, Proceedings of the First International Conference on Simulation of Adaptive Behavior. Eds: Jean-Arcady Meyer and Stewart W. Wilson. MIT Press, Cambridge, Mass., 1991.

Greetings, Hans Blom

Date: Wed Jun 16, 1993 8:52 pm PST
Subject: Another Cycle

[From Jeff Hunter (930617)]

>[re Bill Powers (930616.0900 MDT)]

>

> I say that there is absolutely no evidence for
>collective perceptions, collective reference signals, collective
>error signals, or collective mechanisms for action. I go along
>with Tom Bourbon and Rick Marken: what's your evidence?

Well here's an example of a collective reference signal for you. If a number of women are confined together their menstrual cycles will synchronize. This does not require the conscious knowledge of the women, it seems to happen just by pheromones.

This is pretty remarkable since individual women have different cycle lengths (I'm avoiding the word "periods" :-), and some even have highly irregular cycles. Thus the homonal reference sets both the phase and length of the cycle of each woman.

This seems to be a pretty clear example of a "society" setting the reference signals for individual humans, and it hopefully will be a bit less of a "'tis so, 'taint so" issue.

(Note: I have not claimed that this is an instance of the group containing a control system, rather it merely sets a high-level reference in its component individuals.)

Have fun ... Jeff

Date: Thu Jun 17, 1993 3:27 am PST
Subject: Stuck to the roof of my subjective mouth

From Greg Williams (930617) Bill Powers (930616.1230 MDT)

>>[Oded Maler:]

>>I mean entities like the US, the scientific community, the
>>belief in a single god/theory, the white race, communism, etc.

>So do I. These are ideas in people's heads. They have no
>objective existence. To make them exist, you have to get people
>to AGREE that they exist -- i.e., individually believe it. As
>soon as everyone stops agreeing that they exist, they disappear.
>That makes them quite different things from what we mean by "the
>planet Mars" and "peanut butter."

And, presumably, from "snow" -- but some eskimos believe in several kinds of snow. Wait a minute. I think you are being inconsistent here. Have you given up on the notion that "it's ALL perception"? Don't you think "peanut butter" would disappear, too, if everyone decided that it was better to speak of, say, five "fineness grades" of "chopped peanuts"? If you don't, I think you are apparently smuggling in the worst sort of naive realism: whatever YOU think is independent of belief IS independent of belief.

As ever, Greg

Date: Thu Jun 17, 1993 5:59 am PST
From: tbourbon
TO: * Dag Forssell / MCI ID: 474-2580
Subject: RE: Dag's myths type 2

Dag (direct)

>Tom, please consider responding to the post below. I need help and
>will appreciate any comment you have.

>
>[From Dag Forssell (930615 18.10)]

>
>REQUEST FOR HELP WITH MYTHS TYPE 2:

I didn't get this one. Our server went down for about 28 hours, including the time when you sent that post. None of received mail during that time, and what was here but unread was lost. Can you send a copy direct?

Regards, Tom

Date: Thu Jun 17, 1993 6:25 am PST
Subject: Re: Reconstructing the disturbance

[Martin Taylor 930617 09:50]
(Rick Marken 930616 17:15)

>Why don't you guys just forget about information and information theory
>for a second and just answer these three questions:

Because it is essential to the proper understanding of the control hierarchy (in my view).

>1) Can you, given a sequence of perceptual values, p, that were the input
>to a control system, reconstruct the sequence of disturbance values, d,
>that were present at the time? Yes or no?

No

>2) Can you do it if I also give you the sequence of reference values present
>at the time? Yes or no?

No

>3) Can you do it if I also give you the output function that was in effect at
>the time? Yes or no?

Yes, if you assert that the output function includes the environmental path between the ECS and the CEV to which d is applied, and that the function is defined precisely and is monotonic.

The output function normally cannot be defined precisely, though, both because there is uncertainty about the effect of a particular output signal on the CEV, and because the (uncertain) output function necessarily is time variant in a cross-connected hierarchy. The problem is (at minimum) that the effect of an ECS on its CEV is achieved through the actions of lower ECSs, and their reference signals are derived from multiple and varying sources (the outputs of other higher ECSs).

This variability makes it even more important to note how the information about the disturbance (NOT the disturbing variable) is passed through the perceptual signal. Given that the relation between the output signal leaving the ECS and the effect of that output on the CEV is dynamically variable, the perceptual signal also carries information back to the output about the current relation between the output signal and the CEV. This latter is most important for learning. The essential issues are those of control bandwidth (information rate) and gain, as related to the information rate of the output function itself.

>(The last two are freebees; you would have information that the control
>system itself doesn't have. I agree with Bill that if you take this
>information you are cheating -- but I'm just an easy going kinda guy).

The fact you say this is at the core of why we say you just don't understand what is meant by information. Re-read Allan's posting of last Saturday (Allan Randall 930611.1700). To show that the information about the disturbance was available in the perceptual signal, given the fact of control, all you need is to show that using a given language it is shorter to describe the output function than it is to describe the disturbance time series (or continuous waveform).

Please try to remember: the ability to reconstruct is conclusive proof that the information about the entity is available. We demonstrated that the disturbance could be reconstructed from the perceptual signal, given a known output function. Inability to reconstruct is not proof that there is no information. All an inability to reconstruct can demonstrate is that the available information is not complete.

Martin

Date: Thu Jun 17, 1993 7:48 am PST
From: tbourbon
TO: * Dag Forssell / MCI ID: 474-2580
Subject: RE: Missing post

Dag (direct)

>[From Dag Forssell (930617 direct)]
>Tom, here is a repeat of my post. Thank you for asking. Any help will be
>appreciated. Thanks! Post on CSGnet, of course.

Got it! I will read it after I finish some detail work on a grant proposal this morning and will post a reply soon.

Best wishes, Tom Bourbon

Date: Thu Jun 17, 1993 7:56 am PST
 Subject: Re: Reconstructing the disturbance

From Tom Bourbon (930617.0943)

>[Martin Taylor 930617 09:50] (Rick Marken 930616 17:15)

>>1) Can you, given a sequence of perceptual values, p, that were the input
 >>to a control system, reconstruct the sequence of disturbance values, d,
 >>that were present at the time? Yes or no?

>No

>>2) Can you do it if I also give you the sequence of reference values present
 >>at the time? Yes or no?

>

>No

>

>>3) Can you do it if I also give you the output function that was in effect at
 >>the time? Yes or no?

>>

>Yes, if you assert that the output function includes the environmental
 >path between the ECS and the CEV to which d is applied, and that the
 >function is defined precisely and is monotonic.

But, Martin, doesn't your request for more -- I hesitate to use the word -- information about (in the sense of facts about) the output function vitiate the claim that a modeler can reconstruct the disturbance from the perceptual signal alone? That would be the challenges confronting the organism modeled by the ECS. Or was the claim, not that information in p about d is essential to the organism if it is to control, but that a modeler can reconstruct d using p and other facts, none of which are available to the organism modeled by the ECS?

I don't know if Rick said it, but I imagine the output quantity acted 1:1 on the CEV, but even if it didn't, why would the modeler need to know the output function, if the claim is that the disturbance can be reconstructed from the perceptual signal? I was just getting connected to the net when this particular test was proposed and I missed the details of the offer and of the acceptance. Is the goal to reconstruct the disturbance from the perceptual signal *and* anything else that proves necessary? For the benefit of uninformed observers, can the parties to this demonstration please clarify their understandings of the rules?

>The output function normally cannot be defined precisely, though, both
 >because there is uncertainty about the effect of a particular output
 >signal on the CEV, and because the (uncertain) output function necessarily
 >is time variant in a cross-connected hierarchy. The problem is (at minimum)
 >that the effect of an ECS on its CEV is achieved through the actions of
 >lower ECSs, and their reference signals are derived from multiple and
 >varying sources (the outputs of other higher ECSs).

I don't know which data Rick gave you, but I'll bet they contain only the exactly specified output function he used in the model for that run. If so, the output function and its effect on the CEV are known exactly and the other concerns you express here are not necessary. Again, I am speaking from ignorance of the actual terms of the demonstration, but I suspect Rick gave you data from the simplest possible case, and that he did so in order to eliminate any need for you to be concerned with uncertainties and functions that are time variant in a cross-connected hierarchy. A single-level, perfectly-defined loop, like the one I imagine Rick gave you, should rule out some of those concerns, shouldn't it? (Maybe I should stop imagining what Rick said and intended and wait for him to briefly re-state his offer.)

>This variability makes it even more important to note how the information
 >about the disturbance (NOT the disturbing variable) is passed through
 >the perceptual signal. Given that the relation between the output
 >signal leaving the ECS and the effect of that output on the CEV is
 >dynamically variable, the perceptual signal also carries information back
 >to the output about the current relation between the output signal and

>the CEV. This latter is most important for learning. The essential
>issues are those of control bandwidth (information rate) and gain,
>as related to the information rate of the output function itself.

Again, if my mind-reading of Rick's offer is correct, these concerns are not relevant to the demonstration, are they?

>>(The last two are freebees; you would have information that the control
>>system itself doesn't have. I agree with Bill that if you take this
>>information you are cheating -- but I'm just an easy going kinda guy).
>
>The fact you say this is at the core of why we say you just don't
>understand what is meant by information. Re-read Allan's posting
>of last Saturday (Allan Randall 930611.1700). To show that the information
>about the disturbance was available in the perceptual signal, given
>the fact of control, all you need is to show that using a given language
>it is shorter to describe the output function than it is to describe
>the disturbance time series (or continuous waveform).
>
>Please try to remember: the ability to reconstruct is conclusive proof
>that the information about the entity is available. We demonstrated
>that the disturbance could be reconstructed from the perceptual
>signal, given a known output function. Inability to reconstruct is
>not proof that there is no information. All an inability to reconstruct

Now I am really confused. I thought the claim was an ability to reconstruct the disturbance from the perceptual signal, and that the claim was not yet borne out. The demonstration that the disturbance can be reconstructed from the perceptual signal, given a known output function, must have appeared before I was reliably on the net. Can someone post it again? Was the goal to describe the output function, as you say here? When did the rules change? By "output function" do you mean $o := k(e)$, or are you talking about the time series of values from the output function? Will the parties *please* clarify the rules of the demonstration?

Until later, Tom Bourbon

Date: Thu Jun 17, 1993 8:14 am PST
Subject: Point of view; jumping levels; cells & people

[From Bill Powers (930617.0730 MDT)] Martin Taylor (930616.1716) --

>One CANNOT make a consistent view of information at a place in
>a control system by taking a viewpoint outside the control
>system. It HAS to be applied to the control system FROM ITS OWN POINT OF VIEW.

I know that you've said this, Martin, but the moment you start talking about how a control system perceives, you begin giving that system a very elaborate set of capabilities, or so it seems to me. As you explain how the probabilities are calculated by a perceptual function, I feel that you're describing operations that are pretty unlikely to be going on in the perceptual function itself. Not only do the probability calculations have to be applied TO the proposed control system from its own point of view, they have to be applied BY the proposed control system, without the aid of a human interpreter. The model has to work by itself while you're out having lunch. If your computer program contains any calculations of probabilities or information, that computer program is part of the model, and amounts to a claim that the control system itself is performing those same computations. In my models the computations in the program are intended to represent processes in the perceptual function itself, not computations ABOUT the perceptual function. I don't use any auxiliary computations or meta-computations. Everything the program does (that is part of the model of the behaving system itself rather than its environment or the user interface) represents something the system itself does, or is proposed to do.

We need to talk about this in detail until it's settled. Otherwise we'll be stuck here forever.

Insertion of ECSs:

As a reason for not letting outputs jump levels, you say

>On the other hand, output signals cannot usefully jump levels,
 >even though they may not be prohibited from doing so. This
 >restriction occurs because a level-jumping output would be
 >affecting the same variables (probably the reference signals of
 >low-level ECSs) that would be affected by intervening ECSs, and
 >control by means of those intervening ECSs would be more
 >effective than affecting the low-level references directly.

The reason I have given is a bit more forceful. If a higher level system's output skips a level instead of working through the reference signal of the skipped level, it will be adding a signal to the same place where the output of the skipped level enters. This will simply create a disturbance in the skipped control systems, and the skipped level will, as usual, oppose it. The skipped level will oppose the output from the higher system just as effectively as it would oppose any disturbance arising from outside the system. So the higher system would find it impossible to affect its own perceptions: its output would be cancelled without accomplishing anything.

I'll be interested to see how your ECS-inserting model works. As you're intending to make an actual working model, I won't throw any roadblocks in the way.

 RE: following citations from Nevin:

- > 1. Each cell in an ECS can control itself using intracellular
 >control mechanisms such as ion exchange across membranes.
 >
- > 2. A cell cannot control another cell with these mechanisms.
 >
- > 3. Yet cooperating cells can together constitute a control
 >mechanism of a higher order.
 >
- > 4. This higher order of control is invisible to and does not
 >itself affect the constituent cells.
 >
- > 5. However, chronic error in the higher-order control system
 >is perceived by the cells in the form of environmental factors
 >that are distressing to the cells.

This relates to what I was saying to Hans Blom yesterday. The implication of what Bruce says is that the higher order of control is implemented using the same physical cells that constitute the set of lower order control systems. When I first read this, I happily incorporated it into my own model by assuming it meant that the higher level control systems used the existing lower level ones as an output function, which is how my hierarchical model works. Now I realize that this is not what Bruce was saying (from the context of your citation). Bruce is treating the higher level as if it were simply an emergent property of the existing set of control systems, "invisible" because the very same cells are involved and only their relations to each other express the new level.

This is a concept of hierarchy quite different from mine. In my hierarchical concept, a new level is composed of physically distinct cells, specialized to perform new functions and capable of sending outputs to the old set of cells (through their membranes) that adjust their internal reference signals.

Thus: the liver cells control the sensed concentration of glycogen in their immediate vicinity. As there are many cells, they do interact with one another, and there are probably phenomena of mass action and interaction involved. But liver cells contain no mechanisms for adjusting the desired level of glycogen. That is done by physically different cells that do not control glycogen, but other variables that depend on the level of glycogen. Those cells are far from the liver, in the hypothalamus and perhaps the pituitary. The outputs of the cells in the hypothalamus or pituitary reach liver cells and raise or lower their reference levels for sensed glycogen.

The liver cells are not "distressed" by experiencing different levels of glycogen unless those different levels occur without a change in the reference signal. The liver cells are just as happy to maintain a higher level as a lower level, if the received reference signals so specify. It is not the business of liver cells to judge where the goal for glycogen concentration should be set.

The addition of a higher level of control does not affect existing control organizations at all. It simply makes their reference signals variable.

>If the world's responses to outputs were linear, a one-level
>system of orthogonal ECSs would be perfectly adequate.

Adequate for what? For controlling the rate at which a variable changes back and forth between two values? For adjusting the placement of flowers in a bowl so they have a pleasing look? For raising the garage door before rather than after you back the car out?

You've said before that you think a perceptron using weighted sums of its inputs is adequate for modeling all of perception. I don't think you're considering all the kinds of perceptions there are, and how one kind must be controlled in order to control others because of the way the external world is organized. I know that there's a degrees-of-freedom argument in the background, but there's also a practicality argument: it's just not practical to accomplish all the different higher-level perceptions by means of a single level of transformation, even though mathematically the difference isn't visible. Nonlinearity isn't the only factor. And the mathematics is limited.

Hans Blom (930617) --

I appreciate your clarification; it introduces a nontrivial consideration.

The CROWD program contains ample illustrations of the case you cite.

Suppose we have set up the program so that person A seeks a specific distance from person B, person B from C, and C from D, while D simply seeks the position of a distant target location. The result is that a little procession forms, and when the leader finally arrives at the destination the followers stop in a line trailing behind the leader. This single-file procession is a vivid example of an emergent social phenomenon that is not part of the intentions of any participant, just as much so as the arcs and rings.

It is easy to set up this situation so that about midway during the march across the screen, this processsion encounters a loose group of people traveling on a path crossing at an angle through the procession. A series of disturbances occurs, with the first line of marchers dodging about and deviating far from the original simple lineal order. After the encounter, the single- file march reforms and continues to the goal.

So this does indeed look like a control system that resists disturbances of a social variable, restoring a social relationship to a specific form after a disturbance. But does it pass the Test for the Controlled Variable?

It's easy to forget that resistance to disturbance is only one part of the Test. There are three other parts. One essential part is to make sure that the apparent resistance of the controlled variable is not due entirely to a natural tendency of the controlled variable to resist disturbance. Another part is to interrupt the ability of the supposed control system to perceive the controlled variable, and show that the resistance to disturbance becomes unsystematic -- that control is lost. And the final part is to remove the control system entirely (when possible) and show what the effect of the disturbance on the controlled variable is when the system is not opposing it (there might be some other system acting that we haven't noticed). The effect should be much larger than when the proposed control system is present.

The "control" of the lineal progression passes the first part of the test, but fails all the others. We can't even find a control system that is sensing the state of the proposed controlled variable, the linearity of the formation, so clearly we can't interrupt its perception of that variable. We can't find the variable output by which this control system affects the shape of the formation, so we can't do that part of the test, either. In fact, all we can find is the set of control systems which are each controlling a variable that has nothing to do with the lineal formation. We must conclude that the lineal formation is simply a property of the collection of control systems with their reference signals set as they are, and is not the controlled variable of any superordinate control system. The social behavior we observe is completely explainable through knowing the properties of the individual control systems. No postulate of social control is required.

There's a much simpler example that makes the same point. Suppose that we find a marble in the bottom of a bowl. By applying a known force to the marble, we can make it deviate from the resting position. As the deviation occurs, a counterforce develops; the marble comes to

rest partway up the side of the bowl where the restoring force is equal to the applied force. Is this an example of a control system?

Again, strictly on the basis of opposition to disturbance, yes. But again, none of the other parts of the Test will be passed. We can find nothing sensing the position of the marble. We can find no system that varies its output force in opposition to our disturbance of the marble; the cause of the restoring force, gravity, doesn't vary at all. In fact, the restoring force is produced directly by a physical effect of the disturbance on the marble. No control system is required to explain the observations.

We would require a different explanation if we found that the marble sought a position partway up the side of the bowl, and resisted forces tending to move it from that position. We would have to find the force that is holding it up against gravity, and that force might be coming from a control system. Also, it might not: there might be a magnet off to one side, and the marble might be made of magnetic material. One simply has to explore all the possibilities. Control is not the only possible explanation. We accept the control postulate only when all other simpler explanations have been ruled out.

This is why the Test has four parts, not just one. If all parts of the Test are passed, we can be sure that we have found a control system that works like the model used in PCT.

Best to All, Bill P.

Date: Thu Jun 17, 1993 8:25 am PST
Subject: Re: Another Cycle

From Tom Bourbon (930617.1045) Jeff Hunter (930617)

>

>>[re Bill Powers (930616.0900 MDT)]

>>

>> I say that there is absolutely no evidence for
>>collective perceptions, collective reference signals, collective
>>error signals, or collective mechanisms for action. I go along
>>with Tom Bourbon and Rick Marken: what's your evidence?

>

> Well here's an example of a collective reference signal
>for you. If a number of women are confined together their
>menstrual cycles will synchronize. This does not require the
>conscious knowledge of the women, it seems to happen just by
>pheromones.

> This is pretty remarkable since individual women have
>different cycle lengths (I'm avoiding the word "periods" :-),
>and some even have highly irregular cycles. Thus the homonal
>reference sets both the phase and length of the cycle of each
>woman.

You seem to have ruled out the possibility that synchronization of the cycles is an uncontrolled side effect, much like those Rick discussed in his post "Controlled and Emergent" (Rick Marken 930616.1230). If the pheromone in question acts a disturbance to a variable (as yet unidentified, probably because no one ever looked for it) controlled by all women who become synchronized (not all do, do they?), then you would see exactly the phenomenon described. In that case, there would be no need to assume a change in reference signals. The effect could be like many people in a room full of people salivating in synchrony when periodic blasts of an odorant, associated with their favorite food, enter the room. (By the way, I do not imply that menstruation is nothing more than salivation.)

> This seems to be a pretty clear example of a "society"
>setting the reference signals for individual humans, and it
>hopefully will be a bit less of a " 'tis so, 'taint so" issue.

Maybe it is not so clear as it seemed at first. I'm glad you brought up this example. Someone hurled it at me a couple of years ago and that time I did not think of "pheromone as disturbance."

> (Note: I have not claimed that this is an instance
>of the group containing a control system, rather it merely sets
>a high-level reference in its component individuals.)

See above.

Until later, Tom Bourbon

Date: Thu Jun 17, 1993 9:15 am PST
Subject: Topics for joint paper

[From Bill Powers (930617.1030 MDT)] Tom Bourbon says in a direct post

>What would you like to receive from those of us who are
>pursuing some of the chapters and verses that will be presented
>in the article? As you probably surmised after I posted my
>lists of publications containing odd, erroneous or distorted
>discussions of CT, I have unearthed a few more that were
>missing after our moves. There should be still others. Should
>I review each of them in detail? Or might it be more useful if
>I give a very brief synopsis of the authors' stated goals, then
>lift a few juicy quotes and give a brief critique? It is also
>possible to group the quotes by topics, across articles, for
>topics such as "positive feedback," "feedforward," "control
>systems control their behavior or actions," etc.

The way the paper is developing, it looks as though we should organize the "misinterpretations" by topics as Tom suggests, across articles. I think we should stick to correcting direct statements about control theory and the behavior of control systems that are wrong, without attacking any whole theoretical frameworks. In addition to the topics Tom suggests, there are

Feedback is too slow
Control is a sequence of events
An organism is ONE control system
Feedback doesn't occur unless it is given
Control processes must involve consciousness
Control is homeostasis (goals are fixed)

And probably more.

It would be most helpful right now just to get lists of topics from people, with brief references. As I see the structure of the paper better I can ask for specific citations. Of course if anyone wants to send citations now that's fine; I'll collect everything in one file. Dag Forssell has already send me excerpts from this year's posts, but don't try to avoid duplication. I'll sort that out. We obviously won't be able to have long discussions of every error; I'll consult later about a possible final list.

Best Bill P.

Date: Thu Jun 17, 1993 9:28 am PST
Subject: Collectives, Information

[From Rick Marken (930617.0900)] Hans Blom (930617) --

>You can call the flock a social control system. No single
>individual has a notion of what a flock is: it just controls
>its own movements. Yet a (moving but stable) flock emerges.

The flock itself is not controlling anything; the behaviors of the flock are the result of the controlling done by the individual control systems in the flock. The variable that you call "flockness" has to be better defined. If you mean the average distance between all individuals in the flock, then this is NOT an emergent (uncontrolled) variable; it is a controlled variable; a "piece" of this variable is being controlled by each individual in the flock. The Test you propose would show that "flockness" (my definition) IS a result of the controlling done by each individual.

>You can call the flock a social control system. No single
>individual has a notion of what a flock is: it just controls
>its own movements. Yet a (moving but stable) flock emerges.

"Flockness", defined as net inter-individual distance, is a controlled variable -- controlled, in part by each member of the flock. It is not "emergent" in my sense of the word. One variable that IS emergent, by my definition, is the SHAPE of the flock as seen by an observer. If each member of the flock has about the same reference for the distance between itself and others the default shape of the flock will be a sphere (assuming the individuals are not moving). If you put the flock in a box that is smaller than the default sphere I predict that the flock will assume the shape of the box -- with the individuals in the corners being unable to keep the distance between themselves and their neighbors at their reference level. What will NOT happen (and what shows that the shape of the flock is NOT controlled) is an adjustment of the inter-individual references so that the sphere shape is preserved. If the emergent sphere were really a variable controlled by a "social" control system, then that control system would adjust the references of the individuals in the flock so that the controlled variable (shape) was preserved in the face of the box disturbance; the result would be a smaller sphere that fits in the box. Of course, you would still have to show that "shape" was not a consequence of the controlling done by each individual. And you would also have to identify the sensor and output functions of the "social" control system, if you conclude that shape is a variable that is NOT the result of the controlling done by the individuals in the collective.

Martin Taylor (930616 19:30) --

>I think this discussion of social control might be helped by an insightful
>quote from an old posting by Bruce Nevin...

The problem with Bruce's analogy is that he is not clear about what variables are controlled.

A statement like:

> 2. A cell cannot control another cell with these mechanisms.

is not clear. A cell can only control VARIABLES -- and there certainly may be variable aspects of another cell (those not under control by that cell) that can be controlled by "these [intracellular control] mechanisms".

The rest of Bruce's statement describes something like the PCT model of the reorganization of a cellular control system.

Jeff Hunter (930617)--

> Well here's an example of a collective reference signal
>for you. If a number of women are confined together their
>menstrual cycles will synchronize. This does not require the
>conscious knowledge of the women, it seems to happen just by
>pheromones.

>This seems to be a pretty clear example of a "society"
>setting the reference signals for individual humans, and it
>hopefully will be a bit less of a "'tis so, 'taint so" issue.

Assuming that this is a real phenomenon (I'm prepared to believe that it is) it is not necessarily an example of a "society" setting reference signals for individuals. This would imply that there is a control system called "society" that, in this case, is controlling for the relationship between the periods -- keeping the perception of difference between cycles at zero. I think that it is far more likely that each individual is controlling for the relationship between her period and that of the others (unconsciously, of course). We could test this by keeping the women together but preventing each woman from perceiving the periods of the others. If the synchronicity of the periods is controlled by "society" then it should still occur with this disturbance because the "society" control system can still perceive all the women and can synchronize their periods. If it doesn't occur (as I predict it won't) then the synchronized periods result from the same processes as Hans' flocking -- individual control systems controlling for their own perceptions. The synchronicity is not an emergent phenomenon controlled relative to some "social" reference level; it is a controlled variable that is controlled by the individuals in the collective (just like Tom Bourbon's line pattern in the cooperation experiment).

Martin Taylor (930617 09:50) correctly answered my three questions. But he confuses me once again. In answer to:

>3) Can you do it if I also give you the output function that was in effect at
>the time? Yes or no?

Martin says:

>Yes, if you assert that the output function includes the environmental
>path between the ECS and the CEV to which d is applied, and that the
>function is defined precisely and is monotonic.

This is true. But a control system that worked under these restrictions would be working in a very contrived environment. In fact, it would not even have to be a control system; a calculated output system would do since the outputs that it generates are guaranteed to counter the effects of the disturbing variable. We could call this kind of environment an "artificial intelligence" world since most AI systems ultimately work in this kind of environment (since they are organized to produce preselected outputs, not inputs).

Martin goes on:

>Given that the relation
>between the output signal leaving the ECS and the effect of that output
> on the CEV is dynamically variable, the perceptual signal also carries
> information back to the output about the current relation between the
> output signal and the CEV.

So you are saying that the perceptual signal carries information about the relationship between output and CEV. Yet, in your answer to question 3) you said that you could not reconstruct the disturbance unless you were given information about the relationship between the output and CEV (as part of the specification of the output function).

This is why I want to stop talking about information and start talking about how control systems really work. I don't care if you guys want to go on thinking that information is the greatest thing since sliced bread and that it is the basis of control theory and all that stuff. But it is clear from your post that information in the perceptual signal about the relation between output signal and CEV is of NO USE to you because you cannot use it (as you admitted in your answer to 3) to reconstruct the disturbance. If you can't use it, how do you know that the control system uses it (as it presumably must in order to have effects on the CEV that counter those of the disturbing variable)?

> you just don't understand what is meant by information.

Apparently not; and somehow I feel like I just don't care. I know exactly how control systems operate; I can build control systems that operate in environments where the connection between output and CEV is constantly changing and I can do it without understanding what is meant by information. You're not getting me highly motivated to try to understand information; it seems to me that, whatever information is, it contributes more to mis-understanding than to understanding the nature of control.

>To show that the information
>about the disturbance was available in the perceptual signal, given
>the fact of control, all you need is to show that using a given language
>it is shorter to describe the output function than it is to describe
>the disturbance time series (or continuous waveform).

Ah. A new way to show that there's information in something. Ok Ok There's information in perception. What's it good for? Keeping your descriptions shorter?

>Please try to remember: the ability to reconstruct is conclusive proof
>that the information about the entity is available.

> All an inability to reconstruct
>can demonstrate is that the available information is not complete.

Well, then you can't lose, can you. Congratulations and enjoy all the information in those perceptual signals. May all your descriptions be brief.

Best Rick

Date: Thu Jun 17, 1993 9:35 am PST
 Subject: Control vs entrainment

[From Bill Powers (930617.1100 MDT)] Jeff Hunter (930617) --

> Well here's an example of a collective reference signal
 >for you. If a number of women are confined together their
 >menstrual cycles will synchronize. This does not require the
 >conscious knowledge of the women, it seems to happen just by
 >pheromones.

Why do you say this is a reference signal? I think the word you want is not "control" but "entrainment." When independent oscillators having very low "Q" (little resistance to changing frequency) are loosely coupled together, they will tend to come into synchronization. There is no control action involved; it's just a matter of small timed effects occurring close to the natural resonant frequency. There must be many small influences that a menstruating woman has on other women. Pheromones are one possibility. Perhaps there are others. But as far as I know there is nobody who is watching the various menstrual cycles and adjusting each one so that they all become synchronized, nor is that a shared intention of the women involved (although some have claimed that it is a competitive strategy).

> This seems to be a pretty clear example of a "society"
 >setting the reference signals for individual humans, and it
 >hopefully will be a bit less of a "'tis so, 'taint so" issue.

It looks to me more like an interaction that has an effect on uncontrolled variables. If menstrual cycles were truly under control as to phase and frequency, they wouldn't be that easy to change. I think your example is evidence that they are not under control in those respects.

Best, Bill P.

Date: Thu Jun 17, 1993 10:02 am PST
 Subject: Reply to Tom B.

[From Rick Marken (930617.1100)] Tom Bourbon (930617.0943) --

>But, Martin, doesn't your request for more -- I hesitate to use the word --
 >information about (in the sense of facts about) the output function vitiate
 >the claim that a modeler can reconstruct the disturbance from the
 >perceptual signal alone?

I'll take the liberty of answering for Martin: YES.

> Will the parties *please* clarify the rules of the demonstration?

The rules should have been very simple:

- 1) I send a sequence of 50 p values;
- 2) Martin, Allan, etc send we back the 50 d values that occurred at the same time.

Because of my ridiculously liberal (and flamboyant) nature, I "gave in" to Martin's demands for the reference signal, r, and the output function ($o = k(e)$) despite Bill Powers' protests that this is information that is NOT available to the control system itself.

In April or early may I sent two sets of 50 p values to Allan Randall and Gary Cziko. I also sent the corresponding d values to Gary Cziko. Along with the 50 p values I gave the value of r (I think it was a constant and 0) and the value of k in the output function along with the fact that it was a pure integrator.

I never got back the 50 d values from anyone (except Gary, who properly declared the lack of response a victory for the "no information in perception" crowd).

I think it's clear from the whole go round on reconstructing d that Martin et al have no intention of giving up the idea that information in perception is used by control systems to generate the outputs that control their perceptual inputs. Demos and models and math bounce of these guys like bullets off of Superman. So I'm willing to concede that there must be information in the perceptual input to a control system; it's there because at least two people in the world (Martin and Allan) can perceive it and say a lot of things about it. It is apparently useful, however, only as a basis for conversation; it seems to contribute nothing to our ability to understand or build control systems. So even though you and I can't detect the information in the perceptual input to a control system or understand what it is (I presume you are having the same problems I am) we are apparently OK as long as all we want to do is understand purposeful behavior and build models of purposeful systems.

Best Rick

Date: Thu Jun 17, 1993 10:20 am PST
Subject: Re: Reconstructing the disturbance

[Martin Taylor 930617 13:45] Tom Bourbon 930617.0943)

>But, Martin, doesn't your request for more -- I hesitate to use the word --
>information about (in the sense of facts about) the output function vitiate
>the claim that a modeler can reconstruct the disturbance from the perceptual
>signal alone? That would be the challenges confronting the organism modeled
>by the ECS. Or was the claim, not that information in p about d is
>essential to the organism if it is to control, but that a modeler can
>reconstruct d using p and other facts, none of which are available to the
>organism modeled by the ECS?

There are two separate issues involved.

(1) My claim initially that I could derive the structure of the hierarchic control system from information theory. You pointed out that all sorts of people have claimed to be able to deduce things from information theory that have not proved to be so. I recognized the truth of what you said, and so have held back on that one until I can satisfy myself that the analysis is correct. This I have still not done. But the claim rests on a sub-claim, that to the extent that the disturbance is countered by the output of a control system, information about the disturbance passes through the perceptual signal and is used to form the appropriate compensating output. This is NOT, and never was, a claim that a modeller can reconstruct the disturbance from the perceptual signal alone. The claim is that information about the controlled part of the disturbance is passed through the perceptual signal. Control itself allows the information to be used. The control system does not have to "know" its own output function, because that is handled by bringing the error to zero. The "knowledge" is implicit in the fact of control.

There is also information in the perceptual signal about the part of the disturbance that is not compensated by the control system, but we are not talking about that, since it is accepted to be so by all parties to the discussion.

(2) A demand by Rick, supported by Bill, that we demonstrate the existence of information about the disturbance in the perceptual signal by reconstructing the disturbance waveform given the perceptual signal under the assumption that the reference signal remains zero. We pointed out that success in this endeavour would conclusively demonstrate that the perceptual signal contained information about the disturbance, and Rick agreed that it would. When we showed by example what should have needed no demonstration, reconstructing the disturbance exactly, knowing the output function exactly, Rick then said it was insufficient, and produced a series of examples using different output functions and/or changing reference signals (I forget which). When we ignored this irrelevant challenge, for reasons restated this morning, Rick took that as proof of his position that there is NO information about the disturbance in the perceptual signal.

We have tried to point out in various ways, starting from differently phrased but equivalent bases, that failure to perform a reconstruction was inconclusive about whether the perceptual signal contains information about the disturbance. Despite this, we keep reading that if the reconstruction is imperfect, then THERE IS NO INFORMATION ABOUT THE DISTURBANCE IN THE PERCEPTUAL SIGNAL (capitals courtesy of Rick). The point about the shape of the output function is that if the output function can be described in a finite string,

and is fixed, then eventually its contribution to the information required to reconstruct a disturbance waveform of indefinite length becomes negligible. Only if the output function is changing, so that its description requires a certain amount of information per second, does it become a serious issue in ascertaining the contribution of the disturbance to the information in the perceptual signal.

The two issues, reconstruction and information about the disturbance in the perceptual signal, are separate. Success at reconstruction proves the point that information about the disturbance is available in the perceptual signal. Failure of reconstruction does not disprove it.

>I thought the claim was an ability to reconstruct
>the disturbance from the perceptual signal, and that the claim was not yet
>borne out. The demonstration that the disturbance can be reconstructed from
>the perceptual signal, given a known output function, must have appeared
>before I was relying on the net.

The demonstration is simple. Create a "mystery function" $M(p-r)$, where M has exactly the form of the output function O . To the extent that the output of O is equal and opposite to the disturbance (a claim by Rick and Bill that is simply a statement of perfect control), then the negative of the output of M is the disturbance waveform minus an arbitrary starting constant. It should have been sufficient simply to point this out in order to move the discussion to the next stage, but it wasn't. Rick said it wouldn't work as we thought. We had to (and did) demonstrate that it worked with real data. And then the rules kept changing in ways that confused us as much as they seem to confuse you.

Martin

Date: Thu Jun 17, 1993 10:53 am PST
Subject: Re: Point of view; jumping levels; cells & people

[Martin Taylor 930717 14:10] (Bill Powers 930617.0730)

>>If the world's responses to outputs were linear, a one-level
>>system of orthogonal ECSs would be perfectly adequate.
>
>Adequate for what? For controlling the rate at which a variable
>changes back and forth between two values? For adjusting the
>placement of flowers in a bowl so they have a pleasing look? For
>raising the garage door before rather than after you back the car
>out?

Yes, I think so. Remember that the orthogonality extends across time as well, so that the inputs could be derived from a shift register. Caveat: the "flower" example seems to have an intrinsic non-linearity, in that there is a bifurcation point between a flower going "inside" and the same flower going "outside" the bowl. There are trajectories initially arbitrarily close, one of which ends with the flower in the bowl and one that ends with the flower outside and wilting.

>You've said before that you think a perceptron using weighted
>sums of its inputs is adequate for modeling all of perception.

Not quite. If I said that, I was being clumsy. A multilayer perceptron can categorize any partition of its basis space in an arbitrary way. If you consider only sensory data from a single moment in time, it can categorize any and all kinds of configuration. If you include the possibility of data extended over time, then the categorization incorporates configurations over time (super-transitions, if you like). To the extent that perceptions can be represented as configurations in space-time, a multilayer perceptron whose inputs include the possibility of a shift register can provide the perceptual signals to control them.

>I don't think you're considering all the kinds of perceptions there
>are, and how one kind must be controlled in order to control
>others because of the way the external world is organized. I know
>that there's a degrees-of-freedom argument in the background, but
>there's also a practicality argument: it's just not practical to
>accomplish all the different higher-level perceptions by means of
>a single level of transformation, even though mathematically the

>difference isn't visible. Nonlinearity isn't the only factor. And
>the mathematics is limited.

I agree wholly with the practicality argument. It would not be resource efficient to try to control with a single layer, even if it were possible. I bring in non-linearity, because that makes single-layer control intrinsically impossible (at least a lot of kinds of nonlinearity do).

=====
ECS insertion:

Thanks for the expansion on why the output level-jumping is bad. What will happen, as you say, is conflict, leading to error, leading to reorganization after the insertion of the new ECS. Eventually, the reorganization should lead to reduction or elimination of conflict, and the new ECS will be working more or less orthogonally to the old ones, or will be working in a different region of the perceptual space, which comes to the same thing.

Martin

Date: Thu Jun 17, 1993 11:16 am PST
Subject: Re: Reply to Tom B.

From Tom Bourbon (930617.1253) Rick Marken (930617.1100)

>Tom Bourbon (930617.0943) --

>>But, Martin, doesn't your request for more -- I hesitate to use the word --
>>information about (in the sense of facts about) the output function vitiate
>>the claim that a modeler can reconstruct the disturbance from the
>>perceptual signal alone?

>I'll take the liberty of answering for Martin: YES.

Rick, your answer does not surprise me. By my question, I guess I revealed my agreement with your interpretation of the significance of Martin's request. But I am still eager to see Martin's reply.

>> Will the parties *please* clarify the rules of the demonstration?

>The rules should have been very simple: ...

Apparently my mind-reading of your interpretation of the rules was correct. And they are precisely the rules needed to resolve the question of whether one (any one) can reconstruct d from p.

>In April or early may I sent two sets of 50 p values to Allan Randall
>and Gary Cziko. I also sent the corresponding d values to Gary
>Cziko. Along with the 50 p values I gave the value of r (I think it was
>a constant and 0) and the value of k in the output function along with
>the fact that it was a pure integrator.

Exactly what should be provided -- and then some. You must have been in a mellow mood that day.

>I never got back the 50 d values from anyone (except Gary, who
>properly declared the lack of response a victory for the "no
>information in perception" crowd).

OK. That was about where my connection to the net became firm. I suppose the offer still stands, doesn't it? The test was not the kind that has a time limit. *Any time* someone shows they can reconstruct the 50 values of d from the 50 values of p (even if they must cheat and use the other facts you provided) would prove the claim. I think that *all* future discussions about the importance of information theory in PCT should occur with the fact of the as-yet-unmet conditions of that offer *clearly* stated up front. (Or even at the other end -- just so everyone remains aware of the real state of affairs.)

>... So even though you and I can't detect the information

>in the perceptual input to a control system or understand what it is
>(I presume you are having the same problems I am) we are apparently OK
>as long as all we want to do is understand purposeful behavior and
>build models of purposeful systems.

That's fine by me!

Until later, Tom Bourbon

Date: Thu Jun 17, 1993 11:29 am PST
Subject: You're serious, aren't you?

[From Rick Marken (930617.1200)] Martin Taylor (930617 13:45) --

>We have tried to point out in various ways, starting from differently
>phrased but equivalent bases, that failure to perform a reconstruction
>was inconclusive about whether the perceptual signal contains information
>about the disturbance.

Sounds like a win-win proposition to me.

But tell me, is there some way that I can know in advance whether or not I am going to succeed at my reconstruction? Not that that's important or anything. Oh, by the way, did you know that there is information in a flipped coin about how it's going to land? Yeah. Really. I know this because I can look at films of coin flips and reconstruct what the result of each flip was based on the information in the flip. Sometimes I reconstruct the result correctly -- I say tails when it was tails; and I say heads when it was heads. Of course, I can't tell what happened after some of the flips; for example, I predict tails it turns out to be heads. But this failure to reconstruct the result of the flip based on the information in the flip is certainly inconclusive about whether there is information in the flip about how the coin will land; the correct reconstructions prove that there is information in the flip about how the coin will land.

Oh, and I'm selling my coin flip information detection system for just \$1,000,000 each (nothing compared to the amount you can expect to make with this powerful tool). If you order now, you can get a control system perceptual information detector thrown in for FREE.

>Despite this, we keep reading that if the
>reconstruction is imperfect, then THERE IS NO INFORMATION ABOUT THE
>DISTURBANCE IN THE PERCEPTUAL SIGNAL (capitals courtesy of Rick).

Only because I can't italicize and underline too.

Keep enjoying that old time information.

Best Rick

Date: Thu Jun 17, 1993 12:18 pm PST
Subject: learning

[from Joel Judd 930617]

Following is a short piece somebody left for me last week. It is from
How to Survive in Your Native Land by James Herndon.

JAIL

Of course I have forgotten to tell you again, what you already know. That is that the fundament of the school, even before winners and losers, is that everyone has to go there.

Even if you are rich or have eccentric or far-out parents and go to private schools or invent free schools for yourself, you still have to deal with the public school. But in case most of you are not or have not any of the above. You must go to school.

If kids in America do not go to school, they can be put in jail. If they are tardy a certain number of times, they may go to jail. If they cut enough, they go to jail. If their

parents do not see that they go to school, the parents may be judged unfit and the kids go to jail.

You go to jail. All of the talk about MOTIVATION or INSPIRING kids to learn or INNOVATIVE courses which are RELEVANT is horseshit. It is horseshit because there is no way to know if the students really are interested or not. No matter how bad the school is, it is better than jail. Everyone knows that, and the school knows it especially. A teacher comes into the teacher's room and says happily, I had the greatest lesson today! and goes on to tell the other envious teachers what it was that they hadn't thought of themselves and says, The kids were all so excited! It is horseshit. The teacher has forgotten (as I forget) that the kids have to be there or they will go to jail. Perhaps the grand lesson was merely more tolerable than the usual lesson. Perhaps the kids would have rejected both lessons if they could.

That is why the schools cannot ever learn anything about their students. Why famous psychologists can successfully threaten pigeons into batting ping-pong balls with their wings, but can never learn anything about pigeons.

As long as you threaten people, you can't tell whether or not they really want to do what you are proposing that they do. You can't tell if they are inspired by it, you can't tell if they learn anything from it, you can't tell if they would keep on doing it if you weren't threatening them.

You cannot tell. You cannot tell if the kids want to come to your class or not. You can't tell if they are motivated or not. You can't tell if they learn anything or not. All you can tell is, they'd rather come to your class than go to jail.

Date: Thu Jun 17, 1993 2:11 pm PST
Subject: just a blip

[From: Bruce Nevin (Thu 930617 16:26:24 EDT)]

Lots to say, no time to say it. Probably next week. Easy, however, to quote the following, in case it's of interest:

From: Terry Fogarty <tc_fogar@pat.uwe-brisol.ac.uk>
Date: Tue, 8 Jun 93 10:06:54 BST
Subject: ICGA-93 Workshop on Genetic Algorithms in Control

Call for Participation

ICGA-93 Workshop on Genetic Algorithms in Control

A workshop on the application of Genetic Algorithms to Control Systems Engineering has been proposed for ICGA this year to take place during one of the two workshop sessions. If you are attending ICGA and would like to participate in the workshop would you please let me know as soon as possible.

Terry Fogarty

tel +44 272 656261 ext 3179
fax +44 272 750416
email tc_fogar@csd.uwe.ac.uk
post Dr Terence Fogarty
Faculty of Computer Studies and Mathematics
University of the West of England, Bristol
Coldharbour Lane
Bristol, BS16 1QY
England.

Date: Thu Jun 17, 1993 2:37 pm PST
Subject: Re: Reply to Tom B.

[Martin Taylor 930617 18:10] (Rick Marken 930617.1100)

Rick, please don't start a flame war on this mailing list:

>In April or early may I sent two sets of 50 p values to Allan Randall
>and Gary Cziko. I also sent the corresponding d values to Gary
>Cziko. Along with the 50 p values I gave the value of r (I think it was
>a constant and 0) and the value of k in the output function along with
>the fact that it was a pure integrator.

The last two facts were not included. We explained why that information would be needed.

>Demos and models and math bounce of these guys like bullets off of Superman.

Who was it that refused to believe the results of a demonstration, a mathematical description, and a model that he initially had accepted as being definitive and that would not give the results Allan and I predicted it would? Some loose canon out in California, I think.

Shall we leave it at that?

Martin

Date: Thu Jun 17, 1993 2:38 pm PST
Subject: emergent control

[Hans Blom, 930617a] (Bill Powers (930617.0730 MDT))

Please appreciate that I stipulated quite different conditions in my example from those that you use:

>Suppose we have set up the program so that person A seeks a
>specific distance from person B, person B from C, and C from D,

My control laws for each individual were:

- 1) keep a minimum distance to others;
- 2) keep a maximum distance from others;
- 3) otherwise move at discretion (no specific distance!).

These describe the behavior of a real flock much better, in my opinion. It may be that you do not like the random element, nor the non-linear control law, but let's stick to my example for the moment, shall we?

Now let us look at your other stipulations for The Test to see whether 'flockness' is controlled.

>It's easy to forget that resistance to disturbance is only one
>part of the Test. There are three other parts. One essential part
>is to make sure that the apparent resistance of the controlled
>variable is not due entirely to a natural tendency of the
>controlled variable to resist disturbance.

What do you mean by 'natural tendency'? The same thing as Aristotle when he said that a rock has a 'natural tendency' to strive for the center of the earth? Then it is rhetoric. Please be explicit in what discriminates a 'natural tendency' from control (should you do so, I am tempted to say, I will give you a counterexample :-). When I look at my simulated flock, I may PERCEIVE a 'natural tendency' to form a flock, but that is just my subjective interpretation. To understand (or model) what goes on, I need a law or a formula (and I don't care much whether you call that law or formula a 'natural tendency' or a 'control law').

Excursion: the formulas of quantum mechanics are agreed upon by all physicists, yet there are two vastly different 'interpretations'. Let's try to stick to objective descriptions and try to avoid 'interpretations'.

> Another part is to
>interrupt the ability of the supposed control system to perceive
>the controlled variable, and show that the resistance to
>disturbance becomes unsystematic -- that control is lost.

An individual does not perceive 'flockness', yet 'flockness' is preserved even if individuals are hidden from sight. Seeing only a few individuals preserves flockness. Even seeing only one other individual will do if you accept some funny flock shapes, such as (straight or crooked) lines. Isn't it remarkable how robust flockness is in the face of disturbance?

>And the final part is to remove the control system entirely (when
>possible) and show what the effect of the disturbance on the
>controlled variable is when the system is not opposing it (there
>might be some other system acting that we haven't noticed). The
>effect should be much larger than when the proposed control
>system is present.

Remove WHICH control system? There is no control system for 'flockness'; that hypothetical one cannot be removed. Remove the sensors that measure distances to others, and the flock will disintegrate: a random walk moves each individual away without limit.

Now, does 'flockness' pass The Test for the Controlled Variable?

(Rick Marken (930617.0900))

> The variable that you call "flockness"
>has to be better defined. If you mean the average distance
>between all individuals in the flock, then this is NOT an
>emergent (uncontrolled) variable; it is a controlled variable; a
>"piece" of this variable is being controlled by each individual
>in the flock. The Test you propose would show that "flockness"
>(my definition) IS a result of the controlling done by each
>individual.

Yes, of course.

I define 'flockness' loosely as 'staying together'. Without proximity control, the flock would drift apart without limit. With proximity sensing, the distance of each individual to the common center of gravity remains bounded. Is that good enough? If you want a numeric quantity, take the average distance to the center of gravity of the flock. Flockness is preserved if this quantity stays bounded.

Now 'flockness' is not controlled in the sense that individuals have a sensor for it. Yet 'flockness' IS controlled in the sense that the average distance to the common center of gravity remains within bounds. You are of course right: the individual control laws together 'control' flockness, but in an emergent way, just like a collection of molecules in a gas have a (common) temperature whereas you cannot reasonably talk about the temperature of a single molecule (yes, you can; I just did it :-).

>You can call the flock a social control system. No single
>individual has a notion of what a flock is: it just controls
>its own movements. Yet a (moving but stable) flock emerges.

This kind of simulation is fun. Make one individual different in that you give him/her a NON-RANDOM preferential movement and see what happens. You can endlessly play around with modifications like that and observe some quite remarkable results, some of which can give you a deep insight into notions like 'leadership', 'being afraid to be left alone', and such. Are these human notions just the emergent results of some anomalous parameter settings?

> If each member of the flock has
>about the same reference for the distance between itself and
>others the default shape of the flock will be a sphere (assuming
>the individuals are not moving). If you put the flock in a box
>that is smaller than the default sphere I predict that the flock
>will assume the shape of the box -- with the individuals in the
>corners being unable to keep the distance between themselves and
>their neighbors at their reference level. What will NOT happen
>(and what shows that the shape of the flock is NOT controlled) is
>an adjustment of the inter-individual references so that the
>sphere shape is preserved.

You too assume different conditions from the ones I used. My flock will not be circular, except maybe on average. It does not control for shape (which would require the equivalent of an all-seeing eye if the flock gets large), just for staying together. A box or some such could be called a 'higher level' control law: 'Thou Shalt Stay Within Mine Confines'. Higher level control laws can override or modulate lower level ones, as you show.

(Bourbon (930617.1045))

>You seem to have ruled out the possibility that synchronization
>of the cycles [menstruation] is an uncontrolled side effect ...

This whole discussion indicates to me that there is no common terminology for what control is. To me, a control system is a mechanism that keeps something within specified (but maybe varying) limits or almost so, whereas without that mechanism those limits would not be observed. Therefore, according to my terminology, 'flockness' is controlled, although YOU might consider it an 'uncontrolled' side effect. My home heating system, in your terminology, 'controls' the temperature at its sensor, while the thing that I want, a nice room temperature, is only an 'uncontrolled side effect'. That is not the kind of language that I use, although I can understand it when YOU do. But then, understanding (here I mean: making sense of what someone says) often requires this kind of back and forth 'translation', isn't it?

Greetings, Hans Blom

Date: Thu Jun 17, 1993 2:53 pm PST
Subject: Re: learning and "motivation"

\From Greg Williams (930617 - 2) Joel Judd 930617 [quoting Herndon]

> You go to jail. All of the talk about MOTIVATION or INSPIRING
>kids to learn or INNOVATIVE courses which are RELEVANT is horseshit.
>It is horseshit because there is no way to know if the students really
>are interested or not. No matter how bad the school is, it is better
>than jail. Everyone knows that, and the school knows it especially.
>...

> As long as you threaten people, you can't tell whether or not
>they really want to do what you are proposing that they do. You can't
>tell if they are inspired by it, you can't tell if they learn anything
>from it, you can't tell if they would keep on doing it if you weren't
>threatening them.

> You cannot tell. You cannot tell if the kids want to come to
>your class or not. You can't tell if they are motivated or not. You
>can't tell if they learn anything or not. All you can tell is, they'd
>rather come to your class than go to jail.

This got me thinking about what anyone can tell about ANYONE ELSE who is EITHER acting under perceived threat OR perceived reward. I think it is about as hard to decide whether somebody is REALLY "motivated" (I assume that Herndon means, in some sense, "purely" or "deeply" or "intrinsically" motivated) or not "motivated" EITHER if they've been told they'll go to jail if they don't do something OR if they've been told they'll get to Heaven if they do. And if they believe that they'll go to Hell if they don't and to Heaven if they do... The "power" of social influence via perceived threats or rewards is indeed awesome. Or could be. Herndon's observations could be taken as evidence that the threateners just were a little sloppy -- they MEANT to say "you LIKE going to school, or you go to jail," but it ended up "if you don't SHOW UP at school, you go to jail." I suspect that some PCTers think it is IMPOSSIBLE to enforce an edict such as "you LIKE this, or else," and I am sure that it WOULD be impossible to enforce such an edict ON THEM, and even on MANY individuals in the modern-day U.S. But it might not be so hard to enforce such an edict on many members of, say, an Amish community or an Australian aboriginal group.

As ever, Greg

Date: Thu Jun 17, 1993 2:55 pm PST
Subject: Workman's Comp and Info in Perception

[From Rick Marken (930617.1500)]

Many of you are probably asking yourselves "why does this Marken character get so fired up when the issue of information in perception comes up?" I will try to explain it in terms of Workman's Compensation Insurance.

For the last several years, California (with the help of a bunch of lawyers) has been trying to commit hari kari with Workman's Compensation laws that make it easy for people to get paid big bucks for stress complaints. About all you have to do for a well paid, arbitrarily long vacation is say that something at work "caused" you to have stress. If you have trouble thinking of ways of collecting this largess, all you have to do is call a lawyer - 1 -800- NOSTRESS.

My wife is now a big mucky muck in the personnel department at her place of employment (making big bucks so I can sit here wasting time on the net; but, then, I have to make dinner so it all works out) and she comes face to face with this problem every day; anything that happens to a person at work -- the boss criticizes you, the door isn't open when you get there -- anything is grounds for a stress claim. Obviously, businesses have to pay BIG BUCKS for Workman's Comp so they are motivated to relocate to hell holes like Colorado (just kidding Bill) where Workman's Comp is a lot cheaper.

My wife is currently as fired up against Workman's Comp stress claims as I am against information in perception -- she thinks the former should not exist, I know that the latter doesn't exist. It turns out that these obsessions are related. Workman's Comp stress claims are based on the idea that what happens to people ("stressful information") causes what they do ("have stress"). PCT shows that people control what happens to themselves. So it is as incorrect to blame stress on "stressful information" as it is to "blame" the output of a control system on perceptual information.

A Workman's Comp stress suit is based on the same mistaken concept of behavior as is the notion that there is information in perception. The apparently academic mistake of locating the information for action in perception becomes, with a few word changes, the basis of a destructive myth that is stifling the quality of life in my state. The criteria for stress claims would change significantly if people understood how control works -- it's control OF perception, NOT control BY perception.

It matters. Best Rick

Date: Thu Jun 17, 1993 3:36 pm PST
Subject: Re: Reconstructing the disturbance

From Tom Bourbon (930617.1447)

>[Martin Taylor 930617 13:45]
>(Tom Bourbon 930617.0943)
>

>>But, Martin, doesn't your request for more -- I hesitate to use the word --
>>information about (in the sense of facts about) the output function vitiate
>>the claim that a modeler can reconstruct the disturbance from the perceptual
>>signal alone? That would be the challenges confronting the organism modeled
>>by the ECS. Or was the claim, not that information in p about d is
>>essential to the organism if it is to control, but that a modeler can
>>reconstruct d using p and other facts, none of which are available to the
>>organism modeled by the ECS?

>
> There are two separate issues involved.

>
>(1) My claim initially that I could derive the structure of the
>hierarchic control system from information theory. You pointed
>out that all sorts of people have claimed to be able to deduce
>things from information theory that have not proved to be so. I
>recognized the truth of what you said, and so have held back on that
>one until I can satisfy myself that the analysis is correct. This
>I have still not done.

Thank you for reporting the status of that project.

>... But the claim rests on a sub-claim, that
>to the extent that the disturbance is countered by the output of a

>control system, information about the disturbance passes through the
 >perceptual signal and is used to form the appropriate compensating
 >output. This is NOT, and never was, a claim that a modeller can
 >reconstruct the disturbance from the perceptual signal alone. The
 >claim is that information about the controlled part of the disturbance
 >is passed through the perceptual signal.

Can you tell me (a) how *you* (MT) distinguish between controlled and uncontrolled parts of a disturbance, as they occur in a perceptual signal, and (b) whether, and if so, how, an organism, *modeled* as a PCT system, does so in order to control? I am having a hard time imagining what you mean. Do you mean that information about the part of the disturbance that is *not* controlled at time t "passes through the perceptual signal" and becomes controlled at time $t+x$?

>...Control itself allows the
 >information to be used. The control system does not have to "know"
 >its own output function, because that is handled by bringing the
 >error to zero. The "knowledge" is implicit in the fact of control.

That a control system does not "know" its own output function, I agree. You did not answer my inquiry about your definition of "output function," but I suspect that you mean the time series of actions, not the operator that transforms error signals into actions. Did Rick give you the operator, or the time series? (From what you say later, I infer it was the latter.)

>(2) A demand by Rick, supported by Bill, that we demonstrate the
 >existence of information about the disturbance in the perceptual
 >signal by reconstructing the disturbance waveform given the perceptual
 >signal under the assumption that the reference signal remains zero.

This seems to be a reasonable and modest proposal.

>We pointed out that success in this endeavour would conclusively
 >demonstrate that the perceptual signal contained information about
 >the disturbance, and Rick agreed that it would.

I would agree.

> When we showed
 >by example what should have needed no demonstration, reconstructing the
 >disturbance exactly, knowing the output function exactly, Rick then said
 >it was insufficient, and produced a series of examples using different
 >output functions and/or changing reference signals (I forget which).

But what you did was to construct the disturbance waveform from the output waveform. That was not the demonstration you described in the earlier parts of this paragraph. Had I been actively on the net when Rick gave you the output data, I believe I would have registered a protest. The additional data dramatically changed the nature of the demonstration, away from showing:

>existence of information about the disturbance in the perceptual
 >signal by reconstructing the disturbance waveform given the perceptual
 >signal under the assumption that the reference signal remains zero.

You demonstrated a reconstruction of the disturbance from the output time series. If Rick protested that shift in the focus of the demonstrations, I agree with him, but you shouldn't have been given the chance to use the additional data in the first place.

>The two issues, reconstruction and information about the disturbance in
 >the perceptual signal, are separate. Success at reconstruction proves
 >the point that information about the disturbance is available in the
 >perceptual signal. Failure of reconstruction does not disprove it.

That the two issues are separate, I do not agree -- not if our purpose is to construct better generative, freely-functioning, models of control systems. If you believe discussions *about* control systems are more complete or satisfying or precise if you include the idea that "information about the disturbance is available in the perceptual signal," that is one thing. It seems quite another to claim that a PCT model would behave

more realistically if should you include an identifiable quantity that represents information about the disturbance, as a modeled feature of the perceptual signal. The former preference seems to be from the perspective of an observer, striving for an aesthetically pleasing description. The latter, from the perspective of a modeler who believes the new measure will demonstrably improve the performance of the model. Am I very far from the mark with this interpretation? In which sense do you think we should embrace the idea that "information about the disturbance is available in the perceptual signal?"

Given that your demonstration is not the one originally proposed, I do not see how it confirms your original claim.

As for the idea that a (single) failure to reconstruct does not disprove the original claim, I agree. The claim was about a fact of nature. The failure of a model to reconstruct the fact does not invalidate the fact. (Not any more than a successful reconstruction offers "final" proof that a model is "the correct one.") But a failure to reconstruct does invalidate that particular rendering of the model.

>>I thought the claim was an ability to reconstruct
>>the disturbance from the perceptual signal, and that the claim was not yet
>>borne out. The demonstration that the disturbance can be reconstructed from
>>the perceptual signal, given a known output function, must have appeared
>>before I was reliably on the net.

>
>The demonstration is simple. Create a "mystery function" $M(p-r)$, where
>M has exactly the form of the output function O. To the extent that
>the output of O is equal and opposite to the disturbance (a claim by
>Rick and Bill that is simply a statement of perfect control), then
>the negative of the output of M is the disturbance waveform minus an
>arbitrary starting constant. It should have been sufficient simply to
>point this out in order to move the discussion to the next stage, but
>it wasn't. Rick said it wouldn't work as we thought. We had to (and
>did) demonstrate that it worked with real data. And then the rules
>kept changing in ways that confused us as much as they seem to confuse
>you.

Is M a convenient function for an observer who tries to reconstruct the disturbance from the output function of a PCT model, or is it a feature of a PCT model -- a feature that leads to more realistic performance by the model?

Given that you demonstrated something different from what you said (at the beginning of (2) in your post), I have the impression you speak from the position of an observer when you speak of M, and of the original claim about information in the perceptual signal. Or do I misread you? Do you now suggest that M is a function in the model of an elemental control system?

I am afraid my confusion over your interpretation of these demonstrations has not been allayed. To me, much of my confusion comes, not from interpretations of the meanings of theoretical terms, but from what I see as a dramatic shift in the goal of your demonstration and in the nature of the data you use to perform the reconstructions.

Until later, Tom Bourbon

Date: Thu Jun 17, 1993 4:30 pm PST
Subject: Re: Information

[Martin Taylor 930617 17:35] (Rick Marken 930617.0900)

>Martin Taylor (930617 09:50) correctly answered my three questions.
>But he confuses me once again. In answer to:
>
>>3) Can you do it if I also give you the output function that was in effect at
>>the time? Yes or no?
>
>Martin says:
>
>>Yes, if you assert that the output function includes the environmental
>>path between the ECS and the CEV to which d is applied, and that the
>>function is defined precisely and is monotonic.

>
 >This is true. But a control system that worked under these restrictions
 >would be working in a very contrived environment. In fact, it would
 >not even have to be a control system; a calculated output system
 >would do since the outputs that it generates are guaranteed to counter
 >the effects of the disturbing variable.

A misunderstanding here. All I asked was to know whether the output function was, say $10 * \text{integral}(\text{error})$, or $100 * \text{error}^3$ or whatever. The requirement was that this function be fixed. The assertion that the perceptual signal contains the information about the disturbance comes from the fact of control and the stability (relative to the disturbance) of the output function. And that stability has to extend as far as the CEV. It has to be a control system, not a calculated output system, or none of the argument works (it won't, either).

I agree that it is a most unnatural situation to have a fixed output function. I tried to persuade you of that in the earlier discussion on information, but you insisted that it was not so--that the output function would normally be fixed. My argument, which I think you now accept, was that control can be attained with an output function that is not only nonlinear, but which is also time-variant and uncertain. The information-theory view on this is that the time-variance and the uncertainty of the output function (i.e. the transform between error and the effect on the CEV) has to be considered as information that the control system must deal with, reducing the ability to control against the effects of the disturbance. Numerically, it is subtracted from the controllable information in the disturbance.

>So you are saying that the perceptual signal carries information about the
 >relationship between output and CEV. Yet, in your answer to question 3)
 >you said that you could not reconstruct the disturbance unless you were
 >given information about the relationship between the output and CEV (as
 >part of the specification of the output function).
 >
 >This is why I want to stop talking about information and start talking
 >about how control systems really work.

That "This is why" is the non-sequitur that bugs me. I know I have a problem getting you interested. I have a problem getting Bill P interested, though he has asked a lot of penetrating questions that have helped me refine my concepts. But I hope I will succeed, and not go off into isolation mode, as you are tempted to do in respect of "mainstream" psychologists.

The point is, and always has been, that the "knowledge" of the output function is implicit in the fact of control. It is explicit only in the view of a modeller--an outsider view. It is the outsider who has to know the form of the function in order to do the calculation. The control system does it by correcting for error. If the output function is fixed, the amount of information involved in its specification very soon is swamped by the stream of information from the CEV (i.e. from the disturbance, if the reference is fixed). It vanishes from the calculation, and does not affect the ability of the system to control except insofar as a non-linear output function has a gain that differs for different levels of output. But if the output function is time-varying, information about it has to be factored into the calculation of the ability of the control system to control. That matters. And in a control hierarchy in which reference levels for lower ECSSs come from several independent higher ECSSs, it matters a lot.

>I know exactly
 >how control systems operate; I can build control systems that operate
 >in environments where the connection between output and CEV is
 >constantly changing and I can do it without understanding what is
 >meant by information.

Well, you are one (or more) up on me. I don't know "exactly" how control systems operate, especially if there are lots of them connected in nonlinear ways. I have some notions, however. Can you specify beforehand how such a net will respond, given limits on the perceptual resolution and bandwidth of its various perceptual input functions, and given a variety of behaviours of the multiple disturbances to which the various CEVs will be subjected? Maybe you can. I'm not claiming that information theory will tell you how such a system will operate, but I am claiming that it can provide you with limits on how well the different disturbances can be controlled.

I, too, want to talk about how control systems "really" work. I'm trying to lay the groundwork for doing so from a viewpoint complementary to the one you take, but supportive of it. Didn't thermodynamics help the understanding of many-body systems that were hard to understand using Newtonian mechanics? Newtonian mechanics might be correctly applied, and would give good prediction if the calculations could be carried through (with infinite precision, as we now know). But the thermodynamic approach provided different insights into the behaviour of such systems, insights that would not have been obtained even with an exact Newtonian calculation. In addition, it produced global results that could actually be used, approximate though they might be. I think we are in the same situation.

Martin

Date: Thu Jun 17, 1993 4:53 pm PST
Subject: Re: You're serious, aren't you?

[Martin Taylor 930617 20:30] (Rick Marken 930617.1200)

After this, I'm going to stop direct responses that seem to me to have more the character of a private bicker than of a useful public discussion. I hate 'tis--'tisn't "arguments."

>But tell me, is there some way that I can know in advance whether or
>not I am going to succeed at my reconstruction? Not that that's
>important or anything.

Didn't the answers to your three questions answer that?

Why do you keep focussing on reconstruction? It's the unreconstructible part of the disturbance that leads to the need for more layers of ECSs in the hierarchy. If you can control in one layer, why go for more? The question at issue (for me) is how much is there that needs to be dealt with by higher layers. The control system can control using whatever information it gets through the perceptual system. That much can be reconstructed. What it can't control can't be reconstructed using the perceptual signal, the reference signal and the output function. The questions have the same answer.

The important issues are things like perceptual bandwidth and resolution, and differential system gain at different output levels. Those determine limiting information rates of control, and determine in part what needs to be done beyond the simple ECS, whether it be in parallel ECSs or in higher levels.

Anyway, the general answer is: what is controlled (compensated for) can be reconstructed using the perceptual signal, the reference, and the output function.

Since you know "exactly" how a control system works, I presume that tells you "exactly" what can be reconstructed and what cannot.

Martin

Date: Thu Jun 17, 1993 5:37 pm PST
Subject: Information stuffr

[from Mary Powers 930617.1530] (Jeff Hunter 930617)

What bothers me about your post is

>If a number of women are confined together their menstrual cycles *will* synchronize.

They will? All of them? Always? I suspect that this is one of the gadzillion physiological and psychological *facts* about people in which some statistical probability has been elevated to rock-hard 100% certainty.

I personally lived four years in a female dorm and my 5 week cycle never budged, though I was surrounded by more typical types. I knew 3-weekers who would have been delighted to synchronize, but didn't. Bill's suggestion about entrainment might explain this - in that our oscillations were too far out of synch to be captured. But you claim that even irregular and unusual cycles synchronize too. I suspect that the data behind your assertion was far weaker than you thought it was when you learned that interesting *fact*.

Mary P.

[From Bill Powers (930617.1600 MDT)]

Martin Taylor, Rick Marken, anyone else:

I can see that this discussion of information in the perceptual signal is settling down into a set of fixed misunderstandings. We all seem to be talking about different things.

Martin Taylor (930617 13:45) --)

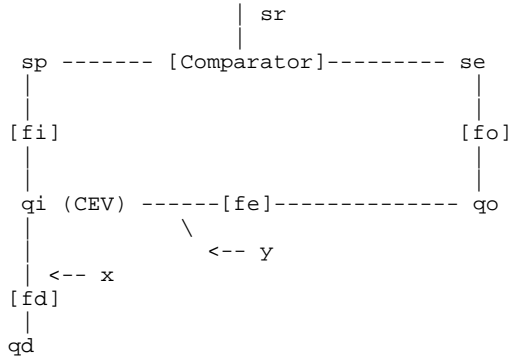
>When we showed by example what should have needed no
>demonstration, reconstructing the disturbance exactly, knowing
>the output function exactly, Rick then said it was
>insufficient, and produced a series of examples using different
>output functions and/or changing reference signals (I forget which).

Rick was holding in reserve a variable feedback function between the output quantity and the controlled quantity. He could give you the perceptual signal, the reference signal, and the form of the output function because he was hiding the (variable) form of the feedback function, rendering the task impossible. You, however, took Rick's later comments this way:

>Yes, if you assert that the output function includes the
>environmental path between the ECS and the CEV to which d is
>applied, and that the function is defined precisely and is monotonic.

This is precisely the part of loop that Rick was making variable without telling you. So Rick continues to assume that you are excluding this part of the loop, and you continue to assume that he is including it. End of that problem. Let's just nail it down.

Here is the diagram:



We've been assuming that fi is a unity transformation involving a change from physical to nervous system units only. So qi is numerically identical to sp. The comparator is a simple subtractor, so that se = sr - sp.

Martin, you've been saying that given sp, sr, fo, and fe, and by default fi, you can reconstruct "the disturbance", whatever that means. I will agree that if by "disturbance" you mean fd(qd), that is, the contribution of the disturbing variable to the state of qi shown as x in the above diagram, you are correct. Anybody could.

Given all functions in the closed loop, plus the reference signal, you can compute "y" in the diagram above, the contribution of the output to the input. As qi = y + x, and given that you know the state of qi from the state of sp and the known input function fi, it is trivial to compute x. qi = sp, and x = y - qi.

In order to reconstruct the state of qd, however, you must also know the form of the function fd, which we've been leaving out. The variable qd is (inverse fd)(x).

When Rick and I have said that you can't reconstruct the disturbance from the perceptual signal, we have meant that you can't reconstruct qd just from knowing sp and sr. Unsaid was "... unless you know the form of all the functions in the loop." But if you know all the functions in the loop, it is trivial to reconstruct qd from sp, because in effect you are

given the states of all the other variables in the loop. Rick allowed you to know (he thought) all the functions in the loop but one, fe. That automatically made the problem unsolvable and the challenge unmeetable. You met the challenge by assuming you knew fe.

Information theory would be interesting indeed if it could provide a way to solve for one of the variables in the loop when insufficient information was given to do so. That is what you seemed to be claiming you could do. And that is why I considered that you were cheating by asking the form of the output function (as you defined it), because being given that information made the challenge trivial.

I hope we can consider that disagreement resolved. From my point of view, however, there still remains a serious problem. It starts with this:

>The claim is that information about the controlled part of the
>disturbance is passed through the perceptual signal.

I have no idea what you mean by "the controlled part of the disturbance." I can't even go on describing my remaining lack of understanding until you tell me what it means. But I will tell you where I am going next.

Way back at the beginning I started asking for an illustration of how information theory would be applied either to data from a real subject or to a PCT model of the subject. I wanted to see what the calculations looked like, thinking that would help me understand the idea better. The next thing I knew, Allan Randall came back with some sort of proof based on carrying out a physically impossible (although apparently mathematically valid) operation. I started getting lessons in why information can't just be calculated, why a whole raft of assumptions have to be made before you can even say what the information in a simple signal is. Basically, I was told "here are the sources, here are the equations, figure it out for yourself."

This is probably reprehensible of me, but I don't want to figure it out for myself; I'm beginning to doubt that I can. What I want is for someone to go ahead and make the necessary assumptions, apply them to the data, and show the calculations, so I can at least begin with one actual example. I don't care what the assumptions are; just make them and say what they are and come up with some sort of result. I want to see how you compute informational measures for all parts of the control loop.

I would be perfectly happy to provide all the data from a running model including everything: disturbance, input quantity, perceptual signal, reference signal, error signal, and the forms of all the functions in the model. I just want to see how the calculations are done by somehow who knows how to do them. I don't want any advanced understanding of theorems and equivalences and principles. I just want to see some numbers, and how the numbers illustrate the things you say. And I want to see the exact assumptions that are necessary to make the calculations possible, so I can make up my own mind about whether to accept them as plausible.

Does this seem unreasonable?

Best, Bill P.

Date: Thu Jun 17, 1993 7:18 pm PST
Subject: Flaming information

[From Rick Marken (930617.2000)]

It's hard for me to edit from this computer so I'll just do this from memory.

First, I'm sorry if it seems like I'm flaming you, Martin. I don't mean to. I'm kibbitzing with you. I thought my coin flip reconstructor was kinda funny. I didn't mean to attack you -- just to make a point about what you seemed to be saying. I've been reviewing documents for the last couple of days -- it gets pretty boring and the net is my only outlet. I guess I've been a little giddy.

Next, Tom and Bill.

Bill Powers hit the nail on the head with his last post, Tom. I was playing cat and mouse with Martin -- letting him have more and more of the control system in order to predict the disturbance, thinking that I would let him get confident and then administer the coup de

grace with a varying, non-linear feedback function and show that, even if you know EVERYTHING about the control system, you still can't reconstruct the disturbance from the perception. I thought it would be obvious that the control system could not conceivably have any knowledge of the feedback function, especially if it was nonlinear AND varying. Obviously, my plan backfired because now Martin is convinced that he can reconstruct the disturbance -- since he can do it if he is given knowledge of everything about the control system except the disturbance itself.

So now there's no more games; if there is information in perception that "gets through" to the control system then somebody ought to be able to find it and reconstruct the disturbance from it. If Martin cannot reconstruct the disturbance from p (as he admits that he cannot) then it's over; there is nothing (call it information, god or whatever) in the input quantity or the perceptual signal that can tell a control system how to produce outputs that keep the perceptual signal at the reference level. Information theory may have something useful to contribute to PCT -- but it is not the idea that there is information about the disturbance in the controlled quantity or the perceptual signal that represents that quantity.

I'll resume the discussion of information in perception when anyone sends me back the disturbing variable values associated with perceptual values that I will post upon request -- no output function, no feedback function. Just perceptual and reference values -- the only two conceivable sources of information coming into a control system. Any takers?

Best Rick

Date: Fri Jun 18, 1993 3:50 am PST
Subject: Blaming who?

From Greg Williams (930618) Rick Marken (930617.1500)

>My wife is currently as fired up against Workman's Comp stress
>claims as I am against information in perception -- she thinks the
>former should not exist, I know that the latter doesn't exist. It
>turns out that these obsessions are related. Workman's Comp
>stress claims are based on the idea that what happens to people
>("stressful information") causes what they do ("have stress"). PCT
>shows that people control what happens to thjemselves. So it is
>as incorrect to blame stress on "stressful information" as it is to
>"blame" the output of a control system on perceptual information.

So how would you assign blame, based on PCT ideas?

>A Workman's Comp stress suit is based on the same mistaken
>concept of behavior as is the notion that there is information
>in perception. The apparently academic mistake of locating the
>information for action in perception becomes, with a few word
>changes, the basis of a destructive myth that is stifling the quality
>of life in my state. The criteria for stress claims would change
>significantly if people understood how control works -- it's control
>OF perception, NOT control BY perception.

How, specifically, do you think the criteria would change?

It seems to me that "stress" claims have more to do with (putative) conflicted or otherwise dysfunctional control systems and (various peoples') stories about the reasons those systems became (putatively) dysfunctional, rather than (putative) discomfort in fully functional control systems. Right off, it appears strictly from orthodox PCT theory that threats based on overwhelming physical force perceived by a person could "trigger" ongoing dysfunction, in the sense that if the threat were not present, the dysfunction would not exist. But I also think (apparently in opposition to some other PCTers) that more subtle triggers are possible.

Rick, if I convinced you that someone near and dear to you had died, and you got terribly worked up and had a serious heart attack, when I told you later that I had just been joking, would you blame yourself TOTALLY for the heart attack ("Gee, Greg, I was silly to jump to unwarranted conclusions -- I should have checked out the facts.)? Or do you think that you might possibly confer with a lawyer about the situation? (Especially if you didn't have insurance to cover your medical bills.)

It would be easy to give examples of subtle triggers of dysfunctional control in the "world of work" (that's what my high school guidance counselor called it). For instance, a boss could determine assembly-line rates which are achievable without worker complaint, then boost the rates to higher levels (as high as the boss can take, with regard to the frequency of worker complaints).

I am not trying to defend the lineal cause-effect thinking of some lawyers. I am trying to express my view that PCT is fully compatible with the notion that blame for the (very real) "pain" of a dysfunctional control system deserves to be shared, at least in some cases, among the individual with the dysfunctional control system and others without whom there would be no dysfunction. I believe that, in at least some cases, it is incompatible with PCT to blame EITHER solely the "victim" OR solely the "perpetrator(s)."

Your wife couldn't be a Republican, could she?

>It matters.

It certainly does. As ever, Greg

Date: Fri Jun 18, 1993 7:43 am PST
Subject: REINVENTING GOV'T - RKC

[From Bob Clark (930618.1115 am EDT)]

Those of you "Netters" who, like Dag, are interested in working with organizations (such as businesses, governments, schools) may find the following book interesting:

REINVENTING GOVERNMENT

How the Entrepreneurial Spirit is Transforming the Public Sector from Schoolhouse to Statehouse, City Hall to the Pentagon David Osborne, and Ted Gaebler Addison-Wesley Publishing Company, Reading, MA

This book offers some rather unusual ways of viewing Social Control Systems. Without using formal "levels of control," it suggests the importance of including and organizing a larger range of individual abilities and initiatives than the usual, strictly limited, authoritarian structure.

I am finding it not only interesting from a theoretical standpoint, but also because it offers some "every day language and concepts" that can help me in my attempts to improve the operations of the Forest Park Government.

I strongly recommend your looking it over.

Regards, Bob Clark

Date: Fri Jun 18, 1993 8:37 am PST
Subject: Facts, Stress, Info, Emergence

[From Rick Marken (930618.0900)] Mary Powers (930617.1530) --

>I suspect that this [menstrual synchronization] is one of
>the gazillion physiological and psychological *facts* about
> people in which some statistical probability has been elevated
>to rock-hard 100% certainty.

I guess we men were all a bit too fast to accept the menstrual phenomenon as "fact". But this does remind me of one topic that should definitely be included in the BBS article (book?). The topic would be something like:

How does PCT explain X?

Where X would be all the purported phenomena that behavioral scientists think are phenomena (usually based on statistical tests) but are not phenomena from a PCT perspective. It looks like menstrual synchronization could be one of the prime examples (maybe Jeff Hunter could post the relevant published reports on this phenomenon). As Mary says, there are probably a gazillion "facts" of this kind to choose from: the social psych and sociology literature would be particularly fruitful, I imagine: obedience, bystander intervention,

etc. All purported facts that are not facts. Not all the non-facts would be statistical. The idea that reinforcement increases the probability of behavior is another non-fact because it is only true in special circumstances (when the animal is starved and it's connection to food input is weak [a demanding schedule]). The "fact" is that animals control food input; the apparent "effect" of that input (reinforcement) on behavior depends on parameters of the feedback loop.

Greg Williams (930618)--

>So how would you assign blame, based on PCT ideas?

I would hope that people could cooperate to achieve their goals in a business. I know that business people often violate this cooperative relationship, setting up contingencies that can hurt people. I know that business people can do bad things -- and I blame them for it. But I also know that employees take advantage of business people by claiming to have been caused stress when, in fact, the stress, if they have it, is not caused. The assumption that the environment causes stress is a problem, not merely because it is false, but because it allows workers to perpetuate a fraud; all they have to do is point to a "stress inducing" stimulus to which they were exposed and the court is required to assume that the result was stress. This is what I object to; using a fake psychological model to make fraud easy.

I am not against workman's comp and there certainly are "stimuli" out there that can cause bad results; if a heavy piece of equipment drops on your foot then it will cause damage. I think a worker should unquestionably be compensated for the effects of this kind of cause. But exposure to "stress information" should not be a basis for assuming that the behavioral result was stress -- because there is no such thing as stressful information.

Your example of lying to me and causing a heart attack is cute; I would be annoyed but I probably would not think of suing you; I'm just not a suing kinda guy. But that is a situation where I might give the "victim" of your "stressful stimulus" a sympathetic hearing -- at least because you did apply that "stimulus" intentionally. I'm sure Linda would have no problem accepting a person's stress claim if it were the result of maliciousness on the part of someone else; I still think the stress (if it exists) was not caused by the "stressful stimuli" applied by the malicious person but that person did, intentionally, create the circumstances where reorganization (and the attendant stress) was possibly required.

The problem with workman's comp is that you can make a stress claim just by showing that you were exposed to stressful information; it is not necessary to show that this information was presented to you maliciously. As I said, if someone left the door to your files locked by accident and you can make a case that a locked file is a stressful event for you, then you're in the money. That's what is no good about the current Workman's Comp law -- it encourages fraud and makes it VERY EASY to carry it out. I am, however, all for protecting workers and think that they should be generously compensated for REAL on the job misfortunes.

>Your wife couldn't be a Republican, could she?

We're both unclassifiable. She, like me, would probably call herself a democrat, but with democrats like Sam Nunn running around, and republicans like Goldwater (who wrote a wonderful, very pro gays in the military editorial recently) I don't know that either of us knows what to be. We both encourage individual freedom tempered by communitarian cooperation. I suppose I'm a libertarian socialist; and Linda's a tad to the right (left?) of that.

Martin Taylor (930617 20:30)--

>It's the unreconstructible part of the disturbance that leads to the need
>for more layers of ECSs in the hierarchy.

This is news. I thought the layers were there to control perceptual aspects of the environment that are not represented in lower level perceptual signals.

>If you can control in one layer, why go for more?

See above.

>The question at issue (for me) is how much is there that needs to be

>dealt with by higher layers.

They are dealing with completely different kinds of perceptions -- not the residual variance of controlled variables at a lower level. At least, in the HPCT model, that's how it works.

>what is controlled (compensated for) can be reconstructed using the
>perceptual signal, the reference, and the output function.

And the feedback function.

Of course you can reconstruct the disturbance if you know all this; then you know everything; All you are saying is that, in an control system, $d = 1/g(o)$ which we are all well aware of. Where does information theory make its contribution? I thought that you were saying that there is something about the information in p and r that would let you compute the outputs that compensate for the disturbance. If you are not saying that, then we have no argument. All we have is a disagreement about the value of info theory for understanding control; you think it has some, I think that it has none.

>Since you know "exactly" how a control system works, I presume that tells
>you "exactly" what can be reconstructed and what cannot.

Yes.

Hans Blom (930617a) --

>You are of course right: the individual control
>laws together 'control' flockness, but in an emergent way

I think we are in violent AGREEMENT. The confusion (probably mine) comes from my understanding of "emergent" as "irrelevant side effect" or "uncontrolled result". Your "flockness" is a variable controlled by the individual members and you agree that this is the case. You go on to call this an emergent variable (meaning, to me uncontrolled) -- I would prefer to call it a controlled result of collective action.

But it is clear from you post that you do not imagine that there is some godlike control system (the "social control system") looking over the flock and keeping flockness under control. That's all that Tom and Bill and I are objecting too -- the idea of an unseen social control system that is "outside" the individuals in the collective. There are variables (that you call emergent) that are controlled by the collective and can only be controlled by the collective of individuals. Perhaps people thought we were denying this. But this is precisely what Tom Bourbon's experiments show. Our point has only been that the only control systems involved in producing these controlled variables are the individual ones in the collective; the collective ITSELF does NOT become a control system; it seems like it does because the collective controlled variable depends on the actions of all the individuals -- but the collective of individuals itself is not a control system because the reference signals for the variables that must controlled in order to controlled the collective controlled variable exist only in the heads of the individual members of the collective. Any change in the individual goals of the members ends control of the collective controlled variable.

Best Rick

Date: Fri Jun 18, 1993 9:21 am PST
Subject: Re: emergent control

From Tom Bourbon (930618.0958) [Hans Blom, 930617a]

In you replies to, Bill Powers (930617.0730 MDT), and Rick Marken (930617.0730 MDT), you seemed to adopt an engineer's definition of control, in which the perceptions of the observer, not of the observed system, are of primary importance. From the perspective of engineers, or users, of a system, variables for which *they* have references are of primary importance, even though those variables might be of only incidental significance to the observed system, or the observed system might be completely unaware that those variables exist, affecting them completely incidentally.

My guess that you speak from the position of a designer, engineer, or user of a system was confirmed when I came to your comments on my post about synchronization of menstrual cycles.

>(Bourbon (930617.1045))

>

>>You seem to have ruled out the possibility that synchronization
>>of the cycles [menstruation] is an uncontrolled side effect ...

>

>This whole discussion indicates to me that there is no common
>terminology for what control is. ...

No terminology common to whom? PCT modelers seem to agree on a definition.

>... To me, a control system is a

>mechanism that keeps something within specified (but maybe vary-
>ing) limits or almost so, whereas without that mechanism those
>limits would not be observed. Therefore, according to my termino-
>logy, 'flockness' is controlled, although YOU might consider it an
>'uncontrolled' side effect.

Fine. That being the case, could you describe the mechanism the "keeps flockness within specified (but maybe varying) limits or almost so?" ("Or almost so?" How close is that?) In the examples presented by Bill, Rick and me, there were specific mechanisms, each controlling particular variables and affecting other variables only incidentally. The incidentally affected variables gave rise to effects that an observer might call "emergent" phenomena. And an oberver might call some of those effects and phenomena "social patterns or organizations."

The conclusion of your comments on my post made your perspective crystal clear:

>...My home heating system, in your terminology, 'controls' the
>temperature at its sensor, while the thing that I want, a nice room
>temperature, is only an 'uncontrolled side effect'. That is not the kind
>of language that I use, although I can understand it when YOU do. But then,
>understanding (here I mean: making sense of what someone says) often
>requires this kind of back and forth 'translation', isn't it?

Your home heating system in fact, not "in your (our) terminology," controls the temperature at its sensor. It is obvious that you do not use the kind of language "we" do, because you are willing to conflate the reference signals of the model or the artificial system with your own intentions. *You* use the heating system to control your perception of a nice room temperature; the heating system does not. I *think* I understand your position pretty well, or am I wrong?

The individuals in GATHER, like the thermostat in your home heating system, lack awareness of, reference signals for, and perceptions of the emergent structures and phenomena that are of interest to you. They do not control those phenomena and structures, although *you* can use them, and their unintended byproducts, to control *your* perceptions.

Until later, Tom Bourbon

Date: Fri Jun 18, 1993 9:22 am PST
Subject: Controlled and uncontrolled outcomes

[From Bill Powers (930618.1000 MDT)] Hans Blom (930617a) --

>My control laws for each individual were:

- >1) keep a minimum distance to others;
- >2) keep a maximum distance from others;
- >3) otherwise move at discretion (no specific distance!).

This can be at least partly implemented with the CROWD program. The "seek another person" control system specifies a distance that is to be maintained from the sought person, neither more nor less (a reference distance). The proximity error gain can be set low so that it seems that a upper and lower limit of desired distance exists. With the destination-seeking control systems all turned off, there will be no tendency to go anywhere in particular. Movement will be maintained by the conflict between collision

avoidance and positive seeking of another person's position. I have tried several setups like this, and they result in a sort of milling around in a general area; the whole group sticks more or less together but may drift randomly in one direction, then another.

I haven't tried this, but by moving other groups through the main one, we can introduce disturbances. I would expect the original group to re-form, because of the seeking of other persons' general vicinities, although I wouldn't expect the centroid to move back where it was. If you have a copy of the CROWD program you might like to experiment with this.

>These describe the behavior of a real flock much better, in my
>opinion. It may be that you do not like the random element, nor
>the non-linear control law, but let's stick to my example for
>the moment, shall we?

The control law in the crowd program is pretty nonlinear already: the controlled variable for collision avoidance, "proximity," is calculated as the sum of the inverse squares of all distances to other persons in the field (divided into left and right proximity relative to the direction of travel), and proximity is also the controlled variable (with a positive but nonmaximum reference setting) for seeking the position of another person. So the interactions actually involve inverse fourth powers.

>What do you mean by 'natural tendency'? The same thing as Aristotle when he
>said that a rock has a 'natural tendency' to strive for the center of the earth?

How about giving me the benefit of the doubt, Hans? The natural tendency of a rock is to maintain its state of motion unless accelerated by an applied force, and then to resist acceleration with a force proportional to its mass. The natural tendency of a board nailed to a wharf is to resist being lifted, not because it's a control system but because it's nailed down.

>Please be explicit in what discriminates a 'natural tendency'
>from control (should you do so, I am tempted to say, I will
>give you a counter->example :-).

Explicitly, it's power gain. A control system has an output function that produces vastly more power output than its inputs provide (drawing, ultimately, on external energy supplies). The physical environment (excluding other control systems) generally imposes a power loss between input and output: work must be done on it to make anything happen. A controlled variable has work done on it by a control system. If the variable is uncontrolled, any potential energy it contains normally simply dissipates and it comes to some equilibrium condition of minimum energy (that's its "natural tendency"). In PCT we speak of "loop gain" which is a measure of power amplification around the loop. The controlling part of the loop is that in which the greatest power amplification occurs: the organism, not the environment.

> When I look at my simulated flock, I may PERCEIVE a 'natural tendency'
>to form a flock, but that is just my subjective interpretation.

It's just a subjective interpretation unless you have analyzed the elements of the flock and can show that the outcome of the interactions among individuals must be a tendency for the whole to remain more or less together. Then it becomes a deduction.

>To understand (or model) what goes on, I need a
>law or a formula (and I don't care much whether you call that
>law or formula a 'natural tendency' or a 'control law').

Well, I care. Using words like "control" loosely results in their not meaning much of anything: have a nice day. The word "control", outside its technical meaning which I ALWAYS intend, is just a slipshod term vaguely indicating the effect of something on something else. We already have plenty of vague terms borrowed from informal usage. Why not refine them for scientific use when we can?

You wouldn't see any semblance of control in your flock if you had ever seen and interacted with a real working control system -- knowingly. You don't arbitrarily perturb a real control system without spitting on your hands first and bracing your feet. When I say that control systems resist disturbance, I mean that they resist it energetically and with as much effort as needed (up to their limits of output). They don't passively wait for controlled variables to be shoved far from their reference levels and then daintily make

polite efforts to encourage them to go back where they belong. They react instantly and strongly and never let the error get large.

If a living control system gives way before a disturbance, it's seldom because it couldn't resist it; it's usually because a higher system is altering the reference signal to prevent disturbance of some higher variable. If you see living control systems drifting around and letting external forces have their way, it's only because the effects of the external forces don't matter much to the control systems (or because they're beyond the capacity of the control system to resist them). The variables that are under control are under CONTROL. There is no doubt about a controlled variable when you find one. You can't simply take hold of a real controlled variable and move it around. If you try you'll have to struggle to get your way. Don't think of a flock of birds. Think of a wild mountain lion. Try picking up the mountain lion and putting it somewhere else to see if it rejoins the group.

>Let's try to stick to objective descriptions and try to avoid 'interpretations'.

I'm the one using an objective description of a control process; you're the one using "control" as a metaphor.

>An individual does not perceive 'flockness', yet 'flockness' is preserved even if individuals are hidden from sight.

"Flockness" is an uncontrolled outcome of an interaction. Our whole problem here is that you don't see the difference between a lawful uncontrolled outcome and an actively controlled outcome. You're refusing to go outside the boundary of phenomena that you are trying to characterize as lawful, to see that there is another kind of qualitatively different phenomenon.

For control in its technical sense to exist, there must be a system that senses the state of the outcome, compares it with a desired state, and turns the error into an output that has a very high-powered effect on the outcome. If flockness were under control by a competent control system, you would be unable to disturb it without using extreme measures.

>Isn't it remarkable how robust flockness is in the face of disturbance?

I don't think you know what "robust" is. If the marble always gets back to the bottom of the bowl after being displaced, you would call its resistance to disturbance robust. If a true control system were keeping the marble at the bottom of the bowl, you wouldn't be able to displace it appreciably in the first place. THAT is "robust."

>Remove WHICH control system? There is no control system for 'flockness'; that hypothetical one cannot be removed.

Correct. If you can't even find a potential control system, the Test is failed before you begin.

>Remove the sensors that measure distances to others, and the flock will disintegrate: a random walk moves each individual away without limit.

The Test would reveal that each control system was sensing and controlling proximity to others (rather weakly). It would not find any system controlling for the outcome of flockness. So the conclusion would be that the individuals are controlling for proximity, but there is no control system for flockness. Flockness is a natural and lawful outcome of the interactions between these individual control systems, but it is not a controlled outcome.

>Now, does 'flockness' pass The Test for the Controlled Variable?

No.

If it were not for the fact that I remain upright while I walk, the air along my path six feet above the ground would not be displaced. The movement of the air follows lawfully from the effects of my remaining upright while I walk, but it is not a controlled variable. It is an uncontrolled, although lawful, outcome of my control processes. A side-effect, as far as I am concerned.

Joel Judd (930617)

Ah, Joel, Herndon understands, doesn't he? And so do you. A society based on coercion gets to the point where it considers coercion a virtue.

Best to all, Bill P.

Date: Fri Jun 18, 1993 11:25 am PST
Subject: Facts, Control

[From Dag Forssell (930618.1200)] Facts: Rick Marken (930618.0900)

> ...But this does remind me of one topic that should definitely
>be included in the BBS article (book?). The topic would be
>something like:
>
>How does PCT explain X?
>
>Where X would be all the purported phenomena that behavioral
>scientists think are phenomena (usually based on statistical
>tests) but are not phenomena from a PCT perspective.... ...As
>Mary says, there are probably a gazillion "facts" of this kind
>to choose from: the social psych and sociology literature would be
>particularly fruitful, I imagine: obedience, bystander
>intervention, etc. All purported facts that are not facts. Not all
>the non-facts would be statistical. The idea that reinforcement
>increases the probability of behavior is another non-fact because
>it is only true in special circumstances (when the animal is
>starved and it's connection to food input is weak [a demanding
>schedule]). The "fact" is that animals control food input; the
>apparent "effect" of that input (reinforcement) on behavior
>depends on parameters of the feedback loop.

Rick, you are seconding my motion of a few days ago:

>Dag Forssell (930615 18.10)
>
>What I began to visualize turn out to be of another kind: MYTHS
>TYPE 2: _Commonly accepted truths which are demonstrably false._
>A clear articulation of these may be much more useful to promote
>PCT. It would be much more upsetting and demanding of attention to
>show that what people think they know is false, than to say that
>something they never heard of, have gotten along fine without, and
>don't care about is right.
>.....
>How about it? I hope for some suggestions. Please post what comes
>to mind as MYTH TYPE 2:
>
> a) Psychological (cognitive, behavioral) science says:
> b) Common leadership understanding:
> c) Perceptual Control Theory says:

I am very interested in a catalog of what Mary wrote about:

>I suspect that this [menstrual synchronization] is one of
>the gazillion physiological and psychological *facts* about
>people in which some statistical probability has been elevated
>to rock-hard 100% certainty.

How can I (we) go about collecting this catalog of *facts* or myths
type 2, make sense of it and see if it appears useful to promote PCT?

Book suggestion? Listing from our minds? Psychiatric reference
works? Psychology textbook?

I am prepared to do some legwork, but PLEASE (all of you sharp
minds out there in CSGnet land) offer some suggestions!

Bill Powers (930618.1000 MDT)

>....The variables that are under control are under CONTROL. There
>is no doubt about a controlled variable when you find one. You
>can't simply take hold of a real controlled variable and move it
>around. If you try you'll have to struggle to get your way.

I once drove Bill to the airport on a Los Angeles freeway. We were in the fast lane with cars all around. Bill asked: "Would you like to see a good, tight control system?" "Of course," I said. Bill reached over from the passenger seat, grabbed the steering wheel and pulled on it, rather hard. I tensed immediately. The car never veered from my intended position in our lane more than a few inches. Bill let go. Again, no deviation of the car's position.

I understood the demonstration of a good, tight control system.

Best to all, Dag

Date: Fri Jun 18, 1993 12:07 pm PST
Subject: Myths

[From Rick Marken (930618.1300)] Dag Forssell (930618.1200) --

>Rick, you are seconding my motion of a few days ago:

I realized it right after I pressed the "Send" button.

>How can I (we) go about collecting this catalog of *facts* or myths
>type 2, make sense of it and see if it appears useful to promote PCT?

I think a current Intro Psychology text would be a gold mine. I have only one old one on the shelf but I'll check it out. Pop psych texts should be good too. I'll take a look through the "10 Habits of effective people" book.

Most health fad books will have tons of stuff -- just about everything they say is a myth. My dad regularly provides me with news of the latest, unpronounceable chemical variable that we are supposed to control (open loop, of course) in order to live to be 250 (my dad, a retired dentist, spends his days varying his eating patterns in order to control variables that he cannot perceive; I sent him a copy of my book but it doesn't seem to have become particularly dog-eared; once an open-looper, always an open looper).

>I once drove Bill to the airport on a Los Angeles freeway.

I'm glad I kept him on surface streets!

Best Rick

Date: Fri Jun 18, 1993 2:30 pm PST
From: Richard Robertson
MBX: urrobert@uxa.ecn.bgu.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: your book suggestion

[From Dick Robertson] (930618.1715)

Hi Dag.

I noticed your suggestion about a book of unfacts that are current in contemp psychology & a comment of Rick's on the subject. I think that there might be enough spinoff from the contributions to Bill's joint-paper to begin something like that. I have been thinking it's about time to put out a new edition of the textbook. So much new knowledge has been developed by the CSG netters (in my opinion) and maybe a book that gave BOTH the current unfacts and what facts ARE known, (and what promising areas are being opened, like Tom's Chuck's and your stuff) would be on the horizon. I have just retired from teaching and haven't yet caught up enough to be able to contribute much to the discussions on the net, but I do keep fairly current. My next objective is to go through all the stuff that Bill and others have already put out on modelling, to try to learn it and then get back to

trying to apply it to higher level behavior that I have been interested in. My clinical work already shows a lot of influence of HPCT but I haven't been able yet to systematize it in a way that I could present to others. But posts like yours inspire me to keep centering on the stuff we already have that ought to be presented to the general public. Keep it coming. Hope to see you in Durango. Hi to Christine

Best, Dick R

Date: Fri Jun 18, 1993 2:30 pm PST
Subject: Re: H(S)

[Allan Randall (930618.1400 EDT)]

Bill Powers (930612.0930 MDT)

>When we say "disturbance" in equations, we mean the STATE OF THE
>DISTURBING VARIABLE:
>...We do NOT mean the amount by which the knot (X) is displaced from
>the target position.

In spite of all the other misunderstandings, I don't think this is a problem. I do not mean the amount by which the knot is moved. I mean the force(s) acting on the knot from the external environment, whether or not the knot actually moves.

Bill Powers (930613.2200 MDT)

> >Bill, I have a real problem with this whole log(D/r) thing...
> I got it from Martin Taylor... I'm just following orders.

Well, all I can say is that this is not the standard definition, which is -log(probability), as I've described many times in the past. I suspect you are using log(D/r) differently than Martin, but I will let him speak for himself. The way you are using it, it is not equivalent to information.

> >If two signals have wildly different scalings of amplitude, but
> >are otherwise identical, then I can write a very short program
> >to convert one to the other...
>
> Suppose you have the relationship $X(t) = Y(t)/10$ No matter
> how you write the program, when you compute X from Y you will get
> a waveform with less relative resolution than there was in Y.
> If you multiply X by 10, you will not get Y back; ...

Not necessarily. It depends on the representation used for the numbers. Eg:

```
00.51928374 / 10.00000000 --> 00.05192837 * 10.00000000 -->
00.51928370  LOSS
```

But:

```
.51928374E+00 / .10000000E+02 --> .51928374E-01 * .10000000E+02 -->
.51928374E+00  NO LOSS
```

So if I scale a sequence of these numbers, I do not have to also scale the resolution at which they were measured nor the resolution at which they are represented in the computer in order to retain information.

I don't want to make a federal case out of this - its not a really important point. I was only trying to illustrate that your statement about log(D/r), while true about THAT measure, is not IN GENERAL true about information. D/r is just the number of possibilities inherent in the numbering system or the measuring apparatus. If you considered this to be the inverse of the probability, then I would have no problem. But this is not how you've been using it. You are tying the whole thing to the way numbers are represented - to the precision at which measurements are made. But many other schemes are also possible - this is why the ideas of computer programs/languages and probabilities/distributions are much better ways to think about it.

Another way of viewing it: D/r *can* be used to compute information on a sequence of numbers *after* they have been compressed to their shortest possible representation in

whatever language we're using. Only then does it make sense to use the precision of the numbers themselves to compute the information. Then your comments would be more or less accurate. But this is a BIG difference. This compressibility has not been addressed in the measure $\log(D/r)$ as you are using it.

```
> >...tell me how to compute log(D/r) for the following sequence S:
>
> > .0 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .5 .2 .7 .7
> >.7 .7 .7 .7 .7
>
> Before I could do that, you would have to tell me the size of r.
> There isn't any inherent "resolution" in an arbitrary sequence of
> numbers like the one above.
```

Sorry - I had thought it implied that the resolution was .1. If r had been .001, I would have written .000 .500 .500, etc.

```
> ...What they mean depends on the
> physical situation they are taken to represent.
```

Sure - and that depends on our *model* of the physical situation. When you say: "...depends on the physical situation they are taken to represent," it sounds like the same thing I mean when I say "...depends on the chosen language."

```
> If this is a
> series of measurements of some physical variable, then we have to
> talk about the measuring device's resolution.
```

If that is the extent of your model of the physical situation, then fine. As I said, if your model or language is restricted to the number system or measuring apparatus, then $\log(D/r)$ will work. However, the "physical situation" the numbers are "taken to represent" is generally more involved than just the measuring resolution. In a human, it is likely to be something like a hierarchy of control systems.

```
> My problem is that I can't find any link between the
> manipulations you talk about and any PHENOMENON.
```

You misunderstand. The arbitrary hypothetical examples are my attempt to explain a mathematical system called information theory. This is not meant to be connected to any particular physical phenomenon. It has nothing to do with control theory or thermal systems or any other physical system. It is a mathematical notion. Now I have ALSO talked about applying this theory to control systems, but in that case, I was not making up arbitrary hypothetical examples - I was using the control hierarchy as the language.

```
> ... the only answer I've received so far is
> "Well, that depends on what you assume."
```

This is not a fuzzy, ill-defined answer. *Given* a language or model to work with, information *is* well-defined. But, you said yourself that what the numbers mean depends on the physical situation we take them to represent. And this *does* depend on what is assumed.

```
> ... How do you
> decide what is the most plausible way to set it up for a control
> system model of a specific example of behavior? So far I'm
> drawing a complete blank on that. And, apparently, so are you.
```

No - I've said many many times that the control hierarchy is the language. The general notion of information *is* completely open-ended as you describe, but as applied to control theory, it is much better defined.

Allan Randall

Date: Fri Jun 18, 1993 3:02 pm PST
Subject: Re: Information Theory (again)

[Allan Randall (930618.1800 EDT)] Rick Marken (930611.2200)

> You can use any kind of control system you like; or any method
> you like whatsoever -- it doesn't even have to be a control system.

No, I need to know what control system to do the reconstruction for, not what to use to do the reconstruction. Correct me if I'm misremembering: You have stated that the simple linear example from Powers' Primer was insufficient due to lack of nonlinearities. Therefore, if I am going to do a demo, I absolutely need to know what kind of control system *will* satisfy you. This seems like a simple requirement.

> I will give you a sequence of numbers, P, that was an actual time
> sequence of perceptual inputs to a control system. ...
> I think I actually gave this data to you (and I gave Gary Cziko the
> actual values of the sequence of D values) but I never got the
> reconstructed values of D.

I thought I made it clear, but I'll repeat it once more: this is an EMPTY challenge. It has nothing whatsoever to do with anything I have claimed about information in the percept or reconstructibility. You are asking for a completely 100% blind reconstruction. This would require 100% of the information about the disturbance in the percept. I never claimed this. You have said in the past that you understood this - and that you were not looking for a complete 100% reconstruction. And yet, here you are again acting like this challenge has some kind of meaning.

> >When did you show this? You showed that it was not possible to perform
> >the reconstruction?
>
> Bill Powers and I both posted data showing that there can be a very
> weak relationship between output and disturbance when there is a
> varying environmental non-linearity between the output and the
> controlled variable.

This is nowhere close to showing that the reconstruction is not possible. You posted correlation figures. As far as I can remember that was about it.

Okay, ONE MORE TIME.

> Without P there is no way at all to reconstruct D except by chance,
> right? What is this fewer additional bits thing?

This "fewer additional bits thing" is the whole notion of information I have been talking about for some time. I am *not* offering to reconstruct the disturbance blind with no additional information. However, if I can use the percept to do the reconstruction using *fewer* additional bits than would be required without the percept, then I have shown that the percept contains information about the disturbance. This is the only reconstruction demo I ever offered to do. When you posted your blind-man challenge, I dismissed it and asked for an example of a sufficiently nonlinear control system. I have not received this from you, so I have not yet done the demo.

To clarify this once more, I am reposting the rules as I see them. I have justified these in the past, so I'm not going to go on at length about them. From a previous post:

The following is what I am assuming as given in any attempt by me to reconstruct D from the percept p:

- the reference r
- the output function O()
- a programming language, such as C

We will all just have to agree that the C programming language has no inherent information about d - that seems like a reasonable assumption to me.

The following are *not* assumed (to do so would be "cheating"):

- the output o
- the environmental feedback function F()
- the disturbance d

> ... I'll give you a string
> of numbers which I swear is P. I'll also give you r, O() and C (I really

> forget why I agreed to giving C if C is values of the controlled variable;

As you can see from above, C was just the C programming language.

> You give me back a string of numbers that is D. You should be able
> to do this because, you claim, there is something (that you call
> "information") in P that let's the control system know what D is.

The idea was to do this in FEWER bits than without P. That is crucial.

Allan Randall

Date: Fri Jun 18, 1993 5:28 pm PST
Subject: Where are the sociologists?

[From Rick Marken (930618.1500)]

After reading the posts from Bill Powers (930618.1000 MDT) and Tom Bourbon (930618.0958) on social control, I have been reconsidering my comments on this topic. I had been assuming that there is such a thing as a "collective controlled variable" which is a variable that is not controlled by any individual but is a result of control by the collective. My example of a collective controlled variable was the three line pattern made cooperatively by the subject's in Tom's cooperation experiment. After thinking about this for a bit I realize that I am wrong -- there is no such thing as a collective controlled variable. A controlled variable ONLY exists if there is a perceptual signal in some system that represents that variable, a reference signal that specifies the reference level of that signal and an output system that can affect the perceptual signal. In Tom's task, the three line pattern is a controlled variable ONLY IF one or the other or both subjects are perceiving and controlling that overall pattern. Any disturbance to the pattern would be resisted (if possible) by the person (or persons) controlling it. The three line pattern would APPEAR to be controlled, however, even if each person were controlling just a piece of it -- the relationship between left and middle line for one subject, the relationship between middle and right line for the other. It might seem like the three line pattern is controlled in this case but it is NOT -- which could be proved if the subjects were allowed to see and control only their two line patterns with the three line result projected on another screen. Disturbances to the three line pattern that did not affect each subject's two line controlled variable would NOT be resisted.

So Hans' and my idea that "flockness" is under control (a "collective controlled variable" by my terminology -- ie. a variable that is controlled by the collective but not by any individual) is incorrect unless there is a control system present that is perceiving and acting to influence this variable.

Where are the sociologists during this discussion of social control, by the way? This is where PCT meets sociology. One of the goals of sociological research, it seems to me, would be to provide formal means for discriminating emergent (uncontrolled) social variables (like the flocking of birds) from controlled social variables (like the three line arrangement in Tom's cooperation experiment--assuming that that IS a controlled social variable). The difference, I think, is that there is no perceptual representation of an emergent social variable in any control system that has an influence on that variable; there is a perceptual representation of a controlled social variable in at least one of the control systems that influences that variable.

Best Rick

Date: Fri Jun 18, 1993 7:00 pm PST
Subject: By the rivers of PCT

[From Rick Marken (930618.1900)] Allan Randall (930618.1800 EDT) --

I said:

> You can use any kind of control system you like; or any method
> you like whatsoever -- it doesn't even have to be a control system.

Allan says:

>No, I need to know what control system to do the reconstruction for,
>not what to use to do the reconstruction.

I gave you, along with the disturbance, EVERY parameter of the control system except for the function relating output to perception (the feedback function). Surely you don't consider the feedback function to be part of the control system, do you? In a tracking task that is the function that relates my muscle tensions (the last output of my nervous system) to the cursor position. Surely you don't consider my desk, the rollers in my mouse, the electronic characteristics of the computer, etc to be part of me as a control system, do you?

> Correct me if I'm
>misremembering: You have stated that the simple linear example from
>Powers' Primer was insufficient due to lack of nonlinearities.


You were able to reconstruct the disturbance when you knew p, r and O(); ie. when I gave you every component of the control system. Then I sent you p along with r and O() but with a feedback function that was not = 1 (and I didn't tell you about the feedback function; but that's not part of the control system so why bother). I thought you would try the reconstruction and find that your computed disturbance was not equal to the actual one anymore. But you didn't do it, so you never found out. You didn't try the reconstruction, you say, because:

>this is an EMPTY challenge. It has nothing whatsoever to do with anything I have
>claimed about information in the percept or reconstructibility. You are
>asking for a completely 100% blind reconstruction. This would require 100%
>of the information about the disturbance in the percept. I never claimed this.

So fine. So there is not enough information in perception from which to reconstruct the disturbance. Great. We are then agreed, except perhaps on the details. You say that the amount of information in perception about the disturbance is >0; I say that it is precisely 0. So all this was wasted effort. Clearly, the control system cannot use the information in perception to produce effects that exactly counter the disturbance because there is NOT ENOUGH information to do it. OK. You can believe that there's a little weensy amount of information in the perceptual signal if you like. Great. What's it good for? What does the control system "use" to produce effects on the controlled variable that are exactly opposed to the disturbance?

Or are you saying that your reconstruction would be only 99% accurate? That would be GREAT. I'd accept a reconstruction as close to the disturbance as was the output of the control system. In fact, why not just reconstruct the output -- forget the disturbance. I'll again give you the perception, reference and output function. Surely there is enough information in the perception from which to reconstruct the output GIVEN THE OUTPUT FUNCTION?

I think it was you, Allan, who said that the input to the control system must have enough information so that it can tell the output of the control system what to do to counter the disturbance. That's why I gave you the reference and output function for your reconstruction -- your concept of control seemed to be:



The perception comes into the system, is subtracted from a reference and the information (i) is extracted. This information is used to tell the "plant" (the output function, O()) what to do. The plant produces an output (o) that will counter the effect of d, resulting in a p closer to r. I believe you came to the conclusion that to believe that control did not work something like the way I diagrammed it above was to believe in magic (or mysticism). Martin said that to believe otherwise would be a "renunciation of logic".

You wanna try reconstructing the output from the perception, reference and output function. Is there at least enough information to do that???

The fact is that you guys just have a wrong conception of how control works -- its a circle, not a chain. Your belief in this information stuff is keeping you from making the small but profound mental change that is needed to be able to visualize how a control system actually works. I'm telling you, you are missing a truly profound experience by

holding on to the life raft of information. Let go, and swim through the refreshing waters of control of perception (if you want to, of course).

Best Rick

Date: Fri Jun 18, 1993 7:55 pm PST
Subject: Re: Randall on information

[From Bill Powers (930618.1930 MDT)] Allen Randall (930518.1400 EDT)]

I said:

>When we say "disturbance" in equations, we mean the STATE OF
>THE DISTURBING VARIABLE:..We do NOT mean the amount by which
>the knot (X) is displaced from the target position.

>In spite of all the other misunderstandings, I don't think this
>is a problem. I do not mean the amount by which the knot is
>moved. I mean the force(s) acting on the knot from the external
>environment, whether or not the knot actually moves.

It's all right to use force as a measure of the disturbing quantity in this example, but before it can be used in the system equations it must be transformed into an equivalent effect on the controlled quantity, the knot position. The effect of a position-independent disturbing force depends on the nature of the connection between the system's output and the controlled variable or CEV. In the rubber-band experiment, the effect on position of a force applied to the knot depends on the spring constant of the other rubber band. If that spring is linear, and assuming that the subject's hand position is under lower-level control so it is unaffected by changes in the pull on the rubber bands, the effect of the disturbing quantity q_d as a force is just q_d/s , where s is the spring constant in units of stretch per unit force. The disturbance function f_d is $1/s$.

When you use force as the disturbing quantity, you're implying that it's unaffected by the position of the knot. If you use the position of the disturbing end of the rubber band instead, as we do for the output quantity at the other end, then the disturbing force does depend on knot position, and the equations become somewhat more complex.

You can't solve the system equations until you've converted the output quantity and the input quantity into terms of effects on the controlled quantity in the same units in which it is sensed. The controlled quantity is a sensed position, not a force. To deduce the applied disturbing force from the perceptual signal, other system variables and functions must be known, particularly (in this case) the nature of the feedback function f_e . The perceptual signal alone gives no clue about the conversion factor $1/s$. In dynamic situations it is also necessary to know the mass of the "knot" which is again not represented in the perceptual signal. The variations in knot position that are represented in the perceptual signal are not directly equivalent to variations in the applied force.

>Bill, I have a real problem with this whole $\log(D/r)$ thing...

>> I got it from Martin Taylor... I'm just following orders.

>Well, all I can say is that this is not the standard
>definition, which is $-\log(\text{probability})$, as I've described many
>times in the past. I suspect you are using $\log(D/r)$ differently
>than Martin, but I will let him speak for himself. The way you
>are using it, it is not equivalent to information.

And then you say ...

> D/r is just the number of possibilities inherent in the
>numbering system or the measuring apparatus. If you considered
>this to be the inverse of the probability, then I would have no problem.

Well, then, we have no problem about that, because I did understand that the probability is simply $-\log(r/D)$ or $\log(D/r)$, at least in simple cases where the divisions of the range into units of r is uniform. The only difference between us is a logarithm.

I'm surprised at your response to my argument about the effect of using finite intervals on scaling. If you have a variable with a range D of 8 units and a minimum increment r of 1

unit, all values of the variable are expressed as whole numbers between 0 and 7. If you then scale this variable down by a factor of 3, you get a new variable with a range between 0 and 2, still with a 1- unit minimum increment. Obviously, if the original variable contained 3 bits, the scaled-down version contains only 1 bit (0 to 1, or 1 to 2). Scaling back up by a factor of 3, you get not the original variable, but a variable with a range from 0 to 5 -- 2 bits and change.

Your conclusion below is simply wrong:

>So if I scale a sequence of these numbers, I do not have to
>also scale the resolution at which they were measured nor the
>resolution at which they are represented in the computer in
>order to retain information.

In the computer, all numbers are represented in binary form. The mantissa has a finite number of bits, which determines the resolution for all computations. If you compute $x/10$, you lose whatever significant digits are unable to be expressed in the conversion from decimal to binary notation and as the quotient of the division. When you scale up again by a factor of 10, in general you get a number different from the original one. You happen to have chosen numbers, in your second example, that come out even in decimal notation. They probably don't come out even in the computer. But even if they did, that would be happenstance, and nothing to base a general principle on. The numbers will come out even only if the scaling division has no remainder in the number base used -- that is, if the dividend is an integer multiple of the divisor.

You have to assume that you will ALWAYS lose resolution when you scale numbers down, given a universe of measurement with a finite resolution. On the average, you will. And if the size of the basic increment doesn't change, the scaled-down number will be divided into fewer intervals of that size than the original number will be divided. This is ridiculously easy to show:

If $I = \log(D/r)$, and we scale down to a new variable G with a range of $D/8$ and the same resolution r , the new information content is $I' = \log(D/8r) = \log(D/r) - 3 = I - 3$. The division is, of course, remainderless, so we can't regain the lost 3 bits by scaling up again. The scaled-down version loses information according to the log of the scaling factor.

This is particularly relevant when you speak of a basic resolution that is much larger than the resolution of your computer. It is not the computer's resolution that counts, then; it is the assumed resolution of measurement of the variable being represented. If r is 1% of D , then the probability of chance occurrence of any value of d between 0 and $99r$ is 1%. If you represent any particular value with more significant digits than 2, the extra precision is illusory. Scaling such a measure down by a factor of 3 and then back up by the same factor, in decimal notation, will lose you nothing if the number is 33, but will lose you 3 digits -- 10 percent -- if the number is 32. $33 / 3 \times 3 = 33$, but $32 / 3 * 3 = 30$. If you scale down by 30 to one, you will lose enormous amounts of precision; on scaling back up you will find just one of four possible numbers: 0, $30r$, $60r$, and $90r$.

>I don't want to make a federal case out of this - its not a
>really important point. I was only trying to illustrate that
>your statement about $\log(D/r)$, while true about THAT measure,
>is not IN GENERAL true about information. D/r is just the
>number of possibilities inherent in the numbering system or the
>measuring apparatus.

It IS an important point. Martin has been implying that the resolution of the input function depends on its RMS noise level, and I completely disagree with this concept of resolution of the measuring apparatus. If the noise is 10% RMS of the range, Martin's interpretation would be that there are only 10 possible values of the perceptual signal with a range of 10 units, 0 through 9. This would make the probability of any one value of perceptual signal 10 percent. In fact, however, the perceptual signal can have a magnitude between 0 and 9 with a precision that depends on how long you observe it: if you observe it 10 times as long, it has 3 times the precision if the noise is Poisson-distributed. And however long you observe, whether for a short time or a long time, the RMS noise does not predict the probability of a specific measure, for a specific measure can have any value in the real number range between 0 and 9. The resolution is infinite, although the precision and repeatability are not.

So the number of possible measures between 0 and 9 is not 10, as Martin's assumption implies, but infinite. The only way in which the probabilities could be correctly computed from r/D would be in the case where D is divided into exactly r intervals with fixed boundaries, so that no measure can be anything but an integral multiple of r . This is not true when r represents the RMS spread of a noise function. In that case, the underlying resolution is far smaller than r -- it must be, if we can even talk about a noise "distribution." If the resolution were equal to the 1-sigma spread in noise amplitude, we couldn't draw the curve of the noise distribution: it would be maximum for all deviations from the mean less than 1 sigma, and 0 for all larger deviations. That is not how real noise behaves.

From what Martin and you have said, and from reading the materials Martin has sent me and that I have looked up, it seems clear that the basic logic of this whole approach was settled many decades ago, and has not been critically examined since. Much more interesting to forge ahead with the complex mathematics. And, after all this time, much more dangerous to fiddle around with the basic assumptions on which so much now depends.

I certainly can't follow the advanced mathematics that have been built up on the base, but I can judge simple basic assumptions as well as the next person. I keep finding sloppiness and ambiguities in them, and even illogic. I also find great pressure to get out of that area, leave the assumptions alone, and concentrate on learning what has been deduced from them. This simply goes against my nature: I can't whip up any enthusiasm for pursuing the complex implications of premises that I still think are probably flawed.

As I said in a previous post, however, none of my objections will mean anything if you can come up with an analysis of an actual control system and show what information theory has to say about it, in numbers. You say

>*Given* a language or model to work with, information *is*
>well-defined. But, you said yourself that what the numbers mean
>depends on the physical situation we take them to represent.
>And this *does* depend on what is assumed.

Fine. So specify a language, state your assumptions, and do it. I'm not going to be impressed by mathematical generalizations. It's too often the case that the mathematical mountain stirs, and gives birth to a mouse. It's also quite frequent that when one gets down to analyzing a specific actual situation, nature pops up with relationships you hadn't anticipated, and you begin to see that there is something wrong with your assumptions. There is no substitute for doing an actual quantitative analysis of real data. If you just want to do abstract mathematical analyses, that's fine, but it just means that I'll have to wait for someone else to demonstrate that the results apply to ANY real cases of control.

As to the "information in perception" argument, I think you're pulling back from the original position, which is OK but you should acknowledge it. Originally, the concept was that the output of the system almost exactly opposes the effects of the disturbing variable on the basis of information contained in the perceptual signal. Now it turns out that your side is admitting that the perceptual signal doesn't contain ALL of the information in the disturbing quantity, but only some part of it. You say

>... if I can use the percept to do the reconstruction using
>*fewer* additional bits than would be required without the
>percept, then I have shown that the percept contains
>information about the disturbance.

Showing that it contains SOME information is not the same as showing that it contains SUFFICIENT information from which to construct the required output. I know you feel that it MUST be sufficient, but that is what remains to be proven, and can't be taken as a premise.

I presume you would agree that the better the control, the smaller the amount of information that gets through to the perceptual signal. So you are now in the position of having to show how an output waveform that almost exactly matches the waveform of the disturbing quantity, and thus presumably contains nearly the same amount of information, can be generated from a perceptual signal that contains only a small part of the information in the disturbing quantity (all in terms of rates).

I will be fascinated to see how you solve this problem and demonstrate that the information in the disturbing quantity is sufficient to produce the output actions. Simply insisting that it MUST be enough isn't acceptable to me: you have to demonstrate that it is, with numbers, before I'll believe it. Martin's mystery function didn't prove this; it simply

showed that if you know all the functions and all but one of the variables, you can solve for the remaining variable. There was no demonstration of the relationship of this fact to information theory. What Martin ended up proving was that algebra works.

You say

>You posted correlation figures. As far as I can remember that was about it.

Rick, how about sending Allen that string of figures again? I also sent some, Allen. Perhaps you just haven't come across them. I'll regenerate mine and send them, too, if you want them.

As far as I'm concerned, you have have all information about all variables and functions. At this point, all I want is to see how you actually do the calculations.

Best, Bill P.

Date: Sat Jun 19, 1993 2:50 am PST
From: Hortideas Publishing / MCI ID: 497-2767

TO: * Dag Forssell / MCI ID: 474-2580
Subject: The trouble is not in your computer...

Hi Dag -

I'm snowed under (thunderstormed under???) right now, working on getting electricity into our new house. I'll comment on your "myth" post as soon as I am able, possibly Monday. Thanks for being patient.

Best, Greg Williams

Date: Sat Jun 19, 1993 8:17 am PST
Subject: (emergent) control

[Hans Blom 930619] (Rick Marken (930618.0900))

>> You are of course right: the individual control
>>laws together 'control' flockness, but in an emergent way

>I think we are in violent AGREEMENT.

That only seems so, for the moment.

(Tom Bourbon (930618.0958))

>My guess that you speak from the position of a designer, engineer, or
>user of a system was confirmed when I came to your comments

Not only. I have been a control engineer for so long (25 years) that I could not miss starting to question its basic assumptions. What I rediscover every time is that modelling is the basis of control. Without an approximately correct knowledge of the system to be controlled it is impossible to achieve any control quality. In Bill Powers' approach the 'knowledge' about the system to be controlled is very limited; most of it is stored in the choice of the parameter that he calls the 'slowing factor'. Choosing an inappropriate slowing factor either results in oscillations or in a system that is too slow to be useful. Bill might protest and say that he needs no model in order to be able to control, but his having to search for an appropriate slowing factor indicates otherwise. Make it a factor of 10 or 100 larger or smaller and chaos or stagnation results. The knowledge incorporated into Bill's models is of unacceptably limited, so other control approaches go further. Using adaptive control, for instance, you would also provide the system with a capability that allows it to discover the best slowing factor. Such adaptive controllers will always ALMOST oscillate (but not quite). The tremors in our limbs may indicate that something like this goes on in humans as well.

But my basic discovery was that each model is an APPROXIMATION of reality, geared toward some goal. A model incorporates the features that you focus your attention on in that one particular application, because you want to study only those, and it does not contain the

features that you deem unimportant in the study. Each model is therefore a human invention, a useful approximation, a tool to study an otherwise too complex reality. It is simply not true that humans ARE control systems; it is true, however, that some human qualities can be MODELLED as control systems. The map is not the area.

> could you describe the mechanism the "keeps
>flockness within specified (but maybe varying) limits or almost so?"

I thought I did: the control laws of all individuals together create the flock implicitly. There is no explicit control for flockness.

> *You* use the heating system to control
>your perception of a nice room temperature; the heating system does
>not. I *think* I understand your position pretty well, or am I wrong?

I think you do. *I* want to control for a nice room temperature. The heating system can only control the temperature at its sensor. So what do *I* do to control the controller? I put the sensor at a location where what the heating system can do approximates as closely as possible what *I* want done. There are two levels: the heating system controls, and I control using the heating system's controller. Could that be a hierarchical control system?

From the perspective of the heating system, which knows nothing of the room temperature except the temperature at its sensor, the temperature at locations other than at its sensor is uncontrolled. From MY perspective, I control the temperature at the position where I sense it, using the heating system. The perspective matters!

(Bill Powers (930618.1000 MDT))

>>Please be explicit in what discriminates a 'natural tendency'
>>from control (should you do so, I am tempted to say, I will
>>give you a counter->example :-).

>Explicitly, it's power gain. A control system has an output
>function that produces vastly more power output than its inputs
>provide (drawing, ultimately, on external energy supplies).

The counterexample (see below for another one): You will agree that your hifi amplifier is a control system. It amplifies the small power at its input terminals into the larger power that the speakers need. Here is the power gain that you stress as the important thing. The feedback is of course needed because the amp's components are not quite ideal, yet we want hifi quality. Now imagine an almost identical situation, where it is the electrical power engineer's task to track the current and voltage of those energy-rich 380.000 Volt overhead power lines on his oscilloscope. What he needs is a tremendous power LOSS. The components of his measurement system are non-ideal as well, and therefore he too needs feedback. Where is the power gain now?

You can, no doubt, think of many more examples in which energy or power must be scaled DOWN rather than up. This example stresses what *I* consider the intent of control: to improve otherwise inadequate results, with or without power gain.

> You don't arbitrarily perturb a real control
>system without spitting on your hands first and bracing your
>feet. When I say that control systems resist disturbance, I mean
>that they resist it energetically and with as much effort as
>needed (up to their limits of output).

This and the next section of your reply make your position much clearer for me. What you talk about is servomechanisms, and you also seem to exclude servomechanisms that have no power gain such as micromanipulators that atomic physicists or surgeons might use to operate on a scale much finer than that of the human hand.

Now servomechanism theory is, indeed, part of control theory, but control theory is much wider. Control theory also studies adaptive control, time-optimal control, and robust control, for instance. In my previous contributions, I have referred to these areas of control theory as well, not knowing that you wanted to limit things to servomechanism theory only. That may be the basis of our mutual misunderstanding about what constitutes control.

>>Isn't it remarkable how robust flockness is in the face of disturbance?

>I don't think you know what "robust" is.

I was talking in terms of robust control systems. Robust control is a new branch of control theory, that you may not be familiar with. It has greatly gained in importance during the last ten years or so. The goal of robust control is to design systems that continue to control well in an environment which may change greatly or with sensors and/or actuators whose characteristics may vary randomly or systematically, as in aging of catalysts or decay of magnetic properties and so on. For example, it is one of the limitations of servomechanisms that they break down when their sensors become very noisy. There are solutions for that problem using other approaches.

>If it were not for the fact that I remain upright while I walk,
>the air along my path six feet above the ground would not be
>displaced. The movement of the air follows lawfully from the
>effects of my remaining upright while I walk, but it is not a
>controlled variable. It is an uncontrolled, although lawful,
>outcome of my control processes. A side-effect, as far as I am concerned.

Again a matter of perspective. If it is my goal to walk, you are right. If it is my goal to displace air (through walking), you are not.

(Rick Marken (930618.1500))

>So Hans' and my idea that "flockness" is under control (a "collective
>controlled variable" by my terminology -- ie. a variable that is controlled
>by the collective but not by any individual) is incorrect unless there is a
>control system present that is perceiving and acting to influence this variable.

We are NOT in violent AGREEMENT now. See how difficult these things are? One moment you see control, the next moment it's gone. That is the best demonstration you, an expert on control, could give me that the concept of control is slippery.

Greetings, Hans

Date: Sat Jun 19, 1993 11:34 am PST
Subject: Power LOSS

[From Dag Forssell (930619 1230) Hans Blom 930619

>I have been a control engineer for so long (25 years)...
>

>...Now imagine an almost identical situation, where it is
>the electrical power engineer's task to track the current and
>voltage of those energy-rich 380,000 Volt overhead power lines on
>his oscilloscope. What he needs is a tremendous power LOSS. The
>components of his measurement system are non-ideal as well, and
>therefore he too needs feedback. Where is the power gain now?

As a mere mechanical engineer, with limited understanding of electronics and control, I am astounded that you offer this as a serious proposition. Perhaps your oscilloscope is *powered* by the 380,000 volt line. I have never heard of one.

Any oscilloscope I have ever encountered works exactly like the HI FI amplifier you describe and *requires* power. The oscilloscope is plugged into the wall, or has a battery to draw on. Typically, The signal from the 380,000 line is reduced by resistors (without distortion) to a suitably low voltage and almost no amperage signal to the oscilloscope input. From there it is AMPLIFIED to drive the cathode ray tube that the human looks at.

Looking at the Niagara Falls is to me quite analogous. There is (paraphrasing you) tremendous power LOSS between the energies in the falling water and the vision in my brain. But there is AMPLIFICATION between the small amount of light that falls on my retina and the image I experience (internal to my nervous system), powered by blood sugar and whatever else from food.

Fortunately for me, I don't have to deal with the full brunt of Niagara, any more than the oscilloscope deals with 380,000 volts and thousands of amperes.

Best, Dag

Date: Sat Jun 19, 1993 3:18 pm PST
Subject: Back to Basics

[From Rick Marken (930619.1400)] Hans Blom (930619) --

>I have been a control engineer for so long (25 years) that I
>could not miss starting to question its basic assumptions.

Given all that experience and questioning of basic assumptions, it seems like you're just
the one to set us straight on the "information in perception" brohaha. Given what must be
your deep knowledge of control theory and its basic assumptions, could you just tell us --
does the sensory input to a control system contain information regarding how "plant"
outputs should vary in order to control the sensed variable? While you're at it, could you
also tell us what variable is controlled by a control system?

If you are lucky, you're answers could get you a section in the BBS "misconceptions about
control theory" paper.

>Bill's models is of unacceptably limited, so other control approaches
>go further.

I feel the uncontrolled side effects of waving hands.

>It is simply not true that humans ARE
>control systems; it is true, however, that some human qualities can be
>MODELLED as control systems. The map is not the area.

I think you can now rest assured that you will have a section in the
BBS "misconceptions about control theory" paper.

>We are NOT in violent AGREEMENT now. See how difficult these things are?

So? It's difficult, sometimes, to tell whether a variable is controlled or not, especially
when 1) you have not built the control system yourself and 2) the variable is the result of
the controlling done by multiple control systems. What's your point? Is the fact that it
may sometimes be difficult to determine whether or not a variable is under control (and
it's probably just more difficult to determine it in the armchair than in the lab) some
kind of indictment of the PCT model of living systems? You got something better? I see you
and raise you.

Best Rick

Date: Sat Jun 19, 1993 4:09 pm PST
Subject: The informative perception

[From Rick Marken (930619.1700)] Bill Powers (930618.1930 MDT)--

>Rick, how about sending Allen that string of figures again?

As the Dred Pirate Roberts said in "The Princess Bride" (a wonderful movie -- my daughter's
favorite) "As you wish". Here are the perceptual values. I believe that it should be
possible to reconstruct the output values from this perceptual signal. Hope those outputs
are nearly the opposite of the disturbance values (which I posted to Gary Cziko). Good
luck.

By the way, the refernce value was a constant -- r = 0
and the output was computed as O := O + 10 * (r-p) * .1
The starting value of O was 0.

What follows is a highly informative perceptual signal.

0 0 0 0 -1 1 0 0 0 1 0 0 0 1 0 1 1 1 1 1 2 1 2
2 2 3 3 3 4 4 4 6 6 7 7 9 2 -8 -12 -11 -10 -9 -7 -5 -5
-5 -5 -3

Best Rick

Date: Sat Jun 19, 1993 7:09 pm PST
Subject: Emergent control

[From Bill Powers (930619.1800 MDT)] Hans Blom (930619) --

I detect the opening of a small window between our worlds. It opened like this:

>So what do *I* do to control the controller? I put the sensor
>at a location where what the heating system can do approximates
>as closely as possible what *I* want done. There are two
>levels: the heating system controls, and I control using the
>heating system's contrller. Could that be a hierarchical
>control system?

This illustration of hierarchical control has been used a number of times by PCTers, but not quite in the manner you describe it. Under normal circumstances, a householder has no choice but to leave the sensor where it was installed. But there is a way to achieve hierarchical control even if the sensor is located so it does a poor job of controlling air temperature in the place where the householder wishes to sit. The knob on the thermostat, which adjusts the reference level, can be turned up or down until the air at some location other than the sensor is at the desired temperature. As long as heat disturbances don't change, the thermostat, by keeping its own sensor at the new temperature, will, as a byproduct, keep the householder comfortable at a different temperature. The householder is a higher-level control system using another control system, through adjusting its reference signal, to control a variable other than the one the lower control system itself is controlling. In this case the difference isn't great -- it's just a matter of where temperature is being sensed -- but the principle remains the same.

Of course the air temperature the householder experiences, being far from the thermostat's sensor, is easy to disturb and indeed can be disturbed without the thermostat's knowing it and correcting the error. So the householder would be best advised to sit near the control unit to avoid having to get up frequently and alter the reference setting of the thermostat. Or, of course, as you suggest, rip the control unit off the wall, solder extensions onto all its wires, and carry it around (as organisms, in fact, do with their own private portable temperature regulators).

Much of the rest of what you say shows that you are actually considering the same kinds of problems with which PCT must some day become involved -- and you're hinting that when that day comes, a lot of the work will already have been done!

>The goal of robust control is to design systems that continue
>to control well in an environment which may change greatly or
>with sensors and/or actuators whose characteristics may vary
>randomly or systematically, as in aging of catalysts or decay
>of magnetic properties and so on. For example, it is one of the
>limitations of servomechanisms that they break down when their
>sensors become very noisy. There are solutions for that problem
>using other approaches.

In order for "robust control" as you define it to be achieved by an autonomous system (i.e., without an engineer standing by with test instruments and a screwdriver), something in the system itself must be able to perceive the quality of control. If a control system threatens to become unstable, for example, something must be monitoring a function of the various system variables so as to report the degree of instability that is present. The desired level may be zero, or it may be a one-way specification: no more than some maximum amount of instability. The resulting error signal has to operate an effector system which then alters the parameters of the control system in the direction that decreases the perceived instability toward the reference setting.

The organization of the stabilizing system is thus just like the basic diagram of a control system in PCT, except for one thing: it does not act by adjusting the reference signals for lower systems, but by adjusting the parameters of their components. It does not perceive a variable that is a report on a higher-level aspect of the environment, but a variable that represents the way the control system is controlling.

I recognized many years ago that this sort of hierarchical control relationship had to exist in the human system. However, I also recognized that modeling it analytically was a very much more complex undertaking than I would be able to carry out. The only step taken in that direction (which Clark and I mentioned in our joint paper in 1960) was the introduction of a reorganizing system, a system which altered parameters at random (anticipating the "genetic algorithm" approach now in vogue). This idea is still necessary, for if there are organized ways of adjusting the parameters of control, the systems that accomplish it must themselves have arisen somehow.

Control systems with random outputs can actually achieve control far more efficiently than one might suspect, by altering the spacing between random changes caused by the output. They explore the total space of possibilities, and so are not constrained to produce any particular type of control system.

But the general topic of control through parameter adjustment, which I believe I mentioned as a research topic even in BCP, is not covered by reorganization. Indeed, the acquisition of systems that control sensed system performance through variation of parameters has the effect, once such a system is working, of greatly reducing the need for random reorganization. The control systems can remain functional even though changes in the environment and their own characteristics go far beyond what a system with a fixed design can accomplish. As you say, such systems are much more robust than are systems with fixed characteristics. So you have my apology for saying you don't know what robust means. I underestimated you. You clearly know more about that subject than I do.

But don't dismiss mere "servomechanisms" too lightly. What I realized long ago was that even if I couldn't model how control systems alter their parameters, I could still study human control systems when they are operating with relatively stable characteristics, in an environment to which they are adapted. In other words, I could ask the question, "How is competent adult human behavior organized when its organization is not being changed?" This is what HPCT is about.

The basic HPCT model, or a PCT model of a specific control system, assumes a system with constant characteristics. As Tom Bourbon has shown with data from tracking experiments, the model parameters that produce the best fit with behavior continue to predict behavior very accurately after a lapse of one year. At our meeting this summer, he will unveil some predictions made 5 years ago, and put them to the same test. We confidently expect very little change in the parameters, and an excellent prediction (3 to 5 percent RMS mismatch between tracings of the real actions and the model's actions). So at least with respect to the particular behaviors we have been able to explore experimentally, the assumption that human control systems retain constant characteristics over some period of time does not seem unreasonable. In fact, it seems correct.

However, it's clear that as we learn how to model higher and higher level control behaviors, that assumption may break down. Then we will have no choice but to include adaptive control in the HPCT model. I don't want to do this until we have a specific problem with modeling that can't be handled with a system having fixed characteristics. But that day will surely come.

When I speak of the excellence and tightness of human control systems, I'm not trying to account for how they got that way. Neither am I trying to explain how they change to maintain that excellence when long-term changes in the environment (or the organism itself) obsolete a formerly competent design. My central goal is simply to get behavioral scientists to start thinking of behavior as a control process, and abandon all the previous bad guesses about its underlying organization. The hierarchical control model with fixed characteristics and control strictly through adjustment of lower reference signals is, I believe, enough to do the job. The sorts of problems you imply with your robustness control are important only after the basic model has been accepted. In fact when people ask me where I think PCT will go in the future, adaptive control is one of the subjects I mention. I also mention that this is a difficult problem and that it will probably be solved by our descendants (it won't, at any rate, be solved by me).

Now to your supposed counterexample concerning power gain:

>Now imagine an almost identical situation, where it is the
>electrical power engineer's task to track the current and
>voltage of those energy-rich 380.000 Volt overhead power lines
>on his oscilloscope. What he needs is a tremendous power LOSS.

In the first place, I'm not sure what you mean by "track." If you just mean "observe," then yes, there has to be a power loss. The sensor used for measuring the line voltage reports the voltage without drawing any significant current, and the current that is drawn (times the voltage drop) is reduced to the level of optical power when the image of the meter face enters the engineer's eyes. [Incidentally, measuring voltage is not measuring power]

However, the power gain I am talking about is power gain measured inside a controlling system. To apply this concept, we would have to have the engineer ADJUSTING the voltage on the 380-volt line. In that case, the engineer would have to have a control knob that allowed the levels of force generated by his muscles to alter the line voltage regardless of the load on it. Coming in to the engineer's senses is a minute scrap of energy that represents the line voltage. Coming out of the engineer is a force that dissipates many orders of magnitude more (muscle) energy than comes into the engineer, and this force results in still greater power amplification in the device that alters the voltage put out by the generator.

The extra power gain in the environment, however, is offset by the power loss in the link from the high-voltage line to the meter and then to the senses of the engineer. There is probably a net power loss in the environment. In the engineer, however, there is an enormous power gain. If we trace gains and losses all the way around the loop, starting anywhere, we will find that the loop power gain is probably on the order of 40,000 (the amplitude loop gain in typical visually-guided motor control processes measures out very roughly at about 200). That assumes that the engineer can control the line voltage as accurately as he would be able to make a cursor track a moving target given a similar manual control knob.

So your counterexample isn't a counterexample, but an example of what I was talking about.

>You can, no doubt, think of many more examples in which energy
>or power must be scaled DOWN rather than up.

Yes, but scaled down where in the loop? All that really counts for control is LOOP power gain, not gains or losses in any one part of the loop. As a control engineer, you surely understand this.

>... you also seem to exclude servomechanisms that have no power
>gain such as micromanipulators that atomic physicists or
>surgeons might use to operate on a scale much finer than that
>of the human hand.

A mechanical micromanipulator is not a control system, although an electromechanical one may be. The electromechanical one has a very high loop power gain if it's a servomechanism. What its actuator loses in amplitude the control system gains back through the sensor detecting the tiny movements, and more power gain is added inside the controlling system. Without considerable power gain, you can't have accurate servocontrol.

Any micromanipulator, even a mechanical one, PLUS an atomic physicist or a surgeon, IS a control system. The amplitude gain that is lost in the reduction of the surgeon's output is regained when he looks through the microscope at the highly magnified movements of the tiny tool. The relationship between hand movements and image movements is about the same as when positioning something directly while looking at the result directly. There is a large power gain going from the optical image to the muscle outputs.

Amplitude loop gain is actually more informative than power gain, because to compute power gain you have to take impedance transformations into account. A frictionless mechanical micromanipulator would actually have a power gain of unity: you trade distance moved for force amplification. But then you have to consider internal losses in the muscles, and it all gets excessively complex. It's easier to look at amplitude ratios, knowing that power gain is simply the square of the amplitude ratio when you measure under the same impedance conditions (as you do by definition when measuring loop gain).

>In my previous contributions, I have referred to these areas of
>control theory as well, not knowing that you wanted to limit
>things to servo- mechanism theory only. That may be the basis
>of our mutual misunderstanding about what constitutes control.

The basic model for behavior has to be, I think, the servomechanism, control-of-input, model. The reason is simple. Organisms are not like commercial control systems; they are

not organized to be used by someone else. The only evolutionary reason for behavior to occur at all is to produce some effect on the organism itself, or to prevent something from happening to it. Open-loop behavior might be practical in the world of devices that are employed by human users, because the human user can monitor the result and alter the commands to or the adjustment of the device to correct any errors. But a behavior that simply affects the outside world with no consequences that ever return to the organism could have no evolutionary basis; as Martin Taylor remarked, that would be a waste of resources. The whole point of behavior is to control the effects the environment has on the organism, keeping those effects in states or conditions or amounts that the organism itself specifies internally. In fact, this ability of control systems to control what happens to them is a strong selection factor favoring the development of closed-loop control rather than open-loop response.

That is the sort of control that adaptive systems in the human organism must maintain, however they do it.

 Finally, I think we have agreed that flockness is not under closed-loop control, but is only under control in the emergent sense. I still have great difficulty in allowing for the latter sense of "control:" it is hard for me to conceive of a kind of control in which the controller never knows about the output and doesn't care what happens to it. My suspicion is that all so-called open-loop controllers are embedded in a closed control loop, and that if they're not, they can't actually control anything.

Best, Bill P.

Date: Sun Jun 20, 1993 12:30 pm PST
 Subject: emergence etc.

[Hans Blom, 930620] Rick Marken (930619.1400)

>could you just tell us -- does the sensory input to a control
 >system contain information regarding how "plant" outputs should
 >vary in order to control the sensed variable?

Sorry, but I decline to be arbiter. As a teacher (my other role) I have to say that having this discussion is by far more fruitful than knowing who is right. If the two of you stop fighting about who is right and who is wrong and start to appreciate that the other may have something useful to say as well, the two of you COMBINED might someday invent optimal control theory :-)

> It's difficult, sometimes, to tell whether a variable is controlled
 >or not, especially when 1) you have not built the control system yourself

YES, YES, YES. This is exactly what I have tried to tell you all along. Savor this thought and let it reverberate for a while!

> I see you and raise you.

I thought, when joining this list, that a scientific discussion was going on, a common search for deeper understanding. That interests me. I am not interested in establishing winners and losers. That is not the game that I play.

>(Bill Powers (930619.1800 MDT))

>In order for "robust control" as you define it to be achieved by
 >an autonomous system (i.e., without an engineer standing by with
 >test instruments and a screwdriver), something in the system
 >itself must be able to perceive the quality of control.

YES, YES, YES. That was a discussion I let loose on you some time ago. Rick Marken shrugged it off as 'the control of Q', but that is what it really is: the control of the *Quality* of the response.

Now this has wide repercussions when you combine it with evolution: just a slightly better quality in how an individual controls its basic evolutionary goal (survival of its descendants) means that slowly by slowly he will start to dominate the population until no others are left. Thus maybe the entire evolutionary process can be seen as emergent control.

It also has wide repercussions for the individual. The slogan "it's all perception" is much too static. It has a connotation of having to act in an a priori circumscribed way given a set of perceptions and a set of top-level (very slowly changing; innate?) reference levels. Blaming what you see for what you do is usually considered a defense mechanism or, if more forceful, for criminality. Optimal control theory applied to humans says that you can fine-tune your actions as well, that you have control over how you control, i.e. that you have SELF-control.

>The organization of the stabilizing system is thus just like the
>basic diagram of a control system in PCT, except for one thing:
>it does not act by adjusting the reference signals for lower
>systems, but by adjusting the parameters of their components. It
>does not perceive a variable that is a report on a higher-level
>aspect of the environment, but a variable that represents the way
>the control system is controlling.

Yes, see how far this reaches when you consider humans. We do not only operate on the outside world but on the inside world as well. We can tune our responses finer and finer, and reach ever higher qualities of response and perception. Control over control is self-control, perceiving your own perceptions is self-perception, consciousness. See how fascinating this approach is?

>I recognized many years ago that this sort of hierarchical
>control relationship had to exist in the human system. However, I
>also recognized that modeling it analytically was a very much
>more complex undertaking than I would be able to carry out.

Chaos theory might show it impossible. A great many processes have been discovered where no prediction (and hence no control) seems possible. In those processes, the only way to know the outcome is to actually run the process and observe how it develops. Yet, even chaos theory develops is (you might say 'emergent') laws, that may provide a little understanding of what goes, and maybe only over short periods of time. That is why ever better simulations are important: even if it is impossible to come to analytical grips with the matter, you might find general tendencies. Maybe only as vague as "flockness", which doesn't really exist (because it isn't controlled for, you might say), but still this kind of vagueness may be preferable to not knowing or misconceptions.

>In other words, I could ask the question, "How is competent adult
>human behavior organized when its organization is not being
>changed?" This is what HPCT is about.

Whereas my focus is more on how human behavior can become even more competent, i.e. more on learning (and evolution as a kind of learning).

>When I speak of the excellence and tightness of human control
>systems, I'm not trying to account for how they got that way.
>Neither am I trying to explain how they change to maintain that
>excellence when long-term changes in the environment (or the
>organism itself) obsolete a formerly competent design.

But those are exactly the things that fascinate me!

>Now to your supposed counterexample concerning power gain

Let me summarize our discussion thus far and maybe make it a little bit more precise. One can, at our level of discourse, see a feedback amplifier or some such device in two very different ways. The first is as a device that transforms an input (voltage, current or power) into an output (voltage, current or power). The feedback is not really relevant here. The (voltage, current or power) gain may be any value; both gains and losses can be similarly realized. The second way is to see the device as a power modulator, where the input modulates the transfer of power from its power supply to its output. The feedback is not relevant here either. Here we will always see a power loss, since the device's efficiency cannot be more than 100%. I often find that people confuse these two quite independent perspectives, especially mechanical engineers :-).

Then, internal in the device, we see the LOOP GAIN, which has nothing to do with power at all, but is the gain in the loop that the SIGNAL travels. A decent control system has a

loop gain much greater than one, although a loop gain of less than one is not unthinkable. In the latter case we might not want to consider the system a CONTROL system.

Loop gain is terminology from the analogue past. In a hifi amp's innards you can still trace the loop with your finger. In a control system that is described in difference rather than differential equations, the 'loop' is functional only at a sequence of infinitely short time intervals. That is why they are also called 'sampled data systems'. Now, to complicate matters even further, in adaptive control systems there are (at least) two loops: one to provide the immediate actions, and another to do the parameter adjustments (learning). These loops interact and hence become much more difficult to trace with your finger. All of this causes the term "loop gain" to slowly fade away.

> All that really counts
>for control is LOOP power gain, not gains or losses in any one
>part of the loop. As a control engineer, you surely understand this.

I understand loop gain, not loop POWER gain.

>The basic model for behavior has to be, I think, the
>servomechanism, control-of-input, model. The reason is simple.
>Organisms are not like commercial control systems; they are not
>organized to be used by someone else. The only evolutionary
>reason for behavior to occur at all is to produce some effect on
>the organism itself, or to prevent something from happening to it.

No. Quality is the thing that counts in evolution, period. Only those with the highest quality leave descendants behind, in the long run, whatever the mechanism. I realize that this usage of the word quality has a circular definition, but that is because this quality emerges, and is not the result of a control process. I realize that this must sound obscure to you, but explaining what I mean would take more time than I can spare at the moment. I have a huge pile of exams to assess. Maybe later.

>Finally, I think we have agreed that flockness is not under
>closed-loop control, but is only under control in the emergent
>sense. I still have great difficulty in allowing for the latter
>sense of "control:" it is hard for me to conceive of a kind of
>control in which the controller never knows about the output and
>doesn't care what happens to it.

Maybe the non-controller DOES care but also knows that he alone could never achieve the emergent thing. Maybe the US is the emergent result of a great number of people who control for living in peace together at the level where they CAN control. I think that the notion is VERY important. Please reconsider your reluctance...

Greetings, Hans

Date: Sun Jun 20, 1993 4:35 pm PST
Subject: Re: WTP SOC SYSTS - RKC

[From Bob Clark (930620.20:15 EDT)]

Sorry to be slow in responding, but my activity with the City of Forest Park is keeping me pretty busy -- attending several Commission and Committee meetings, helping compose a proposed Ordinance, etc.

Bill Powers (930613.2200 MDT) Tom Bourbon (930614.0840)

I guess I inferred too much from Bill's remark:

>This is a fundamental design defect; it creates conflict automatically.

You ask:

>What does not liking them, or who created them, have to do with anything?

"Liking, etc" relates to the vehemence to be expected in discussion. Of course it has nothing to do with the operation of the system considered.

Yet:

>I guess that as an engineer, I am pained when people cooperate in an
>attempt to achieve common goals using means that will almost
>certainly, in the end, frustrate the attempt.

I, too, Bill, am concerned when tasks are not well done -- but I still struggle to avoid having my own "perfectionism" interfere with moving ahead.

Your last few remarks suggest that we are, indeed, in very close agreement, about the location of Social Control Systems. Thus:

>>The only important place where they exist is in the minds of the
>>affected individuals, especially the participating individuals.

>Yes, exactly.

Tom Bourbon (930614.0840).

You also seem to be agreeing with this view where you say:

>Their physical existence is tenuous to the point of non-existence. As you say,

>>"The only important place where thy exist is in the minds of the
>>affected individuals, especially the participating individuals."

Continuing with Bill:

>>"Real" or not, I think we can learn a great deal about how people
>>think about and control their environments by studying their
>>Social Control Systems.

Skipping a paragraph:

>Of course when you say "their" social control systems, you probably
>mean each person's conception of a social control system -- so my
>proposal agrees with yours.

You refer to "traditionally ... in terms of metaphors," in the paragraph "skipped" above. You suggest working with "a model intended to be literally true," yes, of course, that is what I am trying to do.

Comment, in passing: I have been somewhat uncertain about the use of "modelling" in reference to the HPCT. Checking my dictionary, I have no quarrel with this. I am concerned with possible confusion with "metaphors" and "analogies," which are perfectly good words, but defined as differing in some significant way with the object (idea, procedure, system, etc) at hand. I prefer "example" as an equivalent, but somewhat more familiar term for the concept in question.

For example, to me a "Social Control System" is an "example" of a Hierarchical Negative Feedback Control System. It may not work very well, it may have assorted defects, it may be loaded with conflicts, but I think any Social Control System is such an example.

To elaborate: to be considered a "Social Control System," it must be composed of elements each of which is individually an HPC System. Each of the elements must include, somewhere within its hierarchy, the essential Rules, Customs, Procedures, etc defined by the specific SCS the element has accepted. This, in no way, precludes any individual from disagreeing, opposing, etc the system as he understands it. The Rules, etc, defining the SCS need not be written down, but often are. Such writing is, in essence, an extension of individual memories, and can be very helpful.

In terms of the current, "Official," Orders, these Rules, etc. will include some Configurations, some Transitions, some Sequences (or Events), some Relationships, and some Programs. Some Principles may or may not be included, as well as some System Concepts. This is logical and internally consistent. But hard to work with and hard to communicate to ordinary people.

This is among the several reasons I have suggested alternative identifications for some of the Orders. Here, the relevant Order in my suggested sequence is Sixth Order, Interpersonal Relationships. Non-personal Relationships would appear elsewhere. In

addition to fully developed Social Control Systems, Sixth Order would include incidental customs: hand shaking on introduction, sometimes using given names vs formal names, driving on the right hand side (in the US), saying "Please" and "Thank you" at times. Many verbal interactions, theoretical discussions, etc.

Interestingly, Sixth Order and higher Orders would be irrelevant for Robinson Crusoe without his man Friday, because verbal communication would be impossible.

I know this is a sizable departure from the existing naming of the Orders of Perception. But I think it necessary to review these assignments, and I quote, BCP, p 248:

"If we cling too lovingly to the particular structure in this book, we may miss the chance to make major improvements in the near future."

More later, Bob Clark

Date: Sun Jun 20, 1993 4:49 pm PST
Subject: Levl's viewpnt - RKC

[From Bob Clark (930620.08:30 pm EDT)] Bill Powers (930614.0800 MDT)

Referring to Bruce Nevin (930614.0836), you say:

>We need some creative thinking about "the" levels. What would you
>propose makes the difference? For example, what sort of perceptual
>capability is required to distinguish between another person and a
>non-person?

Bill, this is what I'm trying to work out. My post (RKC 921205) suggested a Fifth Order of Skills, primarily muscle skills. Sixth Order, as discussed in recent posts, might become Interpersonal Relationships. The difference would be between inanimate objects and independent active entities. This would include "animals" with people as "persons," above. "Communication" could also be considered for Sixth Order. It would seem to distinguish pretty well between inanimate objects and "independent active entities."

You note, from Bruce Nevin (above):

>Information theory requires a point of view from outside the control
>system; it cannot be applied to the control system itself from its
>own point of view.

Of course, I agree. But more important, if Info Theory has something to contribute, let's hear it. So far I seem to have heard only assertions that it can describe (or do they say "model," or what?) HPCT independently. Let's get away from "Prove it," tests etc. Some kind of working example would at least be interesting.

Regards, Bob Clark

Date: Sun Jun 20, 1993 4:51 pm PST
Subject: Re: emergence, etc.

[From Bill Powers (930620.1600 MDT)] Hans Blom (930620) --

>The slogan "it's all perception" is much too static. It has a
>connotation of having to act in an a priori circumprescribed
>way given a set of perceptions and a set of top-level (very
>slowly changing; innate?) reference levels.

I agree. Perceptions are learned, too. There is a reason for the slogan, however. It's to remind us that each of us sits inside one of these gadgets, and that ALL we know either of the world or of our own actions and inner being consists of perceptions. When we produce a particular carefully crafted action, it is a perception of that action that we know about and control. When we see an effect of the action on the world, it is a perception of the effect that we observe and adjust. When we feel joy or anger at the result, that, too, is a perception. There is nothing else to experience. Our actual outputs, the signals moving in the outward direction, are not part of experience at all. This is of no concern to an

engineer, who looks at his control systems strictly from outside. But in psychology it is the key to understanding what control theory means for behavior.

>Blaming what you see for what you do is usually considered a
>defense mechanism or, if more forceful, for criminality.

True. This is one of the things that PCT is trying to change in psychology: the idea that perception causes behavior. PCT says that the person selects some experience as a goal, and acts to make present-time perception conform to it. Behavior controls perception, not the other way around.

>Optimal control theory applied to humans says that you can
>fine-tune your actions as well, that you have control over how
>you control, i.e. that you have SELF-control.

Again, I agree. Of course you can't fine-tune your actions directly unless you can perceive them; otherwise, you can only reorganize until the perceived result of the actions is what you want, without any direct knowledge of the actions themselves. And it remains true that even when you specifically adjust your actions, it is a perception of the actions you must adjust; the action itself is output, and not sensible.

>Yes, see how far this reaches when you consider humans. We do
>not only operate on the outside world but on the inside world as well.

The distinction between inside and outside is a perceptual classification; both, as far as the brain (or PCT) is concerned, are inside (or both are outside, it makes no difference). Everything the brain can deal with exists in one space, the space we call the experienced world. This world is derived completely from signals generated by sensory receptors; there is no other way to get information about an external world. The nature of that world has to be inferred by the brain from the behavior of the signals and how they respond to attempts to affect them.

>We can tune our responses finer and finer, and reach ever
>higher qualities of response and perception.

I'm not sure how you mean this, but it sounds like one of the concepts we're trying to destroy. "Response" is the conceptual opposite of "control." It implies a blind reaction to an input, and carries overtones of jab-and-jump psychology. If we can "tune responses" we must be sensing something that depends directly on them; all we can actually tune is the sensory consequence, for a pure response (of your own) is not itself experientiable. Whether you intended this or not, this way of speaking about what is learned encourages the old idea that perceptual inputs cause motor outputs -- the very idea that allows people to blame what they see for what they do. This is one of the many basic conventional concepts that stand arrayed against PCT.

>Control over control is self-control, perceiving your own
>perceptions is self-perception, consciousness.

That sounds nice, but I don't believe it. If you diagram a system that senses the stability of a control system and adjusts parameters to control stability, you do not have a system controlling itself: you have a system controlling something about a different system. If you perceive your own perceptions, one subsystem is perceiving the perceptions originating in another subsystem, and most likely interpreting them in a different way. The moment you say "I am thinking," you have denied the statement: the system that is aware of the thinking is not thinking, it is making a statement about a system that is thinking. The "I" of which you speak is never the "I" that speaks.

The only way to make sense of self-reflexive ideas is to treat a person as if that person were solid, like a potato: only the whole person perceives and acts. Only in that way can one say that the referring self is the self referred to. That view is contrary to the modeling approach, in which we try to understand the whole in terms of interactions among its subsystems.

I said:

>>In other words, I could ask the question, "How is competent
>>adult human behavior organized when its organization is not
>>being changed?" This is what HPCT is about.

And you said:

>Whereas my focus is more on how human behavior can become even
>more competent, i.e. more on learning (and evolution as a kind of learning).

I think that my goals have to be reached before yours can be reached (at which point yours would be mine, too). Before you can study how to make the human being more competent, you have to have a way to measure its competence. Psychology has fallen down on that job; nothing it says about behavior can be taken as a clear fact, because its factual statements are riddled with important exceptions and counterexamples. We need a highly predictive and accurately descriptive model of how behavior works when it is not changing. Only then can we measure change in any reliable way, and know whether our attempts to improve competence are having any effect, good or bad.

>One can, at our level of discourse, see a feedback amplifier or
>some such device in two very different ways. The first is as a
>device that transforms an input (voltage, current or power)
>into an output (voltage, current or power). The feedback is not
>really relevant here. The (voltage, current or power) gain may
>be any value; both gains and losses can be similarly realized.
>The second way is to see the device as a power modulator, where
>the input modulates the transfer of power from its power supply
>to its output. The feedback is not relevant here either.

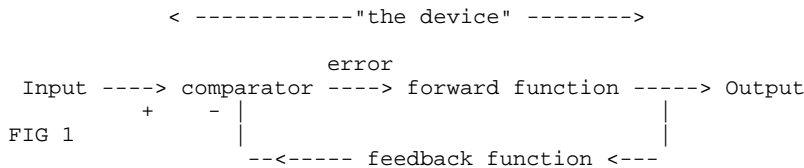
Neither of these concepts is especially relevant to PCT. The first is the usual idea in which the "input" (meaning, really, the reference input) is confused with sensory inputs, so it appears that an input from the environment is causing an output into the environment. In living control systems the reference input does not come from the environment, but from higher systems.

The second applies primarily to the output function. It's not often necessary to draw the power supply of a control system; the behavior of the system is quite insensitive to changes in the power supply, as you know.

>Then, internal in the device, we see the LOOP GAIN, which has
>nothing to do with power at all, but is the gain in the loop
>that the SIGNAL travels. A decent control system has a loop
>gain much greater than one, although a loop gain of less than
>one is not unthinkable. In the latter case we might not want to
>consider the system a CONTROL system.

Loop gain, in PCT, is not "internal to the device." The relevant closed-loop path passes through the environment. I suspect that you haven't yet understood just how the PCT diagram differs from the standard engineering one. Actually, I'm working on this subject right now for the joint paper, but I suppose it won't dilute the writing too much to show you three diagrams that will be in the paper, and preview the discussion. It's important to understand exactly what we have done with the standard diagram.

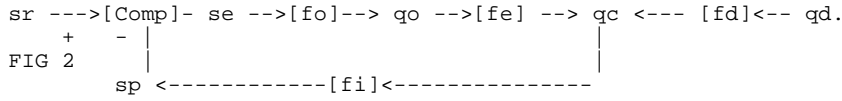
Fig. 1 is what the behavioral sciences have taken from engineering, for the most part:



If you draw a circle around everything under "the device" above, that is everything but the Input and the Output, you have the first case you describe above, as well as your third statement in which feedback is "internal to the device."

Now let's just add a few details without altering the overall appearance: I'll have to use some abbreviations to fit it all in.

< -- "device"----><-- environment ----- >



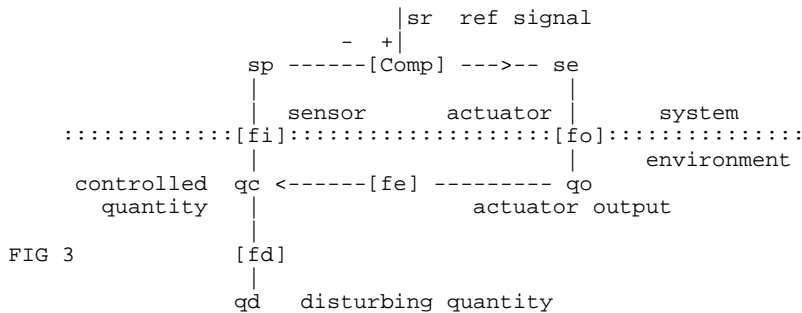
Here sr = reference signal, comp = comparator, se = error signal, [fo] = forward or output function, qo = output quantity (the immediate effector output), [fe] = environmental function (which transforms the effector output into an effect on the controlled quantity), qc = controlled quantity (the physical quantity actually under control), [fi] = input function (which includes the sensory receptors and any computations that immediately follow), and finally, sp = perceptual signal, the internal analog of the controlled quantity. [fd] and qd provide for representing independent disturbances and their influence on the controlled quantity.

I hope you agree that the organization of this model is identical to that of Fig. 1, except for the explicit inclusion of a possible disturbance and insertion of some stages implicit in Fig. 1. I have relabelled the "input" as the reference signal, which does no violence to engineering custom, and the "output" as the controlled quantity, which is also an acceptable alternative in engineering parlance.

However, I have expanded the details at the system's output a bit. I have distinguished between the immediate effector output and the controlled quantity, and introduced an environmental function expressing the dependence of the controlled quantity on the effector output. An example would be a control system that controls a shaft's angular position. The controlled quantity qc is the angle at the end of the shaft where the load or workpiece is; the output quantity qo is the torque output of the driving motor. The intervening [fe] expresses the way torque is converted into shaft position, given the way load resistance depends on angular position (which could be assigned to the disturbing branch).

This separation is always important in detailed control-system design, but especially so when the effector is coupled to the controlled quantity loosely or through complex intervening processes. Then we clearly would expect the effector output to be changing far more than the controlled quantity is changing. Even when we're just talking about output torque and controlled shaft position, the system may have to vary the output torque radically, even changing direction, in order to maintain a specified shaft position, as twisting disturbances are applied one way and the other to the controlled quantity at the end of the shaft. I know you know all this; I'm just making the description complete.

Now consider Fig. 3:



This is organized exactly like Fig. 2. It is simply rearranged. It is actually just like Fig 1., with details added. The plane of separation between system and environment, however, is not the one suggested by the first diagram. To locate it in the first diagram, you would have to draw a line like this:

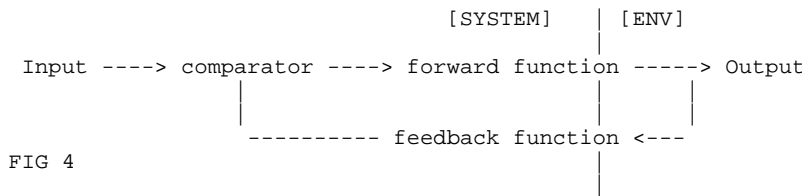


FIG 4

This distinction means little in engineering, but in PCT it is essential for getting the correspondences between the engineering diagram and the physical organism right. In Fig. 3, the horizontal line separates the nervous system of the organism from all that is not nervous system. Sensors and actuators lie on the boundary. Notice that in Fig. 3, there is no chance of mistaking the reference input for a sensory input. The reference signal comes from higher up, inside the behaving system. The sensory inputs are strictly associated with the feedback path through the environment. In living control systems, unlike artificial ones, the reference signals are not accessible from outside the behaving system.

In those control systems that have been traced out in human beings and animals, Fig. 3 corresponds closely to the sensors, intervening cells, and output paths. The reference input corresponds physically and functionally with what are traditionally called "command" signals, signals which carry outputs from systems higher in the brain. Those command signals have been thought of traditionally as carrying commands to the muscles, causing them to contract (the feedback paths are ignored even though they are mentioned in a sort of puzzled way). The control-system diagram, with parameters filled in to make it fit real behavior, shows that the so-called command signal is really a reference signal. Its primary effect is to specify the level to which the perceptual signal will be brought. The actual outputs could be in any state, depending on what disturbances happen to be acting on the controlled variable. The closed loop system varies the output in any way required to make the controlled quantity, and thus the perceptual signal, match the reference signal. It does this without any instructions from the reference signal.

Fig. 3 was drawn as it is with full knowledge of the engineering diagram of Fig 1, for a specific purpose. Almost without exception, behavioral scientists have interpreted the "input" of Fig. 1 to mean "sensory input." When that is done, Fig. 1 becomes nothing but a stimulus-response diagram with an internal feedback loop having no obvious function. Wiener said it "reduced the dependence of the output on the load." This has been taken as cybernetic justification for the old model in which sensory inputs cause behavioral outputs.

The reason for emphasizing the distinction between the actual effector output and the controlled quantity (usually absorbed into a single equation in engineering) is to show the difference between the physical action of the system and the sometimes remote outcome of that action which is actually under control. When we see the controlled variable separated from the effector output, we can much more easily understand that the visible behavior of an organism is really just its actuator output, while the focus of the control action is an effect of that actuator output -- a joint effect, because disturbances act on the controlled quantity, too. Thus, with this diagram, we can point out the specific difference between what we see an organism doing and the controlled outcome of those variable actions.

You might think that this rearrangement would be easy to explain to real control engineers, but that has not always proven to be the case. Control engineers get just as set in their ways as psychologists. Long experience only seems to make matters worse. One old control engineer with whom we went around and around for six months on the net ended by saying that he saw what we meant, but he just couldn't get used to talking about a controlled variable as associated with input. So he bade us farewell, wishing us luck in a gentlemanly way. Of course a much younger one, encountered in a different venue, thought this was terrific, and adopted the PCT model for teaching control theory to graduate students. I guess there are a couple of control engineers on this net who have seen the light. I don't know whether you have or not; it's hard to tell from what you say.

All this volume of output was necessary to explain why I object to your statement

>Then, internal in the device, we see the LOOP GAIN, which has nothing to
>do with power at all, but is the gain in the loop that the SIGNAL travels.

The loop gain in Fig. 3 is the product of the partial derivatives of [fi], [Comp], [fo], and [fe]. That is the gain that determines how tight the control will be. It specifically must include the path through the external environmental feedback function. The control loop in PCT is NEVER "internal to the system." It ALWAYS passes through the environment, no matter what level of control is involved. This is what makes PCT models testable. There may in fact be closed loops totally above the line in Fig. 3, but in behavioral experiments we can do nothing with them. Their effects will simply be absorbed into the basic model of the control system.

And I think that's quite enough of my Sunday and your time to spend on one post.

Best, Bill P.

Date: Sun Jun 20, 1993 4:59 pm PST
Subject: Re: no witnesses - RKC

[From Bob Clark (930620.08:45 pm EDT)] Bruce Nevin (930614.06:57:04)

You refer to the Closed Loop item 3.2:49 about left-handed ball-throwing. You speak of my "analytically replicating the familiar sequence." To me, there was no analysis -- rather I was carefully observing the detailed sequence of movements, including the timing, putting them together to form an identifiable procedure.

Your remark about witnesses and "self consciousness" as usually conceived, is generally pertinent -- but in this case, I had to really concentrate strongly on what I was doing. Any kind of distraction would have, I think, made it impossible.

Try something like this yourself. A simple version is the old bit about rubbing your stomach while patting your head - and vice versa.

Regards, Bob Clark

Date: Sun Jun 20, 1993 6:07 pm PST
Subject: MORE SOC CONT SYST - RKC

[From Bob Clark (930620.09:45 pm EDT)] Rick Marken (930614.1200)

Your position is:

>opposed to those of Bruce Nevin and Bob Clark (who seem to believe
>that social control systems either exist or are useful metaphors)

In earlier posts my remarks were taken to imply that social control systems have some kind of physical existence external to the individuals composing the systems. I think subsequent posts have made it clear that I regard all such systems (including, for that matter, HPCT, etc) as existing only within the "minds" of individuals, and written records reflecting such existence.

Certainly, as you remark, and quote from, Bill, people use these concepts for a variety of purposes. But they can only do these things if both they and those they try to affect have somewhat similar concepts of specific social control systems.

I do disagree with your suggestion that social control systems are:

>just a side effect of the operation of individual control systems.

To the contrary, social control systems are themselves examples of functioning hierarchical control systems. They exist, and function, purely because they exist in the minds of those using them. They may not be very complete, they may not be very effective, they may be loaded with conflicts -- but they exist, at least in the sense that their existence and operation make a difference in the lives of the people involved.

Regards, Bob Clark

Date: Sun Jun 20, 1993 6:25 pm PST
Subject: MORE CELLS/PEOPLE - RKC

[From Bob Clark (930620.1010 pm EDT)] Bill Powers (930615.0900 MDT);

In general, Bill, I agree with your outline of the operation of the cells in a body. The only question in my mind is the extent to which certain cells are directly involved in a feedback control system. Certainly some are, but some are merely passive entities, "doing their thing." This seems to be very much the case with kidney cells, serving as passive filters subject to the "laws of osmosis" and the like.

Over-all, they do not compose a "voluntary" control system -- nor set of such systems. Indeed, the bulk of the physiological systems, including those that function as control systems, seem to be genetically determined.

And this brings up still another topic: which are "built in" and which are "learned?"

Bill Powers (930616.0900 MDT)

After discussing in some detail some possible control relations among cells, media, etc, addressing Oded Maler, you state:

>I think that these apparent collective properties simply emerge from
>the interaction of individual control systems, as measures like
>entropy and pressure emerge from the interaction of individual
>molecules. I think they live in a conceptual space, which exists
>inside your head. You're creating an allegory or a metaphor, not a
>literal model of how things work. You are reading something into
>the world that actually comes from your own imagination. I say that
>there is absolutely no evidence for collective perceptions,
>collective reference signals, collective error signals, or
>collective mechanisms for action.

Nicely said, Bill. I think you thought at first that I conceived some kind of external, physically existing form of social control systems. I think recent posts have now made our respective positions clear -- that they exist exclusively within some portion, or portions, of the hierarchy.

Bill Powers (930616.1230 MDT)

More of the same. Very well said. Yet I fear that I must point out that producing "a stable phenomenon called pressure" in a gas is a bit more complex than just the combinations of their properties and interactions. We need energy, or temperature and gravity etc. The general gas law has it complications.

We probably should develop a group of examples of physical phenomena to illustrate and clarify some of these situations.

Regards, Bob Clark

Date: Mon Jun 21, 1993 2:30 am PST
From: Hortideas Publishing / MCI ID: 497-2767
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Comments

From Greg Williams (930621 - direct)

Comments on your post dated (930615 18.10):

> a) Psychological (cognitive) science says:
> "(Internal) Commands determine behavior"

You already know that I think there is great potential for misunderstanding in this sort of reduction of a complex situation to a single assertion.

> a) Psychological science says:
> Reinforcement selects behavior

Ditto.

>Perhaps a page of 8-12 statements arranged as columns a) and b)
>will allow the reader to identify and agree with the portrayal of
>contemporary wisdom, and see the influence of psychological
>theories on "common sense understanding" and "management practice".

I can sympathize with your need to provide succinct summaries contrasting PCT and other approaches, but I worry about oversimplifications and distortions of the situations, which, as I said above, are not simple. There just aren't monolithic camps out there with the followers in each all singing in harmony. At least, I urge a caveat saying that your summary statements are caricatures drawn for emphasis.

>Each of these accepted "truths" listed above is DEMONSTRABLY FALSE.

>It is very easy to demonstrate that they are false. All that is
>required is to recognize and understand the phenomenon of control.

Oversimplified, they are easily shown to be "false." But I believe it is highly insulting to the reasonable and intelligent people involved with the COMPLEXITIES hidden by the simplifications to act as if PCT ideas have shown what they have been doing worthless. (I'm not optimistic that the core group of PCTers will learn this soon -- and they will continue to constitute just about the ENTIRE group, because of that.)

> c) Perceptual Control Theory says:
> People will always do what they can to satisfy
> internal wants. Reinforcement as such is an
> illusion. Rewards actually serve to reduce
> behavior if they help the person satisfy the
> internal want.

Skinner would have NO problem with the last sentence. The PCTers' MYTH about what reinforcement is ("more reinforcement, more behavior") is NOT what Skinner said. Skinner's main use of reinforcement was in deciding what NEW control structure would obtain following learning.

I don't have any major problems with the rest of your post.

Best, Greg

Date: Mon Jun 21, 1993 2:40 am PST
Subject: No "myth" of stressful info

From Greg Williams (930621) Rick Marken (930618.0900)

>I would hope that people could cooperate to achieve their goals
>in a business. I know that business people often violate this
>cooperative relationship, setting up contingencies that can hurt people.

Why doesn't "setting up contingencies that can hurt [certain] people" amount to the same thing as providing "stimuli" (you can call them "disturbances") which "cause" certain people to experience "stress"? The common meaning of "cause" (with a little attention to rough edges; see, for example, Mackie's THE CEMENT OF THE UNIVERSE for an extended technical discussion) is A, without which B would not follow. In this sense (and adding the further definition of information as "a difference which makes a difference" IN A PARTICULAR SITUATION, there is no "myth" of stress-inducing information. (Of course, in a broader sense, "it's all myth." :->) If it is indeed the case (no fraud) that a boss has provided disturbances to an employee and the employee's control becomes dysfunctional (and the employee feels "stressed"), then, provided that it is reasonable to believe that the dysfunction would not have occurred in the absence of the boss-provided disturbances, I believe it is in full agreement with the normal usage of "causality" to say that "stressful [for this particular employee] information has caused the stress [dysfunctional control itself or a "side effect" of same as perceived by the employee]. It doesn't make it a "myth" that what is stressful information for one employee is not necessarily stressful information for another employee. It doesn't make it a "myth" that there are many other factors in the situation which also could be identified as causal. It doesn't make it a "myth" that the boss cannot predict in advance whether or how an employee's dysfunctional control will manifest itself.

I know that claiming that PCT brands common notions as "mythical" is tempting for PCT revolutionaries. But I don't see any good coming from doing it in this case. It just further polarizes PCTers from everybody else.

>But I also know that employees take advantage
>of business people by claiming to have been caused stress when,
>in fact, the stress, if they have it, is not caused.

Again, the issue of the "myth" of disturbances sometimes causing stress is different from issues of deception, and I can sympathize with your wife regarding the latter problem.

>The assumption
>that the environment causes stress is a problem, not merely because
>it is false, but because it allows workers to perpetuate a fraud; all

>they have to do is point to a "stress inducing" stimulus to which they
>were exposed and the court is required to assume that the result
>was stress. This is what I object to; using a fake psychological
>model to make fraud easy.

I contend that, as per my above comments, the environment of at least some people can cause those people to feel stressed. This is not a "fake" psychological model any more than PCT models are "fake"; it is BASED on PCT models. If the courts too easily accept shoddy evidence for a person feeling stressed and/or too readily classify stimuli as GENERALLY "stress-inducing" for many or most or all individuals, then there are definitely problems -- but those problems aren't due to a "fake" model.

>I am not against workman's comp and there certainly are "stimuli"
>out there that can cause bad results; if a heavy piece of equipment
>drops on your foot then it will cause damage.

You support my contention about the common use of "cause" with this example.

>I think a worker should
>unquestionably be compensated for the effects of this kind of cause.
>But exposure to "stress information" should not be a basis for
>assuming that the behavioral result was stress -- because there
>is no such thing as stressful information.

There is no such thing as the SAME stressful information to EVERYONE, but there is PARTICULAR stressful information for (virtually?) ANYONE.

>Your example of lying to me and causing a heart attack is cute; I would
>be annoyed but I probably would not think of suing you; I'm just not
>a sueing kinda guy. But that is a situation where I might give the
>"victim" of your "stressful stimulus" a sympathetic hearing -- at least
>because you did apply that "stimulus" intentionally.

I think you would be annoyed because you realize that if I hadn't lied, then you wouldn't have had the heart attack. In common parlance, my presentation of certain information CAUSED you to have a heart attack. That certain information was STRESSFUL for you.

>I'm sure Linda
>would have no problem accepting a person's stress claim if it were
>the result of maliciousness on the part of someone else; I still think
>the stress (if it exists) was not caused by the "stressful stimuli"
>applied by the malicious person but that person did, intentionally,
>create the circumstances where reorganization (and the attendant
>stress) was possibly required.

Well, then, you simply sound confused about what "causal" means.

>The problem with workman's comp is that you can make a stress claim
>just by showing that you were exposed to stressful information; it is
>not necessary to show that this information was presented to you
>maliciously.

One more time: I see problems in the courts. But they aren't due to a "myth."

As ever, Greg

Date: Mon Jun 21, 1993 5:56 am PST
Subject: repetition vs. imitation

[From: Bruce Nevin (Mon 930621 09:44:18 EDT)] Bill Powers (930613.2200 MDT) --

> Here's my problem. If I present you with a picture of a grape and
> a picture of an elephant, you can distinguish between them; the
> perceptual input that allows you to perceive the grape does not
> respond to the elephant, and vice versa. So I have established
> that you have two perceptual functions, one for each picture
> (sort of). Am I then justified in saying that you are perceiving
> something called "contrast" between the grape and the elephant?

> Or is the notion of contrast an interpretation by the observer, me?

The analogy obscures the point. The image of an elephant or of a grape is not a behavioral output of another human being that you might likewise produce among your own behavioral outputs. The hypothesis was that you have to produce a phoneme or a word (really or in imagination) in order to perceive one. The phoneme /p/ cannot be said to exist as an object in Boss Reality, independent of our perceptions, in the way that we say that an elephant or a grape exists independent of our perceptions.

The notion of contrast is of course an interpretation by the observer; the phoneme is an interpretation by the observer. The utterance itself, as utterance rather than as mere sounds, is an interpretation by the observer. The contrasts are not only "in" the utterance, they constitute the utterance, they are what make it an utterance rather than mere sounds. Phonemes, words, utterances do not exist in "Boss Reality" as we suppose that elephants and grapes do; they only exist in people's perceptions. Likewise phonemic cues, distinctive features, and so on. They only exist because people, anticipating speech, are looking for satisfiers of a previously known relationship of contrast. (I'm only talking about the perception problem here, not the learning problem. I'll try to avoid getting them muddled again.)

> >The segments are relevant because they represent the contrasts
> >between utterances and locate the points of contrast within
> >utterances.

>

> But isn't it really that the listener's perceptual functions make
> a distinction, rather than that there are objective contrasts in
> the sentences?

Yes, it is the case that the listener's perceptual functions make a distinction, rather than that there are objective contrasts in the sounds that he hears and the articulatory gestures that he perceives. It is also the case that these contrasts, products of the listener's (and speaker's) perceptual functions, are really in the sentences, which are equally products of the same people's perceptual functions.

> But it still is defining "segments" in terms of human perceptual
> functions, not the other way around.

Yes. The thrust of virtually all linguists' work, ignoring Harris, has been to show how segments and contrasts are objective, in Boss Reality, either as physical entities or attributes or indirectly, by virtue of "feature detectors" hard-wired in the human genome for features that are universal across all humanly possible languages. However, these detectors are not for specific features of sound or gesture, rather, for dimensions along which sounds/gestures may be categorized into opposing members of a contrast, with the absolute values used to be determined by the child's experience. Learning, on this view, is parameter setting.

> What grates on my tender
> sensibilities is speaking of the contrast as if it existed in the
> utterance, or pairs of utterances.

Assuming some sort of correspondence of perceptions to Boss Reality, an elephant or a grape exists independently of whether someone perceives them or not, and contrast is a relation between perceptions of an elephant and perceptions of a grape, but not a relation between elephants and grapes. However, an utterance does NOT exist independently of whether someone perceives it or not. In Boss Reality you have only sounds and articulatory movements. The utterance, as such, exists ONLY as controlled perceptions of a linguistically competent control system. An utterance is constituted of contrasts, that is, of differences that make a difference.

(To clear away a possible misunderstanding: the term "utterance" is used instead of "word" or "sentence" because much that is said between pauses in a given language is not a single word, and is more than or less than a sentence. In ordinary, non-technical usage an utterance might also be a non-linguistic sound: "Heathcliff uttered a wierd cry." However, non-language sounds are not utterances in the sense intended. If a tree falls and no one perceives it, the tree (we presume) is still there, in Boss Reality. If a tape recording of the Gettysburgh Address is played in a forest and no one is there, it is just sounds, no utterances at all. The sounds are really there; the utterances are not.)

An utterance is a production in a language which is recognizable by other speakers of the language and repeatable by other speakers such that their repetitions are likewise recognizable as repetitions. As such, utterances literally do not exist as utterances unless they are perceived as such. This is why the contrasts (by which utterances are perceived as utterances) exist in the utterances. If I played a tape of Tagalog or Rwanda conversations a listener who did not know the language on the recording would not be able to repeat any of it (however well she could imitate), would not recognize repetitions, and certainly would not perceive the meanings of the utterances. The listener might suppose that she heard a word repeated in the recording, when in fact those portions of the recording differed by some feature that was distinctive in that language but not in English; or might suppose that she heard different words when in fact the discriminable differences are not contrastive in the recorded language, even though they are so in English.

I have tried to keep distinct two aspects of the problem: the learning or establishing of the phonemes vs. the recognition of words (setting aside for now the limitations of the terms "phonemes" and "words", and ignoring also the evolutionary aspect). We must be careful about this, to avoid getting muddled.

Harris's use of the Test concerns the learning problem; your response concerns the recognition problem. You say that neural nets can find the invariants that are phonemes. In fact, as NNs are "trained" they are given the contrasts, that is, the fact that some inputs are repetition sets whose differences make no difference, and other inputs are not repetitions. The result is that they come up with representations of those contrasts, which they then use (with less than perfect success) to recognize other repetitions. Coming up with representations of contrasts is the learning problem; recognizing other repetitions, given a representation of the contrasts, is the recognition problem. The representations (the "phonemes" or "phoneme recognizers") that one NN or one person comes up with in the learning process may be different from those that another comes up with in a comparable learning process. That does not matter, so long as they are representations of the same contrasts between words (between utterances). The contrasts are the invariants (given a partially shared vocabulary), not the representations by sounds and gestures. It is this that accounts for our ready accommodation to differences in pronunciation, so long as the differences are systemic. If we had a speech synthesis device that achieved naturalistic speech output in various dialects, we could understand it, after a little exposure, no matter which of a range of dialects it was producing. If, however, it was shifting from one dialect to another unpredictably during the course of sentences and words, we could not.

The learning problem is solvable only if the learner is told which utterances are the same and which contrast. This often takes the form of the hearer interpreting utterance XYZ as an instance of utterance XWZ. In other words, the significant pronunciation-differences cannot be learned, as opposed to the pronunciation-differences that are not significant, if the learner has no idea of the differences of meaning between utterances. The least datum about differences of meaning (to which all such differences are reducible) is this: are the two productions repetitions, or do they contrast? Conversely, if pairs of utterances that differ by pronunciation-difference X are consistently different from one another in meaning, they contrast, and X is a significant difference; that is, X is phonemic; that is, pronunciation-difference X is a representation of a phonemic contrast between those pairs of utterances. Kids hear lots and lots of evidence as to what differences make a difference and what differences don't. People misunderstand them (or one another) and correct the misunderstandings, people correct themselves in the child's hearing, people repeat back the words they heard the child say, and words that they heard one another say, and so on.

> If I have two experiences and they actually involve
 > only one perception, I classify this situation as "same." If
 > there is more than one perception, I classify it as "different."
 > But before I can perceive which category of situation it is, I
 > must know already whether one or two perceptions were involved: I
 > must know that both experiences came from this perceiver, or
 > that one came from this and the other from that. If this sort
 > of discrimination hasn't already been made, perceptually, there
 > is no way to decide on the category "same" or "different."

> It's on this basis that I maintain that "contrast" isn't an
 > explanatory term, but only a descriptive one. We don't perceive
 > that two things are the same or different because they ARE either
 > the same or different. We can only make that judgment after the

> discrimination has been made at lower levels. If we perceive two
 > things via two input functions, we conclude that they are
 > different; if both perceptions come from the same input function
 > we call them the same.

At the lower level, we discriminate differences that may be either distinctive or insignificant at the higher level.

The pronunciation of a given phoneme varies from one context to another. The perception cues for distinguishing between similar phonemes at the corresponding places in similar words may be produced with variable gain. For example, the medial consonant in paining/painting/pating (or plating if we can't accept "to pate" as some kind of nonce word). Or consider the flapped r in the Spanish pronunciation of Paraguay. This occurs in American English also. Say "pot of coffee" at an ordinary conversational rate (pot'ocoffee). It occurs in the middle of "latter" as in "he picked up the former and put it on the latter". Now, suppose we have a tool called a former. The same sound/gesture occurs in "ladder", as in "he picked up the former and put it on the ladder". Same input function. In one case, we hear d, in the other we hear t. If the contrast is important, we distinguish t from d more clearly, pronouncing "laTter" or "laDder". (This is the norm in most British and many Canadian dialects.)

With low gain such words may become indistinguishable phonetically. Semantic and syntactic aspects of the context tell us which word to expect (as with outright homonyms like beet/beat, or see/sea/[holy] see/C), and indeed such predictability is a condition for lowering gain on pronunciation of a given word, with the reduction system as a further extension of the same principle. In all of this, it is the contrast between words that is controlled.

When two people converse, they speak more or less different versions of the language. These may be called different dialects if extreme enough, on a continuum out to what we call different but related languages if they are not mutually intelligible (if most of those who speak the two versions cannot converse, each speaking her own version to the other). Assuming their talk is mutually intelligible, they interconvert between the phonemic distinctions made by one and those made by the other. My "kayo" (K.O.) sounds markedly similar to my brother's "cow". (He has lived for many years in Georgia, and we grew up in central Florida, where he identified as a local and I identified as a college-bound kid disaffected with the redneck community around me. This is in reference to prior discussions of Labov's work on social dialect.) Does this mean that people construct some sort of table lookup? Is it not much more parsimonious to suppose that they are controlling the contrasts between words by variable means at the level of pronunciation?

>Second pass:

Actually, this latter part was an earlier pass, and I just gave up on trying to integrate into the "first pass" those points that I wanted to keep.

> what is perceived is a syllable or a word, not a contrast.

What is a syllable? What is a word? You are begging the question here. Syllables and words do not exist in the sounds of speech. This is shown by the fact that speakers of different languages parse the same sounds differently into syllables and words.

> when you speak of contrasts, which contrast do you mean?
 > Every possible discriminable segment differs from all other
 > discriminable segments, simultaneously. If you have 3000
 > discriminable syllables or words in the working set, you have
 > 4,498,500 dyadic contrasts in the set. Does each perceptual
 > function have to search through 2999 contrasts to decide [etc.]

Before children learn phonemic contrasts as a way of representing word-contrasts, their vocabulary is severely limited for just this reason. After they learn to represent word-contrasts by phonemic contrasts, vocabulary takes off. If there are 50 phonemes the search space is feasible. It is greatly reduced by other constraints, such as those we express in terms of what may constitute a syllable. Some of these constraints differ from one language to another, and are learned from "the way people do it" here; others are universal, and reflect physical (physiological, acoustic) properties of the environment.

(Tom Bourbon (930529.0310))--

I asked how would one model repeating as opposed to imitating. What I am interested in, of course, is the difference between repeating a word and imitating a person's pronunciation of sounds that we did not recognize as a word, or where we choose to focus on the pronunciation independently of its recognizability (and repeatability) as a word. We can't model that, so we start with something simpler. For imitation, Tom says:

> There must be something to imitate.

and introduces a triangular waveform for tracking, with a sketch of how it is generated and how it is tracked.

> / \ \ / \ \ / \ \ / \ \ , and so on.

> This is a standard, undisturbed pursuit tracking task, which PCT models with great accuracy. The participant can make the cursor accurately track the target, which is to say, the positions of the cursor imitate those of the target (incidentally, so, too, does the waveform of the participant's hand movements -- the person's actions).

Then for repetition, Tom says:

> For Repeating, there must be something to imitate.

Right away, we're in trouble. I've been trying to show how repetition is different from imitation. Tom is saying repetition is simply imitation of a memory:

> Let that be the person's remembered positions of the cursor as a function of time during the Repeating task. (The target is not on the screen.) Now the model step for the person changes, because the source of the reference signal is different, but the cursor is still determined as it was before.

> For either Repeating or Imitating, the same actions will occur. So, too, will the same positions of the cursor as a function of time.

To get at what I am looking for, we have to introduce some variation. Suppose person B's notion of the pattern is this:

|_ \ | \ | \ | \

Person A's notion is roughly as above.

/ \ / \ / \ / \

Suppose A and B recognize their renditions as having the same intention--the same meaning. Let's say that

A: |_ \ | \ | \ | \ B: / \ / \ / \ / \

A recognizes B's production as a repetition (but not an imitation) of A's, and B recognizes A's production as a repetition of B's.

Now, suppose A and B recognize the following pair of productions as equivalent (as each one's repetition of the other's production):

A: |___ \ |___ \ |___ \ B: / --- \ / --- \ / --- \

And suppose they recognize the following as repetitions:

A: B:

Date: Mon Jun 21, 1993 7:25 am PST
 Subject: Re: By the rivers of PCT

[Martin Taylor 930621 11:00] Rick Marken 930618.1900)

A *slight* inconsistency here, quoted without comment:

(1)

>Surely you don't consider the feedback function
 >to be part of the control system, do you?

(2)

>The fact is that you guys just have a wrong conception of how control
 >works -- its a circle, not a chain.

Do you see the contradiction? We keep trying to tell you that we NEED the circle, and you keep trying to tie us up with the chain. You deny us in your challenges the opportunity to use the necessary circle, and then tell us our analysis won't work because we think we don't need the circle.

=====

I think it was in a very early posting on this information theory bit that I pointed out that the fact of control demanded that the uncertainty about the world remained stable over time. That, in iteself, demonstrated the need for the completed feedback circuit. It would be nice if you could get OUT of your head the notion that an information-theoretic approach deals with an open-ended system.

Bah! Martin

Date: Mon Jun 21, 1993 7:42 am PST
 From: Marken
 MBX: Marken@courier4.aero.org
 TO: * Dag Forssell / MCI ID: 474-2580
 Subject: Personnel discussion

Hi Dag (and Chris)

I told Linda about your interest in talking to her about her new work. She is very reluctant to talk yet because she has just started in personnel and knows virtually nothing about it. This is a completely new field for her-- she was doing regular banking type work until now. All of her ideas about workman's comp, for example, are based completely on first impressions. I think she feels a little insecure about this new position right now; she is not really a mucky muck (she only seems so to me, given my lowly engineer status here).

I think it would be better to just talk about this when we are together at the CSG meeting. Linda was saying that she probably wouldn't feel comfortable taking about this stuff for about 6 months. But I bet she could be encouraged to express some opinions if she feels relaxed at the meeting.

Talk to you soon.

Great post on Power Loss by the way.

Best Rick

Date: Mon Jun 21, 1993 7:46 am PST
 Subject: Re: Power gain, power loss.

[Martin Taylor 930621 11:10] (Hans Blom 930619)

Misunderstandings sometimes are better resolved by a non-combatant.

Bill Powers mentioned power gain as an essential element of a control system. Hans countered with the power LOSS required by an engineer trying to ensure stability in a high-voltage electricity delivery system by observing the images on an oscilloscope.

Bill was talking about the power gain between the error signal and the resulting effector operations that directly affect the CEV. Hans is talking about the power loss between the CEV and the perceptual signal.

It seems to me that just as a power gain is an essential element of the outflow side of a control system, so a corresponding power loss is an essential element of the inflow (perceptual) side. The perceiving of the state of a CEV should not contribute as a disturbance any more than it must (Heisenberg showed that it must, to some extent). Perceiving that state of a CEV should be as power-decoupled from the CEV as it possibly can be. On the other hand, the output power of the control system wants to have maximum effect on the CEV, or as tight coupling and as high power gain as is feasible, given the information limitations on the perceptual side. Output power gain and input power loss are intimately coupled requirements for a control system.

On models, I tend to side with Hans. It is part of the whole information argument. The more information is available within the control system, the less is to be acquired from the CEV through the perceptual apparatus, and the better control can be.

=====

(Bill Powers (930618.1930))

>It IS an important point. Martin has been implying that the
>resolution of the input function depends on its RMS noise level,
>and I completely disagree with this concept of resolution of the
>measuring apparatus. If the noise is 10% RMS of the range,
>Martin's interpretation would be that there are only 10 possible
>values of the perceptual signal with a range of 10 units, 0
>through 9. This would make the probability of any one value of
>perceptual signal 10 percent. In fact, however, the perceptual
>signal can have a magnitude between 0 and 9 with a precision
>that depends on how long you observe it: if you observe it 10
>times as long, it has 3 times the precision if the noise is
>Poisson-distributed.

Gaussian, actually. I'm glad you have got this point. It is the heart of my Gain-bandwidth computation. If you didn't understand that earlier, I'm not surprised you didn't follow the analysis.

I'm sure you sometimes feel a little frustration at people telling you that you said things that are the opposite of what you tried to tell them. So am I, but at least on this point you have come around to the "correct" view.

>And however long you observe, whether for a
>short time or a long time, the RMS noise does not predict the
>probability of a specific measure, for a specific measure can
>have any value in the real number range between 0 and 9. The
>resolution is infinite, although the precision and
>repeatability are not.

The resolution IS infinite under certain conditions, which include infinite observation time and *a priori* certainty that the thing observed does not change over the observation interval. If you cannot be assured beforehand that the thing observed will be unchanging until the end of time, the observation interval can only be as long as it is known to be effectively stable. That's the point of the Nyquist sampling theorem. That's why if there is negative gain in the loop greater than unity, the perceptual sampling rate must be greater than the Nyquist rate for the controlled part of the disturbance (see Friday's postings for the interpretation of that phrase).

Martin

Date: Mon Jun 21, 1993 9:48 am PST
Subject: To loop or not to loop

From Greg Williams (930621) Martin Taylor 930621 11:00

commenting on >>(Rick Marken 930618.1900)

>A *slight* inconsistency here, quoted without comment:

>(1)

>>Surely you don't consider the feedback function
>>to be part of the control system, do you?

>(2)

>>The fact is that you guys just have a wrong conception of how control
>>works -- its a circle, not a chain.

>Do you see the contradiction? We keep trying to tell you that we NEED the
>circle, and you keep trying to tie us up with the chain. You deny us in
>your challenges the opportunity to use the necessary circle, and then
>tell us our analysis won't work because we think we don't need the circle.

Good for you! I would have brought this up, too, but I already know where things would probably head: toward bringing in a conception of the relative "importance" of some components of the loops relative to other components. I've been down that path too many times before with Bill and Rick. When I have told them that assigning such "importances" is not a part of PCT modeling, but an ethically (pragmatically) based adjunct to PCT models, the conversation breaks down. I'll sit this round out... except that my problem with Rick's concept of the "myth" of environmental causation does relate to the relative importance stuff.

Lots of luck, Greg

Date: Mon Jun 21, 1993 9:49 am PST
Subject: Group control processes;stress;contrast

[From Bill Powers (930621.0840 MDT)] Bob Clark (930620.2015 EDT) --

From me:

>>I guess that as an engineer, I am pained when people cooperate
>>in an attempt to achieve common goals using means that will
>>almost certainly, in the end, frustrate the attempt.

You:

>I too, Bill, am concerned when tasks are not well done -- but I
>still struggle to avoid having my own "perfectionism" interfere
>with moving ahead.

"Tasks not well done" does not mean the same as "Tasks doomed to failure," which is what I was talking about. I don't mind joining in tasks that are being done suboptimally. It's the other kind I have a hard time getting enthusiastic about.

We're close indeed about social control systems. I've been thinking lately about how social cooperation toward controlling variables of mutual interest (how's that?) is actually carried out. I was musing about a few halcyon weeks in Boot Camp (now you have me doing it: I mean boot camp) in the Navy when I suddenly became a 17-year-old platoon leader with about 40 people to march around in the compound, practicing close-order drills. Boot camp is basically boring, so we turned this into a form of entertainment. The idea was that the people in the platoon would do EXACTLY what I ordered, and I would suffer the consequences when I got confused. My commands would go like this (MARCH on one foot, execute on the next foot if possible):

FIRST column half-left, MARCH , SECOND column to-the-rear MARCH,
FOURTH column half-right MARCH ...

At this point all the columns (of 10 men each) would be parallel. The outer columns would be moving diagonally left and right, the second column would be heading opposite the original direction, and the third column would be going straight on. The challenge was to get them all back together, marching in four columns again in parallel. This means counting off intervals, and then going through other commands to direct the columns around the field until on the last MARCH they were formed up again. If I goofed it up, the guys would happily crash into each other and fall dead laughing, or continue marching off into the

distance where I had forgotten about them. We got very good at this (we had to practice every day) and it was a heck of a lot more fun than just tromping around in columns. The officers used to come out to watch.

I, of course, got a bang out of having control of all these people and seeing them do what I wanted. But that soon became secondary to thinking up complex series of commands that would leave groups scattered all over the field, and then figuring out a way of getting them back together. It was sort of like programming (which didn't exist then); there is a finite set of operations that have to be used to create an infinite series of different patterns.

It was interesting, however, that the platoon members thought it was great fun, too. We often crashed, but quite often we would go through a really involved set of maneuvers that resolved with a series of consecutive commands into a perfect final formation, and a great cheer would go up: WE DID IT! The best final move was when a series of four commands would be issued, one to each group, without the final "MARCH." Then, at the last second, the command "MARCH!!" would be issued, and on the next step the crash would be averted and order be restored. Usually. After a lot of practice.

In retrospect, it's clear that none of this would have worked if the individuals hadn't had a clear idea of what was going on and even predicted my intentions. After an extended set of maneuvers involving dozens of commands, there's no way that the columns could have remained in proper relationship just by blindly following the orders. The individuals in the groups had to be making fine adjustments all the time, particularly near the end, to keep the lines straight and to be sure that the final step wouldn't occur halfway between the positions of people in other columns. So while I was nominally in control, I was actually delegating detailed control to the individuals.

All this does give some notion of how a social control project could be set up to work. The control isn't arbitrarily imposed the way it is inside a person. It requires all the people involved to understand most of what is going on, at least to the extent that it affects something they can perceive and control by themselves. It was up to me to think up the pattern and the strategy in the marching drill, but it was up to the individuals to make it work.

Also, the speed of control is very slow in such a group effort. The marching cadence and the fixed meaning of the commands, thoroughly understood by each individual and translated reliably into shift of directional reference levels, left a minimum time resolution of four steps, about two to three seconds depending on the pace (one thing that the platoon and I often disagreed about; I ended up following them on that point). The commands had to be issued over a period of at least four counts in advance of the actual execution, even so, to allow for processing time. "First column tatharear HAR!" takes four steps, and is executed on the fifth step (assuming I started on the correct foot for the direction of pivot). Some commands like "First column column-half-left MARCH" were too long to do in four steps and had to be started another step in advance to be understood. At each word, all the people had to (a) recognize that they were in the group to which the command applied, (b) interpret exact technical meaning of the term and turn it into a reference level for an actual physical act, and finally (c) wait for the MARCH command and execute the procedure already set up with a delay of one step (if correctly given) or two steps (if issued on the wrong foot, blush). At the same time, each individual was maintaining a cadence in synchronism with the nearest visible people, maintaining spacing and line straightness, and imagining the move that would result in the correct perception when the MARCH occurred. Furious mental activity in each person.

I'm impressed with how much practice is needed to create a precise group control process, and how simple it has to be to work at all.

Your comments on defining levels of control are noted. I'm not satisfied with them, not because because I disagree that the things you mention are controlled, but because I'm trying to get beneath the surface manifestations to basic controlled variables. When you speak of controlling "interpersonal relations", I immediately wonder what basic capacities of perception and control are required in order to control that sort of thing. I see controlling interpersonal relations as an empirical classification, but not as one that defines the underlying control structure that makes it possible. Your category just doesn't seem general enough. In my attempt to define levels, I tried to define them so it didn't matter what particular things a person was controlling. Control of relationships occurs in all contexts, whether involving other people or not. Ditto, I hope, for all my other levels. Also, each of my levels involves a type of perception, independent of modality, that, when analyzed into its components, proves to be made of elements of the next lower

level. And control of one level of perception REQUIRES varying perceptions of the next lower level. My levels are supposed to meet all these criteria. And they're all supposed to be context-free.

Well, I can't spend all day at the keyboard. The rest can wait.

 Oh, one last thing. I don't think any cells in the body are merely passive participants. Kidney cells control independently for concentrations of many substances in the blood by varying the rate of their removal; one kidney cell contains many control systems. Liver cells control the sensed level of glycogen; raising the circulating glycogen has a negative feedback effect on glycogen production. Liver cells also control for the circulating concentration of cholesterol. If the concentration falls below the current reference level, the liver cells begin manufacturing it. This is why attempts to alter the cholesterol level are so ineffective; it takes rather violent means to change it by as much as 15 percent, and the means have undesired side-effects. If there is a problem, it's in the reference level, not the production machinery or the circulating level. As doctors found out after only about 20 years of effort, trying to alter a controlled variable like circulating thyroxin is no answer to any problem. The body fights any attempt to alter variables it has under control, and in the attempt can damage itself (atrophy of the thyroid gland, for example, when thyroxin is administered over a long period). I haven't yet seen any evidence of cells that are purely passive. It seems safest to assume that everything is under active control.

 Greg Williams (930621) --

I don't want to get into another long war of words, but it seems to me that "stress" is treated as if it's like a measles germ that gets into you and causes internal problems. It seems to me that stress is merely the condition of the system when it's trying to resist a large disturbance. Mental stress, I think, is largely caused by internal conflict. External events can seem to produce it when they bring a person's conflicts strongly into play: when the person is prevented from doing something about an external situation by an internal conflict that prevents any effective action from being carried out.

What seems stressful depends on what your goals are and how effectively you can maintain control. Stressful sexual harrassment for one person is a titillating sexual game for another. Of course in my ethical system, you don't play mind-games with emotional cripples anymore than you go around making life difficult for paraplegics. But the law is a pretty blunt instrument for enforcing ethics.

 Bruce Nevin (Mon 930621 09:44:18 EDT) --

It would help if we could just focus on one issue. I'm raising a question as to whether "contrast" is the right perceptual variable to be thinking of, while you're continuing to miss my question and answer other ones.

Contrast itself can be a perception that is specifically under control, as in adjusting the contrast on a TV set (the difference in brightness between highlights and the darkest areas). Here you are perceiving several different inputs at the same time and perceiving a relationship between them, right now in present time.

When you speak of a contrast between words, you're not talking about that kind of situation. The words aren't being perceived simultaneously, while some higher system judges the difference between them. Only one word is being perceived; its contrast with a rather large set of other words is hypothetical, not actual. The other words are not actually being perceived; there is no basis for detecting a contrast.

I think that this usage of "contrast" is erroneous. The same effect can be obtained from a model in which separate systems emit signals when inputs satisfy their input criteria for recognizing a word. You seem to think that such a perceptual function would require an exact match, so that one and only one set of phomeme signals would result in a perception of a given word. There's no such requirement on an input function. Input functions can be as sharply or broadly tuned as is required to fit the facts.

Your comments on the grape and the elephant seem to assume that the grape and the elephant are distinguished because they REALLY ARE different, whereas phonemes are distinguished because they only SEEM different. This is a spurious argument, because in either case, the perceptions are generated out of lower-order perceptions of sensations. The only difference between the two situations is that for the grape and the elephant, the two responding perceptual functions would probably NEVER respond at the same time except under very

unusual circumstances (a distorted grape with a bit of stem hanging from it that makes the grape look a little like an elephant). In the case of the phonemes, more than one detector can respond to a given set of sensations, and there can be interactions among the detectors that cause the recent presence of one detector's signals to affect the operation of other detectors. The detectors may accept a rather broad range of input sensations as sufficient excuse to produce a perceptual output signal. In no case is it necessary to propose that the detectors are responding to something called "contrast." The idea of contrast in this connection is a dormitive principle: we detect the difference between words because they contrast; a contrast is that which enables us to detect a difference between the words.

In fact, we don't detect differences between words, we detect the words. "Differences" exist only for a difference-detector. In order for a difference or contrast to be detected, the things that differ or contrast must already have been separately discriminated. So difference or contrast has nothing to do with the fact that one system responds maximally to one condition, while another responds maximally to another condition. Both conditions may actually be present at the same time, so we get both signals. Then it's up to higher-level systems to sort them out by context.

Am I getting across at all?

Best to all, Bill P.

Date: Mon Jun 21, 1993 9:52 am PST
Subject: information, stress

[From Rick Marken (930621.1000)] I said:

>could you just tell us -- does the sensory input to a control
>system contain information regarding how "plant" outputs should
>vary in order to control the sensed variable?

Hans Blom (930620) replies --

>Sorry, but I decline to be arbiter. As a teacher (my other role) I
>have to say that having this discussion is by far more fruitful than
>knowing who is right.

But just out of curiosity, who is right? I'm sure we'll go on discussing it anyway, even after we find out. Or is your point that neither side is right? Or that both are right? Or that there is some rightness and wrongness in both positions? You don't have to arbitrate; just tell us which position is right from the point of view of an experienced control engineer -- one or the other, both, or neither. We'll take it from there.

Martin Taylor (930621 11:00) --

>A *slight* inconsistency here, quoted without comment:

>(1)
>>Surely you don't consider the feedback function
>>to be part of the control system, do you?

>(2)
>>The fact is that you guys just have a wrong conception of how control
>>works -- its a circle, not a chain.

>Do you see the contradiction?

Not at all. I was trying to guess why you think there is information in perception. Since the closed loop control system works without there being any information in perception, I figured the only reason you would believe that there is such information is because you imagine that the input must inform the system about what output to produce -- a chain of causality conception of control. Whether you look at control as a chain or a circle, however, is quite orthogonal to whether the feedback function is part of the "box" that contains the actual control system. So, no, there is no contradiction.

>We keep trying to tell you that we NEED the
>circle, and you keep trying to tie us up with the chain.

Ok. You need the entire feedback loop in order to produce outputs that counter the disturbance. Wonderful. Then where does information theory come in -- except as a ridiculously arcane and slipperly way of characterizing the operation of a control system (I say this because we have yet to see a simple, clear [or ANY] demonstration that there IS information in the perceptual signal)? The fact that, give $y=x$ and x you can compute y does not count as proof that there is information in perception -- at least for me. But if this is all you mean by "information in perception" then, yes, there is information in perception.

>It would be nice if you could get OUT of your head the notion that an >information-theoretic approach deals with an open-ended system.

OK. It's out of my head. Now what should I put INTO my head? What do you mean when you say that there is information about the disturbance in perception?

My simple minded interpretation of that statement is that there is something in the input perceptual signal (that you call "information") that is used by the system in order to generate outputs (in the context of the closed loop, OK) that result in control. We have suggested what Bill, Tom, Gary Cziko, myself and some others consider a reasonable test of this assertion: if there is this kind of information in the perceptual signal then you should be able to reconstruct the disturbance from this signal -- knowing ONLY the perceptual signal (which is all the control system knows). But you say you can only do this if you know EVERY function and variable in the control loop -- which the control system itself cannot possibly know. I take this as an admission that there is NO information in perception about the disturbance; the control system's opposing effects are generated as part of the normal operations of the closed loop negative feedback control system. An analysis of this closed loop system (a la PCT) shows that the perceptual signal is a dependent variable; its value is specified by the reference signal. The perceptual signal (when there is high loop gain) has nothing to do with determining the output of the system. The general equation for the output of the system is:

$$o = p^* - 1/g(d)$$

where p^* is the reference signal. Notice that p (the perceptual signal) does not appear on the right side of the equation -- as an independent variable. What does information theory have to say about how a control system works that is not already said by this equation?

Bill Powers (930620.1600 MDT) --

> I suspect that you [Hans] haven't yet understood just how the PCT >diagram differs from the standard engineering one.

This was a wonderful discussion, Bill.

It shows, once again, that PCT is trying to just make a simple point; that is, behavior is the process of controlling INPUT perceptual variables. That is a simple point, but it is basic. If one doesn't fully understand and accept this simple fact about the organization of living systems, there is no chance that their more complex analyses of behavioral phenomena can be worth much because they are based on the wrong premise. I think that many scientists in all fields these days are happy to skip the fundamentals in order to get on with the real interesting, complex modelling. I think you alluded to this, Bill, in your discussion of the information theory material that you had read. It is this apparently irresistible urge to get on with the complex stuff and skip lightly over (or just ignore) the fundamental assumptions that leads to 1) a lack of interest in PCT 2) Karolyan PCT and 3) trendy science. Greg Williams (930621)--

>Why doesn't "setting up contingencies that can hurt [certain] people" >amount to the same thing as providing "stimuli" (you can call them >"disturbances") which "cause" certain people to experience "stress"?

If disturbances cause stress then there is just poor control and that's not the fault of the person (or environment) that produces that disturbance. In order to set up contingencies that can really hurt people (in the sense that they require reorganization) you must use some fairly serious coercion; you must make sure that the only way that a person can satisfy their intrinsic needs is by doing what you want. I'm not sure any employer can really do that. But, if they could, I suppose I would call that a situation that "caused" stress.

>I know that claiming that PCT brands common notions as "mythical" is

>tempting for PCT revolutionaries. But I don't see any good coming from
>doing it in this case.

I don't see any good from not doing it either. Fact is, it's a myth.

>It just further polarizes PCTers from everybody else.

When you're right, you're right -- and you just have to live on that pole; cool but comfortable.

>I think you would be annoyed because you realize that if I hadn't
>lied, then you wouldn't have had the heart attack.

Pardon my french but "shit happens". Knowing you, I might have gotten a heart attack if you didn't say it. We (you and I) can't regulate our behavior to make it seem appropriate to the other person. I suppose I would have preferred that you hadn't played the trick but I can't be sure that that's what you were doing or why I got the heart attack. Besides, someone else might have done the same thing either by mistake or because of a well intentioned effort to make me "tough" or something. I want people to cooperate -- but you can't force them to cooperate. The only thing you can count on is that people will control what happens to themselves -- as best as they can. If we understood this about ourselves and each other, then we would realize that the only way to prevent ton's of interperson conflict is to cooperate. Everybody will still be controlling what happens to themselves -- just a hell of a lot more effectively.

Best Rick

Date: Mon Jun 21, 1993 10:49 am PST
Subject: Re: Where are the sociologists?

Tom Bourbon (930621.1301) Rick Marken (930618.1500)

>


>After reading the posts from Bill Powers (930618.1000 MDT) and
>Tom Bourbon (930618.0958) on social control, I have been re-
>considering my comments on this topic. I had been assuming that
>there is such a thing as a "collective controlled variable" which
>is a variable that is not controlled by any individual but is a
>result of control by the collective. My example of a collective
>controlled variable was the three line pattern made cooperatively
>by the subject's in Tom's cooperation experiment. After thinking
>about this for a bit I realize that I am wrong -- there is no such
>thing as a collective controlled variable. ...

Glad you have seen the light, brother! I thought you were going a little too far in your interpretation of the three-line pattern. It *is* a controlled variable (but not a collective one), some times, but not others. (More on that in a while.)

>... A controlled variable

>ONLY exists if there is a perceptual signal in some system that
>represents that variable, a reference signal that specifies the
>reference level of that signal and an output system that can affect the
>perceptual signal. In Tom's task, the three line pattern is a controlled
>variable ONLY IF one or the other or both subjects are perceiving and
>controlling that overall pattern. Any disturbance to the pattern
>would be resisted (if possible) by the person (or persons) controlling
>it. The three line pattern would APPEAR to be controlled, however,
>even if each person were controlling just a piece of it -- the relation-
>ship between left and middle line for one subject, the relationship between
>middle and right line for the other. It might seem like the three line pattern
>is controlled in this case but it is NOT -- which could be proved if the
>subjects were allowed to see and control only their two line patterns
>with the three line result projected on another screen. Disturbances
>to the three line pattern that did not affect each subject's two line
>controlled variable would NOT be resisted.

For anyone who has not seen a demonstration of the task, or who missed my earlier detailed post, here is a diagram of the connections in Boss Reality (the innards and environs of a PC):



The program steps that determine where each of the three cursors appear on the screen during a particular update of screen information are:

```
Position of Left Cursor := Scaled Position of Left Handle;
Position of Right Cursor := Scaled Position of Left Handle;
Position of Middle Cursor := Sum of Momentary Value of Random Disturbance plus Scaled
Position of Left Handle plus Scaled Position of Left Handle.
```

That is the set of connections in one small part of the world. Nothing in the connections specifies everything that will happen when one or two persons use the handles.

As part of a semi-private discussion off of the net, Rick had said:
=====

[Rick]
>I was trying to read Hans as saying that a controlled variable
>could be the cooperative result of controlling by 2 or more
>control systems. I think that's what is happening in Tom's
>experiment -- is it not? The pattern of the three lines is, indeed,
> a controlled variable, and control of that variable requires that
>the two parties involved control -- and control a particular
>variable each. So there is a controlled variable that is the result
>of collective action. Of course, the two people don't become a
>control system, but their individual controlling controls the
>position of the three lines.

[Tom]
Yes, and they can select another pattern of the three lines, in which case each person will necessarily adopt a different reference perception for part of the new intended pattern.

[Rick]
>I was thinking that, by my definition of flockness as the set of between
>individual distances, flockness is, indeed, controlled. The set of
>distances are controlled, and they are controlled at values which
>depend on the controlling done by each individual.

[Tom]
But, unlike my people, who KNOW they are "making a pattern of three lines," members of a flock need not have such an awareness or intention concerning "being in a flock."

[Rick]
>This is where I might be wrong; maybe by that definition "flocking" is
>still emergent.

[Tom]
See above.

[Rick]
>... Is the line pattern in Tom's experiment a controlled
>variable? Does control of that variable depend on the controlling done
>by a collective? Is flocking a controlled variable by my definition?

[Tom]
What if I do the task with my own two hands? Or what if I use a moused version of the task and do it with one hand? Or a mouse-on-a-stick (I intend to patent and franchise!) and do

it with one joint (this is starting to sound REALLY good!)? Is the line still controlled? Can we draw a line (a division) between any of these permutations and the two-person cooperative task?

In all of the one-person variations, there is no doubt the line is controlled. In the two-person runs, my two people have always known the "superordinate" goal. But what if one person knew the big picture and only told the second person, "just keep this line even with that one." Would the second person be participating in "keeping three lines in a row?"

The REAL social factors begin to pop out, for me, when I think of the *entire set* of permutations. Some are social, some are not. In some, the big picture is controlled, in others, not so for all participants.

=====

Back to the present:

Whether three-in-a-row is a controlled variable, or an unintended by-product of control of other variables, depends on the intentions and perceptions of the control systems. This seemingly innocuous task is rich with implications for several presumably independent areas of study, ranging from social science on one extreme to neuroscience on another.

When one person performs (say the person is me -- TB) "I" control the three lines and keep them in a row - - -, or in some other pattern of "my" choosing:



My actions control two different relationships (middle re right; and middle re left) that are necessary to the production of the "big" pattern. But where do "I" fit in to the picture? Where am I, neurologically? "I" have a clear understanding of what I intend, and of whether I see what I intend. I know that I can achieve the intended result by moving the handles, or the mouse, depending on the version of the program that is running. But by conventional neurological theory, the brain areas responsible for producing the appropriate movements differ across the various one-person implementations mentioned in the post to Rick. Two completely independent control systems can achieve the same results that "I" achieve -- two simple single-loop PCT models can do it. What, and where, "I" am becomes less and less clear to me when I think about these demonstrations. (Hans, I usually go out of my way to say that a person can be represented by two independent PCT models, not that a person is two such systems. The same style applies when there are two people, and two models. The people can be modeled by, or as if they were, two exceedingly simple PCT systems.)

When two people perform the task, the results can be identical to those produced by one person, but for different reasons. Each person controls the relationship between two of the lines, and each might know and try to control a perception of the three lines in a particular relationship. Or each might control only the relationship between two of the lines and neither know or care about the relationship among all three. Or one person might know and control the relationship among the three lines, by directly controlling the relationship between two, and by telling the other person to control the relationship between two -- the second person remaining ignorant of the first person's overall goal.

Only when a person who acts alone, or at least one of two people who act together, perceives and adopts an intended value for three-lines-in-a-particular-configuration does line-ness become a controlled variable. In all other cases, it is an unintended, perhaps unperceived, side effect of the control of other variables.

The same is true of flockness.

Of course, given an organism, or model, or artificial control device with sufficiently complex perceptual functions, a previously uncontrolled variable can be made the object of control. Similarly, given an assemblage of such systems or organisms, a previously uncontrolled (incidentally affected) variable might become the object of collective control. (Hans, if you want to study the individual, social, and evolutionary development of control systems, I suggest this simple cooperative tracking task to you. When you can explain all of its permutations and implications, you should have gone a long way toward your intended end. And along the way, non-engineer PCTers might be able to pick the brains of a control engineer.)

I can offer an example of "emergent" collective control. A few years ago, during a meeting of the American Society for Cybernetics, a group of CSTers (the earlier name for PCTers) crowded together in an elevator, on the way to dinner. Never before had so many CST-PCTers been assembled in any place, much less one so small as the elevator. Someone (Mary Powers, I believe) in the group (the gaggle? the flock?) remarked about how many of us there were. Everyone else agreed -- never before had there been enough of us to fill an elevator. Not long after that, we began to talk about having our own meetings. We did, and CSG was born. And CSG begat CSG-L. And here we are.

Identifying our flockness turned it into a controlled variable, for each of us. From something that had never been a controlled variable, we created one. We continue to do so, in spite of difficulties or complications that might otherwise disturb our individual perceptions of CSG-all-in-one-place. But when it happens, it happens one of us at a time. There is no CSG-in-the-sky that uses each of us to create the perception of the CSG flock, the way any one of us can set the reference signals for the characters in CROWD-GATHER and control our perceptions of characters in a gaggle. (Not unless Bill has installed some pretty interesting equipment on his hill in Colorado! Maybe he uses the net to reset the references every year around this time.) In a few weeks, we will make it happen again. Hans, why don't you come and see it happen?

[Rick]

>So Hans' and my idea that "flockness" is under control (a "collective
>controlled variable" by my terminology -- ie. a variable that is
>controlled by the collective but not by any individual) is incorrect unless
>there is a control system present that is perceiving and acting to influence
>this variable.

I second you on this. This is one of the points I made to Hans in my post (939618.0958) when I said Hans, but not the heating system, controls "a nice temperature" and that Hans, but not the characters in GATHER-CROWD, controls a perception of "flockness." And now when individual CSG people create their own perceptions of the meeting.

[Rick]

>Where are the sociologists during this discussion of social control, by
>the way? This is where PCT meets sociology. ...

I have been wondering the same thing. Has everyone gone on vacation?

Until later, Tom Bourbon

Date: Mon Jun 21, 1993 11:00 am PST
Subject: Re: Power gain, power loss.

From Tom Bourbon (930621.1323)

>[Martin Taylor 930621 11:10] Hans Blom 930619

>Misunderstandings sometimes are better resolved by a non-combatant.

Most of Martin's post was about power gain in a control system. But at one point Martin said:

>On models, I tend to side with Hans. It is part of the whole information
>argument. The more information is available within the control system,
>the less is to be acquired from the CEV through the perceptual apparatus,
>and the better control can be.

Martin, to paraphrase a line from the movie, "Field of Dreams," all I can say is, "If you build it, we will come." Take one example of control, as it is recreated or predicted by PCT models, and show me, in the results of simulations, how making more "information" available "within the control system" improves the recreations and predictions from the model. Then show me that those results generalize, with no further tinkering with the model, to new conditions, with unpredictably different disturbances and targets. That is not much to ask. Just improve on the performance of a single-level, single-loop PCT model.

And please delineate how your ideas in the remark to Hans differ from, say, a plan driven system that relies on information in the form of programs for action, thereby freeing

itself from a need to rely on information about the CEV obtained through the perceptual apparatus. As you stated it, I see no difference.

This is not a put down. It is the only way to do business, if you rely on models to test your assumptions. It is my often repeated plea that you present the evidence, in the form of improved performance of the PCT model. Nothing else will impress us or win us over. You already know that. But I assure you that, if you build it, we will come.

Until later, Tom Bourbon

Date: Mon Jun 21, 1993 11:03 am PST
Subject: of grapes and phonemes

[From: Bruce Nevin (Mon 930621 14:44:33 EDT)] Bill Powers (930621.0840 MDT)

> Your comments on the grape and the elephant seem to assume that
> the grape and the elephant are distinguished because they REALLY
> ARE different, whereas phonemes are distinguished because they
> only SEEM different.

Nope. There are no phonemes there to be distinguished. Phonemes are apparently present only because a perception of contrast is controlled.

Am I getting across at all?

Bruce bn@bbn.com

Date: Mon Jun 21, 1993 11:34 am PST
Subject: Simultaneous vs sequential

[From Rick Marken (930621.1200)] Martin Taylor (930621 11:00)

>Do you see the contradiction? We keep trying to tell you that we NEED the
>circle, and you keep trying to tie us up with the chain.

Greg Williams (930621) --

>Good for you! I would have brought this up, too, but I already know where
>things would probably head:

Well, I think you were wrong about where things would head. But I think there is a better (more accurate) way to characterize the difference between the PCT and information theory of control: rather than circle vs chain I think it is more accurate to characterize it as simultaneous vs sequential.

Both PCT and information theory recognize a circular chain of events occurring in control:

$p \rightarrow (r-p) \rightarrow o \rightarrow g(o) \rightarrow p$

Since p is at the beginning and end of the chain, there is an implicit circle.

What the information theory view seems to miss is that all the variables in this circle are changing simultaneously; it is by simultaneously solving the system and environmental equations that we end up finding that

$o = p^{-1}/g(d)$

with no contribution to the output, o , from perception, p . This is not only because the perceptual variable is part of a causal circle; it is because all variables in this circle have an effect on all others AT THE SAME TIME.

I think the information theorists must imagine that the events in the loop occur sequentially; perception is compared to reference which leads to output which is transformed by the feedback function, $g(o)$, into a new perception; there is sort of an assembly line with the information in perception serving as the blueprint for what should be done to counter the disturbance. This may be why they imagine that something is extracted from p -- the information that is used to produce the next item in the sequence,

o. At least, this is what I think they must be thinking -- otherwise they would see that there is obviously no information in p about d ; p is just one of several simultaneously varying variables that happens (due to the negative feedback, high gain SIMULTANEOUS relationship between all the variables) to be kept equal to one of the variables --r.

Note that $g(o)$ -- the feedback function -- is separate from the control system itself in both the information theory and PCT versions of the control loop. Thus, I fail to see the contradiction that you and Martin seem to see. Help me out here.

Best Rick

Date: Mon Jun 21, 1993 11:55 am PST
Subject: Re: emergence etc.

From To Bourbon (930621.1348) Hans Blom, 930620

Well into your post, you made some comments that Bill Powers has answered in part.

[Hans]

>Then, internal in the device, we see the LOOP GAIN, which has nothing
>to do with power at all, but is the gain in the loop that the SIGNAL
>travels. A decent control system has a loop gain much greater than
>one, although a loop gain of less than one is not unthinkable. In the
>latter case we might not want to consider the system a CONTROL system.
>

>Loop gain is terminology from the analogue past. In a hifi amp's
>innards you can still trace the loop with your finger. In a control
>system that is described in difference rather than differential
>equations, the 'loop' is functional only at a sequence of infinitely
>short time intervals. That is why they are also called 'sampled data
>systems'. Now, to complicate matters even further, in adaptive control
>systems there are (at least) two loops: one to provide the immediate
>actions, and another to do the parameter adjustments (learning). These
>loops interact and hence become much more difficult to trace with your
>finger. All of this causes the term "loop gain" to slowly fade away.

In a reply, Bill addressed your idea that loop gain is internal to the device. (It is not.) Here, I want to point out another of the seemingly large differences between the systems designed by contemporary control engineers, and those we know and love and model as living control systems. In many parts of modern control engineering, it might be correct to say, "Loop gain is terminology from the analogue past," and, "All of this causes the term "loop gain" to slowly fade away." But analogue is what living systems are all about. Neural currents and hormonal fluxes are analogue, through and through. Students of living things do not have the freedom to go with the newest technology, the way a design engineer might, although in the pop science that passes for much of modern psychology and parts of neuroscience, people try to do that by re-stating everything in the latest techno-jargon. But that is all a word game.

Operational amplifiers and potentiometers are pretty good materials for building models of living control systems. Come to think of it, you can use them to build some pretty good artificial control devices, also.

This is not a criticism of design engineers, or a declaration of some kind of superiority in the study of living systems -- just a simple observation of a big difference. It seems to reflect a deeper difference in interests. Some are more interested in building devices that approach optimal control, from their own perspective, no matter how closely the devices approximate the processes and organization of living systems. Those people can and probably should use any technology that fills the bill -- digital, analogue, fractal, you name it. Others are more interested in building models that more accurately approximate the (often other than optimal) actions of living things, by way of processes and organizations that also closely approximate those seen in living things. That means accepting living things as they are, analogue warts and all.

No one on this net has said the twain shall never meet.

Until later, Tom Bourbon

Date: Mon Jun 21, 1993 12:23 pm PST
 Subject: Re: of grapes and phonemes

[From Rick Marken (930621.1230)] Bruce Nevin (Mon 930621 14:44:33 EDT)

>Phonemes are apparently present only because a perception of contrast is
 >controlled.

>Am I getting across at all?

Yes. It is clear that you want language to be a special kind of phenomenon; a perception that ONLY exists because it is controlled. Apparently you have never been to the exploratorium; there you can hear plenty of phonemes generated by nothing more than air rushing through tubes of different shapes -- the air source is not controlled and the tubes are not controlling anything either. I have heard phonemes generated by the wind rushing through trees; by the lunch whistle at places I've worked, by passing trains. Phonemes are just the perceptual signals that come out of phoneme detector neural networks; the external vcause of maximum signal output from these detectors is of no concern to the detector. So you'll hear "she" if the appropriate acoustic signal is produced -- whether it is the controlled result of someone's speech articulators or the "caused" result of steam passing out of a radiator valve.

Best Rick

Date: Mon Jun 21, 1993 12:56 pm PST
 Subject: peek-a-boo

[From: Bruce Nevin (Mon 930621 16:00:33 EDT)]

(Rick Marken (930621.1230)) --

> So you'll hear "she" if the appropriate
 > acoustic signal is produced -- whether it is the controlled result
 > of someone's speech articulators or the "caused" result of steam
 > passing out of a radiator valve.

And there's a rabbit in that cloud up there.

You'll hear phonemes if you're listening for speech. Look at some of the experimental work with synthesized speech. People who are primed to expect speech hear speech; those who are not hear only noises. Look at speakers of languages other than English listening to the same noises in the exploratorium. They don't hear the same phonemes. Why not? One person's language has no contrast between s and sh; another person's language has a 3-way contrast between s, sh, and S (more usually written with a dot under the s). One language has a contrast between i and y (like French "tu") and it sounds more like sy! to them than shi! (your "she!"). You say the phonemes are in the sounds of air rushing through tubes and trees--how can they be different phonemes for different people?

Oh, so they're different phonemes just because each individual is an autonomous control system controlling perceptions independently of the others? Then how come all the native speakers of English hear the same phonemes from the tubes in the exploratorium? (BTW, it doesn't matter whether the sounds recognized as phonemes are coming from a control system, a tape recording, or a mechanism of tubes and air lines with no control system involved, the hearer reconstructs the phonemes she would have intended had she produced those sounds.

The phonemes are not in the noises. They're in your head.

> you want language to be a special kind of
 > phenomenon; a perception that ONLY exists because it is
 > controlled.

It's not a matter of what I want or you want. That appears to be the way it is, by golly. But I'm grateful to you for getting the point, even though you're resisting it as a disturbance to something. Could you specify what?

Bruce bn@bbn.com

Date: Mon Jun 21, 1993 2:17 pm PST
 Subject: Re: Powers on information

[Martin Taylor 930621 16:30] (Bill Powers 930618.1930)

>As to the "information in perception" argument, I think you're
>pulling back from the original position, which is OK but you
>should acknowledge it. Originally, the concept was that the
>output of the system almost exactly opposes the effects of the
>disturbing variable on the basis of information contained in the
>perceptual signal. Now it turns out that your side is admitting
>that the perceptual signal doesn't contain ALL of the information
>in the disturbing quantity, but only some part of it.

I'll simply repeat what I have said several times before. The only argument (initially) was that it was inconsistent to say two things at the same time:

(1) that the output signal exactly mirrors the disturbance signal
(a claim made very strongly by Rick and you in many postings), and

(2) there is no information about the disturbance in the perceptual
signal (a claim made even more strongly, particularly by Rick).

You simply cannot have both (1) and (2) at the same time. They contradict. If you pull back from (1), as you did immediately we presented the "mystery function," then we go along with you, and have done so. To the degree that (1) is inaccurate, to the same degree is extra information beyond the perceptual signal required to reconstruct the disturbance. I don't see any retraction there, beyond to accept that your acknowledgement of the error in the original statements is valid.

>I presume you would agree that the better the control, the
>smaller the amount of information that gets through to the perceptual signal.

Quite the reverse. The better the control, the MORE information has got through (not "through to") the perceptual signal to the output, and the less the perceptual signal retains.

>I will be fascinated to see how you solve this problem and
>demonstrate that the information in the disturbing quantity is
>sufficient to produce the output actions.

"Sufficient" isn't the issue. "Necessary" is. If "sufficient" were the issue, we would be debating S-R versus control. We are not. We start on the basis that there is good control. Without it, there is no basis for our side of the discussion.

Somehow I doubt that you would argue that control is best achieved with no perceptual input (the cognitive planning) position. But logically, that's where you seem to be headed. I expect you to reject that position just as strongly as I reject the S-R position you and especially Rick keep trying to push us into.

=====

Ongoing issue--a slight side-track:

Bill, I simply don't understand how someone who has been involved with real circuits for so much of his life can maintain such a blind spot about the effects of noise, bandwidth, and so forth. You must have known about the time-frequency trade-off, the effects of narrow-band filtering (which is what your long observation interval does), and the like. How is it possible for you to take my discussion of the time-dependent increase in resolution in observations of low-bandwidth signals, and come up with:

>If the noise is 10% RMS of the range,
>Martin's interpretation would be that there are only 10 possible
>values of the perceptual signal with a range of 10 units, 0 through 9.

That violates high-school electricity, let alone what a graduate engineer would think. I am really surprised at you. But I am now less surprised that you did not understand my analysis of the relation between admissible gain and the relative bandwidths of perception and disturbance as related to perceptual resolution, and that to refute the analysis you did a "simulation" in which the perceptual resolution was specified as essentially infinite.

You assert that I gave you D/r as a measure of resolution, ignoring that I pointed out that for continuous signals $(D+r)/r$ is more appropriate, where D and r are power rather than amplitude functions. You return to the assertion that we are talking about quantized signals, even though we had sorted that out before I went to Europe. Quantized signals behave differently, when their variation is smaller than a quantization step.

Control within the quantization step is impossible, UNLESS deliberate variation (noise or shimmy) is added to the disturbance, usually by the output function, in which case control becomes possible over time scales longer than the period of the shimmy or than the inverse bandwidth of the noise. I'm sure you know all this, if you dredge it out of your memory. (Humans do it in several of their control systems, such as ocular pointing, but I think the reason is more subtle than to overcome some intrinsic quantization. I think it is to linearize a zero-point reduction in differential perceptual sensitivity which is both predicted by information theory and observed in psychophysical studies. Oh, I forgot. Psychophysical results are not to be mentioned in the same breath as PCT, because they might say something about perceptual resolution, which we know to be infinite, don't we?)

From time to time you make a posting that makes me think that we are converging on an understanding, and then you go and retract it all again. I don't understand why, except in an abstract PCT sense. Somewhere you are controlling for some value of a perception that could not be maintained if the importance of information obtained through the perceptual signal was accepted as a foundation for the control process. I'm not sure what this perception is, and I don't know how to go up a level to see what can be done to return you to the knowledge of signal theory that you must have used extensively in your professional career. When you were working with noisy signals, you wouldn't have said things like:

(Bill Powers 930422.0030)

>Martin, however, has said that the white noise disturbance he
>assumed must have uniform power over its frequency range. This
>implies multiplying the FT entries by $1/f^2$, because power
>increases as the square of frequency for constant amplitude.

The power of a sine wave is independent of its frequency. It is

$1/x$ (integral from $x=0$ to $x=t$ ($\sin^2 fx$) dx) (omitting the odd 2π factor)

which, provided there is an integer number of cycles in the range $x=0$ to $x=t$, is independent of f . The answer is $1/2$. (The indefinite integral is $(x/2) - (1/4f)\sin(2fx)$.)

Strictly speaking, before the notion of "power" is correct, the sine wave has to be the amplitude of some entity such as voltage applied across an impedance, but "power" is often used colloquially for the square of amplitude averaged over time.

>Martin agreed that varying the phase angle as I propose is the
>proper procedure (not realizing, apparently, what this implies in
>terms of the power distribution).

Corrected, now? No f^2 ?

>But he also recommended that
>the amplitude of each entry be given a Gaussian distribution. I
>can do that by using the formula $1/(2\pi)\exp(-x^2/2)$, choosing x
>with a random number generator

Quite adequate Gaussian numbers can be generated by adding 12 random numbers (12 because, if I remember rightly, that gives unit variance if the numbers are chosen from a unit rectangular distribution. Don't quote me, though.)

>and adding or subtracting the
>result from the amplitude of each entry. However, as I planned
>how to do this I realized that something is awry here.

>

>Why should the disturbance have to be so carefully selected? And
>particularly, why should a Gaussian variation in the amplitudes
>be needed? The only reason I can see for adding the Gaussian
>factor is to provide a value for r in the expression $\log(D/r)$.

The reason is that the Gaussian variation provides the maximum uncertainty for a given power. If you do not have the Gaussian variation, the signal can be translated into another signal of lower power but the same resolution, that can convey the same information (equally, it can be converted into one of lower bandwidth but the same power per unit frequency or one of shorter time but the same power per unit time).

Any other disturbance can be considered "carefully selected," but the Gaussian cannot. It is an extremum.

In the same posting:

>What this implies is that talk about the information content of
>physical variables is empty. For all practical purposes, their
>objective information content is infinite. How much information
>we can obtain from them depends strictly on the dynamic range of
>the sensor.

And its resolution. Yes, you have come here to recognize one of the ideas I tried often to get across. The concept of information depends on the receiver, and cannot be determined *a priori* by an outside observer. It is at the heart of your complaint (in some other posting) about being told "it depends on what you assume to be known." It does, and the piece I just quoted shows that you know that to be so.

To emphasize that point, I make just one comment about your statement that the Perceptual Input Function just "is," and has no prior information about the incoming signal--have you ever heard of a filter?

Martin

Date: Mon Jun 21, 1993 2:23 pm PST
Subject: Phomeme perception

[From Rick Marken (930621.1500)] Bruce Nevin (Mon 930621 16:00:33 EDT)

> Look at
>speakers of languages other than English listening to the same noises in
>the exploratorium. They don't hear the same phonemes. Why not?

They have developed different phomeme detector perceptual functions.

>You say the phonemes are in the sounds of air rushing through tubes and trees

Never said that. The phomemes are in the brain (as perceptual signals); the causes of the phomemes are the sounds in the tubes and trees.

>how can they be different phonemes for different people?

Different phomeme detectors.

>Oh, so they're different phonemes just because each individual is an
>autonomous control system controlling perceptions independently of the others?

No. Different phomeme detectors.

>Then how come all the native speakers of English hear the same
>phonemes from the tubes in the exploratorium?

They have developed similar phomeme detectors -- if they didn't, they would not be able to perceive the phonemic distinctions that make it possible for them to communicate with each other -- assuming that they WANT to communicate with each other, of course, which they apparently do.

>(BTW, it doesn't matter
>whether the sounds recognized as phonemes are coming from a control
>system, a tape recording, or a mechanism of tubes and air lines with no
>control system involved, the hearer reconstructs the phonemes she would
>have intended had she produced those sounds.

"Reconstructs" may be a bit much. The PCT perceptual model just says that there is a perceptual signal in the neural output of a phomene detector.

>The phonemes are not in the noises. They're in your head.

That's right -- as the outputs of perceptual functions (phomene detectors).

I said:

> you want language to be a special kind of
> phenomenon; a perception that ONLY exists because it is controlled.

You say:

>It's not a matter of what I want or you want. That appears to be the
>way it is, by golly. But I'm grateful to you for getting the point, even
>though you're resisting it as a disturbance to something. Could you
>specify what?

It is a disturbance (though not a very big one, I must admit) to my understanding of how the PCT model works -- at least, how the perceptual aspect of the model works. I see no reason for language sounds to be treated any differently than any other sounds. The model of perception in PCT is pretty simple. Here it is, specially tailored for phoneme detection:

sounds ----> PD --->p

Sounds out in boss reality are inputs to ALL phoneme detector (PD) perceptual functions. One is shown here. The PD is a neural network that transforms a sound signal into a neural signal, p. The bigger the neural signal, the more the sound matches the input that that PD is designed to detect. An "ah" PD puts out a big p signal when the sound is like "ah" (actually, like a waveform that we experience as being like "ah"). The particular set of PDs we have depends on what phomemes we have had to learn to discriminate.

This is the PCT model of ALL perception. Note that there is no "reconstructing" involved. The perceptual fuctions (PDs) don't care how the input (sound) was generated; they are just transducers -- complex transduces but transducers nevertheless.

So what you are saying is a disturbance to my concept of the perceptual side of PCT. You seem to be saying that the perception of a phoneme (or other speech sound) requires that the perceiver actually be controlling the phoneme perception or reproducing (in imagination) the actions that were used to control it. I don't believe this model but I would be more interested in it if I could at least see how it works (a diagram, perhaps). It just seems unnecessary to me. I can recognize ALL the voices in a Bach fuge without making the finger movements that I would have made in order to create these perceptions (I couldn't create these voices with my own fingers anyway; I am not skilled enough as a pianist). I can recognize the sounds of different objects even though I can't make those sounds. I guess the disturbance to me about your position is this "mystique" about language. I also think you may be mistaking PCT for an "analysis by synthesis" version of speech recognition. Basic PCT is "analysis by transduction".

Language perceptions are important and useful and uniquely human; but so are Bach perceptions. I bet it all works the same way; perceptual signal outputs of neural network perceptual functions.

Best Rick

Date: Mon Jun 21, 1993 2:23 pm PST
Subject: assorted replies, and some more

[Hans Blom, 930621] Bill Powers (930620.1600 MDT)

>>The slogan "it's all perception" is much too static. It has a
>>connotation of having to act in an a priori circumprescribed
>>way given a set of perceptions and a set of top-level (very
>>slowly changing; innate?) reference levels.

>I agree. Perceptions are learned, too. There is a reason for the
>slogan, however. It's to remind us that each of us sits inside
>one of these gadgets, and that ALL we know either of the world or
>of our own actions and inner being consists of perceptions.

Agreed. This is the important point that CSG makes and that I fully agree with. In places, I disagree with some of the details of your models, however (you notice how each time I stress the 'optimal' usage of what is perceived), but that does not detract from the importance of this basic point.

>>We can tune our responses finer and finer, and reach ever
>>higher qualities of response and perception.

>I'm not sure how you mean this, but it sounds like one of the
>concepts we're trying to destroy. "Response" is the conceptual
>opposite of "control." It implies a blind reaction to an input,
>and carries overtones of jab-and-jump psychology.

I am talking in terms of the response of a control system to, say, a step change of a reference level or of a disturbance. As you know, the response might be overly damped or oscillatory. By a high quality response I mean one that is approximately critically damped. Maybe you show a great deal of self-control when your response is critically damped :-)

>>Control over control is self-control, perceiving your own
>>perceptions is self-perception, consciousness.

>That sounds nice, but I don't believe it. If you diagram a system
>that senses the stability of a control system and adjusts
>parameters to control stability, you do not have a system
>controlling itself: you have a system controlling something about
>a different system. If you perceive your own perceptions, one
>subsystem is perceiving the perceptions originating in another
>subsystem, and most likely interpreting them in a different way.
>The moment you say "I am thinking," you have denied the
>statement: the system that is aware of the thinking is not
>thinking, it is making a statement about a system that is
>thinking. The "I" of which you speak is never the "I" that speaks.

I do not believe in a monolithic "I", and neither do your models. Just look at that "society of mind". There is a lot of parallel processing going on in that brain of ours, and it is quite possible that one part talks about what another part is thinking. In a mentally healthy person, I assume that the parts are more or less in mutual contact; in multiple personality disorder they are not.

>The only way to make sense of self-reflexive ideas is to treat a
>person as if that person were solid, like a potato: only the
>whole person perceives and acts.

Demonstrably false. In split brain patients, one hand of the patient may caress his wife while the other hand hits her. In my experience, many people (all?) show such (psychological?) splits to some extent. Inconsistencies, due to 'conflict'.

>Loop gain, in PCT, is not "internal to the device."

In a hifi amp, it is. That was my example.

> I suspect that you haven't yet understood just how the PCT diagram
>differs from the standard engineering one.

Yes, I think I do. But sometimes I need to make a principle clear using the simplest example that I can think of, and I think that this kind of reductionism is the basis of science. If you cannot clarify a principle using a simple example, you will certainly be unclear when things get complex. Anyhow, thanks for the repeat. One comment to this section:

>This separation is always important in detailed control-system
>design, but especially so when the effector is coupled to the
>controlled quantity loosely or through complex intervening
>processes. Then we clearly would expect the effector output to be
>changing far more than the controlled quantity is changing.

That would be bad in any control system, especially if the load can vary. An optimal control system needs to know as exactly as possible what the effects of its OUTPUTS are. This can be done by SENSING THE ACTIONS OF THE EFFECTORS (by sensors in or near the effectors) and OBSERVING THE REACTION OF THE OUTSIDE WORLD to the actions of the effectors (in any sensory modality). I therefore hypothesize that human effectors have must have built-in sensors. These should sense what the effectors do, not through a (changing) environment as your model shows, but as intimately and directly as possible.

Note that this directly solves the adaptation (learning) problem as well: correlating the effector's action with the world's reaction 'identifies' ('systems identification') the world. This is how adaptive control systems work, and this is how I suppose a human works. This is also where control theory meets information theory.

Our effectors do in fact have built-in and/or built-on sensors, as you know. In my opinion, they have this additional function above the one that you describe, the comparator function.

>When we see the controlled variable separated from the effector
>output, we can much more easily understand that the visible
>behavior of an organism is really just its actuator output

Not when the actuator is loosely coupled to its environment (through some layers of skin tissue, for example). An analogy is a battery with a large internal resistance. Its 'visible' voltage greatly depends on its load.

(Martin Taylor 930621 11:10)

>It seems to me that just as a power gain is an essential element of
>the outflow side of a control system, so a corresponding power loss
>is an essential element of the inflow (perceptual) side. The
>perceiving of the state of a CEV should not contribute as a disturbance
>any more than it must (Heisenberg showed that it must, to some extent).

How true. Very familiar, too. In my blood pressure controller, the arterial pressure must be measured. Imagine what a thick needle in a thin blood vessel can do to disturb the circulatory system and thence the pressure measured!

> On the other hand, the output power of the control system wants to
>have maximum effect on the CEV, or as tight coupling and as high power gain
>as is feasible, given the information limitations on the perceptual side.

Feel that micromanipulator carefully tear a single cell away from its surrounding tissue as your hand squeezes on the macroscopic counterpart of the micro-pincer. Power GAIN? No! Very carefully scale the power down! Almost Virtual Reality...

(Rick Marken (930621.1000))

>>could you just tell us -- does the sensory input to a control
>>system contain information regarding how "plant" outputs should
>>vary in order to control the sensed variable?

>Hans Blom (930620) replies --

>>Sorry, but I decline to be arbiter. As a teacher (my other role) I
>>have to say that having this discussion is by far more fruitful than
>>knowing who is right.

>But just out of curiosity, who is right? I'm sure we'll go on
>discussing it anyway, even after we find out. Or is your point that
>neither side is right? Or that both are right?

My experience says that servomechanism theory and information theory are orthogonal. Neither has any use for the other. In the field that is called 'optimal control theory', however, where you deal with noisy sensors, noisy actuators and a noisy and/or a changing environment, it becomes critically important to try to identify the characteristics of the environment's reactions to your actions as accurately as possible. Optimal control theory is very much an extension of the combination of servomechanism theory and information

theory. And therefore not very attractive for people who refuse to think in terms of stochastic differential equations and their ilk, I might say :-)

>It shows, once again, that PCT is trying to just make a simple point;
>that is, behavior is the process of controlling INPUT perceptual
>variables. That is a simple point, but it is basic.

Yes, but not only. See above.

(Tom Bourbon (930621.1301))

>Identifying our flockness turned it into a controlled variable [CSG- L],
>for each of us. From something that had never been a controlled variable,
>we created one.

Nicely put! Yet, no single individual can 'control for' CSG-L. But where two or three are gathered in the name of CSG-L...

(Tom Bourbon (930621.1323))

> Take one example of
>control, as it is recreated or predicted by PCT models, and show me,
>in the results of simulations, how making more "information"
>available "within the control system" improves the recreations and
>predictions from the model.

Let me give you a practical example from my blood pressure controller. The arterial blood pressure decrease Δp due to an infusion flow rate i can be modelled as

$$\Delta p = \text{sensitivity} * i$$

where sensitivity is a constant that describes an individual's sensitivity for the drug used. It can vary by a factor of 80. That is just the static response, that is obtained as soon as a stable pressure is established. The dynamics of the process can be modelled by a delay time of about one minute and a dominant time constant of about one minute as well, in case you're interested. I will leave out other gruesome details, such as an often (but not always!) very pronounced non-linearity of the response and frequent sensor malfunction due to drawing blood through the same line that the pressure is measured with.

In this controller it was crucial to estimate the the individual's sensitivity (not necessarily very accurately; within a factor of two was good enough) to obtain a step-response that did not oscillate yet was fast enough to satisfy the anesthesiologist. Note that the sensitivity was not known and could not be established beforehand and had to be established while controlling. I dare you to design a servo-mechanism that can control under these circumstances. I tried but failed miserably.

>Then show me that those results generalize, with no further tinkering
>with the model, to new conditions, with unpredictably different
>disturbances and targets. That is not much to ask. Just improve on
>the performance of a single-level, single-loop PCT model.

We did. With slight modifications (different choices for sensitivity range, delay time, time constant), we got the system going with another drug. That was also the control of blood pressure. Now we are working on a very similar 'robust' control system for muscle relaxation, with excellent prospects.

(Tom Bourbon (930621.1348))

>In a reply, Bill addressed your idea that loop gain is internal to
>the device. (It is not.)

In a hifi power amp, it is. See above.

> But analogue is what living
>systems are all about. Neural currents and hormonal fluxes are
>analogue, through and through.

No, they are not. They are MODELLED as such. Neural currents come in units called action potentials, and hormones come in units called molecules. Using 'fluxes' in a MODEL is fine, as long as you do not forget that a model necessarily is a simplification of reality.

Greetings, Hans

Date: Mon Jun 21, 1993 2:42 pm PST
Subject: Two can play this game, Rick

From Greg Williams (930621 - 3)

>Rick Marken (930621.1000)

>When you're right, you're right -- and you just have to live on that pole;
>cool but comfortable.

I agree. I assume you are alluding to me being right, since you have offered no arguments, only assertions, to contradict my claims.

Thanks for the confidence,

Greg

Date: Mon Jun 21, 1993 3:02 pm PST
Subject: Re: Simultaneous vs sequential

[Martin Taylor 930621 18:50] (Rick Marken 930621.1200)

>I think there is a better (more accurate) way to characterize the difference
>between the PCT and information theory of control: rather than circle
>vs chain I think it is more accurate to characterize it as simultaneous
>vs sequential.

I don't think you can say this either. We (I, anyway) take for granted that all the variables are changing simultaneously, but unlike you, we do recognize that there is a construct "loop delay." If you try

>simultaneously solving
>the system and environmental equations that we end up finding that
>
>o = p*-1/g(d)
>
>with no contribution to the output, o, from perception, p.

you are working with no loop delay, which is unrealistic. You have to include in your analyses not just p or o or d, but p(t), o(t), d(t). If you do that, you get part-relations like $o(t_1) = f(e(t_1 - \Delta t))$ and the like. As you work your way around the loop, you then run into problems of stability at critical frequencies, so you have to do some sort of spectral compensation (e.g. integration in the output function). There's a lot that happens when you don't allow infinities in the analysis, such as infinite speed and infinite resolution.

But we do agree that all this is happening simultaneously.

Martin

Date: Mon Jun 21, 1993 3:06 pm PST
Subject: Re: assorted replies, and some more

Tom Bourbon (930621.1748)

Only time for a brief reply to part of Hans' post:

>[Hans Blom, 930621]

>(Tom Bourbon (930621.1348))

>> But analogue is what living
 >>systems are all about. Neural currents and hormonal fluxes are
 >>analogue, through and through.
 >
 >No, they are not. They are MODELLED as such. Neural currents come in
 >units called action potentials, and hormones come in units called
 >molecules. Using 'fluxes' in a MODEL is fine, as long as you do not
 >forget that a model necessarily is a simplification of reality.

Yes, they are. Don't forget, action potentials neither originate at nor travel to synapses. At the synapse, slow potentials, correlated highly with fluxes in the synaptic space, are the order of the day. I don't use fluxes in a model just to be using them; rather, I use them because the concepts of flux and current come a lot closer to describing the synaptic business of nervous systems than do the "impulses" moving along single axons. Even using the action potential as a measure, it is action potentials across axons across time that comprise the affairs of the nervous system -- back to currents and fluxes, again. The nervous system seems not to be very discreet -- oops, discrete -- in its affairs.

Until later, Tom Bourbon

Date: Mon Jun 21, 1993 3:51 pm PST
 Subject: I've seen the light

[From Rick Marken (930621.1600)]

Greg Williams (930621 - 3), Martin Taylor (930621 16:30)

Ok, you guys. I give. I know when I've been licked.

Yes, I admit it:

Stressful situations cause stress.

There is information about the disturbance in perception.

I'll even throw in "loop gain is internal to the living system" .

These are fascinating new facts about perceptual control systems. I can hardly wait to tell a Bill and Tom.

Best Rick

Date: Mon Jun 21, 1993 8:05 pm PST
 Subject: Collective Controlled Variables

[from Gary Cziko 930622.0323 UTC]

Rick Marken said:

>After thinking about this for a bit I realize that I am wrong --
 >there is no such thing as a collective controlled variable. ...

Tom Bourbon said:

>Glad you have seen the light, brother! I thought you were going a little
 >too far in your interpretation of the three-line pattern. It *is* a
 >controlled variable (but not a collective one), some times, but not others.

But don't collective controlled variables exist? I can understand why the above example doesn't count as one, but what about the situation when two or more people have a common reference level but neither alone can satisfy it. For example, my wife and I both want to see the picnic table in the shade (it is now in the sun), but neither one of use alone can get it there (too heavy and big). In this case we both have the same reference level but we don't have adequate output functions unless we put them together. Shared reference levels, combined output functions.

Or what about the warehouse worker helping the truck driver back up to a narrow loading dock? Both of them want to see the truck docked, but neither can do it alone. Here it has to do with the driver making use of the perceptual functions of the warehouse worker. Again, there are shared reference levels here. But this time the two individuals are sharing their perceptual functions.

Then there is the case when both my wife and I want to have the doors locked before we go to sleep. If she goes up to bed first, I lock the doors. If I go up first, she locks them. Here we have again a shared reference level, but we share neither output functions nor perceptions to lock the doors.

--Gary

Date: Mon Jun 21, 1993 8:37 pm PST
Subject: Miscellany

[From Bill Powers (930621.1600 MDT)] Martin Taylor 930621 11:00

>(To Rick:) We keep trying to tell you that we NEED the circle,
>and you keep trying to tie us up with the chain.

I'm curious about this. I haven't seen anything yet in the information theory stuff you've sent to me that requires a closed loop for anything. Got any examples?

>I think it was in a very early posting on this information
>theory bit that I pointed out that the fact of control demanded
>that the uncertainty about the world remained stable over time.
>That, in itself, demonstrated the need for the completed feedback circuit.

Your argument was philosophical and verbal, not mathematical. I'm still waiting to see information calculations done with a closed-loop system.

It isn't the uncertainty that has to remain stable, is it? If control is taking place under conditions of 10% uncertainty about the state of the controlled variable, it won't disappear if the uncertainty decreases to 1%, will it? It should get better even if the amount of uncertainty changes.

RE: Power gain

>Bill was talking about the power gain between the error signal
>and the resulting effector operations that directly affect the CEV.

No, I was talking about the power gain between any variable in the loop and that same variable after a complete trip around the loop. There can be power gains and losses at various points in the loop; what matters is that there be a net power gain around the WHOLE CLOSED LOOP. If you're measuring at a point where signals carry low power, then the returned signal (with the circuit broken) must have a much larger amplitude than the original signal. Where you're measuring power at the same place you're measuring amplitude, power ratios are just the square of amplitude ratios.

My point was that if we are to designate a particular part of the loop as the control system, the power gain in that part of the loop must be much larger than the power gain in the remaining part. The details of gains and losses within these parts are irrelevant.

>It seems to me that just as a power gain is an essential element of the
>outflow side of a control system, so a corresponding power loss is an
>essential element of the inflow (perceptual) side.

Not so. Control systems with no power loss at the input work quite well. However, if you're going to measure the state of a variable without significantly perturbing it (not really important in a control system because of the presence of much larger disturbances from elsewhere), you'll want to draw as little power from it as possible.

>>Poisson-distributed.

>Gaussian, actually.

No, Poisson. The Poisson distribution is appropriate for measuring signals composed of discrete impulses averaged over many parallel independent channels (electron flow, for

example). In the Poisson distribution, the RMS noise level (σ) is proportional to the square root of mean signal amplitude.

>I'm sure you sometimes feel a little frustration at people telling you that
>you said things that are the opposite of what you tried to tell them. So am
>I, but at least on this point you have come around to the "correct" view.

It gets even greater when people think I have come around to the correct view when I'm expressing things I've known since I was 18. Everyone in electronics has to learn about signals and noise and how measurements are affected by noise, integration time, and so on. All this can be expressed quite accurately without information theory, and much more simply. Maybe it's not as philosophically correct, but it gets to the same answers.

>The resolution IS infinite under certain conditions, which
>include infinite observation time and *a priori* certainty that
>the thing observed does not change over the observation interval.

The resolution is infinite if the measuring device is an analog instrument, period. Resolution is simply the minimum difference between two readings that can be indicated. Whether it indicates a real difference of the measured variable is beside the point; that is a question of precision or accuracy, not to mention epistemology.

If you take a snapshot of pure noise using an analog meter, you will get a reading on the instrument. If you take another snapshot, you will get another reading. The two readings can differ by any amount whatsoever all the way down to zero. The number of possible different readings is infinite. If you can't see any difference by eye, get out an electron microscope. There is no point on the scale at which the needle could not come to rest.

Resolution is finite for a digital measuring instrument. It is the least change in the readout that can occur: one Least Significant Digit. The digital meter simply cannot indicate any reading between two consecutive numbers on its readout. Looking at the world through a digital measuring device, one sees a discrete universe in which nothing can change by an amount smaller than one least significant digit, and in which measurements are always an integer multiple of that least digit. There are no values at all between the indicated readings.

In the nervous system, all signals pass through a pure analog stage, where the signal amplitude is represented as a chemical concentration in a cell body. The least amount by which the chemical concentration can change is by one molecule inside the cell. If you measure in terms of synaptic potentials, the least difference is even smaller than that, because it also depends, then, on the positions of molecules within the cell. For all practical purposes, therefore, the resolution of a neural signal measured in impulses per second is infinite. There is no signal magnitude that cannot be represented by a neural frequency; two signal magnitudes can differ by as little as the amount represented by a single molecule inside the cell body, or less, which is far smaller than the amounts by which natural neural signals fluctuate -- for all practical purposes, infinitely smaller.

Changing the observational interval and the integration time and all those other factors relates to the precision and repeatability of a measurement, not to its resolution. I think you're just using the word resolution with a nonstandard meaning.

Greg Williams (930621) --

>I have told them that assigning such "importances" is not a
>part of PCT modeling ...

I've forgotten the original context, and have lost most of my interest in finding out, but just for the record:

"Importance" in many contexts can be translated to "error sensitivity." That is, the ratio of effort to error. If keeping a variable at a specific level is "important" to someone (and importance is always importance to someone), then only a small error is sufficient to result in a large corrective effort. A variable of less importance to the person will show a lower ratio of effort to error.

Rick Marken (930621.1000 PDT) --

>I was trying to guess why you think there is information in perception.

If information is defined as $\log(R/r)$, where R is the range of fluctuation of a signal and r is the least discernible change in the signal, then there is information in the perception. It can be calculated in bits per second. This measure does not represent knowledge; it is not "about" something. It is just information, a quantity calculated from measurements of R and r that can be represented by a single scalar number. Do not confuse the word information with the word information. Those are two quite different words with entirely unrelated meanings. We just happen to spell and pronounce them the same way. Perhaps we should capitalize one of them to remind us of the difference, like coulomb and Coulomb, the first being a unit of charge and the second a person. Information is a measure of the range and resolution of any signal, regardless of its causes or detailed behavior through time. On the other hand, information deals with relationships to other signals and is concerned with the meaning of a signal's behavior through time: what the behavior of one signal tells us about the behavior of another one. The most we can get from a measure of Information is a relationship between the range of one signal and the range of another, in terms of the least measurable unit. From knowledge only of the range and resolution of a signal, we cannot of course reconstruct its detailed behavior through time.

v

Martin and Allen are saying there is Information in the

v

perceptual signal. You are saying there is no information in the perceptual signal. You are both right. You cannot get information from Information.

 Bruce Nevin (Mon 930621 14:44:33 EDT) --

Bill:

>> Your comments on the grape and the elephant seem to assume
 >>that the grape and the elephant are distinguished because they
 >>REALLY ARE different, whereas phonemes are distinguished
 >>because they only SEEM different.

Bruce:

>Nope. There are no phonemes there to be distinguished.
 >Phonemes are apparently present only because a perception of
 >contrast is controlled.

Try again; we'll get there yet. If you control and perceive contrast, then what you perceive is contrast. You say, "Wow, there's some contrast." You may decide that it's too much, and do something to make it less, or that it's too small, and do something to make it greater. But when you're done, you still have only what you started with: contrast.

There's a very important principle of modeling lurking here, so let's not give up on this.

Best to all, Bill P.

Date: Mon Jun 21, 1993 8:37 pm PST
 Subject: Here one is - where are you - another list of truths

FROM CHUCK TUCKER 930622

I have not read the body of any post since 930617. I was in Baltimore playing with my grandchildren (and Clark's grandchild) for the weekend. Clark has been out of town (and is out of town as I write). I think that Danny and Kent are teaching Summer school (like I am) but I can't account for the other sociologists. When I looked a my screen today and saw 73 posts my first thought was "Call a 'time out' so those of us who work can catch up" but I doubt that I could influence anyone to stop for a moment. So until I can read everything and formulate a statement which will demonstrate that all of you are into a semantic swamp up to your eyebrows I offer these statements:

THE FIVE CONCLUSIONS OF SOCIOLOGY

People influence each other

When people influence each other they generate/create/construct culture (i.e., language)

When people influence each other they generate/create/construct relationships (i.e., groups)

When people influence each other they generate/create/construct sets of relationships (i.e., organizations)

When people influence each other they generate/create/construct relationships among sets of relationships (i.e., systems or structures)

A SHORT QUOTE ON ORGANIZATIONS

Karl Weick says "Organizations keep people busy, occasionally entertains them, gives them a variety of experiences, keeps them off the streets, provides pretexts for storytelling and allows socializing. They haven't anything else to give."

More later, Chuck

Date: Tue Jun 22, 1993 1:20 am PST
Subject: Re: Cells and people: emergent phenomena vs control

[From Oded Maler (930622)]

[Bill Powers (930616.1230 MDT) and other anthropocentrist entities :-)]

* >I mean entities like the US, the scientific community, the
* >belief in a single god/theory, the white race, communism, etc.
*
* So do I. These are ideas in people's heads. They have no
* objective existence. To make them exist, you have to get people
* to AGREE that they exist -- i.e., individually believe it. As
* soon as everyone stops agreeing that they exist, they disappear.
* That makes them quite different things from what we mean by "the
* planet Mars" and "peanut butter."

Does a "living cell" has an objective existence?

*
* >A perceptual signal in your theory is a huge abstraction step
* >-- no molecule in the body has any "evidence" concerning the
* >existence of perceptual signals, not to mention whole
* >persons/hierarchies.
*
* Not so. A perceptual signal in my theory is in principle, and
* often in practice, an observable signal in a physical pathway. It
* is defined as a perceptual signal because of the way it is
* generated. The subsystem itself doesn't know about perceptual
* signals, of course; that knowledge ABOUT the subsystem must
* always be formulated from an external point of view. But that's
* what a theory is, as various contributors to CSG-L have been
* pointing out: an external view of a behaving system (but residing
* in the brain of each viewer).

A perceptual signal is observable in your theory because of the type of measurements instruments you have, and because of your own set of conceptual categories. If instead of electrodes that can record nerve pulses at some temporal and spatial granularity, you only had instruments that could detect properties of single atoms and molecules for very short periods, you would never deduce the existence of reference and perceptual signals, and not even the existence of cells.

I think that as a human you are in a similar situation toward social and historical entities.

Concerning emergence, You admit, I think, a transition from dead matter to living systems. Some organization of matter, although still governed by physical and chemical laws, made possible the emergence of such things as (chemical) reference and perceptual signals. Further developments, such as neurons, provided for more complicated emergent phenomena, etc., culminating in human beings. There is no reason, in principle, except the fact that we, the theorizers and observers, are humans, that further developments such as agriculture, language, cars or e-mail, furnish the way to higher order organisms, for which human beings are like (complicated) molecules inside a cell. I don't say such organisms are

control systems in the PCT sense; The very fact that we are what we are, prevents us from giving a scientific "evidence" for their existence.

Concerning collective reference signals, I think some participants don't get the point. Take a CROWD-like demo that will always settle in a circle. The circle IS a reference signal for the higher-order systems. The circle is meaningless for each individual who can only perceive its relation to his neighbors (this point is also related to the resistance of some to the idea of distributed representation of signals). If you deny this, and try to be consistent from the bottom up, you will deny the existence of "normal" reference signals, because when you look close enough, they constitute a set of purposeless measurable quantities. Try to think of the simplest instance of a control system, than look one level (*) below and you will see that the theorizing step you perform is not better than the step involved in hypothetizing collective socio-historical signals.

Best regards --Oded

(*) "Level" is the sense of organizational level: Atoms - molecules - cells - organs, not in the sense of the hierarchy.

p.s.

Rick said somewhere that Zoroastrian religion is not practiced anymore (btw, I think it is somewhere in Iran) and that is a an evidence that the existence of higher-order entities depends completely on the free will of individuals. This is not so, due to the following reasons:

- 1) I didn't say that such entities, (like other living systems..) cannot die. They can.
- 2) The limitation of human language and categorization system prevents us for seeing the continuity when the name and the details change. For example the, the middle-age knighthood can be transformed into the 20th century motor-bike skinheadhood. Catholic inquisition into scientific program committees. And the idea of people chosen by god to show the light to the rest of the world, can also be transformed into various forms...

Date: Tue Jun 22, 1993 6:26 am PST
Subject: Re: Collective Controlled Variables

From Tom Bourbon (930622.0817) Gary Cziko 930622.0323 UTC

>
>Rick Marken said:
>
>>After thinking
>>about this for a bit I realize that I am wrong -- there is no such
>>thing as a collective controlled variable. ...
>
>Tom Bourbon said:
>
>>Glad you have seen the light, brother! I thought you were going a little
>>too far in your interpretation of the three-line pattern. It *is* a
>>controlled variable (but not a collective one), some times, but not others.
>
>But don't collective controlled variables exist?

In a word, no. Not in the sense raised so often on the net in recent days.

>...I can understand why the
>above example doesn't count as one, but what about the situation when two
>or more people have a common reference level but neither alone can satisfy it.

They don't have a common reference level. Each has a reference level for a reference perception and neither alone can satisfy his or hers (or its -- I have become a model rights activist).

>... For example, my wife and I both want to see the picnic table in the
>shade (it is now in the sun), but neither one of use alone can get it there
>(too heavy and big).

See what I mean? Each of you has a reference perception that pertains to an object and some relationships in your shared environment. The environment comes the closest to being something you "have" in common. "Common reference signals" really are individual reference signals that bring people into mutual influence on common features in the environment. And when people act in concert to satisfy their personal references, the fact of their (the references') individuality usually becomes clear -- "Right here." "No, over here." "What do you mean over there? I thought you wanted it in the shade!" And all of the other verbal and nonverbal attempts by each person to get her or his perceptions "just right."

>... In this case we both have the same reference level
>but we don't have adequate output functions unless we put them together.
>Shared reference levels, combined output functions.

See my comments immediately above.

>Or what about the warehouse worker helping the truck driver back up to a
>narrow loading dock? Both of them want to see the truck docked, but
>neither can do it alone. Here it has to do with the driver making use of
>the perceptual functions of the warehouse worker. Again, there are shared
>reference levels here. But this time the two individuals are sharing their
>perceptual functions.

Take my eyes, please!

>Then there is the case when both my wife and I want to have the doors
>locked before we go to sleep. If she goes up to bed first, I lock the
>doors. If I go up first, she locks them. Here we have again a shared
>reference level, but we share neither output functions nor perceptions to
>lock the doors.

Yes? And?
=====

Try to portray each of your examples in a diagram showing each of the persons as a single-loop PCT model, carefully draw in the various CEVs and the paths to them and from them, for each modeled person, and draw in the paths of communicative deliberate interactions between the modeled people. I'll bet the results look more like some of the social diagrams I posted a few days ago than like a system with a common (single) reference signal, set to the same reference level for all people.

I think this artistic exercise would be especially helpful were you to make another pass at the example of the truck driver who "uses the perceptual functions" of the assistant! ("Excuse me. Can you show me where these optic nerve thingies plug in?")

Until later, Tom Bourbon

Date: Tue Jun 22, 1993 1:23 pm PST
Subject: Information (big i), social control

[From Rick Marken (930622.0830)]

I said:

>I was trying to guess why you think there is information in perception.

Bill Powers (930621.1600 MDT) replies:

>If information is defined as $\log(R/r)$, where R is the range of
>fluctuation of a signal and r is the least discernible change in
>the signal, then there is information in the perception. It can
>be calculated in bits per second. This measure does not represent
>knowledge; it is not "about" something.

Yeah. That's the kind of information there is in perception. In fact, there's all kinds of things like that crammed into perception. There is not only Information (big i) in perception, but there is average, median, variance, kurtosis, etc.

>Martin and Allen are saying there is Information in the

>perceptual signal. You are saying there is no information in the
>perceptual signal. You are both right. You cannot get information
>from Information.

Yes. This could have been solved long ago if Martin had just sent back the value of $\log(R/r)$ associated with the perceptual signal I sent. Then I would have seen the Information in that perception. It would have been something like 2.643856 (assuming \log_2).

Tom Bourbon (930622.0817) --

This post should be required reading for everyone (like myself) who was ever tempted to believe in "collective controlled variables". Very clear. I think this "collective controlled variables" topic is extremely important. The appearance of a "common reference level" in cooperative control is extremely seductive -- as seductive as the idea that there is information (small i) in perception.

Best Rick

Date: Tue Jun 22, 1993 1:23 pm PST
Subject: Re: power gain and loss

[Martin Taylor 930622 11:30] Hans Blom, 930621

>>It seems to me that just as a power gain is an essential element of
>>the outflow side of a control system, so a corresponding power loss
>>is an essential element of the inflow (perceptual) side. The
>>perceiving of the state of a CEV should not contribute as a disturb-
>>ance any more than it must (Heisenberg showed that it must, to some
>>extent).

>

>How true. Very familiar, too. In my blood pressure controller, the
>arterial pressure must be measured. Imagine what a thick needle in a
>thin blood vessel can do to disturb the circulatory system and thence
>the pressure measured!

>

>> On the other hand, the output power of the control
>>system wants to have maximum effect on the CEV, or as tight coupling
>>and as high power gain as is feasible, given the information limita-
>>tions on the perceptual side.

>

>Feel that micromanipulator carefully tear a single cell away from its
>surrounding tissue as your hand squeezes on the macroscopic counter-
>part of the micro-pincer. Power GAIN? No! Very carefully scale the
>power down! Almost Virtual Reality...

You are mixing up different control systems, I think. The perceptual energy levels must be greatly reduced from levels that would tear away that single cell; the corresponding output must come back up to a level that permits the cell to be torn away. You have to look at both in the same environment. The output must have enough energy to dominate any random effects on the CEV, and enough to counter whatever disturbance there might be (which normally would be substantially greater than the random effects). The perceptual input, on the other hand, must be based on sensors that work below the level of the random effects on the CEV, or at worst at the same energy level. There has to be a substantial power gain on the output side that (to an order of magnitude) compensates for the power loss between CEV and sensor.

The fact that in your example there are several intervening control systems that work at very different power levels is irrelevant. There may be great power gains and losses on both sides of the multi-level circuit, but in the end, they must more or less compensate when measured between the CEV and the perceptual signal that corresponds to the CEV.

Martin

Date: Tue Jun 22, 1993 1:43 pm PST
Subject: Social control

[From RICK Marken (930622.0900)] CHUCK TUCKER (930622) --

> Karl Weick says "Organizations keep people busy, ..."

PCT says "People keep organizations (as perceived by each individual) busy".

Oded Maler (930622)--

>I think that as a human you are in a similar situation toward social
>and historical entities [as cells are to the control systems of which
they are a part].

>The very fact that we are what we
>are, prevents us from giving a scientific "evidence" for their
.[higher level control systems] existence.

And you are certainly free to think that. But since, as you admit, there is no way to obtain scientific evidence about these systems, your thoughts about them (as Bill pointed out) are simply mysticism. Nothing wrong with mysticism -- it's just not what we do in PCT.

>The circle IS a reference signal for the higher-order systems.

I think you mean that the circle is the reference level of a controlled perceptual variable; the reference level is specified by the reference signal in the purported higher-level system. Bill already explained why the circle fails The Test for the Controlled Variable (even if it resists disturbances there is no sensor or output device to be found outside the group of individuals itself). You might want to read about The Test in BCP. But somehow, I have a feeling that this means of determining whether or not a variable is under control will not suffice to convince you that social variables (like the circles in CROWD) are not controlled by an invisible control system in the sky (with a white beard and flowing robes?). As long as they aren't giving public money to fund research on these invisible, higher level control systems, I have no problem with it.

Best Rick

Date: Tue Jun 22, 1993 3:23 pm PST
Subject: Information; controlling effects

[From Bill Powers (930622.0730 MDT)] Martin Taylor (930621.1630 EDT) --

>I'll simply repeat what I have said several times before. The
>only argument (initially) was that it was inconsistent to say
>two things at the same time:

>(1) that the output signal exactly mirrors the disturbance
>signal (a claim made very strongly by Rick and you in many
>postings), and

>(2) there is no information about the disturbance in the
>perceptual signal (a claim made even more strongly,
>particularly by Rick).

These arguments are indeed inconsistent; when they are stated without qualification they contradict each other. The cause of the difficulty, for which both Rick and I are responsible, is that these qualitative statements are derived from a quantitative analysis in which a parameter, $G/(1+G)$ where G is loop gain, is allowed to go to the mathematical limit of 1 to yield an approximate solution (i.e., G goes to infinity). Furthermore, even the approximate solution depends on an assumption that had been stated but tends to get left behind: that the bandwidth of the disturbance is far less than the cutoff frequency of the various components in the control loop. Strictly speaking, the above two statements are true only of a system with infinite loop gain presented with a constant disturbance.

The conditions you assumed in showing that these statements are inconsistent do so precisely by violating to the maximum degree possible the assumptions that led to our approximations. Instead of using a constant disturbance, you used a step-function to show that the disturbance must have an effect on the perceptual signal. A step function contains an instantaneous transition, which is guaranteed to be far outside the bandwidth of any physical control system.

Our statements and your counterstatements depend on adopting assumptions that go to opposite extremes and then making qualitative statements about the consequences. If the arguments then proceed on that qualitative basis, as they have done, the only possible result is a barrage of attacks and denials. In fact, your argument is unable to handle the case of a constant disturbance, while ours is unable to handle the case of a high-frequency disturbance. There is no way to reconcile conclusions based on such extreme opposite assumptions.

 To say that there is NO information about the disturbance in the perceptual signal, whether one defines information formally or informally, is clearly wrong. After a step function, the change in the perceptual signal represents the change in the disturbance at least for some short time after the step. The degree of resemblance depends on the rise-time of the input function and the time-scale on which we observe the signal (the rise-time of the measuring instrument).

If there were no feedback connection, the perceptual signal would simply rise to a new value and stay there, the new value being proportional to the disturbance (in a linear system) and accurately representing the magnitude of the disturbance. However, with the loop closed and an integrator in the output function, the initial value of the perceptual signal will exist only for an instant, after which the signal will decline exponentially to zero (we assume zero reference signal). The speed of the exponential decline will be determined by the constant multiplier in the output integration (or by the amplification of the error signal prior to the integration). This exponential decline of the perceptual signal clearly does not represent what the step-disturbance is doing; the disturbance remains constant while the perceptual signal disappears, returning eventually to the same value it had prior to the step. After the step has occurred, therefore, the perceptual signal rapidly loses its similarity to the disturbance.

This is a time-domain representation of the difference between our positions. Yours is based on the state of affairs immediately after occurrence of the step. Ours is based on the state of affairs long after the step. In your case, the change in the perceptual signal is a reasonably faithful rendition of the amplitude of the step and is essentially unaffected by the reference signal. In our case, the perceptual signal bears no relationship to the amplitude of the disturbance and represents only the setting of the reference signal.

So whether the perceptual signal in any sense contains information about the state of the disturbing variable depends on when, in relation to the time of occurrence of the step, you observe the perceptual signal. The longer you wait to observe the perceptual signal, the less able you are to deduce the state of the disturbance from it.

 All of this suggests a way to apply information theory that has a point of contact with PCT. Instead of using a step-function, use an impulse disturbance. With an impulse disturbance and no feedback connection, the perceptual signal will rise and fall in a way typical of the impulse response of the input function. The error signal, which we can assume to be instantaneously computed, follows that waveform. With the loop open in the environmental path, the integrator output will show a transition from an initial steady level to a new steady level. When the loop is closed, the perceptual signal will show the impulse-response of the closed-loop system.

Now applying superposition (this is a linear system), the dynamic response of the system can be built up by representing any arbitrary disturbing waveform as a series of impulses of appropriate amplitude and sign, and superimposing the smooth responses to the successive impulses. Therefore if you can analyse the behavior of the loop for a single impulse disturbance, you can sum over the set of overlapping impulse responses to get the continuous behavior of the system. This is a well-known method called convolution, with which I'm sure you are at least acquainted. It bears a strong resemblance to your approach that assumes sampling, but it straddles the pure "sampling" concept and the pure continuous-variable concept. While the input is treated as a purely discrete sample, the output is a continuous waveform. The waveform of the perceptual signal after an impulse disturbance of the integrating control system would be of the form $\exp(-kt) * [1 - \exp(-kt)]$: the product of a rising and a falling exponential. It seems to me that this would provide a much smoother transition from the world of strictly discrete calculations to the kinds of continuous analysis we use in PCT. And I would surely understand it better than I understand Shannon. An added advantage is that I already have a nonanalytical method for directly deriving the impulse response of a real person's control system from the continuous record of performance (assuming linearity).

>How is it possible for you to take my discussion of the
>time-dependent increase in resolution in observations of low-
>bandwidth signals, and come up with:

>>If the noise is 10% RMS of the range,
>>Martin's intepretation would be that there are only 10 possible
>>values of the perceptual signal with a range of 10 units, 0 through 9.

Easy. You said that the equivalent of r (in D/r) for a continuous system with noise would be the RMS noise level. If the RMS noise level is 10% of D, then D/r is 10, and the probability of any one measure would be r/D or 1/10. In the discrete case, D/r represents the number of possible values of d. Thus there are D/r = 10 possible values of d (or some signal) in the analog case, too, if you're computing probability the same way.

>That violates high-school electricity, let alone what a
>graduate engineer would think ...

Yes, but because of what I know about electricity, not because of what I know about information theory. I can already arrive at a correct conclusion from electricity theory; I'm trying to see how you arrive at it from information theory. So far, the information theory I know seems too elementary to come up with the right conclusion. No doubt I'm applying the calculations incorrectly.

>You assert that I gave you D/r as a measure of resolution,
>ignoring that I pointed out that for continuous signals (D+r)/r
>is more appropriate ...

So there are 11 values instead of 10 when the RMS value of r is 10% of D. Isn't that rather a quibble?

>Oh, I forgot. Psychophysical results are not to be mentioned in
>the same breath as PCT, because they might say something about
>perceptual resolution, which we know to be infinite, don't we?)
>engineer would think.

Now, now. I understand resolution to be the least possible difference between two readings. How do you understand it?

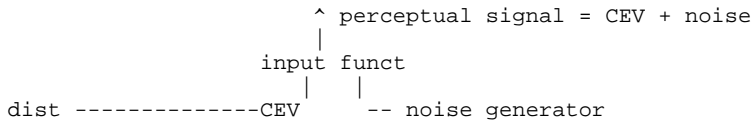
The power of a sine wave is independent of its frequency. It is

1/x (integral from x=0 to x=t (sin^2 fx) dx) (omitting the odd 2 pi factor)

You're right. Sorry. This means I can take the inverse Fourier transform of a set of equal amplitudes in an artificial transform and get the required waveform, right? I can make it random just by picking random phase-pairs for the real and imaginary parts.

>The reason is that the Gaussian variation provides the maximum
>uncertainty for a given power.

I'm getting confused here. The disturbing variable is an arbitrary waveform. If we add a Gaussian noise to it, we will still have just an arbitrary waveform, won't we? I thought the place to add the noise would be in the input function, so we have



Is there any point in putting noise into the environment? I can see using a Gaussian distribution for the noise generator, but why for the disturbance?

The output of the system has to come near mirroring the waveform of the disturbance. The disturbance waveform is just whatever it is. If you make the disturbance itself Gaussian, the output then has to be nearly the same waveform, with the same Gaussian distribution, if control is to be near-perfect. I think maybe the above diagram is really what you're getting at. The above diagram can be implemented easily with Simcon.

>If you do not have the Gaussian

>variation, the signal can be translated into another signal of
>lower power but the same resolution, that can convey the same
>information (equally, it can be converted into one of lower
>bandwidth but the same power per unit frequency or one of
>shorter time but the same power per unit time).

Have you been following my discussion with Randall? If you translate a variable with a range of 8 units and unity resolution into a signal with a range of 4 units _and_ unity resolution_, you will lose information, won't you?

Also, isn't the resolution you have to use the resolution with which you express the Gaussian distribution? If you use a 1-sigma resolution, you won't end up with a Gaussian distribution around the mean. You'll end up with a noise that changes amplitude in 1- sigma jumps.

Hans Blom, (930621)--

All right, we are getting some apparent disagreements out of the way. There are some left, but maybe they will go away, too.

>you notice how each time I stress the 'optimal' usage of what is perceived...

It's hard to define "optimal" in a massively parallel hierarchy where a jillion different goals are being satisfied all at once -- and some are in conflict. There's no evolutionary reason for optimal control, unless in "optimal" you include quitting the improvement of a control system when it's just good enough to assure survival. My view is that living control systems have just barely enough "quality" to enable them to survive to the extent that they do, and not a bit more. An engineer who has to satisfy only a few performance criteria in an artificial system can hone the behavior of the system to a very high degree of precision, speed, efficiency, and so forth. I don't think that happens in many living control systems.

>I am talking in terms of the response of a control system to,
>say, a step change of a reference level or of a disturbance. As
>you know, the response might be overly damped or oscillatory.
>By a high quality response I mean one that is approximately
>critically damped.

OK.

>I do not believe in a monolithic "I", and neither do your models.

Good.

>>The only way to make sense of self-reflexive ideas is to treat
>>a person as if that person were solid, like a potato: only the
>>whole person perceives and acts.

>Demonstrably false. In split brain patients, one hand of the
>patient may caress his wife while the other hand hits her.

Whoops. I was saying that the only way you can talk about a self describing the SAME self (that's what I meant by self-reflexive) is to assume just a single self. Otherwise you're talking about one subsystem describing another one -- and specifically not the subsystem doing the describing. Von Foerster's idea of "recursive" consciousness is what I was talking about, where the recursion involves just one system doing things to itself. His model is mathematical recursion: $\sin(\sin(\sin(\dots x)))$, in which the very same function is applied to its own value, recursively. That's what I'm arguing against.

Maybe David Goldstein will some day favor us with a description of his application of PCT "levels" to a multiple-personality patient.

>>Loop gain, in PCT, is not "internal to the device."

>In a hifi amp, it is. That was my example.

That depends on where you draw the system-environment line. The last active component in a hifi amplifier is the output power stage. You can draw the line immediately after that, for the output side. On the input side, the feedback from the transformer secondary (if used)

or from the voice coil voltage, is the sensory input from the environment. The music-signal input voltage is the reference signal in a hifi amplifier; the whole system makes the feedback from the output equal to the reference signal, and keeps it matching with a very wide bandwidth. This is no different from the way you would describe such an amplifier, but it's rearranged to show the parallels with the PCT model.

>>This separation is always important in detailed control-system design, but especially so when the effector is coupled to the controlled quantity loosely or through complex intervening processes. Then we clearly would expect the effector output to be changing far more than the controlled quantity is changing.

>That would be bad in any control system, especially if the load can vary.

On the contrary, it's a necessity if you want good control. When you design a regulated power supply, the voltage sensor does not measure the output right at the terminals of the power supply. Instead, the sensing lines are moved out to where the load will be applied, sometimes several feet away. The reason is that controlling the immediate output from the filter capacitor will allow a voltage drop along the line, and the voltage will not be well regulated at the load. Moving the sensor to the load makes THAT voltage remain constant, while the output at the capacitor varies up and down to compensate for varying losses in the connecting wires.

Likewise, if you want to control shaft speed, you don't put the tachometer right where the motor is attached to the shaft. If you do that, the shaft will "wind up" and the speed at the other end, which can be meters away, will fluctuate up and down as the load varies. You put the tach at the other end, and let the motor end advance and retard in phase as loads come and go.

You don't put the bimetallic element of a thermostat in the duct leaving the furnace. You don't put the sensor for a nuclear reactor's coolant flow on the motor that drives the valve, or on the button that turns the motor on (unless you're the designer of the Three Mile Island coolant control system). You put the sensor where the controlled variable is, and make sure it is sensing exactly the thing you want controlled. Sensing the physical output of the effector simply puts that output under control; it doesn't put the downstream consequences of that output under control.

Engineers who build control systems know this. They just redefine "output" to mean whatever variable the sensor is attached to, regardless of how far it is from the effector.

>An optimal control system needs to know as exactly as possible what the effects of its OUTPUTS are.

You put the emphasis on the wrong word: that should read

>An optimal control system needs to know as exactly as possible what the EFFECTS of its outputs are.

It doesn't give a damn what its OUTPUTS are. The outputs automatically adjust to make sure that the EFFECTS are those that are wanted. That's what PCT is about: controlling EFFECTS, not controlling OUTPUTS.

>This can be done by SENSING THE ACTIONS OF THE EFFECTORS (by sensors in or near the effectors) and OBSERVING THE REACTION OF THE OUTSIDE WORLD to the actions of the effectors (in any sensory modality).

If that's how you always design control systems, you're missing a good bet. If you can observe the reaction of the outside world, you already have a sensor for the reaction. Put that sensor in a feedback loop and you can control the reaction directly, out there at the end of the causal chain. You don't really want to control the actions of the effectors, do you? What you want to control is the effect they're having. Sometimes you can't sense the effect in time, but usually you can. When you can sense the effect, the result will always be better. When you can't, you're not going to get great control by any means -- you can't correct changes in the result due to unforeseen disturbances acting directly and independently on it, downstream from your carefully-controlled effector outputs. If you sense the result directly, letting the effector outputs vary as required, you can counteract any disturbances, even those that can't be sensed before they happen.

>Note that this directly solves the adaptation (learning)
 >problem as well: correlating the effector's action with the
 >world's reaction 'identifies' ('systems identification') the world.

But it's not as good a solution as detecting the world's reaction, comparing that against the reaction you want, and converting the difference into the direction of output that will make the difference smaller. Decidedly smaller. When you can do it that way, you don't need any correlations.

>Our effectors do in fact have built-in and/or built-on sensors,
 >as you know. In my opinion, they have this additional function
 >above the one that you describe, the comparator function.

Yes, they do. That allows us to control sensed force and very roughly, velocity and position of a limb. The comparator functions for these loops are the spinal motor neurons. But higher loops use these controlled effectors as outputs in systems that sense the EFFECTS of the motor outputs and use higher-level comparators. Nothing runs open loop in an organism, with only a few trivial exceptions. Every consequence of action is directly sensed and controlled, and consequences of those consequences are ALSO sensed and controlled. Nothing is left to change.

When necessary, living control systems at some levels can run open-loop for short periods of time. If the lights go out, you can keep walking across the room, maintaining balance kinesthetically and approximating the actions required to get to the light switch. But you won't hit the light switch on the first try, and if you have to walk too far you'll start running into things. The trouble with open-loop behavior is that it works only in an absolutely constant environment, which is hard to find. Open-loop systems have to extrapolate on the basis of the last known situation; there's no way they can deal with the effects of unpredictable disturbances (unless you build them like locomotives). And they are terrible at producing extended sequences of action in which the starting point for one action is the ending point of the previous one. The only way to make that work is with stepper motors and guiding rails -- and strict attention to preventing disturbances.

>>When we see the controlled variable separated from the
 >>effector output, we can much more easily understand that the
 >>visible behavior of an organism is really just its actuator output.

>Not when the actuator is loosely coupled to its environment
 >(through some layers of skin tissue, for example). An analogy
 >is a battery with a large internal resistance. Its 'visible'
 >voltage greatly depends on its load.

In the hierarchy, the "actuator" for higher-level systems consists of lower-level control systems. Thus what we call actions are generally actually controlled inputs of lower systems, and so don't show the loose coupling that is actually there. Only at the lowest level are the actions hidden, inside the muscles. However, when we're looking at one "action" that is really a controlled variable, we miss the variable that is being controlled by the action, so unless you understand control you'll still miss the real controlled variable. In truth, practically everything we call "behavior" is actually a controlled outcome of some lower-level action. We name behavior by the outcome, not by the action that produces it. If we see a person varying the position of a pencil held in a hand, we don't name the behavior as "moving a pencil against varying frictional resistances by altering muscle tensions." We say that the person is "writing", because what is produced by these actions is seen to be writing. Yet in the next breath we can treat the "writing" as an action, and say "He's writing a letter apologizing to his sister." Of course he isn't (at another level): he's writing words, folding paper, licking a stamp, and so forth.

I'm going to leave power gain for another time.

Best to all, Bill P.

Date: Tue Jun 22, 1993 5:28 pm PST
 Subject: Re: Collective Controlled Variables

[from Gary Cziko 930622.2000 UTC] Tom Bourbon (930622.0817) said:

>Try to portray each of your examples in a diagram showing each of the
 >persons as a single-loop PCT model, carefully draw in the various CEVs and

>the paths to them and from them, for each modeled person, and draw in the
>paths of communicative deliberate interactions between the modeled people.
>I'll bet the results look more like some of the social diagrams I posted
>a few days ago than like a system with a common (single) reference signal,
>set to the same reference level for all people.

I never meant to claim that people could share a common (single) reference signal (having two separate brains gets in the way). But they can come to have very similar separate reference signals nonetheless, such as the way you and I (but not always Rick Marken) spell English words.

I will have to go back and look at your diagrams again. I think the ones you asked me to consider drawing have already been drawn by McPhail and Tucker in their American Behavioral Scientist "Purposive Collective Action" paper.

>I think this artistic exercise would be especially helpful were you to
>make another pass at the example of the truck driver who "uses the
>perceptual functions" of the assistant! ("Excuse me. Can you show me where
>these optic nerve thingies plug in?")

You should have seen how they renovated the inside of our football stadium last year. Huge cranes were used reaching inside the stadium with the operators seated outside the stadium. Inside the stadium were guys with radios saying things like "you're getting close, slow down, a bit more to the right, take it away." Seems to me there would have been big trouble for the crane operators if the inside guys lost their eyesight all of a sudden.--Gary

Date: Tue Jun 22, 1993 5:29 pm PST
Subject: Collective Controlled Variables?

From Kent McClelland (930622)

Tom Bourbon (930622.0817) and Gary Cziko (930622.0323 UTC) have been involved in an exchange with Rick Marken about the possibility of "collective controlled variables." As a sociologist, I've been staying on the sidelines of the recent "social control" debates, partly because I've been out of town till recently, and partly because I'm still too busy from the remnants of last semester to give much attention to the net. But the points being raised are interesting enough that I'd like to jump in.

I'm willing to concede Tom and Rick's point that each human control system has its own reference levels, and that no super-ordinate social-control system exists to impose reference signals on independent individuals. Two individuals cannot literally share the same reference level. However, Gary's examples of cooperatively accomplished tasks do suggest some ideas that get lost when we insist too strongly on the independence of each social actor. Bill Power's recent post about close-order drill in the Navy (930621.0840 MDT) illustrates the issue even more sharply.

Even if two people can't share the same reference levels, social life is only possible because people can find reference levels that are similar enough to allow effective cooperation. For "practical purposes" two people's reference levels can be the same--good enough for government work, in the case of Bill's Naval comrades. I'm talking about control at the higher levels of the hierarchy, the programs in one's head for shared cooperative action, or the system concepts that imply coordination of efforts, or the rules for producing and interpreting language (to take a page from Bruce Nevin's book).

The term I like to use to designate this practical similarity of reference levels is "alignment"--a concept that I think is key for doing any kind of social analysis with PCT. This summer I'm once more trying to rework my paper on power to argue that in talking about social power we're almost always talking about the alignment of reference levels of large numbers of people. To allow cooperation, people's reference levels do not have to be identical, but they do have to be aligned. In most cases people who cooperate need to be perceiving themselves as working together on the same project, and their programs for carrying out the project have to be similar enough that their joint efforts aren't simply disturbances to each other.

Several interesting questions arise when we think about social interaction in terms of alignment:

1. How do people succeed in getting their reference levels into alignment? Sometimes it must be imitation of what the other person seems to be doing; sometimes verbal exchanges are obviously used; a lot of the time it must be trial and error, trying something and seeing if the other person resists, a practical application of "the Test".
2. How good is good enough when it comes to alignment? If the alignment isn't perfect (which according to PCT it can't ever be), when do cooperative efforts turn into a tug-of-war? (By the way, Tom, have you ever set your experiments up so that the human control system and the computer model are simultaneously trying to do the same thing, both controlling the same cursor? What happens? How does it feel? How similar do reference levels have to be for the two to be doing the "same" thing?)
3. What are the consequences when lots of control systems, not just two, are effectively aligned in some coordinated effort? Do they become indistinguishable from one "super-powerful" control system? How is the collectivity different from a single high-gain system?

There's obviously more to say on this, but I'll go back to trying to put it into a paper.

Kent

Kent McClelland	Office: 515-269-3134
Assoc. Prof. of Sociology	Home: 515-236-7002
Grinnell College	Bitnet: mcclel@grinl
Grinnell, IA 50112-0810	Internet: mcclel@ac.grin.edu

Date: Tue Jun 22, 1993 6:36 pm PST
Subject: Re: Miscellany

[Martin Taylor 930622 18:00] Bill Powers 930621.1600

>>(To Rick:) We keep trying to tell you that we NEED the circle,
>>and you keep trying to tie us up with the chain.
>
>I'm curious about this. I haven't seen anything yet in the
>information theory stuff you've sent to me that requires a closed
>loop for anything. Got any examples?

The analysis of gain-bandwidth-resolution depended entirely on the loop, didn't it? I don't see why you are "still waiting to see information calculations done with a closed-loop system." How could the analysis have worked otherwise (whether or not you accept its correctness)?

>>Bill was talking about the power gain between the error signal
>>and the resulting effector operations that directly affect the
>>CEV.
>
>No, I was talking about the power gain between any variable in
>the loop and that same variable after a complete trip around the
>loop. There can be power gains and losses at various points in
>the loop; what matters is that there be a net power gain around
>the WHOLE CLOSED LOOP. If you're measuring at a point where
>signals carry low power, then the returned signal (with the
>circuit broken) must have a much larger amplitude than the
>original signal. Where you're measuring power at the same place
>you're measuring amplitude, power ratios are just the square of
>amplitude ratios.

Correction noted. I see the point. But I think it reinforces what I said to Hans Blom, doesn't it?

>>It seems to me that just as a power gain is an essential
>>element of the outflow side of a control system, so a
>>corresponding power loss is an essential element of the inflow
>>(perceptual) side.
>
>Not so. Control systems with no power loss at the input work
>quite well. However, if you're going to measure the state of a

>variable without significantly perturbing it (not really
>important in a control system because of the presence of much
>larger disturbances from elsewhere), you'll want to draw as
>little power from it as possible.

I think a little misunderstanding. Your power loss is presumably between the sensor and the rest of the perceptual system. I was referring to what I think Hans was talking about, the different power levels of the CEV itself and that taken by the sensor apparatus detecting the CEV state. Your second sentence reflects this. If the disturbances are "much larger" than the binding energy levels of the CEV, they are likely to destroy rather than disturb the CEV. I think we must be talking at cross-purposes in some way I don't understand here.

>>>Poisson-distributed.

>

>>Gaussian, actually.

>

>No, Poisson. The Poisson distribution is appropriate for
>measuring signals composed of discrete impulses averaged over
>many parallel independent channels (electron flow, for example).
>In the Poisson distribution, the RMS noise level (sigma) is
>proportional to the square root of mean signal amplitude.

Different conditions. Gaussian distribution for maximum uncertainty. Poisson for impulses that have a small finite probability of occurring in any fixed time interval. Neither for neural impulses that have a pretty regular pulse rate, though either or neither may be approached by averaging over many channels, depending on how the different channels relate to each other. And later in the posting you insist that what we should be talking about are continuous signals, anyway, which is the domain of interest to me.

>The resolution is infinite if the measuring device is an analog
>instrument, period. Resolution is simply the minimum difference
>between two readings that can be indicated. Whether it indicates
>a real difference of the measured variable is beside the point;
>that is a question of precision or accuracy, not to mention
>epistemology.

A real difference in what we are talking about. I treat a measurement as being *about* something. All information (or Information, if you want) is *about* something. Sure, the analogue instrument has a potentially infinite resolution. But no observer of the analogue instrument knows what its reading is to that precision. If the analogue instrument is reading something that is about an interesting value in the world, its reading is not infinitely precisely related to the actual (unknowable) value. It is an uninteresting fact that the analogue instrument has a continuum of possible values, all along the real number system. The interesting fact is to what degree the observable reading on the analogue instrument reflects the value of the thing in the world that it is measuring. That determines the amount of information about the interesting value that is obtainable from the measuring instrument. I don't see that as a question of epistemology. It is a question of how you can build a control system in a real world, as opposed to a simulated world in which arbitrary precision can be attained.

>Changing the observational interval and the integration time and
>all those other factors relates to the precision and
>repeatability of a measurement, not to its resolution. I think
>you're just using the word resolution with a nonstandard meaning.

Well, it is the standard I've always been familiar with. But I'm happy to change the word, if only we can come to an understanding of what really happens in a sensory-perceptual system that is part of a control system. Words don't matter, provided we agree on what actually happens.

>If information is defined as $\log(R/r)$, where R is the range of
>fluctuation of a signal and r is the least discernible change in
>the signal, then there is information in the perception. It can
>be calculated in bits per second. This measure does not represent
>knowledge; it is not "about" something.

Information can't be redefined as $\log(R/r)$, and it always has to be "about" something. Under some special and unusual conditions, it may take $\log(R/r)$ as its value.

>...

>Information is a measure of the range and
>resolution of any signal, regardless of its causes or detailed
>behavior through time. On the other hand, information deals with
>relationships to other signals and is concerned with the meaning
>of a signal's behavior through time: what the behavior of one
>signal tells us about the behavior of another one.

Information is a measure of the relative probability distribution of some value (possibly multidimensional) given some other facts. It is about relationships. Under some special circumstances it is related to (not a measure of) the range and resolution of a signal. Information is about how much one signal tells us about another signal. Information is ALWAYS about something, and it depends entirely on the presuppositions and prior structure of the recipient. It cannot be measured in the abstract, as an absolute physical quantity, though in many circumstances extrema can be determined for information rates or the mutual information between signal sources (not between signals). Information and meaning are intimately related, though not identical.

Perhaps we should start compiling a list of myths about information, parallel to the list of myths about control. Like the PCT myths, many of the information myths are widely believed. They can be (and obviously are) equally misleading.

=====

>Bruce:

>>Nope. There are no phonemes there to be distinguished.
>>Phonemes are apparently present only because a perception of
>>contrast is controlled.

>

>Try again; we'll get there yet. If you control and perceive
>contrast, then what you perceive is contrast. You say, "Wow,
>there's some contrast." You may decide that it's too much, and do
>something to make it less, or that it's too small, and do
>something to make it greater. But when you're done, you still
>have only what you started with: contrast.

I have refrained from commenting on this question, but I have been rather surprised that you and Rick have had problems with Bruce's point that contrast is the important issue in language. Try listening to your own speech, to see whether you say a particular phoneme in the same way in different contexts, or the same word the same way in different dialogue contexts. If you do this analytically, using a tape recorder, you find that the first time a new major content word is used in a dialogue, its phonemes are pronounced in a more extreme way than on later occurrences--each phoneme contrasts more with the neutral central form early rather than later--and you may even find that the word, if complex, loses whole or partial syllables if it is repeated often enough in the dialogue. The whole structure of language is based on the ability to discriminate, whether it be phonemes, words, or argument structures.

Perhaps we do have phoneme detectors, of a form. I'm not greatly inclined to that view, given the evidence that the ability to perceive (consciously) phonemes depends on whether the listener has learned to read in an alphabetic rather than a syllabic script. Phonemes seem to be a construct of linguists as much as of language users. I'm more inclined to think of parallel feature value detectors (sophisticated filters) rather than category detectors (which phoneme detectors would be). Granted, there has been moderate success in making speech recognizers that generate phoneme probability distributions as an early stage of recognition, but much of their success has been in handling phonetic context properly to take account of the vast acoustic differences among instances of the "same" phoneme, and of the great overlap acoustically across instances of "different" phonemes. It depends a great deal on the word (in French) whether a particular sound might be l, r, s, or f (I think I got those right), and (in English) l and o are often quite alike.

I should be surprised if there are not structured perceptual input functions that give higher outputs the closer the input is to their "preferred" form. I should be surprised also if these "filters" are unaffected by prior context.

It is sometimes hard to recognize how much of language differs from one language family to another, and how much of what we take for granted in our own language simply doesn't apply to another. Linguists (and more particularly psycholinguists, since we are talking about people rather than mathematics) disagree, and I'm sure that a lot will disagree about what I wrote above. But it does come from a professional association with the field, not from

personal or folk observation. (I know that in this group, that makes what I said suspect. Too bad.)

I hope that I will continue to ignore this thread in the near future. I, like most of the posters here, have too much else to do.

Martin

Date: Tue Jun 22, 1993 7:35 pm PST
Subject: Re: assorted replies, and some more

From Tom Bourbon (930622.1157) Hans Blom, 930621

```
>(Tom Bourbon (930621.1323))
>>
>>control, as it is recreated or predicted by PCT models, and show me,
>>in the results of simulations, how making more "information"
>>available "within the control system" improves the recreations and
>>predictions from the model.
```

That remark was in as post directed to Martin Taylor. It was a reiteration (the third or fourth) of an offer I first made a couple of years ago. Rather than assert that Martin's claims that information theory can contribute to PCT (no, that PCT derives necessarily from information theory), I suggested that Martin and anyone who cares to join him take one published example of prediction by a PCT model and show that procedures or measures from information theory improve the predictions. When that demonstration is achieved, there will be no need to debate whether information theory can contribute to PCT; it will already have done so.

The offer remains in effect. The additions, which must be information theoretic, must improve the correlations between predicted and human performance beyond the present .996 or .997. Stated another way, the additions must "explain" or "account for" more than the .991-plus of the variance accounted for by the present PCT model.

```
>Let me give you a practical example from my blood pressure controller.
>The arterial blood pressure decrease delta_p due to an infusion flow
>rate i can be modelled as
>
>      delta_p = sensitivity * i
>
>where sensitivity is a constant that describes an individual's sensi-
>tivity for the drug used. It can vary by a factor of 80. That is just
>the static response, that is obtained as soon as a stable pressure is
>established. The dynamics of the process can be modelled by a delay
>time of about one minute and a dominant time constant of about one
>minute as well, in case you're interested. I will leave out other
>gruesome details, such as an often (but not always!) very pronounced
>non-linearity of the response and frequent sensor malfunction due to
>drawing blood through the same line that the pressure is measured
>with.
>
>In this controller it was crucial to estimate the the individual's
>sensitivity (not necessarily very accurately; within a factor of two
>was good enough) to obtain a step-response that did not oscillate yet
>was fast enough to satisfy the anesthesiologist. Note that the sens-
>itivity was not known and could not be established beforehand and had
>to be established while controlling. I dare you to design a servo-
>mechanism that can control under these circumstances. I tried but
>failed miserably.
```

Nice, but not an example of the results from an application of the simplest PCT model to human performance. We are not attempting to be control engineers. Horrors, that would mean we might make a living out of our work!

```
>>Then show me that those results generalize, with no further tinkering
>>with the model, to new conditions, with unpredictably different
>>disturbances and targets. That is not much to ask. Just improve on
>>the performance of a single-level, single-loop PCT model.
```

>We did.

Did you? Where are the results of using the PCT model to recreate or predict an instance of human performance?

> With slight modifications (different choices for sensitivity
>range, delay time, time constant), we got the system going with
>another drug. That was also the control of blood pressure. Now we are
>working on a very similar 'robust' control system for muscle relaxa-
>tion, with excellent prospects.

Again, that is nice to know. It reaffirms my genuine belief that control engineering is a thriving enterprise and that it can lead to useful, many times important, products. I know that many aspects of our lives have been influenced by devices crafted by control engineers. But I am a reformed psychologist who recently renounced that professional title. I am trying to understand control by living things. Even with my severely limited skills, I have been able to apply the PCT model to instances of control by individuals, pairs, and groups of four. The results amaze me, coming as I do from an academic and research background in which statistical trash passes as the currency of the realm. I am eager to see the PCT model added to and improved in any way possible -- so long as the addition or improvement works -- and my criterion for "working" is improvement in prediction, either by producing closer fits for previously studied phenomena, or by predicting phenomena not previously modeled. When I ask for demonstrations, that is what I mean, nothing more.

Until later, Tom Bourbon

Date: Tue Jun 22, 1993 7:36 pm PST
Subject: contrasts as phonemic elements

[From: Bruce Nevin (Tue 930622 13:16:16 EDT)] Bill Powers (930621.1600 MDT)

> Bruce Nevin (Mon 930621 14:44:33 EDT) --
> Bill:
> >> Your comments on the grape and the elephant seem to assume
> >>that the grape and the elephant are distinguished because they
> >>REALLY ARE different, whereas phonemes are distinguished
> >>because they only SEEM different.
>
> Bruce:
> >Nope. There are no phonemes there to be distinguished.
> >Phonemes are apparently present only because a perception of
> >contrast is controlled.
>
> Try again; we'll get there yet. If you control and perceive
> contrast, then what you perceive is contrast. You say, "Wow,
> there's some contrast." You may decide that it's too much, and do
> something to make it less, or that it's too small, and do
> something to make it greater. But when you're done, you still
> have only what you started with: contrast.

And that's what phonemes are. They are phonemic distinctions.

You can't make the contrast between phonemes too much. There are physiological limits to how far you can extend the active articulators, there are acoustic limits on the effect of some articulatory changes, and (most importantly) there are limits on combinatorial possibilities for contrastable articulations that have contrastive acoustic effects, and for contrastive acoustic effects that result from contrastable articulations.

The contrast might be made less for various reasons. You might not need to distinguish t from d for recognition of the words in "it edited it", for example. The contrast might be made more, for example if the person didn't recognize the intended words first time, in noisy conditions, etc. An increase in contrastiveness is traditionally called "fortis" articulation for consonants, and a decrease in contrastiveness is called "lenis" articulation. This pair of terms does not apply to varying contrastiveness of vowels, though the terms "peripheral" vs. "centralized" often apply (referring to distribution in the "vowel triangle" in the acoustic space defined by making F1 and F2 the axes of a

Cartesian graph). However, making contrast less is generally called "lenition" and increasing it is sometimes called by the parallel term, "fortition".

If I say "This one is butting out" with the alveolar flap [D] that is ambiguous between t and d, and you say "where are the buds" (and are not just making a pun), my clarification is "I said buTTing out, not buDDing out." Even if I omit the second phrase, it is implicit and understood. I would never say "I said buTTing out, not bu[D]ing out" (where [D] reverts to the lenis articulation that confused you), and you would never understand that as the implicit juxtaposition were I to omit saying the second phrase. Clearly, what is controlled here is the contrast and not the individual sound segment or gesture-combination segment.

The notion that there must be something like the letter t or the letter d there, occupying a sequential position among other such elements making up the word "butting" or "budding", is an artifact of our alphabetic way of representing speech by letters. Visual perceptions do not align well with the combination of acoustic and articulatory perceptions that constitute speech. It works because the graphical objects that are letters are taken as representing the contrasts. Ignoring the problems of English orthography (which arise in large part because the language has changed since the writing was standardized), t in the string "butting" is a representation of the contrast between "butting" and "budding", "bunning", "bussing", "buzzing", etc.; or, more compactly stated, between the sound/gesture features that differentiate those words from one another and the corresponding features (perceptions) in "butting".

It works, too, because there is precedent in how language is acquired. In a pre-speech infant, there are sound/gesture combinations that are just part of its "play," producing behavioral outputs and taking in their perceptual consequences. At first, the infant can produce little more than things we might write "ma", "ba", and "pa" (or m@, b@, p@ with some indistinct, centralized sound for @) and "nga", "ga", "ka", because the larynx has not descended and there is little room for movement of the tongue in the mouth independent of the mandible. (Simians never get much past this, that's an important reason they don't have speech.) But these are just accidental cooccurrences of gestures with the lips, the tongue (undifferentiated), the velum, and the larynx, in various combinations. Later, as the larynx descends and there is more space for the tongue to move, other possible combinations appear. During all this time, the parents and other adults are recognizing syllables and words much as we recognize horses and bunnies in the shapes of clouds. "She said 'mama,' I know she did!" Sooner or later, the child comes to recognize some of the adult productions as repetitions of her productions, and begins to produce utterances as deliberate repetitions of adult words, with a growing sense of their significance.

However, there are still no phonemes. Every word is an event with no systematic relation to other words. At some point there comes control of the contrasts between similar words, and control of the contrasts themselves--

larynx open (voiceless) vs. larynx at the critical degree
of closure for voicing

velum up vs. velum down (nasalization)

wide opening vs. mid opening vs. critical closure (for turbulence) vs.

eclosure
complete closure for oral articulators (tongue and lips)

tongue tip vs tongue blade vs. tongue back (vs. root in many languages)

teeth vs. alveolar ridge as passive articulator with tongue (e.g.
s-sh contrast in English, part of w-v contrast)

It is not the specific gestures that the child controls--she's been doing that all along--its the contrasts between them, the significant differentiating of them along orthogonal dimensions for contrast. (The specific parameters and values I have listed are roughly those used at Haskins Laboratories.) The sound/gesture combinations that the child has been controlling all along in babbling play now become exploited as ways of representing the contrasts between words. Whereas before there was a continuum of sound/gesture possibilities, now a polarization is imposed.

You could argue that the child controls "for" an ideal sound/gesture for each phoneme-segment, which just happens to contrast with each other sound/gesture

phoneme-segment in the language, and that control is imperfect because gain varies and because control of nearby segments disturbs control of the current segment ("coarticulation effects"). This leaves the ready recognition of divergent pronunciations, and the systematic variation of one's own dialect according to social context, still to be explained and somewhat puzzling. If the child is controlling contrast, you get that explanation for free. Changes from one dialect to another are systematic, preserving all or most of the contrasts (sufficient to enable word recognition, or else by definition the dialects would have diverged into distinct languages).

Incidentally, Chomsky and others never understood this aspect of Harris's work. The attack on "taxonomic phonemics" by Chomsky and Halle in the early 1960s applies to many other linguists' work, but not to Harris's. For the others, contrast was to be defined by analysis of "the data of pronunciation." For Harris, contrast was given as an observational primitive by the Test. This makes a very great difference.

Bruce bn@bbn.com

Date: Tue Jun 22, 1993 10:08 pm PST
Subject: The emperor's old theory

[From Rick Marken (930622.2200)] Martin Taylor (930622 18:00)--

> Information

>is about how much one signal tells us about another signal. Information
>is ALWAYS about something, and it depends entirely on the presuppositions
>and prior structure of the recipient.

(NB. Bill Powers. Looks like Martin et al have been talking about information (small i) all along.)

So, I sent one signal (p) along with the "presuppositions" (r) and prior structure (O()) of the recipient. How much does that signal tell you about another signal, d, that was the one (along with o) that determined the value of p?

Let me strongly echo the sentiments expressed by Tom Bourbon (930622.1157):

>Rather than assert that Martin's claims that information theory can
>contribute to PCT (no, that PCT derives necessarily from information
>theory), I suggested that Martin and anyone who cares to join him take one
>published example of prediction by a PCT model and show that procedures or
>measures from information theory improve the predictions. When that
>demonstration is achieved, there will be no need to debate whether
>information theory can contribute to PCT; it will already have done so.

How about it, IT fans. How about one (count them one -- a mere one half bit of information worth) example of a way that IT can help us understand purposeful behavior. An improved prediction of control, a la Tom's suggestion above, would sure be convincing. Or perhaps an example of an illuminating experiment based on information theory; perhaps one that hones in on the perceptual variables that are actually being controlled in some situation -- IT had a big impact on perceptual psychology; perhaps it's in that area where we can find the IT "beef". As it sits, even after having caved in and accepted the fact that there is all this information in perception, I still don't know how to find it, measure it, test for it, or put it into models. I feel a bit like the emperor who has just been handed a nice new set of clothes that everyone seems to think are fantastic, but look rather transparent. Other than the fact that everyone says "information theory is great", why shouldn't I go with my perception that there's nothing there? Or is that what makes it so great -- the fact that there really is nothing there?

Best Rick

Date: Tue Jun 22, 1993 11:25 pm PST
Subject: transitivity of information

[From: Bruce Nevin (Tue 930622 10:19:25 EDT)] Martin Taylor 930621 16:30

> The only
> argument (initially) was that it was inconsistent to say two things at

> the same time:

> (1) that the output signal exactly mirrors the disturbance signal
> (a claim made very strongly by Rick and you in many postings), and

> (2) there is no information about the disturbance in the perceptual
> signal (a claim made even more strongly, particularly by Rick).

> You simply cannot have both (1) and (2) at the same time. They contradict.

It follows from (1) that there is information about the disturbance in the output. For (2) to follow, it seems to me that you have to demonstrate some kind of link between the output and the perceptual input, such that there is information about the output in the perceptual input, and (furthermore) that the information about the disturbance that is in the output (for an outside observer) is preserved amid the information about the output that is in the perceptual input (for the control system). Have you demonstrated this link, or something like it, and I missed it?

Bruce bn@bbn.com

Date: Wed Jun 23, 1993 12:54 am PST
Subject: phoneme detectors

[From: Bruce Nevin ()] Rick Marken (930621.1500)

It seems that we are in violent agreement about at least one issue here. I said:
> >The phonemes are not in the noises. They're in your head.

You agreed:

> That's right -- as the outputs of perceptual functions (phomene detectors).

The input to a phoNeme detector is not a phoNeme; the output is.

I only have a problem with your notion of lineal causality here:

> The phomemes are in the brain (as perceptual signals); the
> causes of the phomemes are the sounds in the tubes and trees.

"Cause" is a misleading word. First (in temporal order), the sounds are not the causes of the forms of input functions of phoNeme detectors as a child learns or "acquires" the phoNemic contrasts of a given language. OK, so we've put aside the learning problem and we're talking only about the perception problem after the phoNemes are set. Secondly, after the child has learned the phonemic distinctions of the language, the same sounds may cause more than one phoNeme detector to fire: "It's a t!" "It's a d!" "It's an r!" (this for the middle consonant of the first word in "latter-day saints", in a U.S. pronunciation--the same sound is an r in some pronunciations of "throw"). I believe you would want to say that only one phoNeme is actually there, right? Don't just read the letters, say the following sentence several times, in an ordinary, conversational way:

It edited it.

You will probably hear four occurrences of this "flapped r" between the vowels i_e_i_i_i, followed by one released t at the end.

I asked:

> >how can they be different phonemes for different people?

You said:

> Different phomeme detectors.

I asked:

> >Then how come all the native speakers of English hear the same
> >phonemes from the tubes in the exploratorium?

You said:

> They have developed similar phomeme detectors -- if they didn't,
> they would not be able to perceive the phonemic distinctions that

> make if possible for them to communicate with each other -- assuming
 > that they WANT to communicate with each other, of course, which
 > they apparently do.

How does this "social agreement" between their phoNeme detectors come about? What is going on here, and how can we model it? I have been trying to get beyond hand-waving.

You say:

> "Reconstructs" may be a bit much. The PCT perceptual model just says that
 > there is a perceptual signal in the neural output of a phomene detector.

And farther on:

> you may be
 > mistaking PCT for an "analysis by synthesis" version of speech recognition.
 > Basic PCT is "analysis by transduction".

I guess you missed it when we talked about reasons for believing that some form of analysis by synthesis was necessary for language. I guess also you weren't watching this thread closely when I cited evidence that there are no reliable cues in the acoustic signal for identifying phonemes. Bill's response was essentially "that can't be true, because the simplest version of PCT demands that phonemes be recognized from the acoustic image, so the people whose work you cited must all be incompetent." Or maybe you were there and just tacitly agreed with Bill. Let me ask you, then: is that your position too?

Identical acoustic input is perceived as one phoneme in context A and as a different phoneme (or even combination of phonemes) in context B, where the change in context may be different neighboring sounds, different word (identified by neighboring words making up the syntactic/semantic environment), or shift from one person's speech to another person's speech. Pick any acoustic feature that you think is essential to perceiving a given phoneme. It can be excised, and the phoneme is recognized on the basis of what remains. (Excise enough, of course, and you get a signal that is ambiguous over a class of two or more phonemes.) I certainly agree with your basic conception, as far as it goes:

sounds ----> PD --->p

The input function of PD is where the interest lies. The input function of PD includes input from parallel and higher levels, in addition to "sounds" from below. (Higher levels by the imagination loop.) The same input function and the same PD controls your own production of the phoneme as well as your perception of the other person's phoneme. Where actual input is lacking, it is supplied by imagination: as I speak, the imagination loop fills in what you are supposed to be hearing; as I listen, the imagination loop fills in what I would be saying if I were producing those sounds. The diagram for all this parallel control and the imagination loops is more than I can undertake. But it doesn't seem PCT-heterodox to me, and it seems to account in a very direct way for the experience of recognizing words and repeating words in a recognizable way.

You said:

> you want language to be a special kind of
 > phenomenon; a perception that ONLY exists because it is
 > controlled.

And farther on you said:

> I guess the disturbance to me
 > about your position is this "mystique" about language.

This is not unique to language. You recognize the sounds of different objects being manipulated in a sound-producing way because you have manipulated them or objects like them so as to produce sounds, or you have seen it done. No experience, no memory, no imagination, then no recognition. The association of sounds with visual perceptions and tactile, kinesthetic, and other perceptions, is something that we've been working on since infancy. And since infancy we've worked especially hard at the association of speech sounds with intended speech gestures, and then with intentionally distinct speech gestures, distinctions that make the difference between one word and another, similar word. This is important to us, so that we can know what words the other person intends us to hear, and so that we can produce words in such a way that the other person recognizes the words that we intend them to hear, despite the degeneracy of the acoustic signal.

This is not unique to language. It is just that language is what I am interested in and what I know most about. I asked an analogous question about the Gather program: how can we model individuals coming to have the same setting for a perception of appropriate

inter-individual distance? Presently, we just set those values from outside the model. This surely is not how it happens with people.

When you recognize the sounds of objects you do so by imagining the way in which the object produces the sound. When you recognize a phoneme, you do so by imagining the way in which a phoneme is produced. It happens that you yourself produce phonemes, and it happens that it is important for your phonemes and the other person's phonemes to be mutually recognizable by you and she. These two added conditions make the difference between phoneme recognition and recognizing what kitty knocked off the shelf in the next room. You discriminate between the voices in a performance of a Bach keyboard piece, but discrimination is not the same as recognition, no? You discriminate between the voices conversing in the next room, and you recognize your wife's voice and her friend's voice. But this general recognition of voice quality is scarcely as specific and detailed as word-recognition. Nor is it socially set. By that I refer to the fact that I could recognize the words if you sent me a tape, but I would not recognize the voices. Sure, I don't have a "Mark's-wife-detector" MWD and you do. Nor is it important or particularly useful to me to (however important and useful it may be to you!). The two conditions just mentioned above make all the difference. We don't model the development of a PD or a MWD. We just take it as given that they exist. That's the learning problem.

Bruce bn@bbn.com

Date: Wed Jun 23, 1993 5:49 am PST
Subject: Re: Collective Controlled Variables

[From: Bruce Nevin (Wed 930623 08:47:27 EDT)]

If language is based on control of perceptions of contrasts or distinctions, I think that I no longer need to suppose that speakers of the same language share similar settings of similar reference perceptions. The similarities across individuals fall out from maintaining distinctions in a recognizable way.

Bruce Nevin bn@bbn.com

Date: Wed Jun 23, 1993 6:32 am PST
Subject: Re: Collective Controlled Variables

From Tom Bourbon (930623.0850) Gary Cziko 930622.2000 UTC

>Tom Bourbon (930622.0817) said:
>

>>Try to portray each of your examples in a diagram showing each of the
>>persons as a single-loop PCT model, carefully draw in the various CEVs and
>>the paths to them and from them, for each modeled person, and draw in the
>>paths of communicative deliberate interactions between the modeled people.
>>I'll bet the results look more like some of the social diagrams I posted
>>a few days ago than like a system with a common (single) reference signal,
>>set to the same reference level for all people.

>
>I never meant to claim that people could share a common (single) reference
>signal (having two separate brains gets in the way). But they can come to
>have very similar separate reference signals nonetheless, such as the way
>you and I (but not always Rick Marken) spell English words.

I really didn't think you did, but you were posting on a net where some participants have come to believe in unicorns and social control systems. I was only encouraging you to remember *they* are watching and often take such things literally. (I tried to indicate my tongue in cheek attitude by the silly remarks about eyeballs and optic nerves.)

>I will have to go back and look at your diagrams again. I think the ones you
>asked me to consider drawing have already been drawn by McPhail and Tucker in
>their American Behavioral Scientist "Purposive Collective Action" paper.

Again, that was a hint more for the watchers than for you. But you *did* leave a nice opening.

>>I think this artistic exercise would be especially helpful were you to

>>make another pass at the example of the truck driver who "uses the
>>perceptual functions" of the assistant! ("Excuse me. Can you show me where
>>these optic nerve thingies plug in?")

>

>You should have seen how they renovated the inside of our football stadium
>last year. Huge cranes were used reaching inside the stadium with the
>operators seated outside the stadium. Inside the stadium were guys with
>radios saying things like "you're getting close, slow down, a bit more to
>the right, take it away." Seems to me there would have been big trouble
>for the crane operators if the inside guys lost their eyesight all of a sudden.

No doubt! But that has nothing to do with sharing someone else's perceptual functions --
not in the sense some people on the net probably would take it. It would be a lot more
like my diagrams, which differ in some important respects from Clark's.

Until later, Tom

Date: Wed Jun 23, 1993 7:42 am PST
Subject: PCT and other theories

[From Bill Powers (930623.0700 MDT)]

Yesterday I got a total of 35 posts. If I took an average of 15 minutes to consider each
one and reply to those about which I had something to say, I would have devoted 8 hours and
45 minutes to the mail. This has been going on for many months, and I'm not the only one
who has been swamped. What is it that's been taking up so much of our time?

By and large, it's the same thing that has caused PCT to take 40 years to get where it is:
arguing with people who have other points of view to promote. I am coming close to
concluding that this is a waste of time.

Consider the information-theoretic thread. Does it really matter what information theory
has to say about control theory? To those interested in information theory, obviously it
does. Does it really matter how PCT fits in with various (competing) concepts in
linguistics? How PCT compares with existing sociological theories? With personality
theories? With theories of education? With cognitive theories? With chaos theory, or
optimal control theories, or action theories? To those whose primary interest is in
promoting those theories, or whose expertise lies in those areas, it is clearly important
to discuss how control theory relates to, or can be absorbed into, each of those theories.
But is that important to the development of PCT?

In each case the answer is no. The development of PCT rests on putting it to experimental
test, discovering its flaws, and working out ways to correct the flaws by improving the
theory. Trying to work out its relationship to all the other theories of behavioral
organization that exist is a waste of time, because by the time PCT has developed enough to
deal with the kinds of problems implied by these other theories, assuming it ever does,
NONE OF THESE OTHER THEORIES WILL EXIST ANY LONGER. By the time PCT has developed enough to
handle the same subject matter, the subject matter itself will have taken on a completely
new appearance. What will PCT have to say about intelligence, or aggression, or you name
it? Nothing, because those words won't mean anything any more.

A little light went on in my head this morning. Why, I thought, are PCTers asked so often
what PCT has to say about this fact or that theory or the other explanation of behavior?
Obviously, because the askers have not worked out any answers for themselves. They are
throwing up a challenge and waiting for someone else to meet it. In many cases, they are
waiting for an answer to knock down. But why do they need to ask?

Obviously, because they don't understand PCT well enough to answer the question themselves.
But if they don't understand PCT well enough to answer the question, how are they going to
understand the answer? In fact they generally don't: the next word they utter following the
answer is "But ..." The answer makes sense within the context of PCT, but not within the
context of whatever other theory they're coming from. So they are left with the same
problem: not understanding PCT.

Of course those who are trying to teach other theories have exactly the same problem: an
information theorist offers an explanation, and the PCTer replies "But ...". This shows the
futility of trying to use one theory as the medium for understanding another, particularly
when they deal with orthogonal subject matters (for which expression we may thank Hans
Blom).

In teaching PCT, the object is for the learner to grasp its principles and structure to the extent that they have been developed, so that when a question comes up regarding observations of behavior or other proposals concerning behavior, the learner can apply PCT to working out an answer -- or, given enough advancement in the art, so the learner can realize that there is something wrong with the theory that needs fixing. When a student is learning how to solve the quadratic equation, the point is to learn the method of solution and the principles involved; it's not to deal with every new algebraic expression by going to the teacher and saying, "I solved the last problem, but this one is different -- how do you solve this one?" This is not learning how to solve the quadratic equation; it's avoiding learning. You learn by trying to apply the principles and the method, and when you fail, asking about the principles and the method, not about the answer. The answer has no importance at all; when you understand the principles and the method, you'll know when the answer is right.

One of the giveaways concerning the degree of understanding is in the reaction to PCT "slogans." We say "It's all perception." We say "Control systems control their own inputs." We say "Control systems control outcomes, not outputs." To a person who doesn't understand PCT, these are indeed just slogans; they sound like the basic dogma of PCT. But to a person who does understand PCT, these statements are simply succinct summaries of deductions that rest on a closely-reasoned analysis of the way all control systems work: they're conclusions, not premises. They're where you end up, not where you start.

A lot of the verbiage I've churned out on the net has been by way of offering answers to questions. That's fun; it shows how much I know, or think I know. But what good does it do anyone who doesn't understand how I got to the answers? I don't have any magical tricks or any genie whispering in my ear. All I do is apply PCT. If I can do that, so can anyone else. There's nothing secret about it. Or there isn't any secret once you have assimilated the theory. That's what's so gorgeous about PCT: no secrets. No arbitrary assumptions "for the sake of the argument." No appeals to philosophical generalizations. No references to 200 authorities you must have read before you can understand the answer (although reading an algebra book would help).

As a way of opening a conversation, answering a question about some aspect of behavior isn't a bad move. As soon as the "But..." occurs, however, the next thing to do isn't to try to convince the listener of the truth of the answer. It's to ask "Do you want to know how I got to that answer?" If the reply is "no" the conversation is finished: all the person wants is for his conclusion to win. If we want to teach PCT, what we have to teach is PCT. It's tempting, I can tell you, to play guru and just come up with a series of mysteriously penetrating answers. But what's rewarding is to see a person (like a lot of people I could name in the CSG), START to ask a question, and then say "Oh, never mind, I see."

I am definitely not saying that PCT should go its own way and ignore everyone else's ideas. That's not the point. The point is that we shouldn't keep on pitting the conclusions of PCT against other people's conclusions. Well, I shouldn't. All that does is prolong the agony. If PCT has something to say about another person's field of expertise, the best person to work that out is the other person.

I've found myself pretending to be a linguist, a mathematician, a neurologist, a biochemist, a sociologist, a psychotherapist, and God knows what else, when all I really am is a control theorist saying how I would think about problems in those areas. If I were talking to a biochemist-control-theorist, or a sociologist-control-theorist, and so on, I would be talking about PCT, not about the other person's discipline. And the other person would be saying all the same things I would say and a lot more, because that person would know his own field AND PCT.

All of our problems come from trying to talk with people who don't understand PCT. The solution is to get them to understand PCT. The rest will take care of itself.

Best to all, Bill P.
Date: Wed Jun 23, 1993 8:33 am PST
Subject: Martin-Tom: #1

From Tom Bourbon (930623.0936)

Yesterday, I lost access to CSG-L for several hours. During that time, Martin Taylor and I pursued a direct discussion about information, plans and control. The discussion began on CSG-L and we thought we were still on the net, but were not. We agree that parts of our

private discussion might be of interest to others following the threads on those topics; consequently, we will post the discussion. It will come in two installments -- Martin-Tom #1 and #2.

The discussion began on the net when I read a remark by Martin to Hans Bloom:

=====
Subject: Re: Power gain, power loss.

[Martin Taylor 930621 11:10] Hans Blom 930619

On models, I tend to side with Hans. It is part of the whole information argument. The more information is available within the control system, the less is to be acquired from the CEV through the perceptual apparatus, and the better control can be.

=====

From Tom Bourbon (930621.1323)

>[Martin Taylor 930621 11:10] (Hans Blom 930619)

>Misunderstandings sometimes are better resolved by a non-combatant.

Most of Martin's post was about power gain in a control system. But at one point Martin said:

>On models, I tend to side with Hans. It is part of the whole information >argument. The more information is available within the control system, >the less is to be acquired from the CEV through the perceptual apparatus, >and the better control can be.

Martin, to paraphrase a line from the movie, "Field of Dreams," all I can say is, "If you build it, we will come." Take one example of control, as it is recreated or predicted by PCT models, and show me, in the results of simulations, how making more "information" available "within the control system" improves the recreations and predictions from the model. Then show me that those results generalize, with no further tinkering with the model, to new conditions, with unpredictably different disturbances and targets. That is not much to ask. Just improve on the performance of a single-level, single-loop PCT model.

And please delineate how your ideas in the remark to Hans differ from, say, a plan driven system that relies on information in the form of programs for action, thereby freeing itself from a need to rely on information about the CEV obtained through the perceptual apparatus. As you stated it, I see no difference.

This is not a put down. It is the only way to do business, if you rely on models to test your assumptions. It is my often repeated plea that you present the evidence, in the form of improved performance of the PCT model. Nothing else will impress us or win us over. You already know that. But I assure you that, if you build it, we will come.

=====

Inadvertently, the next exchanges were private, starting from Martin:

=====

Subject: Re: Power gain, power loss.

[Martin Taylor 930621 18:00] Tom Bourbon 930621.1323

>And please delineate how your ideas in the remark to Hans differ from, say, >a plan driven system that relies on information in the form of >programs for action, thereby freeing itself from a need to rely on >information about the CEV obtained through the perceptual apparatus. As you >stated it, I see no difference.

The difference is in those words "relies on" and "freeing itself from." I have no concept of either. Change them to "uses" and "reduces its need for" respectively, and I have less of a problem.

>Then show me that those results generalize, with no further tinkering with
>the model, to new conditions, with unpredictably different disturbances and
>targets.

What I assume would be in the model doesn't have much to do with the disturbances and targets, but with what Bill has labelled f(e)--the effect of a particular output change on the CEV. Reorganization is one way of building a f(e) that conforms to a predetermined model, which has the characteristic of being monotonic, as steep as can be constructed, and leads to negative feedback. That "model" needs no explicit form. It works with little information about the environment (which, Bill, incorporates all the lower-level ECSs, not just the part of the world outside the skin envelope) other than that the sign of the feedback is constant and the environmental gain stays adequately high. Bats, on the other hand, seem to adjust their perceptual input filters according to the expected time and frequency of the (doppler-shifted) echo. They need the model to distinguish the very low-power but precisely determined echo from whatever else is going on in their acoustic world.

How could such a model work? In the neural-net world, one rather powerful form of node is called a sigma-pi node. It does summation and multiplication, and can be used as a variable filter. It would be quite reasonable, I think, for a perceptual input function to contain the pi part of the sigma-pi, in addition to the sigma that is generally acknowledged to be there. The input to the pi could come from the output signal, changing the relative sensitivity of different elements of the PIF, and thereby changing its prior uncertainty about the expected signal. That's just one way it could work.

I'm not committed (yet) to internal models in general. I can see their potential usefulness, but they add a complexity to the ECS with which I am not happy. In the syntax predictor that Allan is developing for me, we do not include (yet) any internal model. We hope we will not need to include one to achieve good prediction. We are starting by relying on perceptual input functions that include differentiation. Nevertheless, when we get to noisy, smoothly changing representations of the syntax, I am at least open to the idea that we will have to incorporate models.

As I said, it's a question of the required information rate from perceptual signals. If you are among those who consider it an uninteresting quantity, you will not be interested in the possible value of an internal model as a component of an ECS.
=====

Here's a counter-challenge to the skilled modellers. I think it is fair, because we have not yet developed our own model, so we can see whether anyone, ourselves included, can solve the problem.

Define a formal grammar (say a BNF grammar) with 3 levels between the root and the leaves. Assert for each leaf symbol a description consisting of a location in an arbitrary 3-space (by analogy, think of phonetic feature values for phonemes). Let a control system "see" the succession of locations defined by the successive symbols output by executing the grammar with predefined probabilities of taking the different branches. The output of the control system is a location in 3-space. The three "intrinsic variables" that the control system must maintain are the difference between the locations of the output symbols and its own three dimensional output. The control system may be designed or it may learn (ours will learn).

Obviously, if the grammar output moves very slowly, any 3-D control system will work. Our problem is to get the control system to move to the right place as early as possible, preferably in synchrony with the motion of the grammar output point, which is moving quickly.

So far, we have not defined a challenge grammar or specified its rate of output, but we assume that the output point will have to stay stable for at least two compute cycles for the control system to have any chance of learning. We think that our control system will learn to have about as many levels as there are in the grammar, but that remains to be seen (it will grow by inserting ECSs between the "intrinsic variable" control ECSs and the top perceptual layer, as discussed last week).

=====
=====

Subject: Re: Power gain, power loss.

From Tom Bourbon (930622.1231)

>[Martin Taylor 930621 18:00] Tom Bourbon 930621.1323

>

>>And please delineate how your ideas in the remark to Hans differ from, say,
>>a plan driven system that relies on information in the form of
>>programs for action, thereby freeing itself from a need to rely on
>>information about the CEV obtained through the perceptual apparatus. As you
>>stated it, I see no difference.

>

>The difference is in those words "relies on" and "freeing itself from." I
>have no concept of either. Change them to "uses" and "reduces its
>need for" respectively, and I have less of a problem.

Fine. Change the words. Now, please because I still do not understand, tell me how the model implied in your remarks to Hans differ from, say, a plan driven system that "uses" information in the form of programs for action, thereby "reducing its need for" information about the CEV obtained through the perceptual apparatus. As you stated it, I see no difference.

>>Then show me that those results generalize, with no further tinkering with
>>the model, to new conditions, with unpredictably different disturbances and
>>targets.

>What I assume would be in the model doesn't have much to do with the
>disturbances and targets, but with what Bill has labelled $f(e)$ --the
>effect of a particular output change on the CEV. Reorganization is
>one way of building a $f(e)$ that conforms to a predetermined model, which
>has the characteristic of being monotonic, as steep as can be constructed,
>and leads to negative feedback. That "model" needs no explicit form.
>It works with little information about the environment (which, Bill,
>incorporates all the lower-level ECSs, not just the part of the world
>outside the skin envelope) other than that the sign of the feedback is
>constant and the environmental gain stays adequately high. Bats, on
>the other hand, seem to adjust their perceptual input filters according
>to the expected time and frequency of the (doppler-shifted) echo. They
>need the model to distinguish the very low-power but precisely determined
>echo from whatever else is going on in their acoustic world.

>

>How could such a model work? In the neural-net world, one rather
>powerful form of node is called a sigma-pi node. It does summation
>and multiplication, and can be used as a variable filter. It would
>be quite reasonable, I think, for a perceptual input function to contain
>the pi part of the sigma-pi, in addition to the sigma that is generally
>acknowledged to be there. The input to the pi could come from the
>output signal, changing the relative sensitivity of different elements
>of the PIF, and thereby changing its prior uncertainty about the expected
>signal. That's just one way it could work.

>

>I'm not committed (yet) to internal models in general. I can see their
>potential usefulness, but they add a complexity to the ECS with which I am
>not happy. In the syntax predictor that Allan is developing for me, we do
>not include (yet) any internal model. We hope we will not need to include
>one to achieve good prediction. We are starting by relying on perceptual
>input functions that include differentiation. Nevertheless, when we
>get to noisy, smoothly changing representations of the syntax, I am
>at least open to the idea that we will have to incorporate models.

>

>As I said, it's a question of the required information rate from
>perceptual signals. If you are among those who consider it an
>uninteresting quantity, you will not be interested in the possible
>value of an internal model as a component of an ECS.

Please, all I asked was:

>>Then show me that those results generalize, with no further tinkering with
>>the model, to new conditions, with unpredictably different disturbances and
>>targets.

Of course, in the original I asked to see a generalization of the results of simulations by the model you suggested. That is all I need to see, for you to convince me that what you say about information theory *does* translate into improvements in the performance of the PCT model. In the demonstration, you are free (encouraged) to assume the model in its fully developed and informed state. You need not simulate evolution, conception, birth, maturation, learning, social control processes, or enlightenment. Simply take an extant PCT model, add to it the features or measures you believe must be there for it to be an information theoretic PCT (IT PCT) model, and let it run. I described my criteria for improvement in other posts long ago, and in one addressed to Hans Bloom a few minutes ago. A demonstration like that would clear the air of gigabytes of "I said," "You said," "We said," and the like. And it would focus the discussion on the real issue -- does the PCT model work and, if so, can it be improved?

>Here's a counter-challenge to the skilled modellers. I think it is fair,
>because we have not yet developed our own model, so we can see whether
>anyone, ourselves included, can solve the problem.

This is another kind of "challenge" entirely. In fact, my offer is not a challenge. I am merely saying that we know the PCT model works for certain instances of control by humans. We know the model can be and should be improved. We are eager to enlist the support of anyone who wishes to join in that endeavour. The criteria for demonstrating improvement in the model are simple and direct. Have at it. We have even published and posted the PCT model (all two lines of it, if you include the environment) many times, so you can avoid the need to develop your own model. Please, use ours as a testbed for your ideas. (I am completely serious -- no attempt by me to be cute, clever or condescending.)

This is not a contest in which we try to prove prowess and skill -- not for me it isn't -- I have neither of those "attributes." My skills are limited. I would like to see people with skills and resources superior to my own devote some of their time and creativity to working on our project.

>Define a formal grammar (say a BNF grammar) with 3 levels between the
>root and the leaves. Assert for each leaf symbol a description consisting
>of a location in an arbitrary 3-space (by analogy, think of phonetic
>feature values for phonemes). Let a control system "see" the succession
>of locations defined by the successive symbols output by executing the
>grammar with predefined probabilities of taking the different branches.
>The output of the control system is a location in 3-space. The three
>"intrinsic variables" that the control system must maintain are the
>difference between the locations of the output symbols and its own
>three dimensional output. The control system may be designed or it
>may learn (ours will learn).
>
>Obviously, if the grammar output moves very slowly, any 3-D control
>system will work. Our problem is to get the control system to move
>to the right place as early as possible, preferably in synchrony with
>the motion of the grammar output point, which is moving quickly.
>
>So far, we have not defined a challenge grammar or specified its rate
>of output, but we assume that the output point will have to stay stable
>for at least two compute cycles for the control system to have any chance
>of learning. We think that our control system will learn to have about
>as many levels as there are in the grammar, but that remains to be seen
>(it will grow by inserting ECSs between the "intrinsic variable" control
>ECSs and the top perceptual layer, as discussed last week).

You got me there, Martin. Congratulations. I sure can't do that, but then I never claimed to be a skilled modeler. Now, can I interest you in helping us figure out how to improve the PCT model for something as mundane and trivial as stick wiggling?

```
=====
=====
From Martin Taylor, 22 June 1993
Subject: Re: Power gain, power loss.
```

Tom,

You responded to me personally, so I do the same. Did I mail to you, rather than posting to CSG-L? I don't remember.

As I understand it, the problem with outflow plans is that they cannot work because the world both is disturbed by influences unknown to the pseudo-control system and because the impact of the output on the thing to be "controlled" is not always the same. If the impact were defined with high probability, and if the thing to be controlled were effectively isolated from disturbance, outflow planning would work.

Do we agree so far?

[Tom -- present time -- 23 June]

Yes.

[Martin]

Now:

[Martin]

>>The difference is in those words "relies on" and "freeing itself from." I
>>have no concept of either. Change them to "uses" and "reduces its
>>need for" respectively, and I have less of a problem.

[Tom]

>Fine. Change the words. Now, please because I still do not understand, tell
>me how the model implied in your remarks to Hans differ from, say, a plan
>driven system that "uses" information in the form of programs for action,
>thereby "reducing its need for" information about the CEV obtained through
>the perceptual apparatus. As you stated it, I see no difference.

[Martin]

If you use "relies on" and "freeing itself from", you are talking about a plan-driven outflow system. If you say "uses" and "reduces its need for" you are talking about a normal control system that probabilistically anticipates outputs that might affect the world usefully, but remains based on the present difference between current perception and current reference.

[Tom - now]

I am not sure how a normal control system "probabilistically anticipates outputs that might affect the world usefully," unless you mean something like evaluating programs in imagination mode, then using one that seemed to work sufficiently well in imagination. If that is the case, the program is for perceptions, not actions, otherwise there is no difference between the wordings you compared: there is not much difference, if any, between "plan-driven outflow" and "outputs that might affect the world usefully."

[Tom, previously]

>Of course, in the original I asked to see a generalization of the
>results of simulations by the model you suggested. That is all I need to
>see, for you to convince me that what you say about information theory
>*does* translate into improvements in the performance of the PCT model.

[Martin]

You use the example of a sawtooth tracking task quite often in your demonstrations. What do people do if after many cycles of the sawtooth you stop the target at the mid-point and leave it there. Doesn't the tracker overshoot before coming back to the target? But does the tracker overshoot when the target reverses direction at the peaks of the sawtooth? I think not, at least not after the first few cycles.

I haven't tried this, but here is a place where I make a prediction that you have plenty of data to test. I am guessing that a simple ECS model of one level tuned to a best fit to the human data will fail in two specific places: (1) There will be a reversal in the sense of the predictive miss at the peak between (a) the first one or two peaks, and (b) peaks late in the sequence; (2) if you stop the target motion at a peak after many tracking cycles, the human will reverse and come back, whereas the model will not. The sawtooth should be fast enough that the human tracks well, but measurably imperfectly.

Is this right?

[Tom -- now]

In examples I post on the net, I often use sawtooth targets, but only because I can draw them in ASCII. I haven't figured out how to draw random target paths. But since you mentioned them, events will occur as you described them in the hypothetical demonstration -- if you stop the target, the person will overshoot, then come back. (And will also reveal the presence of other levels, by looking at you, or at the computer, and saying things like, "it stopped working," or "what happened?")

However, the "overshoot" probably is not a result of "predictive movements." The situation looks very much like the one described by Rick and Bill in their chapter, "Levels of Intention in Behavior," in Wayne Hershberger's book, *Volitional Action: Conation and Control*, (1989, North-Holland, for those who do not know the book). What they studied, and what you describe, is the result of differences in time constants for different levels in a hierarchical control system. The "overshoot" you describe in the demonstration, like that they describe in their chapter, occurs when for a brief time a lower level system continues operating on the "old" reference signal from above. It takes time for the higher level to detect error which produces the new reference signal for the lower level. During the time between onset of changes in the rules in the environment, and alteration of the error signal qua reference signal from the higher level, the lower level keeps on with the originally "correct," but now "incorrect," reference signal -- it keeps working perfectly.

[Martin]

I'm actually rather surprised to think that the "world-model" idea might work at this low level. I had thought of it as being more useful above the category level, and particularly at the program level and above. But the theoretical background makes no such distinction of levels, and if it actually does work as I presume, the result should abort a lot of future fruitless discussion. The question will (I hope) turn to how the world model information actually is implemented in an ECS. Is it in the output function, the perceptual input function, or (as I presume) in the imagination loop?

[Tom -- now]

A world model (in imagination mode?) must affect the performance of the system at *every* level, once the system begins to act on the model.

[Martin]

As for the word "challenge," I agree with your comments. It is an ill-chosen word. However, much of what has been going on has the flavour of challenge, and the word came quite naturally. I would be much happier if it didn't. And I'm not the modeller. That's not my skill, either. I'm more of a theoretician (as you doubtless have observed). Allan is doing the work.

=====
=====

[Tom -- now]

Fine! I like the idea of cooperation or collaboration, perhaps with a bit of heat to liven things up, better than challenges. (Care for a round of stick wiggling?)

That has us almost caught up. I will send a few of today's posts a little later.

Until later, Tom Bourbon

Date: Wed Jun 23, 1993 9:04 am PST
Subject: Re: PCT and other theories

From Tom Bourbon (930623.1246) Bill Powers (930623.0700 MDT)

I won't reproduce your eloquent post, Bill. Those are the ideas that, in much less articulate form, are behind my pleas that others put aside some of their more lofty goals

and try using PCT to model a few deceptively simple-looking phenomena. But that just doesn't seem to be in the cards.

"Well put," and thanks for saying it, Bill.

Tom Bourbon

Date: Wed Jun 23, 1993 9:56 am PST
Subject: PCT and linguistics

[From Bill Powers (930623.0930 MDT)] Bruce Nevin (930622.1316)

RE: phonemes, words, etc.

In the PCT model, a perception is a signal; it is a function either of signals generated directly by sensory endings, or of signals that are outputs of lower-level perceptual functions. Some of the signals in lower-level perceptual functions are under feedback control; some are not.

These functions, at the lower levels, are primarily analogue in nature. That is, they handle continuous variables, making one output variable a function of some set of input variables. The following discussion, however, applies to any kind of function.

A variable y is a function of a set of x 's: $y = f(x_1, x_2, \dots, x_n)$. Y represents the magnitude of the perceptual signal output by a perceptual function, the x 's represent the magnitudes of perceptual signals collected from lower perceptual functions and serving as input to, or arguments of, the function.

For any such function, there are many sets of values of the x 's that will produce the same value of y . All sets of x 's that produce the same y are therefore perceived as the same state of the same variable. If the x 's vary in combinations that leave y unchanged, the changes in the x 's are not perceived at the level of y . Any other combinations of changes in the x 's result in a changed value of y , and the perception y is then affected by the changes in the x 's. The entity perceived as y is then perceived in a new state. If the function is continuous, the states may be indicated on the real number scale. If the function is discrete, the states may be indicated on an integer scale, including the binary scale.

Two perceptions, y_1 and y_2 , may be derived through two different functions f_1 and f_2 , from the same set of x 's, $x_1 \dots x_n$. If the functions are different, then combinations of changes in the x 's that leave y_1 constant may alter y_2 , and vice versa. Other changes may alter both y_1 and y_2 . Still other patterns of change in the x 's may leave both y_1 and y_2 the same. This general relationship between x 's and y 's holds true for any number of y 's, although the maximum number of y 's cannot exceed the number of x 's if independent changes in the y 's are to result.

If we could plot the dependence of y on its corresponding set of input x 's in as many dimensions as desired, we could draw "surfaces of indifference" in this space passing through all the values of the x 's that left the value of y unchanged. Any or all x 's could change without altering y if the values lay in the same surface of indifference. Changes in y result from the x 's changing in a way that moves the result to a new surface of indifference: a change normal to the surface of indifference in hyperspace.

In two dimensions, if the functions are $y_1 = x_1 + x_2$ and $y_2 = x_1 - x_2$, we would have lines of indifference. If x_1 decreased by the same amount that x_2 increased, the value of y_1 would remain unchanged but y_2 would change. If x_1 and x_2 increased or decreased by the same amount, y_2 would remain constant and y_1 would change. Any function of x_1 and x_2 would create surfaces of indifference: for $y_3 = x_1^2 + x_2^2$, the surface of indifference would be a circle. The perceptual signal y would thus represent only the radius of the circle, and would be indifferent to changes in x_1 and x_2 that represented different points on the same circle.

If x changed uniformly from 1 to -1 while y changed uniformly from -1 to 1, we would find that y_1 remained constant throughout, y_2 changed uniformly from -2 to 2, and y_3 declined from 1 to 0 and then increased to 1 again. Three different aspects of the behavior of x_1 and x_2 would be perceived simultaneously, and all the behaviors through time would be different. Only two of these aspects could be controlled at the same time; the other would simply follow as a consequence.

Each x is derived in a similar way from a set $w_1...w_n$ at a lower level (if one exists). This is the basic hierarchical relation between levels of perception in the hierarchical PCT model.

When some physical change occurs that affects sensory receptors, it has a systematic effect on many perceptual signals. Different kinds of physical changes will have different patterns of systematic effects on individual perceptual signals. Some perceptual signals will be affected more than others; some will be increased while others will be decreased.

These patterns of change can be resolved into two components for any one higher-level perceptual function. One component alters the set of lower-level perceptions parallel to a surface of indifference. That component has no effect on the higher perception. The other component is normal to the surface of indifference; it results in a change in the output perceptual signal. The magnitude of the output perceptual signal thus represents a particular aspect of the external event that altered the lower-level perceptions. What that aspect is depends entirely on the organization of the perceptual function involved.

In general, any external event that alters a set of perceptual signals at one level will have differential effects of either sign on perceptual signals at the next level. Some signals at the next level will not be changed at all. Some will change slightly. Others will change greatly. Each perceptual function defines its own set of surfaces of indifference in the space of the lower-level perceptions (each lower-level perception representing one dimension of the space). A path orthogonal to the surface of indifference created by any function represents the range of the variable being perceived, and the position along that path represents the magnitude of that variable. The common external event is thus first represented as a set of states in the lower-level collection of x 's, and is then rerepresented at the higher level by a new set of signals indicating many aspects of that event, each aspect being perceived in some state between zero and maximum.

In passing upward through the hierarchy of perceptions, this representation and rerepresentation occurs over and over, new types of functions being introduced and new aspects of the lower-level world being created and signified by perceptual signals in the next higher level. Some functions may compute weighted sums. Some may compute spatial or time derivatives. Some may compute logical relationships. In HPCT, an attempt has been made to define the functions that are typical of particular levels, where the functions are only suggested by the classes of perceptions that they relate. The primary task of perceptual research in HPCT is to discover, through finding what appear to be hierarchically-related types of perceptions, the kind of functions that must be performed to create the observed relationships between adjacent levels. Perceptions can be defined through control tasks.

The research question is thus not so much whether configurations are derived from sensations, but what kind of perceptual function would be required to create that apparent relationship. Of course it is also important to establish the actual existence of specific pairs of hierarchically-related perceptions. The two questions must be answered together.

This is what PCT has to say about perceptions. Applying this model to the perceptions of phonemes and words is a matter of proposing appropriate perceptual variables and showing that one type depends on other types. Identifying surfaces of indifference amounts to finding changes in lower perceptions that make no difference to higher ones, and defining the particular function of the lower perceptions that is orthogonal to those surfaces of indifference. Identifying perceptions is a matter of setting up appropriate control tasks and applying the test for the controlled variable. Go to it.

Best, Bill P.

Date: Wed Jun 23, 1993 11:03 am PST
Subject: There are people, Bill....

from Ed Ford (930623:1100) (Bill Powers 930623.0700)

>...all of our problems come from trying to talk with people who don't
>understand PCT. The solution is to get them to understand PCT. The
>rest will take care of itself.

A few minutes before I read your PCT And Other Theories post, I had been talking with Eric Newton, Assistant Superintendent, Sweetwater School District No. 2, Green River, Wyoming, where I have done extensive training. With his permission, I have used his name as a

reference when soliciting other school districts about my program and I was calling to ask if he had heard from anyone. He said he hadn't but not to worry, if anyone called he would certainly be glad to tell them about my program and its benefits, especially PCT.

He then just came out with the following remarks which should warm the heart of our friend in Durango.

"Once you understand control theory, you can apply it in a school district. There is no sense trying to make any meaningful changes in a district unless you really understand control theory and then with that understanding, begin to make changes that help children. Understanding that theory really does make a difference in how you deal with children as well as support staff, teachers, and parents. You need to internalize this as a way to deal with people. You have to get out of the theoretical stage and say this is the way things are happening within us. And then you can use. As long as you say 'if this is the way things are happening, then...' you'll never get any where. You have to say this is how things work inside the brain and go from there. You have to apply it. Get away from the theoretical to the practical.

The world seems to be stimulus response world judging from our experiences but then after you learn control theory and begin to use it, then you realize that the world is much different, it's all choices. You have to look at your world through control theory eyes and then you begin to see a whole new world and then you can make effective changes that benefit children within your school district."

For those who like to criticize public schools, they should visit Green River and get to know people like Eric Newton and the staff and teachers of Sweetwater School District #2 in Wyoming. It gets real cold there in the winter time, but they've got a lot of very warm and intelligent hearts.

Best, Ed

Date: Wed Jun 23, 1993 12:13 pm PST
Subject: Re: Collective Controlled Variables?

From Tom Bourbon (930623.1306) Kent McClelland (930622)

>Tom Bourbon (930622.0817) and Gary Cziko (930622.0323 UTC) have been
>involved in an exchange with Rick Marken about the possibility of "collective
>controlled variables."

>...

>I'm willing to concede Tom and Rick's point that each human control system
>has its own reference levels, and that no super-ordinate social-control
>system exists to impose reference signals on independent individuals. Two
>individuals cannot literally share the same reference level. However, Gary's
>examples of cooperatively accomplished tasks do suggest some ideas that get
>lost when we insist too strongly on the independence of each social actor.
>Bill Power's recent post about close-order drill in the Navy (930621.0840
>MDT) illustrates the issue even more sharply.

>

>Even if two people can't share the same reference levels, social life is only
>possible because people can find reference levels that are similar enough to
>allow effective cooperation.

Kent, I don't think effective social life, with cooperation as an example, requires that two or more people find, accept, adopt, or create similar reference signals. What matters is that the participants adopt reference signals (that they adopt goals which we model as reference signals -- Hans?) that result in each acting in a way that produces a match between personal goals and present perceptions. The same constraints apply to my two hemispheres when "they cooperate" to produce one-person performance of what can also be a two-person cooperative task: they need not adopt similar reference signals; all that is necessary is that each adopt reference perceptions that result in the intended perceptions. Of course, cooperation between hemispheres, whether in one head, or two, means actions that affect variables controlled by another entity. (When applied to events inside one head, the language becomes a bit weird.) Cooperation can, probably always does, entail some necessarily different reference perceptions in the participants.

>...For "practical purposes" two people's reference
>levels can be the same--good enough for government work, in the case of

>Bill's Naval comrades. I'm talking about control at the higher levels of the
>hierarchy, the programs in one's head for shared cooperative action, or the
>system concepts that imply coordination of efforts, or the rules for
>producing and interpreting language (to take a page from Bruce Nevin's book).

Here, to be sure, for people to knowingly cooperate, each must agree to cooperate. But in itself that kind of agreement says nothing about the specific contents of the participants' heads. The socially affected and approved options all reside in individual heads, as products of individual interpretations or understandings. Necessary and important, to be sure; the same in each person, probably not.

>The term I like to use to designate this practical similarity of reference
>levels is "alignment"--a concept that I think is key for doing any kind of
>social analysis with PCT. This summer I'm once more trying to rework my
>paper on power to argue that in talking about social power we're almost
>always talking about the alignment of reference levels of large numbers of
>people. To allow cooperation, people's reference levels do not have to be
>identical, but they do have to be aligned. ...

I understand your meaning and I appreciate the practicality, for you, of employing this terminology in some settings. For me, the problem comes when people who speak of aligning references or perceptions begin to believe that is what really happens. What we are talking about is seeing people act so that they all perceive what they intend and that (in the applications you describe) they should all intend certain kinds of experiences of the world. (And a lot of the "intending" is in the head of someone who is trying to modify, or create, a particular social interaction or system so that it suits their own ends.)

>... In most cases people who
>cooperate need to be perceiving themselves as working together on the same
>project, and their programs for carrying out the project have to be similar
>enough that their joint efforts aren't simply disturbances to each other.

By definition, for them to cooperate they *must* do that (perceive themselves working together); but their programs could be wildly different -- the various roles assigned to and adopted by members of the crew of a large ship, plane, or other installation come to mind. If everyone had similar programs, they would fail. I know that is not the sense in which you mean similar programs, but that is a meaning that falls out very easily, and some people who "use" control theory as a "framework for viewing" social phenomena appear to think of it that way.

In the cooperative tracking tasks, each person or chunk of brain is best modeled as a system that independently controls part of the "big picture." At the level of effective control, each system is independent. That means the operative reference perceptions in each must be different, at that level of control. This fact raises some interesting and vexing questions about which persons, or which brain parts, must know what during effective interactions. Your next questions also acknowledge that fact.

>Several interesting questions arise when we think about social interaction in
>terms of alignment:

>
>1. How do people succeed in getting their reference levels into alignment?
>Sometimes it must be imitation of what the other person seems to be doing;
>sometimes verbal exchanges are obviously used; a lot of the time it must be
>trial and error, trying something and seeing if the other person resists, a
>practical application of "the Test".

I will say more about imitation in a later post, in which I will reply to both you and Bruce Nevin.

>2. How good is good enough when it comes to alignment? If the alignment
>isn't perfect (which according to PCT it can't ever be), when do cooperative
>efforts turn into a tug-of-war? (By the way, Tom, have you ever set your
>experiments up so that the human control system and the computer model are
>simultaneously trying to do the same thing, both controlling the same cursor?
>What happens? How does it feel? How similar do reference levels have to be
>for the two to be doing the "same" thing?)

I have. So has Rick. And so has a thesis student of mine -- the very last one, Michelle Duggins, who will be at CSG to talk about her thesis work on "helping." More on this then.

>3. What are the consequences when lots of control systems, not just two, are
>effectively aligned in some coordinated effort? Do they become
>indistinguishable from one "super-powerful" control system? How is the
>collectivity different from a single high-gain system?

The answers to all of your questions depend on from whose perspective you define alignment and effectiveness. If from the vantage point of a social engineer who wants to create certain patterns of social interaction (a position analogous to that of the control engineer who designs an artificial control system), the answer can be quite different than from the position of the persons, or PCT models, who do the acting.

In none of the cases I have modeled did a quartet take on super powers, different from those seen with a dyad, or with a two-fisted human controller --or with two or four PCT models. In all cases, the different combinations simply got the job done. Under various conditions, the job could only be done by one, two, or four systems acting in concert, but that was because of the way I parceled out access to the physical devices that could influence the different independently controllable variables in the experiment.

Until later, Tom Bourbon

Date: Wed Jun 23, 1993 1:14 pm PST
Subject: Re: Information; controlling effects

[Martin Taylor 930623 15:00] Bill Powers 930622.0730

Your time-domain representation of the way the control system works is exactly the way I see it and have seen it since before the beginning of the information theory argument. No problems. If you remember, I used a colloquial phrase to which you objected, about the information about the disturbance in the perceptual signal being "bled off" over time into the output signal, so that after a while there was almost none left in the perceptual signal, if the perception was effectively being controlled.

>Instead of using a constant disturbance, you used a step-function to show
>that the disturbance must have an effect on the perceptual
>signal. A step function contains an instantaneous transition,
>which is guaranteed to be far outside the bandwidth of any
>physical control system.

When the disturbance moves very slowly compared with the control bandwidth, the phase lag is very small, so there is at any time very little information that has not been "bled off." But if you integrate that "very little" error over time, it turns out to be exactly what is required to create the output signal (assuming a simple integrator output function). The changing disturbance, however slowly it changes, does create an error that is corrected over time (exponentially as you say given the simple output function). In my "silver tongued" posting, to which you refer, I argued that the same function applied as the magnitude of the step approached zero, in the classic way that differentials are taught as the infinitesimal end-point of differences. There are no peculiarities in this system about the approach to zero, and the sum of step function effects approaches the integral of continuous effects. I don't think we are talking about different extreme conditions in any way that alters the argument.

=====

>>How is it possible for you to take my discussion of the
>>time-dependent increase in resolution in observations of low-
>>bandwidth signals, and come up with:

>

>>>If the noise is 10% RMS of the range,
>>>Martin's interpretation would be that there are only 10
>>>possible values of the perceptual signal with a range of 10
>>>units, 0 through 9.

>

>Easy. You said that the equivalent of r (in D/r) for a continuous
>system with noise would be the RMS noise level. If the RMS noise
>level is 10% of D , then D/r is 10, and the probability of any one
>measure would be r/D or $1/10$. In the discrete case, D/r
>represents the number of possible values of d . Thus there are D/r
>= 10 possible values of d (or some signal) in the analog case,
>too, if you're computing probability the same way.

I still don't understand this. 3.01572 is different from 3.01583, even if the RMS error of measurement is 300. The readings differ, even if they don't usefully differ when referred to the thing being measured. There is no finite number of possible readings when we are dealing with continuous systems. There is a finite amount of information that you can get from a reading about the thing being measured. That's quite a different kettle of fish.

>I understand resolution to be the least possible
>difference between two readings. How do you understand it?

You are an old radio astronomer. I take the resolution of a telescope to be roughly the distance at which two point-spread functions meet at their 3 dB points. I've forgotten the name, but isn't the point spread function called the Airy disc or something like that? The centre points of the star images can be anywhere, but a another star closer than the diameter of the Airy disc cannot be resolved. That's what I mean by resolution. (Sorry if I have the name totally screwed up; the concept is what counts).

The resolution is referred to the star field, not to the film grain.

I think your diagram looks OK for the added noise. The control system can see the value of the CEV perturbed by the added noise, but can affect only the value of the CEV. What it will presumably do is to affect the CEV in such a way that the perceptual signal matches the reference, and therefore in such a way that the CEV does not have the desired value. The real-world effectiveness of control depends on the reliability with which the perceptual signal reflects the value of the CEV (you might drive off the cliff if you couldn't tell where you were on the road to better than 20 metres). So if you want to control the CEV's value (as seen by a neutral and infinitely precise observer) to within $\pm X$, you had better be sure that your perceptual signal is reliably different when the actual value of the CEV changes by $2X$. You can do that, whatever the RMS noise level, by averaging over a long enough time.

>Have you been following my discussion with Randall? If you
>translate a variable with a range of 8 units and unity resolution
>into a signal with a range of 4 units _and_ unity resolution_, you
>will lose information, won't you?

Yes, if it is quantized resolution you are talking about, or if you are talking about the same observation conditions when determining the value of the resolution.

I appreciate your problems with too much irrelevant mail. You put information theory among the "Yes, but..." class of stuff that wastes your time. I think of it as neither orthogonal nor competitive, but as a foundational aspect of PCT that has been ignored to date.

The discussion above does not use the word "information" with or without a capital "I" (unless I slipped up), but it is mostly about the importance of information. So is the example that Tom Bourbon kindly reposted to the net, about the characteristic differences between a one-level simple ECS and human behaviour in the sawtooth tracking task. Whether these characteristic differences depend on a world model that forms part of an always active imagination loop, or depend on shifting reference signals (of a form I do not understand) from above (as Tom suggests), they represent information used by the human, derived from the redundancies found in the history of the task, that helps the human track.

A one-level ECS that tracks as well (or badly) as the human on average, with its parameters fitted to correlate 0.997 with the human, does not overshoot when the sawtooth motion of the target stops at the centre point, and does not reverse direction when the target stops at a peak of the sawtooth. The human does, however, because of expectations about the future movement of the target, wherever those expectations are stored. The model ECS uses only information from the sensory input. The human uses that in conjunction with information from the history of the movement, and takes longer than the model to respond to changes in the parameters of the movement. The model would not care if the movement changed from sawtooth to random. But would the same parameter settings of the model provide the optimum fit both to the human tracking the sawtooth and to the same human making the transition to tracking a randomly moving target of the same maximum slew rate?

=====
Answering Hans Blom (related to the previous section)

>>Note that this directly solves the adaptation (learning)
 >>problem as well: correlating the effector's action with the
 >>world's reaction 'identifies' ('systems identification') the
 >>world.

>
 >But it's not as good a solution as detecting the world's
 >reaction, comparing that against the reaction you want, and
 >converting the difference into the direction of output that will
 >make the difference smaller. Decidedly smaller. When you can do
 >it that way, you don't need any correlations.

But you get the best of both worlds by using the correlations, because so long as the world remains consistent in its reactions to output, you can reduce the phase lag of control. The penalty is increased error when the effect of output on the CEV changes.

Martin

Date: Wed Jun 23, 1993 2:34 pm PST
 Subject: Re: Information; controlling effects

From Tom Bourbon (930623.1619)

>[Martin Taylor 930623 15:00] Bill Powers 930622.0730
 ..

>I appreciate your problems with too much irrelevant mail. You put
 >information theory among the "Yes, but..." class of stuff that wastes
 >your time. I think of it as neither orthogonal nor competitive, but
 >as a foundational aspect of PCT that has been ignored to date.

>
 >The discussion above does not use the word "information" with or without
 >a capital "I" (unless I slipped up), but it is mostly about the
 >importance of information. So is the example that Tom Bourbon kindly
 >reposted to the net, about the characteristic differences between a
 >one-level simple ECS and human behaviour in the sawtooth tracking
 >task. Whether these characteristic differences depend on a world
 >model that forms part of an always active imagination loop, or depend
 >on shifting reference signals (of a form I do not understand) from
 >above (as Tom suggests), they represent information used by the human,
 >derived from the redundancies found in the history of the task, that
 >helps the human track.

>
 >A one-level ECS that tracks as well (or badly) as the human on average,
 >with its parameters fitted to correlate 0.997 with the human, does not
 >overshoot when the sawtooth motion of the target stops at the centre
 >point, and does not reverse direction when the target stops at a peak
 >of the sawtooth. The human does, however, because of expectations about
 >the future movement of the target, wherever those expectations are stored.
 >The model ECS uses only information from the sensory input. The human
 >uses that in conjunction with information from the history of the movement,
 >and takes longer than the model to respond to changes in the parameters
 >of the movement. The model would not care if the movement changed from
 >sawtooth to random. But would the same parameter settings of the model
 >provide the optimum fit both to the human tracking the sawtooth and to the
 >same human making the transition to tracking a randomly moving target
 >of the same maximum slew rate?

Martin, apparently I did not (I will not yet say cannot) urge you strongly enough to examine the chapter by Rick and Bill. The point was that a two-level PCT model duplicated the "overshoots" that occurred when a person, who was performing a simple tracking task, encountered various alterations of the conditions in the task -- sudden reversal of the direction the mouse moved the cursor, a signal to change reference perceptions for where the cursor should be relative to the target, and a couple of others. The point was that the models, and presumably the persons, "kept on tracking" using the old reference signal. What they did for a few hundred milliseconds was "wrong," relative to the new condition, but exactly "right," relative to the old reference signal.

Your remarks in the following section of your post lead me to believe you really should read that chapter:

>...Whether these characteristic differences depend on a world
>model that forms part of an always active imagination loop, or depend
>on shifting reference signals (of a form I do not understand) from
>above (as Tom suggests), they represent information used by the human,
>derived from the redundancies found in the history of the task, that
>helps the human track.

The effect occurred, for the 2-level model, as a direct result of its hierarchical connections, with error from level 2 as the reference signal for level 1. There were no world models, or imagination loops, or redundancies found in the history of the task, or stored expectations about the future movement of the target. Not that those things never occur or are never important; but in the tracking task, a simple bare-wire two-level PCT model reproduced phenomena that you believe require those things.

Please, read the chapter.

Until later, Tom Bourbon

Date: Wed Jun 23, 1993 4:32 pm PST
Subject: PCT, resolution, etc

[From Bill Power (930623.1545 MDT)] Martin Taylor (930623.1500) --

>Your time-domain representation of the way the control system
>works is exactly the way I see it and have seen it since before
>the beginning of the information theory argument. No problems.
>If you remember, I used a colloquial phrase to which you
>objected, about the information about the disturbance in the
>perceptual signal being "bled off" over time into the output
>signal, so that after a while there was almost none left in the
>perceptual signal, if the perception was effectively being
>controlled.

If I had an objection, it was to using "information" where all you meant was "signal." If this image helps do the informational calculations, however, by all means use it. I still want to see the calculations, unless they're just the same ones we would use in control theory with different names.

>I still don't understand this. 3.01572 is different from
>3.01583, even if the RMS error of measurement is 300. The
>readings differ, even if they don't usefully differ when
>referred to the thing being measured..

If all you know about the thing being measured is contained in the reading, as is the case for living organisms, you have only two choices: treat the (300-unit) nonsystematic variations that you measure (to six significant digits) as real and accurately measure the fluctuations in the measured variable, or assume that some noise source has entered the situation, and average readings together to measure the average value of the (assumed) non-noisy external variable to six figures. If you guess wrong, of course, you will average out the perfectly meaningful variations in the external variable.

>I take the resolution of a telescope to be roughly the distance
>at which two point-spread functions meet at their 3 dB points.
>I've forgotten the name, but isn't the point spread function
>called the Airey disc or something like that? The centre
>points of the star images can be anywhere, but a another star
>closer than the diameter of the Airey disc cannot be resolved.

As two star images get closer together, more and more precise measurement of the intensity distributions is required to distinguish a double star from a single one. There is no theoretical minimum detectable separation other than that set by the grain of the emulsion or the spacing of CCD pixel elements, and in principle that can be overcome by projecting the image to as large a scale as necessary. The only absolute limit is set by available observing time and the assumption that the source of the image remains constant over the observation. The resolution that matters in astronomy is that of the detector, and always

refers to the least possible separation of adjacent measurements, not to noise level (temporal or spatial). The resolution of a noise-limited measurement can always be improved by reducing the noise (for example by cooling) or extending observation time. Instrument-limited resolution can't be improved regardless of observation time or noise reduction. When you can see the pixels or the grains, that's it unless you magnify the image more.

To some extent, pixel-limited resolution can be improved by combining multiple images. This can effectively increase the number of available pixels. But whether this results in improvement depends on being able to position one image relative to another within a known fraction of a pixel. That implies some other means of resolving the image to less than one pixel on the detector. Using the entire picture to get an average match can help, but only if the representation is precisely linear in space. With great labor, spatial distortions in the image can be mapped to a fraction of a pixel, but again that requires some other way of measuring positions of point-sources with fractional-pixel accuracy. And you must assume that the map remains valid from one exposure to another.

>The resolution is referred to the star field, not to the film grain.

No. The "resolution" of the star field is infinite. It is only the detector grain that ultimately limits the resolution of the picture. That resolution may be expressed in seconds of arc, but nobody is under the impression that this means the stars themselves are fuzzy. If anyone had thought that, there wouldn't have been much to be gained by putting high-resolution detectors on the Hubble space telescope.

Star images can be located relative to others within 0.002 sec of arc with an 18-inch telescope which has an Airey disk of about 0.25 sec of arc, under seeing conditions averaging about 1 sec of arc. I did about 100,000 of the film measurements (61 Cygni).

>I appreciate your problems with too much irrelevant mail. You
>put information theory among the "Yes, but..." class of stuff
>that wastes your time. I think of it as neither orthogonal nor
>competitive, but as a foundational aspect of PCT that has been
>ignored to date.

I appreciate your problems with my refusal to see how information theory is "foundational" to control theory. What is foundational to the notion of control theory, in my opinion, is the logic of closed loops of causation. I haven't yet seen how information theory would lead to this notion if it hadn't already been stated. It would be impossible to understand any control system, or even know that one was possible, without the idea of the closed causal loop. But it has been possible to understand, build, and predict the behavior of control systems without any use of information theory. Along with Tom Bourbon, I have yet to see anything coming out of information theory that will help us do our modeling any better. Perhaps some day you will demonstrate to us something that will accomplish this. Until then, I have no basis for accepting your "foundational" claim. Nor do I have any reason to doubt it. The case is not proven.

>A one-level ECS that tracks as well (or badly) as the human on
>average, with its parameters fitted to correlate 0.997 with the
>human, does not overshoot when the sawtooth motion of the
>target stops at the centre point, and does not reverse
>direction when the target stops at a peak of the sawtooth.

The first statement is incorrect. In a tracking model with an integration constant of about 6/sec and a perceptual delay of 0.16 sec, a sudden stop of the target results in an overshoot that is then corrected. The real person may overshoot a bit more.

Your second point, however, is correct. A multiple-level system is needed to take advantage of any regularity in the disturbance. As we have not yet modeled systems that can take advantage of regularities, we use unpredictable disturbances. Whether a human being is able to take advantage of the regularity of a sawtooth movement depends very much on the frequency of the sawtooth. If it is too slow, tracking isn't improved because human estimates of frequency tend to wander over longer periods of the wave. If it is too fast, the lower-level control systems are on the verge of losing control anyway, so the higher-level control simply prevents control from deteriorating as fast as frequency rises. The main change between unpredictable and predictable disturbances is that the delay becomes smaller and more variable for a predictable disturbance.

>The human does, however, because of expectations about

>the future movement of the target, wherever those expectations
>are stored.

We will not need to include "expectations" or "predictions" in this model for a long time. A system that matches a predictable waveform doesn't have to predict it. All it has to do is remember it. Actual prediction is something else, although the appearance of prediction can easily be created by a control system that has dynamic stabilization using derivatives in the perceptual function. I don't think we will run into true predictive control until we get to the cognitive systems.

>The model would not care if the movement changed from sawtooth
>to random. But would the same parameter settings of the model
>provide the optimum fit both to the human tracking the sawtooth
>and to the same human making the transition to tracking a
>randomly moving target of the same maximum slew rate?

So far there is little difference between human and model behavior in either case -- not enough to measure reliably. Do you agree, Tom?

Best, Bill P.

Date: Thu Jun 24, 1993 7:07 am PST
Subject: Re: PCT, resolution, etc

From Tom Bourbon (930624.0854) >Bill (930623.1545 MDT)] Martin (930623.1500)

[Martin]

>>A one-level ECS that tracks as well (or badly) as the human on
>>average, with its parameters fitted to correlate 0.997 with the
>>human, does not overshoot when the sawtooth motion of the
>>target stops at the centre point, and does not reverse
>>direction when the target stops at a peak of the sawtooth.

[Bill]

>The first statement is incorrect. In a tracking model with an
>integration constant of about 6/sec and a perceptual delay of
>0.16 sec, a sudden stop of the target results in an overshoot
>that is then corrected. The real person may overshoot a bit more.

>

>Your second point, however, is correct. A multiple-level system
>is needed to take advantage of any regularity in the disturbance.
>As we have not yet modeled systems that can take advantage of
>regularities, we use unpredictable disturbances.
>Whether a human being is able to take advantage of the regularity
>of a sawtooth movement depends very much on the frequency of the
>sawtooth. If it is too slow, tracking isn't improved because
>human estimates of frequency tend to wander over longer periods
>of the wave. If it is too fast, the lower-level control systems
>are on the verge of losing control anyway, so the higher-level
>control simply prevents control from deteriorating as fast as
>frequency rises. The main change between unpredictable and
>predictable disturbances is that the delay becomes smaller and
>more variable for a predictable disturbance.

>

[Martin]

>>The model would not care if the movement changed from sawtooth
>>to random. But would the same parameter settings of the model
>>provide the optimum fit both to the human tracking the sawtooth
>>and to the same human making the transition to tracking a
>>randomly moving target of the same maximum slew rate?

[Bill]

>So far there is little difference between human and model
>behavior in either case -- not enough to measure reliably. Do you
>agree, Tom?

[Tom -- now]

I am looking at two plots of data, one from a person, the other from a one-level PCT model with no perceptual delay and with k (the integration factor) estimated as described in our publications. The target was the variable-rate triangular wave Bill and I used in *Models and Their Worlds*. The target did not stop during these runs, but at the maximum of every excursion up or down, the increment of change in the target for the next excursion was randomly selected from a set of three possible increments.

I see segments of the data where the person's cursor lagged behind the target; the model's cursor did the same. Where the person's cursor ran ahead of the target, so did the model's. At some extrema, the person "overshot" the peak, then moved rapidly back on target; the model did the same. And at some extrema, the person "undershot" the peak, reversing direction too soon; so did the model.

Whether the target is triangular with variable rates, triangular with fixed rate, or a smoothed random function, these are the results. Some time ago, I ran (but of course never published) several trials in which the target switched between a triangular wave and a smoothed random function. The results were the same for the person and the model, even though the model's parameters (p^* and k) were estimated from the person's earlier performance while tracking a triangular wave.

I have not systematically matched slew rates of targets that follow either a triangular wave or a smoothed random function, as Martin suggests. I have done runs in which the target stops: both the person and the model continue moving -- they overshoot, and they both come back to the motionless target. The person overshoots farther than the model. The person looks at me quizzically and says something like, "It stopped working," or "It's broken," or "What happened;" the model does none of those things. (By the way, the people are right. Unlike traditional behavioral scientists, whose human participants are described as ignorant and unable to appreciate what is happening to them in an experiment, control theorists seem to find people who understand when things are going as they want, or not, and when things are working, or not.)

Rather than stopping the target, if I "disconnect" the control device, by eliminating it from calculations of the position of the cursor, then during tracking with no disturbance on the cursor, the cursor stops moving but the target continues on its way. Nearly all persons perform some vigorous handle movements, often accompanied with remarks implying that the device is broken or disconnected (They really are smart!), then they stop moving the handle; the model keeps on cranking out huge simulated movements and never understands.

Until later, Tom Bourbon

Date: Thu Jun 24, 1993 10:31 am PST
Subject: Simcon 4.5

[From Bill Powers (930624.1200 MDT)]

Sorry for the delay, but Simcon version 4.5 is now on Bill Silvert's server at "biome.bio.ns.ca" for access via ftp. Change directory after login to

pub/csg/simcon.

Included are simcon45.exe, a self-extracting compressed BINARY file, the C source code for compiling with Borland Turbo C 2.0 or later, a UUENCODEed version of the self-extracting file, and a .ASH version (Greg and Pat Williams' program for converting to ascii files). The .UUE and .ASH files can be downloaded as ASCII files.

Also included are four segments of the Primer series. More are in the works.

Contact me by email directly if you need any help with this.

Bill P.

Date: Thu Jun 24, 1993 12:44 pm PST
Subject: Re: model test

[Martin Taylor 930624 13:40] Tom Bourbon 930624 0854

Tom has proposed privately that we cease this interchange until we agree on the appropriate conditions to test my prediction. I agree, and will try to get the chapter by Rick and Bill. He has described data that may or may not conform to the required conditions. I can't tell.

Here's are the conditions that I think are needed for a test. There may be others I haven't thought of, but for now these seem adequate:

(1) a tracking setup in which the target moves with some regularity that can be detected by the subject. A regular sawtooth of constant speed and amplitude is fine.

(2) The target motion must be fast enough that the human has appreciable difficulty in tracking accurately. In other words, the cursor-to-target correlation must be fairly low, at least below 90%, and the lower the better (0.5 to 0.7 would be nice), provided that

(3) The single ECS model predicts the human behaviour with good fidelity (say over 0.99 correlation between model prediction and cursor movement). The model must have the same difficulty as the human in performing the tracking.

Having determined the parameters of the ECS for this set of conditions, change condition 1 to

1a) The target moves with no regularity that can be detected by the subject.

To match the intrinsic tracking difficulty, the maximum slew rate of the target should not change. Previously, I suggested a random white noise motion, but perhaps better would be a motion in which the target is always moving at the same speed, but changes direction at random times (i.e. Poisson distributed). I would assume that a single ECS model could be made to fit the human performance in this condition, degraded though that performance is likely to be. I doubt that the parameters of this model will be the same as those computed for the regular conditions.

If the random change condition degrades performance too much, the experiment might work by doing the 1a condition first, adjusting the target speed to permit the human to do at least better than randomly wiggling the control. We would want a correlation of, say 0.5 or 0.7 between cursor and target in one or other of the conditions.

The key point is to make it difficult for the human to track. If you don't have that condition, you won't see any effect.

I'll ask our librarian to get the book. Meanwhile, I think that if the results are available, they could be posted, and if not, Rick's Hypercard stack could probably be modified quite easily so as to do the experiment. Maybe, if I can find the time, I will try it. But time is not easily come by at present.

Martin

Date: Thu Jun 24, 1993 7:19 pm PST
Subject: IT prediction?

[From Rick Marken (930624.2000)] Martin Taylor (930624 13:40)

>Tom has proposed privately that we cease this interchange until we agree
>on the appropriate conditions to test my prediction.

I haven't been following this thread very carefully so I probably missed something but what is the prediction that is being tested?

>Here's are the conditions that I think are needed for a test.

[... the conditions]

>The key point is to make it difficult for the human to track. If you
>don't have that condition, you won't see any effect.

What effect?

Is this an effect that is predicted by information theory? If so, could you post the calculations or model equations?

Thanks Rick

Date: Fri Jun 25, 1993 2:50 am PST
Subject: paper suggestion

[Avery Andrews 930625.2000]

I'm hopefully emerging from the heap I've been buried in, so here's a suggestion for the paper (v 4 being the one I'm looking at):

Reduce the introductory section (above the centered heading `Control') to a single paragraph, along lines to the effect that although control has been on the scene in the behavioral sciences for a long time, its nature and potential have been seriously misrepresented, and that the purpose of the paper is to document this fact. E.g.

Feedback control has been seen as a central concept in the behavioral sciences for several decades. But its actual nature has been widely misunderstood, and for this reason, its potential and significance have been seriously underestimated, especially since the mid-to-late seventies, with the development of motor programming theory [assorted Schmidt et. al. refs] and Action Theory [assorted Kugler et. al. refs]. We begin by establishing some terminology, especially a technical usage of the term `control', which is normally used rather loosely.

E.g. no apologies or explanations for why people have trouble with it (I think the tone of Bill's intro is rather defensive, which would get readers off on the wrong foot). And I don't think I'd go on about thermostats for so long - I suspect that everybody really does understand them pretty well.

More later, hopefully.

Avery.Andrews@anu.edu.au

Date: Fri Jun 25, 1993 3:52 am PST
Subject: PCT revolution

I see an opportunity to make another supporter of PCT in the UK. The only thing needed is to make this person to fully understand PCT, and my guess is that it is possible (probable, perhaps?). He already realises the existence of control mechanisms at all levels of organisation of animal behaviour.

Professor George Houghton, University College, London, is giving a course on "Mechanisms of Cognitive Control: Insights from Neural Networks and Neuropsychology". I am attending his course in the beginning of September in Edinburgh, and if it proves to be what it looks like, I'll invite him to come to our European CSG Meeting of next year. I'll take my copy of BCoP with me to Edinburgh, in case he has never heard of Powers.

The abstract of his course reads: "This course will discuss the problem of control mechanisms in cognitive functioning from a connectionist/ neuropsychological point of view. Cognitive control mechanisms may be broadly defined as those involved in the selection and sequencing (including timing) of dominant perceptual, behavioural and mental states. Control issues arise at all behavioural levels, from the global organisation of coherent, adaptive social behaviour down to, for instance, the control of finger movements in typing.

Ideas regarding control in cognitive psychology have been influenced by the processing characteristics of the von Neumann serial computer, and the control structures provided by conventional programming languages. However, it is likely that control mechanisms in biological systems have little in common with such devices. For instance, biological systems appear to depend heavily on parallel rather than serial processing. Control of this processing may be based on a variety of inhibitory mechanisms, which have no analogue in conventional architectures. In this course I will review some recent theoretical and

empirical work in neural networks and neuropsychology which attempts to develop more biologically plausible models of control. Specific examples considered will include such phenomena as selective attention, serial order, planning and control of willed action."

There are a few things that hurt my ears, and I hope he is receptive.

Regards, Marcos.

Date: Fri Jun 25, 1993 6:05 am PST
Subject: Re: IT prediction?

From Tom Bourbon (930625.0834) Rick (930624.2000) Martin (930624 13:40)

[Martin]>

>>Tom has proposed privately that we cease this interchange until we agree
>>on the appropriate conditions to test my prediction.

[Rick]>

>I haven't been following this thread very carefully so I probably
>missed something but what is the prediction that is being tested?

[Me]

I'm not sure there is a quantitative prediction, Rick. I have suggested that Martin and I suspend our discussion about the need for memory, storage, history and the like, until after he reads the chapter you did with Bill, in Wayne's book. Then, after he is up to date on some of the things a simple 2-level PCT model can accomplish without the traditional "explanations," we will continue.

I believe Martin wants to use the model in a condition where direct control of cursor relative to target position will become poor, then, in a way not yet specified in our discussion, introduce imagination, internal models or the like. This is only my interpretation; Martin can tell you for himself. He will be at CSG, where I think the loops will be tighter and the chances for us to demonstrate the modeling will be better.

Until later, Tom Bourbon

Date: Fri Jun 25, 1993 9:15 am PST
Subject: Re: paper suggestion

From Tom Bourbon (930625.1145)

>[Avery Andrews 930625.2000]

>

>I'm hopefully emerging from the heap I've been buried in, so here's
>a suggestion for the paper (v 4 being the one I'm looking at):
..

>E.g. no apologies or explanations for why people have trouble with it
>(I think the tone of Bill's intro is rather defensive, which would get
>readers off on the wrong foot). And I don't think I'd go on about
>thermostats for so long - I suspect that everybody really does
>understand them pretty well.

The wrong food? Well put, Avery, even if it *was* by accident! We certainly don't want to give them any more reasons to start the feeding frenzy than we absolutely must!

As for everyone understanding thermostats pretty well, if they did, they would not use them inappropriately -- wrongly -- when they try to show us how well they understand. If we lined up all of the people who say they understand thermostats (as control systems) and asked for all who believe thermostats (and A/C systems) control the temperature of the air in the room to raise their hands, nearly all of them would raise their hands. The result would be the same were we to ask for a show of hands by all who believe thermostats control the outputs of A\C systems, or that they control to make a nice comfortable room. On the other hand, if we asked for all who believe thermostats control only their own input signals to raise their hands .. you know the rest. I am afraid a significant proportion of the group who would agree with that statement resides on this net.

Until later, Tom Bourbon

Date: Fri Jun 25, 1993 12:05 pm PST
Subject: CSGrev? (Electronic Journal)

[from Gary Cziko 930625.1934 UTC]

The information I recently received (attached below) after subscribing to EDPOLYAR caused me to consider the possibility of starting an electronic journal for CSG as well.

We already have CSGnet, the electronic forum which corresponds to EDPOLYAN below. We also, of course, have the hard copy CLOSED LOOP, edited by Greg Williams which includes coherent threads from CSGnet and which has begun to include articles as well. What would people think of starting CSGrev, which would disseminate peer reviewed articles to interested subscribers? What are some of the advantages and disadvantages of an electronic journal over the hard copy version to which we seem to be moving? Are there reasons for both?--Gary

P.S. I'm not sure I could set up all the fancy indexing and retrieval services that EDPOYAR enjoys, but I imagine that we could set this up somewhere, perhaps relying so more on Bill Silvert's help and file server.

P.P.S. One problem with electronic journals is their current "hostility" to figures (graphs). There is, however, a new standard for Internet graphics taking shape called MIME. Also, CSGnetters have shown amazing ability to put together informative ASCII-character graphics.

=====

Dear Subscriber:

Welcome to the EDUCATION POLICY ANALYSIS ARCHIVES, a peer-reviewed electronic journal accessible as a LISTSERV under the name EDPOLYAR at ASUACAD.BITNET. The journal is edited by Gene V Glass of the College of Education at Arizona State University to whom inquiries may be addressed: BITNET address Glass@ASU, INTERNET address form Glass@ASU.EDU; voice phone number is 602-965-2692 and ordinary mail address is College of Education, Arizona State University, Tempe AZ 85287-2411.

As articles are published by the ARCHIVES, they are sent immediately to the EDPOLYAR subscribers and simultaneously archived in two forms. A running notebook of all published articles in a given year is compiled and can be retrieved for any year by addressing the following letter to LISTSERV@ASUACAD.BITNET: GET EDPOLYAR LOG93 F=MAIL. To retrieve the journal for 1994, for example, you would send an e-mail letter to LISTSERV@ASUACAD.BITNET and make the sole contents of the letter read GET EDPOLYAR LOG94 F=MAIL. Articles are also archived on EDPOLYAR as individual files under the name of the author and the Volume and article number. For example, the article by Stephen Kemmis in Volume 1, Number 1 of the ARCHIVES can be retrieved by sending an e-mail letter to LISTSERV@ASUACAD.BITNET and making the single line in the letter read GET KEMMIS V1N1 F=MAIL. For a table of contents of the entire ARCHIVES, send the following e-mail message to LISTSERV@ASUACAD.BITNET: INDEX EDPOLYAR F=MAIL, that is, send an e-mail letter and make its single line read INDEX EDPOLYAR F=MAIL.

You may cancel your subscription to EDPOLYAR at any time by sending an e-mail letter to LISTSERV@ASUACAD.BITNET and including in the letter the single line SIGNOFF EDPOLYAR. You may wish to consider subscribing to an open forum where issues like those discussed in the articles published in the ARCHIVES are freely discussed. This forum operates as the LISTSERV known as EDPOLYAN and is located at ASUACAD.BITNET as well. You may subscribe by sending an e-mail letter to LISTSERV@ASUACAD.BITNET and making the sole contents of the letter read SUB EDPOLYAN <your actual name>. The Internet address LISTSERV@ASUVM.INRE.ASU.EDU will serve as well for transacting LISTSERV subscriptions and other business.

To receive a publication guide for submitting articles, send an e-mail letter to LISTSERV@ASUACAD.BITNET and include the single line GET EDPOLYAR PUBGUIDE F=MAIL. It will be sent to you by return e-mail. General questions about appropriateness of topics or particular articles may be addressed to the Editor.

Date: Fri Jun 25, 1993 1:05 pm PST

Subject: Contagion!

From Tom Bourbon (930625.1359)

Looking for a little light reading for the weekend? A little science fiction, perhaps? Try the latest issue of *Current Directions in Psychological Science*, vol. 2, no.3, published by the American Psychological Society. (I am allowed to comment -- I am a charter member of APS.) CD is, "the bimonthly journal that provides a timely source of information spanning the entire spectrum of scientific psychology and its applications. CD publishes brief (2,000-2,500 words), scholarly reviews that focus on emerging trends, controversies, and issues of enduring importance to the science of psychology." (from the inside, front cover) Sounds good, doesn't it? (I wonder if they would publish an article ... Nah!)

The article of choice for PCTers is:

Elaine Hatfield, John T. Cacioppo, and Richard L. Rapson,, Emotional Contagion, pages 96-99. There is something here for everyone -- language, mimicry, information, social control, moment-to-moment tracking (of other people's emotions), physiology, neuroanatomy, feedback, commands, consciousness, unconsciousness, emotion, and a slew of "tendencies" and "these results may be important in" Whew!

page 96:

"We define emotional contagion as

the tendency to automatically mimic and synchronize expressions, vocalizations, postures, and movements with those of another person's and, consequently, to converge emotionally. (pp. 153-154)

page 96:

"Because the brain integrates the emotional information it receives, each of the emotional components acts on and is acted on by the others." (The "components" include, but I presume are not limited to, the following: "conscious awareness; facial, vocal, and postural expression; neurophysiological and autonomic nervous system (ANS) activity; and instrumental behavior.")

page 96:

"Theoretically emotions can be caught in several ways."

By design and editorial policy, the article contains no original data; it is a review. But I cannot tell if *this* scholarly review focuses on an emerging trend, a controversy, or an issue of enduring importance to the science of psychology. Mostly it seems to be filled with assertions that "people do so and so," backed by brief summaries of research -- summaries that include numerous statements like, "tend," and "often," and "sometimes."

page 97:

As for "feedback:"

"Subjective emotional experience is affected, moment to moment, by the activation of and feedback from facial, vocal, postural, and movement mimicry."

"Theoretically, emotional experiences are influenced by the central nervous system commands that direct such mimicry and synchrony in the first place; the afferent feedback from such facial, verbal, or postural mimicry and synchrony; or conscious self-perception processes, wherein individuals make inferences about their own emotional states on the basis of their own expressive behavior. Given the functional redundancy that exists across the neuraxis, all three processes may (sic) operate to ensure that emotional experience is shaped by facial, vocal, and postural mimicry and expression."

And so it goes. The authors all are reputable and well known. True to the intent if APS, the article is representative of the state of the behavioral, social, psychophysiological sciences.

One in a while it is good to take a close look at real science.

Until later, Tom Bourbon

Date: Fri Jun 25, 1993 2:53 pm PST
Subject: Speech, thermostats

[From Rick Marken (930625.1500)] Bruce Nevin ()

>I guess you missed it when we talked about reasons for believing that some
>form of analysis by synthesis was necessary for language. I guess also you
>weren't watching this thread closely when I cited evidence that there
>are no reliable cues in the acoustic signal for identifying phonemes.

I had a brief fling with analysis by synthesis in the early 1970s, before I understood (or knew about) PCT and when I had only a nodding acquaintance with modelling. I think the best way to cure yourself of (or strengthen your commitment to) your belief in analysis by synthesis is to try to model it -- and compare that model to the performance of a transducer model. I believe that there is "synthesis" (in the sense of feedback effects of articulatory outputs on perceptual inputs) involved in the production of your own intended speech sounds. But what I don't see is how this imagined synthesis can help you analyze (recognize) what was said by another person. Presumably the synthesis would aid in telling you whether what you hear could be what you would have produced if you were trying to produce phomeme a,b, c...etc. The problem is that the synthesis must be done in imagination -- it doesn't go through the external environmental feedback path (or are you suggesting that people actually "mirror-talk" out loud in order to understand what other people are saying?). So there is no actual comaprison of what you hear to what might have been the result of such and such a pattern of articulation -- unless you imagine that sounds are deterministic results of particular articulation patterns. Even if you imagine this (surely false) fact about speech production, you still don't need synthesis to analyze what you hear anyway because the only way you can check what your imagined articulation would produce is against the template of the sound input itself -- which you claim is not a reliable indication of which speech sound was intended anyway.

The point is (to echo Tom Bourbon's wonderful reference to Field of Dreams) " build it (a working analysis by synthesis model) and we will come". I understand "control by synthesis" but I need a little help on the analysis part.

Tom Bourbon (930625.1145) to Avery --

>As for everyone understanding thermostats pretty well, if they did, they
>would not use them inappropriately -- wrongly -- when they try to show us
>how well they understand.

Absolutely right, Tom. Everybody says they understand how that tired old thermostat example works -- and yet VERY FEW really do. The idea that everyone understands the behavior of a thermostat is the biggest myth of them all, perhaps. The thermostat controls an electrical signal -- which is its perceptual representation of the controlled variable (temperature) . It is easy to prove that it is this electrical signal (not the variable it represents) that is controlled -- and one can do the appropriate experiments right in one's own home (if one is prepared to tamper with the thermocoil or the wires leading from it). The thermostat (like all systems that exist in a negative feedback relationship to their environment) controls its perception. If people won't stop and try to understand that simple little point then how can they go on to the next step that they so eagerly want to get to. Since they don't really understand how the simple little thermostat works, the next step for these erstwhile cyberneticists and control theorists is typically "input -- processing -- output" (or TOTE) land.

By the way, the paper you referred Martin to (by Bill and me) is, of course, in Mind Readings -- I think it's called "Levels of intention".

Best Rick

Date: Fri Jun 25, 1993 9:18 pm PST
Subject: Thermostat

[From Dag Forssell (930625 2215)] Rick Marken (930625.1500)

> The thermostat controls an electrical signal -- which is its
>perceptual representation of the controlled variable (temperature)

I think it would be more correct to say that a (typical, bi-metal spring action) thermostat controls the position of the end of the bi-metal spring.

The reference is the position of a contact point, set by turning a thumbwheel or such. The temperature at the thermostat is converted into a position by the bi-metal spring (by different thermal expansion of two metals). The mechanism is designed so that when contact is made, an electrical (error) signal is sent to the furnace, turning on the burner and fan. As the air temperature eventually rises in the room (at the thermostat) the spring breaks contact and the furnace stops. As the temperature in the spring drops, contact is restored and an error signal is again sent to the furnace.

With this, more precise analysis, it is clear that the dimension of the reference signal (position) can be compared apples and apples with the perceptual signal (position).

Best, Dag

Date: Sat Jun 26, 1993 12:17 am PST
Subject: thermostats

[Avery Andrews 930626]

Re Rick Marken and Tom Bourbon on knowledge of thermostats, I'd say that for the audience at which the attempted BBS paper is aimed ('theorists' in swear-quotes), you should assume that everybody really does know how thermostats work, tho they might not have very good ways of conceptualizing it, or of applying their knowledge to other situations. At any rate, they all *think* they know how thermostats work, and therefore *will not read* long explanations of them. The longer the explanation of thermostats, the more likely such readers are to a) skim it b) stop reading the article.

The way to smuggle in illumination of thermostats is on a 'hook' provided by some prevalent and identifiable misunderstanding, such as Abbs & Winstein's 'technical definition' wherby feedback is supposed to be restricted to correcting errors 'at the site where the error is introduced'. Thermostats can be used to illustrate the incoherence of this idea, as well as its divergence from actual practice (errors introduced at the front door being corrected in the basement).

Avery.Andrews@anu.edu.au

Date: Sat Jun 26, 1993 8:54 am PST
Subject: thermostats, conversions

[From Rick Marken (930626.1000)] Dag Forssell (930625 2215)

>I think it would be more correct to say that a (typical, bi-metal
>spring action) thermostat controls the position of the end of the
>bi-metal spring.

Thanks -- even though that means I will have to re-write the thermostat section in my (still in progress) book about perceptual control theory; unless there are thermostats that control an electrical representation of the thermocouple. I hope so; it would make comparison to the brain much easier. Do you know of thermostats that work by controlling an electrical voltage or current?

Avery Andrews (930626)--

I was agreeing with Tom about the lack of understanding that it is an input variable that is controlled by a thermostat. But I heartily agree with your review of version 3.0 or whatever it was of the BBS paper; introducing PCT with thermostats is definitely the wrong approach -- for the reasons you gave. The current version (that I have just scanned, Bill) is a significant -- very big -- improvement. Nice work Bill (and all the reviewers who have contributed to it -- which definitely leaves me out).

Marcos Rodriguez (sp?) --

>I see an opportunity to make another supporter of PCT in the UK. The only
>thing needed is to make this person to fully understand PCT, and my guess
>is that it is possible (probable, perhaps?). He already realises the
>existence of control mechanisms at all levels of organisation of
>animal behaviour.

Boy, are you in for the surprise of your life, Marcos. I can hardly wait to meet you at the EPCT conference and listen to your stories of this almost certain non-conversion.

Here's some tips if you really want to try this conversion: Make sure that you are both talking about the same thing when you use the word "control". If you manage to get that far, then see if he will agree to one, basic little fact about control -- that what is controlled is perceptual input. If you get that far then you have done more than our small flock has ever been able to do with any behavioral scientist, control engineer or cyberneticist. Then give the final exam question:

What is the main goal of the study of living control systems?

If he answers "To discover the perceptual variables that they control and how they control them" then you have found yourself a new PCTer.

Good luck!

Best Rick

Date: Sun Jun 27, 1993 6:07 am PST
Subject: Re: thermostats, conversions

[Martin Taylor 930626 18:00] (Rick Marken 930626.1000)

>Then give the final exam question:

>

>What is the main goal of the study of living control systems?

>

>If he answers "To discover the perceptual variables that they

>control and how they control them" then you have found yourself a new PCTer.

>

There is only one correct answer? I thought one of the tenets of PCT was that you are always satisfying multiple goals.

More to the point, since it is necessarily true that the perceptual variables people control change from moment to moment, what you are stating as the goal is to track individuals over time, which sounds really like a political-ethical problem. Check their bank statements, monitor their movements and their contacts, eh?

No, I prefer a goal like "To discover how perceptual control works, and to find its limits and possibilities for aiding the healthy survival of individuals, species, and ecosystems." It can't be of much interest what perceptual variable someone is controlling for at any particular moment, unless you are, for your own purposes, interested in affecting the feedback loop that the person is using, either to aid or to thwart the control of that particular perception. Science is about how things work, not about momentary or personal specifics.

>see if he will agree to one, basic little

>fact about control -- that what is controlled is perceptual input. If

>you get that far then you have done more than our small flock has ever

>been able to do with any behavioral scientist,...

A slight hyperbole, perhaps?

Martin

Date: Sun Jun 27, 1993 12:11 pm PST
Subject: Re: resolution and time

[Martin Taylor 930617 15:40] (Bill Powers 930623.1545 MDT)

>>I still don't understand this. 3.01572 is different from

>>3.01583, even if the RMS error of measurement is 300. The

>>readings differ, even if they don't usefully differ when

>>referred to the thing being measured..

>

>If all you know about the thing being measured is contained in

>the reading, as is the case for living organisms, you have only
 >two choices: treat the (300-unit) nonsystematic variations that
 >you measure (to six significant digits) as real and accurately
 >measure the fluctuations in the measured variable, or assume that
 >some noise source has entered the situation, and average readings
 >together to measure the average value of the (assumed) non-noisy
 >external variable to six figures. If you guess wrong, of course,
 >you will average out the perfectly meaningful variations in the
 >external variable.

There are, of course, other choices, but the principle is right. It is one of the things I have tried to get at over these months. You can improve the resolution of your measurement (not your reading) by averaging. (The other choices depend on what a prior knowledge you have about the noise sources and the external variable, such as their spectrum or temporal variability or internal correlation ...).

 >>I take the resolution of a telescope to be roughly the distance
 >>at which two point-spread functions meet at their 3 dB points.
 >>I've forgotten the name, but isn't the point spread function
 >>called the Airy disc or something like that? The centre
 >>points of the star images can be anywhere, but a another star
 >>closer than the diameter of the Airy disc cannot be resolved.
 >

>As two star images get closer together, more and more precise
 >measurement of the intensity distributions is required to
 >distinguish a double star from a single one. There is no
 >theoretical minimum detectable separation other than that set by
 >the grain of the emulsion or the spacing of CCD pixel elements,

You forgot the diffraction limit imposed by the telescope aperture

>and in principle that can be overcome by projecting the image to
 >as large a scale as necessary. The only absolute limit is set by
 >available observing time and the assumption that the source of
 >the image remains constant over the observation.

Yes, that says in another way what I have been saying. The resolution of a measurement improves over time if the data in successive time intervals is less than 100% correlated.

>The resolution of a
 >noise-limited measurement can always be improved by reducing the
 >noise (for example by cooling) or extending observation time.
 >Instrument-limited resolution can't be improved regardless of
 >observation time or noise reduction. When you can see the pixels
 >or the grains, that's it unless you magnify the image more.

The eye actually does better under some conditions. It can't help in resolving as distinct two point sources near to each other, but it can resolve non-collinearity (vernier acuity) and related resolution like aspects of vision to more than an order of magnitude better than the pixel grain (cone separation). It does this by oscillating the eye at a frequency faster than the visual system can deal with (around 60 Hz). This allows it to move the point spread function across the boundary of the individual cones (which are separated, in the fovea, by about the width of the diffraction-limited point spread function). Hence the time-averaging can provide a more subtle representation of where things are in the visual field than is permitted by the grain of the receptor field. If the visual question at issue is where things are relative to one another, and the eye can average over linear edges, then the answers one gets tend to be around 1 second of arc rather than 1 minute of arc.

>To some extent, pixel-limited resolution can be improved by
 >combining multiple images. This can effectively increase the
 >number of available pixels. But whether this results in
 >improvement depends on being able to position one image relative
 >to another within a known fraction of a pixel. That implies some
 >other means of resolving the image to less than one pixel on the
 >detector.

Well, yes and no. If it is possible to make some assumptions about the form of the thing observed, such as that it consists of patches with relatively smooth sharp edges, then one can get sub-pixel placement. The same can be done by shifting the sensor array with respect to the image. In visual astronomy, the atmosphere does that for us, which makes for a problem with static film, but allows dynamic CCD arrays to provide the data for later deblurring computations.

>Using the entire picture to get an average match can
>help, but only if the representation is precisely linear in
>space. With great labor, spatial distortions in the image can be
>mapped to a fraction of a pixel, but again that requires some
>other way of measuring positions of point-sources with
>fractional-pixel accuracy. And you must assume that the map
>remains valid from one exposure to another.

Yes, I've done this with satellite radar images in multiple bands. You have to assume that the real-world scene contains corresponding points and that distortion from image to image is smooth on a scale of many pixels. It does not require that the images be linearly related to each other, though that helps if it can be guaranteed.

In view of the context of this discussion, it is interesting to note that the currently favoured techniques for improving resolution in ways related to this come under the general name of "maximum entropy" methods. The idea is that the most likely representation of what is "out there" is at an informational extremum. It works pretty well most of the time. (The actual maths is a bit over my head, I'm sorry to say). I remember it being demonstrated first at the Montreal World's Fair in 1967.

>>The resolution is referred to the star field, not to the film grain.
>
>No. The "resolution" of the star field is infinite. It is only
>the detector grain that ultimately limits the resolution of the picture.

I said that it is "referred to" not that it is "limited by." One is talking about where in the sky stars are, not where on a film plate or imaging device their images are. The resolution of whatever you are measuring must be referred to the thing being measured, wherever it might be limited. The point of a measurement is to know more about the thing being measured than you did before you took the measurement.

My point in all this has been that the resolution of a measurement in noise can be improved by extending the measurement over time, the limit being (to quote you paraphrasing me) "set by available observing time and the assumption that the source of the image remains constant over the observation." The measuring instrument can gather information at some set rate. Over time, the information gathered increases linearly, provided that the autocorrelation function of the measurement goes to zero fast enough after some finite time. Your discussion of my telescope example agrees with this, so I presume that we can continue from a mutual acceptance of that point?

Martin

Date: Sun Jun 27, 1993 3:47 pm PST
Subject: Re: Resolution and time

[From Bill Powers (930627.1630 MDT)] Martin Taylor (930617 15:40) --

I don't want to get into an argument about who knew what and when, but the fact is that the tradeoffs between observing time and resolution (both amplitude and spatial) have been well-known in the observational sciences for a long time, including the mathematical relationships involved; information theory is not required to derive them. By citing things I learned during my experiences with astronomy in the 1960s, I have been gently trying to say that you are not "reminding" me of these relationships, nor am I "paraphrasing" you in describing them. There are other ways to get at the same results without ever mentioning the terms "information" or "entropy" and without converting any quantities to logarithms and giving them special status because of that elementary transformation.

Many of the probability calculations may be the same, but as I found out in developing low-light-level imaging systems, the probability calculations are very approximate anyway. Actual noise distributions seldom fit any convenient form such as Gaussian or Poisson, and real bandwidths are neither uniform nor truncated at some convenient frequency. The

statistical distribution of electrons leaving a photocathode is different from that of electrons leaving a secondary emitter, both depend on the solid angle of emission, both are variable with temperature and magnetic (focussing and deflection) fields, and neither distribution fits the standard forms particularly well.

All of these real-life considerations mean that the more elaborate the mathematical scheme that is imposed on the real measurements, the less likely it is to be of much empirical help or predictive value. It's nice to be able to guess within a factor of ten what the quantum efficiency of the overall system will be, photocathode to display screen, but when you plug in a new image isocon you have to start over. It's a lot faster just to plug it in and turn it on.

When you say that quantity of information is linearly dependent on observing time, this is not a RESULT from information theory; it is the basic PREMISE by which information is defined. Information is a construct deliberately chosen to have the convenient feature of growing linearly with time in certain narrowly-defined situations, so that instead of having to deal with products one can deal with sums.

I did not forget the diffraction limit of the telescope. In fact, in citing an instance of determining relative stellar positions to 0.002 sec of arc, I referred to the diffraction "limit" of the 18-inch telescope as being about 1/4 sec. arc.

But you are right on one score: to separate very close double stars, you have to assume that you are dealing with two point-sources. By eye and without detailed intensity measurements, all you can say for sure is that an elongated image is not made by a single point source. I doubt, however, whether either the telescope or the film makes those assumptions. What is recorded is what is recorded. The assumptions come later during the cognitive analysis of the data.

As to visually detecting smooth sharp edges with second-of-arc resolution, I don't believe it. I think this phenomenon is a mass measure resulting from visually detecting the invariance of a line image with respect to translations along it. In college I heard about line curvatures being detected with such impossible accuracy. It took just a few minutes to work out that a comparison of the angles of the two ends would reveal a very small departure from a chord as a rather large change in angular orientation of the ends. It's very tempting to accept such gee-whiz findings at face value, but before you do you had better think of all the alternatives that don't violate physics so severely. And also before you rush in with an ad-hoc explanation based on all sorts of assumptions that will later prove irrelevant, as the phenomenon is actually explained in a different and much simpler way.

Best, Bill P.

Date: Sun Jun 27, 1993 4:43 pm PST
Subject: Goals of PCT

[From Rick Marken (930627.1730)]

I said that the main goals of the study of living control systems (from a PCT perspective) were:

> "To discover the perceptual variables that they
>control and how they control them"

Martin Taylor (930626 18:00) replied:

>No, I prefer a goal like "To discover how perceptual control works,
>and to find its limits and possibilities for aiding the healthy survival
>of individuals, species, and ecosystems."

That's fine -- and I did include the question of "how" perceptual variables are controlled. But if you don't try to figure out what people are controlling, what do you have to explain?

> It can't be of much interest what perceptual variable someone is
>controlling for at any particular moment,

But that's what it's all about, isn't it. People control perceptual variables; in order to determine how this is done (how it works), you have to have examples of perceptual control to study. Demonstrable examples of controlled perceptual variables are the archival data on which PCT must be built. Without such data, what do you have to explain? With such data, you can see if there are relationships between the kinds of variables that are controlled; see if the variables seem to clump into types (as implied by the hierarchy), see if some types of variables are always used to control other types -- or not, etc. Even if you just want to study the parameters of control of one particular variable, you still have to know what variable is actually being controlled, and, possibly, other variables that are being controlled in order to control that variable. It all comes down to determining what perceptual variables an organism is controlling -- that's the basic phenomenon that PCT aims to explain.

> Science is about how things work, not about momentary or personal specifics.

So the momentary, personal specific rate at which a particular ball is rolling down a particular point of a particular inclined plane is not what science is about? Maybe not -- once you have a theory that allows you to deduce that data. But science (at least the kind I like to do) starts with data. Then we make up models in the hopes of finding the underlying explanation of why we get the data we get. I think that the success of this modelling process has led many scientists to forget about the importance of observation. That's part of the reason for the lack of interest in PCT -- behavioral scientists are more interested in fancy models (whether they explain any observations PRECISELY or not) than in phenomena. Thus we get "trendy science": strange attractors, complex systems, information theory, etc.; theories that are loved more for how they sound than for their ability to explain phenomena.

I said:

>see if he will agree to one, basic little
>fact about control -- that what is controlled is perceptual input. If
>you get that far then you have done more than our small flock has ever
>been able to do with any behavioral scientist,...

Martin replies:

>A slight hyperbole, perhaps?

Nope. Precisely true. The only people who have ever REALLY understood and agreed to the idea that behavior is the control of perceptual input variables have come to that understanding and agreement all on their own.

Best Rick

Date: Mon Jun 28, 1993 8:08 am PST
Subject: Re: thermostats, Goals of PCT

From Tom Bourbon (930628.0945) [Avery Andrews 930626]

[Avery]

>Re Rick Marken and Tom Bourbon on knowledge of thermostats, I'd say
>that for the audience at which the attempted BBS paper is
>aimed ('theorists' in swear-quotes), you should assume that everybody
>really does know how thermostats work, tho they might not have
>very good ways of conceptualizing it, or of applying their knowledge
>to other situations. At any rate, they all *think* they know how
>thermostats work, and therefore *will not read* long explanations
>of them. The longer the explanation of thermostats, the more likely
>such readers are to a) skim it b) stop reading the article.

[Tom]

I agree, so long as we also place swear-quotes around "really does know how thermostats work."

Rick Marken (930626.1000) >Rick replied to Avery Andrews (930626)--

>

>I was agreeing with Tom about the lack of understanding that it
>is an input variable that is controlled by a thermostat. But I

>heartily agree with your review of version 3.0 or whatever it was
 >of the BBS paper; introducing PCT with thermostats is definitely
 >the wrong approach -- for the reasons you gave. The current version
 >(that I have just scanned, Bill) is a significant -- very big --
 >improvement. Nice work Bill (and all the reviewers who have contributed
 >to it -- which definitely leaves me out).

Agreed, on both counts.

=====
 [Martin Taylor 930626 18:00] (Rick Marken 930626.1000)

[Rick]
 >Then give the final exam question:
 >
 >What is the main goal of the study of living control systems?
 >
 >If he answers "To discover the perceptual variables that they
 >control and how they control them" then you have found yourself a new PCTer.

[Martin]
 No, I prefer a goal like "To discover how perceptual control works, and to find its limits and possibilities for aiding the healthy survival of individuals, species, and ecosystems." It can't be of much interest what perceptual variable someone is controlling for at any particular moment, unless you are, for your own purposes, interested in affecting the feedback loop that the person is using, either to aid or to thwart the control of that particular perception. Science is about how things work, not about momentary or personal specifics.

[Tom - now]
 So, mainstream psychologists are right, after all: psychology is and ought to be the science of "hypothetical nomothetic androgynous persons," or statistical fictions -- the science of a misapplied "method of relative frequencies," to use Phil Runkel's phrase.

I don't really think you think that, Martin, but your remark to Rick brought flashbacks of my days as a psychologist, in a department of psychologists who believed in their bones that the old ways are the best ways in the science of psychology. Phil's elegant discussion of differences between the method of relative frequencies and the method of specimens reveals the seemingly paradoxical nature of the science of living control systems: careful study of a *single* instance of control achieved by a *single* "specimen" reveals facts that apply to virtually *all* members of the "natural kind" of which the specimen is -- well, a specimen. In that light, Rick's characterization of the goal of PCT-science is correct, even if it is carried to the level of studying one specimen; a limit Rick never mentioned. When it comes to control, the generalities are in the specifics.

In contrast, endless "net casting" through misapplications of the method of frequencies (statistical testing of traditional hypotheses of linear causation) leads to a science in which it is impossible to predict the specific actions of even one specimen, with any meaningful degree of certainty or precision.

As you say, a science of living control systems is about more than, "what perceptual variable someone is controlling for at any particular moment," but that is precisely where the science begins, and that is what it should explain.

Until later, Tom Bourbon

P.S.
 For anyone who is new to the net, I was referring to the following (excellent) book comparing research methods in psychology (and beyond):

Phillip J. Runkel (1990). Casting Nets and Testing Specimens: Two Grand Methods of Psychology. New York: Praeger.

Highly recommended by nine out of nine perceptual control theorists.

Date: Mon Jun 28, 1993 8:58 am PST
 Subject: Re: thermostats, Goals of PCT

[Martin Taylor 930628 12:30]
 (Tom Bourbon 930628.0945 and Rick Marken --didn't save the date stamp)

I guess I should conform to the practice of putting smiley-faces on postings that are partly tongue-in-cheek :-).

Now, do I believe that or not? And how partly deserves what part of a smiley?

Martin

Date: Mon Jun 28, 1993 2:29 pm PST
 Subject: CSGrev (Electronic Journal)

[From Rick Marken (930628.1430)] Gary Cziko (930625.1934 UTC) --

> What would people think of starting CSGrev,
 >which would disseminate peer reviewed articles to interested subscribers?

Sounds great. But who are the peers (of the PCT realm?) and what would they have to review? A lot of opinions are tossed about on CSGNet but I'm not aware of a lot of research projects that could be the basis for PCT type journal articles (the kind of research which involves the collection of actual DATA from living systems -- not just data from models of those systems). I'd be happy to send in my oft rejected "Hierarchical behavior of perception" paper to kick things off. Maybe if I were a peer (Lord Richard?) I could finally get the thing accepted. Then I could spend more time working on my spelling.

Best Rick

Date: Tue Jun 29, 1993 1:17 pm PST
 Subject: Natural Life Too Limiting?

[from Gary Cziko 930629.1935 UTC] Tom Bourbon (930628.0945) said:

>So, mainstream psychologists are right, after all: psychology is and ought
 >to be the science of "hypothetical nomothetic androgynous persons," or
 >statistical fictions -- the science of a misapplied "method of relative
 >frequencies," to use Phil Runkel's phrase.

Just after reading Tom's remark, I received an announcement for the new journal Artificial Life (MIT Press). I know Tom will love the stated rationale for this journal.

"Biology is the scientific study of life - in principle, anyway. In practice, biology is the scientific study of the only kind of life that has been available to observe - life on Earth based on carbon-chain chemistry. Thus, theoretical biology has long faced the fundamental obstacle that IT IS IMPOSSIBLE TO DERIVE GENERAL PRINCIPLES FORM SINGLE EXAMPLES [emphasis added].

"In order to derive general theories about life, we need an ensemble of instances to generalize over. Since it is quite unlikely that alien life forms will present themselves to us for study in the near future, our only option is to try to create alternative life forms ourselves - ARTIFICIAL LIFE - literally 'life made by man rather than by nature.'"

So, it look like the biologists are bored with being limited to carbon-based life forms and have been unable to "derive general principles" therefrom. Perhaps if they were to entertain the possibility that life involves and requires perceptual control, they wouldn't be so bored.

And what kind of alternative life forms do ALifers create? I've seen lots of neat computer simulations, but aside from the Powers-McPhail-Tucker "Gather" program, I've yet to see one that includes a simulation of environmental disturbances. All the other simulations I've seen exist in alien world in which commands are transformed unerringly into consequences. And the ALifers will probably go on to use this imaginary environment in order to show how organisms function without perceptual control.

Perhaps this kind of artificial life should be called instead "imaginary, alien life."--Gary

Date: Tue Jun 29, 1993 1:55 pm PST
 Subject: Re: Natural Life Too Limiting?

From Tom Bourbon [930629.1551] Gary Cziko 930629.1935 UTC

>>So, mainstream psychologists are right, after all: psychology is and ought
 >>to be the science of "hypothetical nomothetic androgynous persons," or
 >>statistical fictions -- the science of a misapplied "method of relative
 >>frequencies," to use Phil Runkel's phrase.
 >
 >Just after reading Tom's remark, I received an announcement for the new
 >journal Artificial Life (MIT Press). I know Tom will love the stated
 >rationality for this journal.
 >
 >"Biology is the scientific study of life - in principle, anyway. In
 >practice, biology is the scientific study of the only kind of life that has
 >been available to observe - life on Earth based on carbon-chain chemistry.
 >Thus, theoretical biology has long faced the fundamental obstacle that IT
 >IS IMPOSSIBLE TO DERIVE GENERAL PRINCIPLES FROM SINGLE EXAMPLES [emphasis
 >added].

What a gem!

I amend my original post:

Psychology is and ought to be the science of "hypothetical nomothetic androgynous
 non-carbon-based units," and furthermore, psychologists are obligated to advocate the
 elimination of all carbon-based units, which, after all, are to blame for the fact that
 theoretical psychology has long faced the fundamental obstacle that IT IS IMPOSSIBLE TO
 DERIVE GENERAL PRINCIPLES FROM SINGLE EXAMPLES.
 Signed, V'ger

First psychologists blame people for the messiness of their data and in retaliation they
 create nomothetic persons. Now biologists blame nature for the messiness of their theories
 and in retaliation they create -- who knows what -- nonlife? Ain't real science grand?

Watching out for my carbon-based backside, I remain,
 Tom Bourbon Department of Neurosurgery
 University of Texas Houston Medical School Phone: 713-792-5760
 6431 Fannin, Suite 7.138 Fax: 713-794-5084
 Houston, TX 77030 USA tbourbon@heart.med.uth.tmc.edu

Date: Tue Jun 29, 1993 6:52 pm PST
 Subject: Generalizing; statistics vs generative models

[From Bill Powers (930629.1845 MDT)]

Gary Cziko (930629.1935 UTC), Tom Bourbon (930629.1551)

>IT IS IMPOSSIBLE TO DERIVE GENERAL PRINCIPLES FROM SINGLE EXAMPLES.

What an event! Theoretical biology commits suicide before our very eyes! For if it is
 impossible to derive general principles from single examples, then it is altogether
 impossible to derive general principles. There is only one universe.

 I've been preaching for quite a few years, not always coherently, that generalization is
 simply not basic science. Generalizations are what you create when you have no theory of
 the regularities underlying specific phenomena. When you don't have any concept of how
 natural phenomena are created from beneath, all you can do is observe what happens and try
 to guess at some superordinate general rule of which specific events are only examples.

Many scientists seem to confuse this search for general rules with the search for
 mechanisms underlying phenomena. Some, of course, are perfectly clear about the difference,
 and come down solidly on the side of generative modeling. Tom, can you pick out that

citation (I think it was in "Worlds") and post it again -- the one about how the laws of motion would have been treated under the strategy of generalization?

This is what is behind my stubborn rejection of "probability" as an operative factor in nature. The essence of probability is to observe what has happened and to compute the chances of its happening again, with no basis for predictions except experience at the level of phenomena. All kinds of generalizations can be made about probabilities, single, jointly, and in bunches. Mathematical theorems can be developed to bring out interesting properties of probabilistic calculations. But behind all these calculations, there is the simple fact that we must use probabilities mainly when we have no model of the underlying order behind phenomena.

To characterize anything in nature as being intrinsically probabilistic is simply to give up the search for a generative model. Sometimes this is the only choice: we must treat pressure and temperature as probabilistic phenomena because it is simply not practical to keep track of all the molecules and compute their individual interactions from basic principles. Of course the world of supercomputing is changing that picture; what was once wildly impractical is now done routinely with quite large assemblages of molecules, or stars. Fluid dynamical calculations on a Cray are not probabilistic because the behavior of each molecule (or packet) is computed from an underlying generative model of physical interactions. Such models don't yield just general mass measurements; they yield the behavior of all the particles involved, in detail. Knowing what actually happens, one can then create generalized measures at leisure.

If we were to construct a generative model of a pair of dice, complete with shape measurements and coefficients of elasticity, hardness, mass distribution, and so on for both the dice and the surface against which they are thrown, we could no longer use the throwing of dice as an example of probabilities. True, there would still be certain conditions that would require more precision of prediction than obtainable, as the dice land in orientations near a bifurcation on an edge or corner, but between these conditions, as the model was refined, there would be broader and broader regions in which the outcome was deterministic with a high degree of accuracy. For almost all selections of initial conditions of position, velocity, and spin, predicting the outcome would simply become a matter of running the model. Given the initial conditions, we could no longer say that the probability of a certain number coming up on one die was $1/6$.

A probabilistic treatment of throwing dice depends on the assumption that all ending orientations would occur unmpredictably and with equal frequency over an infinite number of throws. This assumption is not, of course, warranted for any particular pair of dice or for all initial conditions, as dice cheaters know. But more important, the underlying assumption is that the final orientations ARE NOT IN FACT DETERMINED BY ANY REGULARITY IN NATURE. If there were underlying laws, if generative models could be stated, then using probability measures might become a matter of convenience or necessity, but would not reflect anything fundamental in nature. The fundamental laws of dice-throwing would be the laws of physical mechanics, not probability.

There is always a temptation, I think, to look on the unpredictable noise in an experiment and decide that this, at last, is the level at which nature itself is governed by chance. Quantum physicists of at least one school appear to have made this decision. But so did psychologists, when they decided that the variability of behavior, the unpredictable variations that made their causal models fit the data so poorly, was inherent in organisms. In the case of the psychologists, this was the wrong decision; they were simply using a model that was insensitive to most of the actual regularities. I think this is always the wrong decision; at least it is never possible to prove it is the right one. All that's required to refute the hypothesis of inherent randomness is to come up with a generative model than makes the phenomenon quantitatively predictable.

The problem with the decision to choose randomness as an inherent aspect of nature is that one then stops looking for an underlying generative model. One starts to invent all sorts of reasons for not looking, but they all rest on the assumption that there is no point in looking because there are no further regularities to be found. I think this decision has been taken over and over, at various levels of sophistication, throughout human history. It's much the easiest to give up on explaining why plants grow better in one kind of soil than in another, and throw it all in the laps of the gods.

The concept of science as a process of generalizing really rests on giving up on generative models (or on never having heard of them). If you have no model that can explain phenomena

in terms of underlying mechanisms, the best you can do is to sort the phenomena by subjective similarities, classify them, classify the classes, and look for empirical rules that apply to the classes. If the rule fails too often, you can always form a more general set of classes still and look for rules at the new level of abstraction.

This leads to the myth of the Grand Master Generalization, the rule that is so abstract that it applies to everything. Finding such a generalization is like learning the Secret Names of all things; one is filled with a sense of understanding everything, but is incapable of actually doing anything. Ursula Le Quin's Wizard of Earthsea, through learning those secret names, was able to command supernatural forces that transcended the laws of ordinary nature. But the Wizard of Earthsea could never have calculated an orbit. The Wizard of Earthsea lived in an imaginary universe where anything becomes possible if you say it is possible.

Probabilistic calculations are sometimes the only practical approach to prediction. But they must always give way to calculations based on generative models when such models exist, because the generative models predict specific events, while the probabilistic approach is inherently incapable of doing that. The generative model disproves the assumption that there is no regularity behind the fall of the dice. And as I said, that reduced the statistical interpretation from a fundamental insight into nature to a mere practicality.

Best, Bill P.

Date: Wed Jun 30, 1993 4:43 am PST
Subject: Bugging Gary

From Greg Williams (930630) Gary Cziko 930629.1935 UTC

>And what kind of alternative life forms do ALifers create? I've seen lots
>of neat computer simulations, but aside from the Powers-McPhail-Tucker
>"Gather" program, I've yet to see one that includes a simulation of
>environmental disturbances. All the other simulations I've seen exist in
>alien world in which commands are transformed unerringly into consequences.

You forgot about Dr. Beer's artificial cockroach program. In our IBM port of this program, you can alter the positions of barriers and food at any point in a run; it wouldn't be much work to make the barriers and food move as the program runs. In BOTH "Gather" and "BeerBug," "commands are transformed unerringly into consequences*," but the consequences do not always (immediately or sometimes even ever) result in reducing the errors of control loops to zero. I.e., the artificial cockroach can get stuck in a cul-de-sac and "die" when its "energy" reserves run out.

*"low-level" consequences, such as moving a certain foot forward at a certain speed

As ever, Greg

Date: Wed Jun 30, 1993 6:09 am PST
Subject: First European CSG Meeting

[From Marcos Rodrigues (930630.1455 BST)]

The prospects for our 1994 meeting are looking good. 19 people have already indicated their intention to attend: 17 as you can see above (3 to be confirmed: Ian, Jon, and Hans) plus Franz and Hettie Plooi from The Netherlands with no e-mail address.

For reasons outside our control, we have to postpone the meeting (one day). Instead of commencing on Wednesday evening, we have to schedule it from Thursday evening, 23 June 1994 to Monday morning, 27 June 1994. As before, Thursday is arrival day and Monday is departure day, with no sessions planned neither on Thursday nor on Monday.

I am sorry for the change and hope it will not bring any inconvenience to anyone. However, if for any reason you still want to come on Wednesday, you are much welcomed; we will be able to accommodate a few people with no problems. If the change of dates means that you will not be able to come, please let me know as soon as possible so that we can think of alternative arrangements.

In addition, if you intend to bring any member of your family who is not attending the meeting, please let me know before Christmas. We have booked Gregynog in for 30 people; in principle, the bed and dining facilities of the place is larger than that, but we have to know the exact numbers well in advance to avoid disappointment.

 To Mary Powers --

I've received the HFSP application forms. Thanks a lot. Our meeting definitely fit the criteria and I'll be applying soon. Hope we can get it.

Best regards, Marcos Rodrigues mar@aber.ac.uk

Date: Wed Jun 30, 1993 10:27 am PST
 Subject: Science

[From Dag Forssell (930630 0910)]

I noticed the following review in the paper. I am now reading the book and find it delightful and useful. While I am presently only 1/4 of the way through, I am already considering the counter-intuitive aspects of PCT and formulating analogies with inspiration from the book.

 Book review Los Angeles Times June 4, 1993

A BIOLOGIST'S SPIRITED DEFENSE OF SCIENCE, REASON

By LEE DEMBART, special to The Times

THE UNNATURAL NATURE OF SCIENCE: Why Science Does Not Make (Common) Sense by Lewis Wolpert
 Harvard University Press \$19.95, 191 pages

If Shakespeare had never been born, we would not have "Hamlet," "Macbeth" or "The Merchant of Venice." No one else would have written them.

But if Isaac Newton had never been born, we would still have the law of gravitation. Sooner or later, we believe, someone else would have discovered it.

These statements cannot be proved. Shakespeare and Newton were born, and we have no Shakespeareless and Newtonless world to compare with our own to see whether "Hamlet" was nonetheless being performed and gravity was unknown.

But the idea that artists create something unique while scientists discover laws that are already there is deeply embedded in our collective notion of reality.

This idea has come under powerful attack in the post-modernist age, where the "social construction of reality" is the rallying cry of those who claim that readers are more important than authors, that reality is malleable and that scientists, like artists, create the world more than discover it.

For example, some feminist scholars have written that modern science embodies a male, sexist reality in which forces "act on" objects, DNA is a "master molecule" and evolution is driven by a "struggle" to survive. A bias-free world, they say, might require redefining objectivity, rationality and the scientific method.

Into this argument steps Lewis Wolpert, a biologist at University College, London, who offers a spirited defense of traditional science, objectivity and reason. The world and its laws are independent of us, and scientists try to figure out those laws, which are timeless, universal and independent of human wishes. Scientific truth is public, testable and verifiable. In assessing a scientific statement, the only thing that matters is its correspondence with reality and its ability to predict the results of experiments.

"Science always relates to the outside world, and its success depends on how well its theories correspond with reality," Wolpert writes on Page 2 of "The Unnatural Nature of Science," a series of essays that uncompromisingly defends science as fully entitled to the special and privileged position that it enjoys. "Science provides the best way of understanding the world," he says.

Wolpert is especially withering in his response to the relativists, philosophers of science, sociologists of science and Kuhnians (they of the "paradigm shifts"), who, in his view, create conundrums where there are none.

He probably speaks for all working scientists when he asserts: "My own position, philosophically, is that of a common-sense realist: I believe there is an external world which I share with others and which can be studied."

He acknowledges that "authority, fashion, conservatism and personal prestige play important roles" in science. "But it is misleading to think, as some have claimed, that science is really nothing but rhetoric, persuasion and the pursuit of power."

If a theory "does not conform with the evidence, If it is not internally consistent, it does not provide an adequate explanation, the authority and all the other social factors count for nothing: It will fail."

To the relativists and post-modernists, Wolpert issues a simple and elegant challenge: What alternate theories do you have to explain the data? Is there a physics that is not based on a set of basic forces? Is a biology possible that is not based on cells and DNA? Is there a chemistry that does not include the periodic table?

There is much more to this wonderful book. Wolpert explains the difference between science and technology. Science is about generalization and abstraction. In itself, it is useless. Technology, on the other hand, produces usable objects.

Many civilizations had technologies, some of them, like the Chinese, quite advanced. But only the Greeks sought understanding. As a result, we trace science to them. No matter that almost all of the Greeks got the wrong answers. They asked the right questions.

Amid all of this fanfare for science and the scientific method Wolpert's prose is measured and thoughtful. He acknowledges the sources and strengths of other views, and he makes no claim that science can know everything. In fact, he makes clear that it can't.

Moral and political problems are outside its purview, as are justice, happiness, love and ultimate values. These matters are in principle beyond the range of science.

Furthermore, some systems, such as human behavior and society as a whole, are so complex that "knowledge in these fields is barely at the stage of a primitive science." We may some day be able to predict the stock market, but that day is neither foreseeable nor imaginable. So much for the social sciences.

In an age when fundamental ideas about the nature of truth are assailed, when scientists are derided as madmen who threatened the world with nuclear weapons and genetic engineering, it is a pleasure to read a clear, level-headed and persuasive defense of the scientific enterprise. We are better off with knowledge than without it.

Best, Dag

Date: Wed Jun 30, 1993 12:53 pm PST
Subject: Philosophy Discovers Control (Again)

[from Gary Cziko 930630.1900 UTC]

van Gelder, Tim. (1993). Pumping intuitions with Watt's engine. *Cog Sci News*, *_6_(1)*, 4-7 (Cognitive Science Program, Lehigh University, Bethlehem, PA)

In this article, an Indiana University philosopher discovers control theory via Watt's steam engine governor and makes some quite appropriate observations such as:

"The centrifugal governor keeps the engine speed at a constant speed, and it does so by adjusting the throttle valve, but it never COMPUTES how much that valve should be adjusted."

"Input, internal operation and action do not take place at distinct times; rather they are all continuous and simultaneous."

"The best way to think about the computational governor is by means of an algorithm--an hierarchically organized series of instructions to be followed, one after another. These instructions correspond, of course, to the sequence of discrete operations in the machine. But you can't think of the centrifugal governor in these terms ; there simply aren't any discrete, ordered steps taken, and so nothing for the instructions to specify. To rigorously characterize the centrifugal governor's behavior we have to use tools that philosophers and cognitive scientists are typically not so familiar with: differential equations and dynamical systems theory."

Can the realization that control systems control their input be far behind?

I hope van Gelder will be pleased to learn that there has been considerable theoretical development and empirical research on applying control theory to human behavior and cognition--I plan to let him know. We may be able to add another philosopher to CSG (I think Hugh Petrie gets lonely at times)!

=====
 Cog Sci News is distributed free of charge. You may request a copy of Vol. 6 No. 1 by sending and be placed on their mailing by sending a message to that effect to jbg1@lehigh.edu (Bill P., I've already asked that a copy be sent to you).

=====
 In his references, van Gelder lists:

van Gelder, T., & Port, R. (Eds.) (to appear 1993). *_Mind as motion: Dynamics, behavior, and cognition.* Cambridge, MA: MIT Press.

--Gary

Date: Wed Jun 30, 1993 2:05 pm PST
 Subject: Re: assorted replies, and some more

Rick Marken,

I am sorry to take so long to respond to your response (930614.1400) to my questions about modeling the phenomenon of understanding (930614). I have been off the net for two weeks due to a family crisis, APS convention, etc.

I appreciate your post. As a way of helping you see where I am coming from, I thought it might be best to send you parts of a poster presentation I did at APS (American Psychological Society). I am intending to submit something along this line for publication.

I hope you would take the time to go through it and help me out with your PCT input. I am aware that its not reporting The Test with higher level functioning, but to me its a step in that direction.

My questions about understanding are rooted in this paper but are also an attempt to go beyond this present work.

Title--Computer-based Drill Performance Predicted by Feedback Processing: Explained by Perceptual Control Theory

Abstract

The theoretical concern is explanations of higher level human functioning which are rooted in Perceptual Control Theory (Powers, 1973) The applied concern is the identification of the source of differing performance levels. Data are reported from over 150 hours of computer-based drills. Fifty-four subjects were tasked to learn name labels for 27 combinations of simultaneous displays of graphic, aural, and iconic information. After each response subjects were asked to judge the certainty of their response correctness (a form of metacognitive estimate). Then response sensitive feedback information was displayed. The highest performing subjects showed an inverse relationship between certainty and feedback frame latency. This could indicate that subjects had a goal (or reference level) for understanding the associations between stimuli and the correct label. They also displayed a significant increase in feedback time when responses were incorrect, indicating an additional goal of responding correctly to the stimuli. The middle performing subjects did not demonstrate a goal of understanding but tended to exhibit only a goal for responding correctly (and that with lower magnitude). The lowest

performing subjects exhibited little consistent relationship between feedback time and correctness, nor between feedback time and certitude. We interpret their poor performance to be due to a lack of goals (reference levels) for correctness or understanding. This work demonstrates an empirically based and theoretically driven explanation of an enduring practical problem in training and education.

Introduction

Our proposed solution to the enduring problem of understanding instructional feedback (e.g. Kulhavy & Stock, 1989; McKendree, 1990) for adaptive instruction, is to view humans as hierarchical perceptual control systems. Three aspects of PCT or Perceptual Control Theory (see Bourbon, 1990; Cziko, 1992; Marken, 1986; Powers, 1973, 1978, 1990) are pivotal to our understanding of how humans use instructional feedback: 1. a human always seeks to maintain or control perceptual input so that it matches internal perceptual references or goals, in order that error signals or discrepancy will be minimized; 2. a human's behavior is driven by discrepancy/output, and as such behavior is simply the means of controlling perceptual input; 3. a human is composed of a hierarchy of levels of control, such that the output of higher level control systems control the internally driven goal state of the lower level control systems (For the most detailed explanation of PCT, see Powers, 1973.)

Thus when a student in a class or a subject in an experiment perceives a feedback message, that person's use of the feedback is to be explained in terms of the following: 1. the subject views the message, compares perceptions stimulated by the message to the internal goals that are operative at that moment, and then 2. continues to spend time with the message (probably executing processing programs) as long as the message continues to appear helpful in reducing the magnitude of discrepancy. And the crucial point in this paper is that 3. the references that will be operative as the subject spends time with the feedback message are determined by the output of higher level control systems related to the task at hand.

The basic concern here is that subjects should have various references or goals which drive the lower level programs and the use of the feedback message and thereby affect the performance. That is, subjects could have the following higher level goals: a) getting done with this boring task/perceiving the task as finished; b) getting responses correct/perceiving correct feedback messages; c) understanding why a response is correct/perceiving distinct connections between category labels (correct responses) and the output of programs for perceiving the associates of the categories. These types of concerns are not only theoretically noteworthy but are also important to a trainer who has goals that the students would be correct and/or understand. We propose that a means of testing a subject's higher level control is by observing his/her use of the feedback message. By looking at the variance of feedback frame latencies across levels of metacognitive judgments of response certainty (e.g. Hancock, Stock, Kulhavy, 1992), and between response correctness (right or wrong), we hope to determine whether each subject's higher level control systems are efficiently related to learning to respond correctly. Thus, we are attempting an empirically based and theoretically driven understanding of an enduring practical problem in training and education.

We assume that a subject who rates his/her certainty of responding correctly is thereby providing evidence of the amount of discrepancy that persists in control systems related to responding correctly. For example, a subject who rates 100% certainty of being correct is presumably sensing no discrepancy related to responding correctly, while a subject who rates 25% is experiencing substantial discrepancy. And secondly, we assume that a subject who spends more time with a feedback frame is using the information in the frame to reduce discrepancy, and thereby the performance should be improved as well. Based on these assumptions, we would hypothesize the following: 1. Those subjects who have stronger goals for understanding, will have feedback frame times that increase as certainty decreases. 2. Those subjects that have stronger goals for responding correctly, will have feedback frame times that are longer for incorrects than for corrects. 3. Those subjects who have stronger goals for something other than understanding or responding correctly (such as getting finished), will have feedback frame times which bear little consistent relationship to either correct/incorrect differences or to certainty. In this context, we hypothesize that the objective correct performance will be ordered according to hypothesized higher level references: understanding > correct responding > something else.

Method

The basic experimental design was a completely within subjects factorial design--5 levels of certitude estimate, and 2 levels of response correctness. The predicted value was the feedback time. For the group analyses the subjects were ordered by mean correct response--high, middle, low. The ultimate concern was not just grouped data, but the trends within individuals. Fifty-four university undergraduates participated. The subjects were randomly assigned to the experimental conditions. The stimuli consisted of 27 separate items each presented as a screen of information. Each item consisted of three separate but simultaneous displays of graphic, aural and iconic information. In addition, the graphic, aural and iconic information could be presented at one of three levels, yielding a total of 27 combinations. The right side of the screen also displayed 27 names. Each display item was associated with one and only one of these names. The subject's task was to identify the display item by clicking on the appropriate name with the computer's mouse. The certainty rating screens included a certainty rating scale: How certain are you that your response is correct? 100% certain, 75% certain, 50% certain, 25% certain, 0% certain. There was a radio button to the left of each level of certainty.

Each feedback screen displayed response sensitive feedback information.

The procedure during each of the five approximately 40 minute sessions was as follows:

- 1. view stimulus item and click on a name button;
- 2. view a certainty rating scale and select a rating;
- 3. view the feedback screen, and press a continue button;
- 4. view the next stimulus item, etc.

Results

First, in casting a net (Runkel, 1990), performance (ability/ achievement) groups were determined by dividing the subjects into thirds according to mean number of correct responses. All group and individual analyses of variance were performed on log2 feedback frame times. Inspection of the graphic plots of the data indicated that the trends were as hypothesized.

The partial R squared was calculated for the relation between correctness and feedback time, adjusting for certitude group, .570; middle group, .185; low group, .058. And the R squared was calculated for the relation between certainty following corrects and feedback time group, .305; middle group, .173; and low group, .070.

With the same analyses conducted separately on each subject there were the following trends: high group every subject yielded a significant effect ($\alpha = .01$) for correctness and 10 of the 16 had significant effects for certainty; middle group 11 of 16 significant effects for correctness and 6 of 16 for certainty; low group 3 of 16 significant effects for correctness and 3 of 16 for certainty.

Date: Wed Jun 30, 1993 3:52 pm PST
Subject: Science & PCT

[From Rick Marken (930630.1400)] Bill Powers (930629.1845 MDT)

>I've been preaching for quite a few years, not always coherently,
>that generalization is simply not basic science.

Well, you sure got it coherent this time. What a great post!

Dag Forssell (930630 0910)--

>THE UNNATURAL NATURE OF SCIENCE: Why Science Does Not Make (Common)
>Sense by Lewis Wolpert Harvard University Press \$19.95, 191 pages

Great find. Wolpert gets it right about science and then seems to fall apart when he gets to behavioral science.

>"Science always relates to the outside world, and its success
>depends on how well its theories correspond with reality,"

Correct.

>To the relativists and post-modernists, Wolpert issues a simple and

>elegant challenge: What alternate theories do you have to explain the data?

We know how they answer to that one: change the subject.

>Moral and political problems are outside its purview, as are
>justice, happiness, love and ultimate values. These matters are in
>principle beyond the range of science.

Wrong on that one. Why do scientists think that? Especially ones who have held brains in their hands. Is it just that human behavior is the last place where materialists can clutch at the remnants of a lost spirituality?

>Furthermore, some systems, such as human behavior and society as a
>whole, are so complex that "knowledge in these fields is barely at
>the stage of a primitive science."

Then why are they were treated as though they were simpler than a ball rolling down an inclined plane? Perhaps the people who are dealing with these purportedly complex phenomena just haven't got a clue about what they are looking at -- and none of the tools to understand it.

It seems to me that many people WANT to believe that human behavior is complex. This is another reason PCT has problems; PCT says that, underneath it all, human behavior is NOT complex; it is simply the process of controlling perceptual input. The basic law of behavior (control of input) is simple -- I think this offends many people. For some reason, people can accept the fact that a simple law (inverse square) underlies all the physical complexity they see; but the idea that a simple law underlies all the behavioral complexity just seems like hubris or something.

Gary Cziko (930630.1900 UTC) commenting on an article by van Gelder asks:

>Can the realization that control systems control their input be far behind?

To which I say: What do you mean by "far behind"? A year? 50 years? A millenium? The realization that control systems control their input is probably more than 50 years but not more than a millenium behind. In that range.

Best Rick