

Date: Tue Sep 01, 1992 6:21 am PST
Subject: Re: Agre's Cooking problem

I think Avery almost has it. Most formal planning starts with an attempt to fit a known template over a current set of facts (Pick A Model). Too many times this results in forced ignorance of precisely those things which make the current case unsuitable for analysis by the model chosen. "Best fit" means, very regularly, that certain facts and contingencies must be lopped off the decision tree.

What is attractive about PCT for the planner is that one is forced to start with what is, and not what one must observe in order to find a best model fit.

Date: Tue Sep 01, 1992 7:07 am PST
Subject: Promoting PCT;Science Musings

[From Hank Folson (920901)]

Dag Forssell (920830)-promoting PCT:

>We need to reach people who are dissatisfied with what they can accomplish,
>people with a yearning for something better. A dissatisfied person will be
>open to suggestions and interested in trying a different solution.

And later Dag says:

>in collaboration with Hank Folson,
>whom I am proud to have sponsored onto this net).

I am an anecdotal example of what Dag is talking about: I won't embarrass Dag by listing all the different things he has tried to get me interested in over the years, but there were many. I never acted on any of them because I had no error signals that would be reduced by his proposals. But when Dag first told me about PCT, I was looking for a general theory of business management to help me run my business better. I had an error signal, a dissatisfaction with what I saw in my business experience and in the literature. PCT offered the hope of reducing my error signal, and I have become a convert, although I won't hound people at airports for donations, yet.

Dag Forssell (920830)-About Science:

>Verbal exercises can be carried on indefinitely without any tests ever
>possible.

Dag, if people are controlling for intellectual stimulation, more interested in the hunt than the kill, they will happily do aerobic verbal exercises forever.

A person who is interested in the kill and is controlling to get to the meat of the matter can still appear to be willingly participating in a never-ending verbal exercise. This is the result of not having an effective strategy or techniques for communicating.

The first instance is not a serious problem on the net. The delete key has been suggested before. The second case is harder. A useful thread on the net would be the application of PCT principles to communication. Teaching ourselves more effective PCT based communications skills could really speed up the scientific process.

>I am suggesting that those who have a soft science background and are
>wrestling with PCT may have a greater personal challenge than those with
>a hard science background.

If you come from a world where soft narrative science is the norm, hard science creates an error signal, which you will control to reduce. PCT says you are asking a lot of someone to make that degree of change.

>You need to experience PCT, in hard experiments and/or in your own life.

PCT does not give easy answers. It tells us why and how the world works, but it doesn't tell us directly how to apply this understanding. Don't feel bad, Bill. I know that $E=MC^2$, and I don't know what to do with that either. My personal bias is to put more emphasis into developing techniques for applying PCT.

Hank Folsom, Henry James Bicycles, Inc.
704 Elvira Avenue, Redondo Beach, CA 90277
310-540-1552 (Day & Evening) MCI MAIL: 509-6370
Internet: 5096370@MCIMAIL.COM

Date: Tue Sep 01, 1992 11:49 am PST
Subject: Re: Misunderstandings & miscellany

(ps 920901.1200)

[From Bill Powers (920826.2000)]

Penni Sibun (previous posts) --

>From all the comments by others, I suspect that I misunderstood your last post in most of the possible ways. My apologies. As Avery said, we are ALL arguing against somebody else's point of view without really understanding it. You keep your temper remarkably well, which should be a lesson to all of us.

well, thanks. i'm not so sure about that though....

[From Bill Powers (920827.0800)]

--
Penni Sibun (920826) --

I'm being cautious here to avoid more podiphagia.

You say

>the interactionist argument is in general deemphasizes individuals
and >emphasises their interactions.

Does this mean that interactionists believe that interactions can be
EXPLAINED without reference to the properties of individuals, or that
they can be DESCRIBED without such reference?

i don't think i mean either. we fall into our language problems here though.
maybe we can say that interactionists think that you can't describe or explain
the individual w/o also the interaction.

Does the interactionist
viewpoint differ from, say, that of conventional sociology, in which
it is held that there are laws of social systems that exist apart
from, and indeed unconnected to, laws of individual behavior?

yes, indeed. a primary interactionist field, ethnomethodology, came about
explicitly in reaction to sociology. rather than asking ``how can we describe
there here institutions,' it asks, ``are there institutions, and if so what,
and if so how constituted.''

I have a copy of a '92 paper for AAAI (whatever that is)

american association for artificial intelligence.

by Agre and
Horswill called "Cultural support for improvisation." In it he

``they'!!!

the theory is certainly primarily agre's, but horswill is the one who wrote
toast.

Before this section, the authors discuss the philosophy of their
approach, which depends on using the properties of objects in the
environment to replace a great deal of computation. Here's my
interpretation of what they say; make allowances for the fact that I'm
trying to swim through a sea of unfamiliar terminology. Gratuitous
comments by me are in square brackets.

Objects in the world are characterized by the roles they play in
customary routine activities. [The playing of roles is attributed to
the objects. A lot of the language that follows could be cleaned up by
inserting "are perceived to be" or some such proviso at frequent
intervals].

well, ``playing a role'' is a well-understood way of talking about these
things. the advantage of talking about roles instead of perceptions, is that
you don't have to ask perceived by whom.

A given object type is characterized by a state graph: nodes are
states and arcs are labeled as operations that change the state from
one node to another. [These are apparently thought of as properties of

the objects, but operations are defined in terms of their outcomes, not in terms of the actions that bring them about. In HPCT terms, the operations could be re-phrased as the setting of goals specifying a new state which lower systems then bring about, thus moving the object state to the new node].

An operation can be a label that applies to several arcs leaving a given node, thus being "nondeterministic." [This baffles me. Do they mean "stochastic?" Do they mean that under various circumstances, a given act can have different outcomes?

they defined it: ``in the sense that it labels several arcs leaving from the same state.'' the terminology is straight from computer science theory: a machine is deterministic if there is only one way to get from one state to the next; it's nondeterministic otherwise.äthe point is that there's more than one way to get from one state to another.

A world state is a set of instances and a mapping from each of these instances to a state in its type's graph. A history follows a set of instances through a series of changes. [This would seem to amount to a model of the environment's behavior]

no, a history is what happens between one point in time and another point in time.

An action type is a sequence of pairs made of an object type and an operation from that type's state graph.[So an action type is a set sequence of operations that produces a set sequence of states.

it's not a set sequence of states if the operations are nondeterministic. it's a set of possible sequences.

A goal for an object is satisfied if the world state contains some instance of the object type that is in the indicated state.(So one cannot assert that two instances of a given type must be in a given state). [I don't see the advantage of this rule. It says, apparently, that the goal of breaking a given egg is satisfied if the egg being operated on remains intact, but another egg breaks.

no, it says that the goal of breaking an egg is satisfied if an egg (any old egg) is broken. it doesn't have to be egg-47 or egg-13.

So the only egg that can break is the one being operated on by the break operation and there's no need to keep track of which egg is broken.]

that's true. why do you find this contradicts the above?

Summing up, the authors give this description of the algorithm:

While there are unsatisfied goals,
Choose arbitrarily an unsatisfied goal.
Let G be the specified goal state.
Find an instance M of the specified material type.

Determine its current state S.

Let A be the entry for the S and G in M's action table.
Choose tool instances to fill each of A's participant slots,
preferring instances which are already in their normal state.

If some chosen tool T is not in its normal state, then perform the
action specified in T's action table for its current state. Otherwise
perform A using M and the chosen tools.

-

This last summary is ever so much simpler than all the verbiage that
led up to it, because the attempt to present a general case has been
reduced to a specific case.

in what way is it specific?

One change I would suggest is the explicit addition of an error graph
for each operation. The main weakness of this algorithm as it stands
is that every operation must produce a predictable state. By using an
error graph this limitation could be removed.

yes, i agree. it would at least be interesting.

it would certainly be interesting if you (or someone wanted to try to
incorporate your suggestion into toast. it would be instructive for everyone,
i'm sure. i think you can get the code from horswill, though i would guess
it's all in lisp.

cheers. --penni

Date: Tue Sep 01, 1992 1:47 pm PST
Subject: Cognition and Control

[From Rick Marken (920901.1430)]

I have a new question, particularly for those who support some version of
"cognitive psychology". It is related to the recent discussion of behavior
manipulation. The question is:

What does cognitive psychology have to say about the control of behavior?

As far as I can tell, behavior control does not seem to be much of an issue
for cognitive psychologists. Why? There are some cognitivists who have spoken
to this issue. Chomsky, for example, seems to think it's possible but not nice
to do. Neisser seems to think it is impossible because people are too smart. I
don't know if AI types deal with the issue at all, though they seem to believe
in "reinforcement" (see Minsky's "Society of Mind"). So what is the cognitive
(or, in general, the "non-behavioristic") position on control of behavior? Can
the behavior of an organism (viewed as a "cognitive" system) be controlled?
Why or why not? Please, when answering this question, assume that "control"
means "to bring behavior to a state selected by the controller and to maintain
it in that state against disturbance".

I would really like an answer to this. I think this discussion could help us articulate the difference between cognitive and PCT views of behavior. I, of course, would also love to see how an interactionist deals with this topic (control of behavior). What do interactionists say about shaping and operant control?

Thanks Rick

Date: Tue Sep 01, 1992 3:28 pm PST
Subject: Promo, Martin, Science

[From Dag Forssell (920901-1)]

Hank Folson (920830) on promotion:

>I won't embarrass Dag by listing all the different things he has tried
>to get me interested in over the years, but there were many.

Would it embarrass me? Was I trying to interest you (focus on you), or just bubbling along about what interested me at the time (focus on me)? I got no error signal from your refusal to take my expert advice, did I?

Most of us are much more concerned with ourselves than with others. Control of self comes before control of others?!

Martin Taylor's Paris paper:
PRINCIPLES FOR INTELLIGENT HUMAN-COMPUTER INTERACTION.
(Version control - My copy printed July 30, 1992)

I have read your 50 page paper, with emphasis on the PCT parts. Very interesting to learn some of your personal background and to see many of the issues you have brought to the net in your context.

A few, nitpicky, comments:

In fig 2.: "The ECS has as its purpose to match the Perceptual Signal to its Reference Signal which comes from a higher level ECS."

As I see it, the Reference Signal IS the purpose. The matching happens automatically because the ECS is a control system. It has no purpose of its own.

Since you are jumping in with both feet into the (PCT concept) ECS anyway, you might consider inserting (distributed, local) memory (BCP p. 218, 221) as a source for the reference signal. Then say that this is implied in subsequent simplistic representations of ECS's.

Rick's paper: "The behavior of perception" deals with this and shows clearly (as I recall) that it is memories of perceptions that "behave." In our daily lives, most of our behavior is simply the repetitious recreation of past perceptions. A capable person has many past perceptions to choose from. These memories become purposes when they are specified by higher level ECS's. I

witnessed a discussion between Ed Ford and Bill P. a year ago about why it is that Quality Time works. In a broken marriage, the spouses have no fresh recollections of pleasant perceptions together. So they cannot recreate pleasant experiences in the present either. They do not know what to do. By asking them what they individually enjoy, then prescribing the common denominator as a joint activity - quality time - Ed coaches the couple to create recent perceptions of each other with a pleasurable content. They can then keep on recreating these perceptions, restoring affection, love and marriage.

The reference signal from memory can of course be quite complex in a human, made up of temporal and spacial patterns and several sensory components. Equivalents for the computer Claude that I imagine are "subroutines," "algorithms," "matrixes" and many other things computer scientists talk about and I don't know much about.

Fig 3.: Putting words to the levels of perception so the concept is intuitively grasped without explanation is hard. This looks like a good start.

Kent McClelland provides a good narrative in his paper (several levels of cruise control), which I plan to graph out in the next several weeks.

P. 9: "But the central point is that the only aspects of the world that can be perceived are those that can potentially come under control."

I vaguely recall not paying attention as you discussed this on the net. Offhand, it does not make any sense. Look at the stars! Perhaps a reformulation can convey what you are trying to say. The following two sentences make sense, but do not add up to the one above.

Fig 4 & 5: I can see that the mirror diagram makes intuitive sense. You have left the control diagram out of the boxes in the environment, which is a nice, subtle touch. (Should they be put back in Fig 9).

Fig 9: This shows me what you are driving at with this paper, right? With the other intelligent entity being a computer, you demonstrate the utility of matching the computer's layers perceptually with those of the human it is interfaced with. The human is existing, the computer must conform. The implication that computer designers should look at the design of humans is compelling.

P 15: "There is very little to choose... I don't follow the english or meaning of this sentence.

Fig 14: The imagination loop would make much more sense with the recognition of memory references fed into perception, as shown in BCP. You can show the address signal and memory box in the loop shown.

P 32: I shall remember the principle of reorganization shown here, with autumn leaves accumulating where they are left alone; moving at random where they are not sheltered. Thanks!

It is nice to have a fellow student of PCT addressing an application and audience where this ingenious insight can be welcomed and put to good use. Keep up your hard work.

Hope this is helpful.

Bill Powers (920830.1545)

Thanks for your nice comment. Glad you liked articles. Scanning it is. A flatbed scanner and Omnipage character recognition software do a good job fast. Spell check does the rest. I am not patient to type this quantity of stuff without very good reason.

Re: Soft and Hard science.

It occurs to me that a more useful terminology for the distinction between soft and hard might be Descriptive and Explanatory.

With "Descriptive," there is no expectation of prediction.

With "Explanatory," testing of predictions is implied.

Best to all, Dag

Date: Tue Sep 01, 1992 3:34 pm PST
Subject: Re: Agre's Cooking problem

(ps 920901.1600)

[From Jeff Hunter (920831)]

 This was (if you pardon the pun) a cooked example. None
of the subtasks interfered with any other in any major way.
 To show why, here is an example of weekend goals:

- wear sunhat
- get new car tires
- spend weekend at cottage

 And these are fed into Agre's weekend planner (main mechanism
recapped below):

i'm quite certain that neither agre nor horswill would dream of suggesting using toast for planning a weekend. (and *i* certainly wouldn't use a planner for my weekend!) is it unreasonable to assume that different sorts of activities can use different sorts of planning, in both the technical and the colloquial sense?

 Any decent planner can find the best ordering for these
weekend goals in nearly the same time that Agre's one would
find a lousy one.

substantiation?

 The "planning weenies" have considered a lot of cases of

conflicting tasks, and their work should not be discarded lightly.

yes, the planning weenies have spent a lot of time worrying about subgoal conflicts (sometimes called ``goal clobbers brother goal''). one of agre's points is that usually you just don't have to worry about it. and making breakfast is (claimed to be) a good example of just that.

I think I see why the planner died. At least part of it is that Toast is implicitly told that any pat of butter is the same as any other pat, whereas the SNLP planner was not.

Why are there 15 pats of butter? Why so that the planner must consider more than 23 factorial orders of placing butter and eggs in the frypan.

have you ever seen a kitchen w/ one pat of butter in it? in fact i bet you've never seen a kitchen w/ **any** pats of butter! i assume there are fifteen pats of butter becasue that's roughly what you'd get out of a stick and they didn't want to get into individuating the stick into pats so they assumed it came that way. i think both toast **and** the snlp would have a tough time simulating slicing off a piece of butter!

De apibus semper dubitandum est. Winni Ille Pu

;-} --Penni

Date: Tue Sep 01, 1992 4:52 pm PST
Subject: Planning and Agre's Cooking Problem

[from Avery Andrews (920902.0914)]
(Thomas Baines 92090909)

> I think Avery almost has it.

Spurred on by this, I'll dump the rest of this file. Before Penni introduced me to Interactive AI, but after I started getting CSGNet, it struck me that there was something very wierd about the discussions of the `frame problem' in my AI books, because they seemed to be concerned with the problems faced by programs that attempted to imagine everything that might happen in some domain that they had no practical experience in.

By contrast, real planning, as carried out with some modicum of success by real people (we have 4 people, 1 car, 7 hard commitments & 3 wishes -- how are we going to get thru Saturday?) is carried out in domains where people have a lot of practical experience. Consider a typical operator (= step in a plan). In addition to its desired effect, it will have various side-effects, which can be grouped into three types:

- a) irrelevant (when I execute DRIVE(OWEN,BELCONNEN), the place on the driveway where the car is normally parked will get wet if it rains).
- b) relevant, but routinely controllable (as a result of the driving, the fuel-level of the car will decrease). One of the effects of culture is to increase the routine controllability of side effects.

c) relevant, but uncontrollable. (while I am doing the driving, there won't be any car at home, so no other person can do any driving).

So the real-world planner can find out from hearsay and experience what the type (c) consequences of his/her actions are, and construct 'clean' sets of operators that have few such consequences. The general intractability of planning means that real plans have to be quite short, so that experienced and competent people would have large libraries of operators, which might be individual quite complicated (like the tricks that people make up to do the last layer of the Rubik's cube).

In terms of everyday live, it strikes me that errors are very often made when a side-effect of an operator switches from type a) to type c). E.g., when a family with two drivers is reduced from having two cars to only one, they are constantly making plans which presuppose that another car will be available when one has been driven off somewhere. It took My wife and I a tremendously long time to stop making up these kinds of errors after we had the use of a friend's car for six months (tho we did usually manage to spot our mistake before anyone actually drive off to somewhere).

And, as Penni has been saying to me recently, a lot of the info about useful operations doesn't have to be figured out at all - you can pick it up by watching and listening, because it's just floating around in the culture.

Avery.Andrews@anu.edu.au

Date: Tue Sep 01, 1992 4:52 pm PST
Subject: Kent's paper

[From Dag Forssell (920901-2)] For Kent McClelland:

Read your paper on the plane two weeks ago, Draft 7/21/92, which I picked up in Durango.

Very nice. Really a compelling presentation. One of the best, comprehensive introductions to PCT we have.

A few nitpicks, of course.

2.2.4.4. : I have never known wether my foot is at 60 degrees or any other angle. I doubt that your readers will identify with this and say: "of course." I would suggest that the perceptual system controls pressure on the pedal, plain and simple, which translates nicely into tendon tension. All I ever do when I drive is squeeze the accelerator, and let up, respectively. I do not care what the angle is, and would guess poorly if asked. The angle becomes a consequence of the properties of the environment (car design, uphill, wind resistance etc.).

2.3.2.5. : You bring in emotions and I don't understand "invites negative emotions." I do not know what insight of Rick's you refer to. If you want to deal with emotions, I would think a summary of Ed Ford's presentation in Freedom from Stress would be more appropriate and easier to relate to. This could be its own section. LCS II has the missing (from

BCP) chapter on feelings.

Between Rick and Bill and all, there is much talk about loop gain. But you are not using equations with constants all over the place. Considering your audience, you could choose to not mention it, deemphasize it or use / substitute more illustrations.

Personally, I find it intuitively easier to talk about how a control system directs energy and or resources to influence the variable.

The nerve signal gives instructions to the muscle fiber, which contracts, using the energy / resource of sugar to create an effect far stronger than the energy content of the signal itself. (This is amplification). What size biceps?

Your cruise control presses the gas pedal, which provides more fuel to the engine. What size engine? (This is strength).

When you talk about conflict, you can talk about two control systems of equal or different resources, strength, evoking the image of armwrestling or such and make the same points about fatigue, reorganization, lack of satisfaction.

3.2.8. : There [are] some ..

5.2. : ... To put it bluntly[: I]f...

This is a marvelous paper. What are your plans for it? Surely it deserves publication in a highly visible place.

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

PS Kent, please acknowledge. Believe you are on net. If not, I will mail.

Date: Tue Sep 01, 1992 5:04 pm PST
Subject: Delayed; misc subjects

[From Bill Powers (920901.1000)] Dag Forssell (920831) --

What happened here? Otherwise looks good.

If the associate wants the shareholder..... Associate satisfaction is all there is. Associates is all you have to work with. Associates must want and perceive many things and take effective action for the company to function well.

Martin Taylor (920831) --

>... when we are trying to be precise, and communicate exactly, then we
>have to go back to separating the controlled percept from the CEV,

>which is what I was trying to do. We do, easily but sloppily, use
>third-stage language most of the time.

But doesn't this say that there "is" a CEV with its own definition Out There in the environment, and that the percept simply represents it?

I'm trying to make the percept the central concept, because that's what the control system knows and what it controls. The CEV is found (by somebody who knows more than we do) by working backward through the perceptual function and finding what combinations of external variables would be sufficient to produce that percept via that input function. In general there are many combinations of external variables that will lead to the same percept -- that's the other side of the concept "invariant."

We have to distinguish between a CEV as seen by the controlling system and as seen by a different person, a different perceptual hierarchy. This will not in general be the same CEV. The only completely correct definition of a CEV is $f(p)$, where f is the input function in the controlling ECS. Because the inverse of a many-to-one function is always ambiguous or multiple-valued, all we can say about the CEV as deduced from the inverse is that it consists of ANY ONE of the combinations that will yield that value of that percept. We have to consider many different CEVs in order to pin down the environment to a less ambiguous definition of each CEV. So a single CEV really doesn't exist as a single independent unambiguous entity when considered alone.

Through joint efforts like physics, we can come to agree on definitions for certain CEVs like mass, energy, shape, color, and so on. But for each of us, these definitions remain private; we can't go further than being satisfied that other people seem to use the agreed-on terms in situations that we find consistent with our own perceptual meanings of the terms. What a given person controls when controlling for "energy" is precisely the perception to which this word points. But that perception is known only to that person.

>As you put it here, ALL CEVs are and always will remain misdefined, >inasmuch as there probably exists no construct whatever in the outer >world that behaves precisely as the perceptual input function that >defines the CEV does.

Yes, that's another way to say it. But you can also say that for each person, the CEVs are precisely defined, but coincide only in part with the the CEVs that other people also precisely define. The ideal of defining a CEV so that all people will agree (operationally) is not reachable. And of course the idea that there are objective CEVs independent of our ways of perceiving them is a question we would all like to have answered. But there seems to be, at present, no hope of answering it.

Your concept of X-control seems to be just a way of describing the control process that one person attributes to another. The Test reveals X-control. We hope it bears some resemblance to P-control, and keep looking for ways to make that hope more realistic.

RE: transients.

>You can't have it both ways. Either there is an error signal that
>results from the disturbance and generates an output signal that

>resists the disturbance, or there isn't.

We're having a quantitative misunderstanding here, resulting from using words in a qualitative way. There are degrees of error and degrees of effect of a disturbance on a variable. If the variable is under control, a disturbance will affect it a lot less than if the same variable were not under control. I'm drawing the line between "effect" and "no effect" between these two cases.

In a control system with a loop gain factor of -30, a disturbance that would have 100 units of effect on the variable when uncontrolled will have only 3 units of effect when control is present. So yes, the disturbance does disturb the variable, but in comparison with the effect that would occur with no control the amount of effect is essentially nothing. It is, in fact, so small that it can be neglected. If that were not true, the control system would have continued to reorganize until it had a gain of 300, or 3000, or whatever was required to make the deviations unimportant to it.

And if the transient response delay of the control system is 1 second, then disturbances that last less than that amount of time will be unimportant to the control system. If that were not true, the control system would have evolved to respond in 1/2 second, or 1/4 or 1/8 second, whatever is required (within the limits of the possible).

When I spoke of a transient effect on the controlled variable, I meant an essentially uncontrolled effect -- in other words, an effect that is just as large as it would have been with no control. In a control system whose output rise time is 0.2 seconds, a disturbance that lasts only 0.05 seconds will alter the variable just as if no control existed. But the control system will not respond to that change, so that brief a disturbance can't be used to control the action of the system. You must maintain the disturbance long enough for the output to come to equilibrium with it if you want to control the output.

You understand me correctly in saying "It means that the organism will survive, behaving this well, and that no excess resources will be applied to ensure better behaviour than that."

We also agree, I think, about manipulation:

>No difference when we "control" another person. We keep trying different
>"influences" until the person does what we want or until we give up.

Right. Manipulation, if goal-directed, always comes down to control. And when you try to control another control system IN ANY RESPECT THAT IS ALREADY UNDER CONTROL BY THAT SYSTEM, you generate conflict, a contest of brute strength. Such contests can be fun when done for the purpose of developing skills. And when you know the game will be over pretty soon. As a way of life, they stink.

>No, control of another has nothing whatever to do with setting
>reference levels for another. Control of another may be achieved
>easily if it so happens that the other's reference levels come to lead
>to actions that are perceived as appropriate to the reference levels in
>the "controller." But the other's reference levels are a red herring
>in the discussion of whether one person can control another.

Right on. Either you control other people in the trivial sense that you elicit error-correcting actions from them that happen to suit your own purposes, without serious inconvenience to the other, or you wrestle with them over who gets to control something that can't be in two states at once. So far our civilizations have been founded on the latter mode of interaction, which accounts for their violent history.

I do want to cavil about control existing when one person happens to control something the way somebody else wants it controlled. I said this to Greg Williams but it's worth repeating. Control requires a closed loop. If one person is controlling a variable, it's highly unlikely that another can control it at the same time without generating conflict. So the spectator may be satisfied with how the other person is doing the controlling (as the spectator sees the CEV), but the spectator can't DO anything about it if an error appears. The spectator is not controlling, because controlling requires action to close the loop. The passenger can jam down on an imaginary brake, but had better not jam down on the real one if conflict with the driver is to be avoided.

Mark Olson (920831) --

Tell your friend to hang on; I'm working on an extensive writeup for Little Man Version 2. I'll announce it when it's done. I'm going to sell this program for \$100, but have decided to give an 80% discount to CSG members (Dag note). And of course it's basically shareware because people will rip it off anyway. I don't really mind if people who can't afford it or can't finagle department funds just copy it. But I am trying to pay for my new computer, and want to build some experimental apparatus, and get Frans Plooijs to the next meeting, and all that sort of stuff, so the money is really needed.

Martin Taylor (920831.1720) --

>"All behaviour is the control of perception." But is it?

Rick said it: all control behavior is control of perception. The problem here is again language. "Behavior" can mean anything we see happening anywhere, in organisms or otherwise. The behavior of a rolling marble. The behavior of a thunderstorm. It can also mean irrelevant side-effects, especially if you don't know anything about control theory.

The problem is that historically there's been no understanding that the behavior of an organism while controlling something is different from any other sort of activity we might notice. In one sense, behavior has been used to refer to events almost as if they were intentional when in fact they aren't. To speak of the behavior of a marble rolling around in a bowl carries a subtle idea that the marble has something to say about how it's going to move -- we attribute the behavior to the marble almost as if the marble is an agent. In fact the path of the marble is due entirely to external forces acting on its spherical shape and mass which are the same no matter what the marble is "doing."

>Exploratory behaviour is "what if" behaviour. One is not controlling

>for any particular reference (at a higher level one is controlling for
>a percept of knowing more than one does, but that's outside this
>point). One is just acting, more or less randomly at the level at
>which exploration is being done. Of course it isn't random at lower
>levels, and probably not at higher ones.

Exploratory_behavior_ -- meaning observable actions -- is just like any other kind of behavior. If you see someone sharpening a pencil, you can't tell whether he's sharpening it because he doesn't know what will happen and wants to find out, or because he does know what will happen.

What's different about exploratory behavior is the goal, not the behavior. At some level, as you say, the action must differ from ordinary action, in that it's more or less random (although when you do a search for a lost child, you organize to do systematic sweeps of the territory, so all that's exploratory about this is not knowing whether the error will be corrected or not).

The goal-directedness of exploratory behavior becomes obvious when you ask what will make it stop.

>In the Layered Protocol formulation, before I heard of PCT, I argued that
>the purpose of casual conversation was for the partners to develop models
>of each other that would permit them to communicate effectively when the
>need arose. In the same vein, I suggest that the purpose of casual
>exploration is to develop organizations that permit effective control
>when the occasion arises. Exploratory behaviour, then, is NOT the
>control of perception, but the discovery of how to control perception.

Tsk, tsk, Martin. Are you saying that the goal is to learn how to control, but without any perception of whether you've learned to control? You forgot to go up one more level.

Jeff Hunter (920831) --

Thanks for the further info on the Toast approach.

Planning certainly happens. If people do it, it belongs in a model of people. The only question I have is how elaborate it really is, and at what level it is really applied. From what I've seen of planning programs, they include a lot of things from lower levels of organization that could be done a lot more simply just by control systems. By confining the planning aspect of the model just to the levels where planning alone is the problem, you can make the planning tasks much easier and also do a better job of explaining behaviors that really don't have any planning aspect to them. In a hierarchical model you can take care of lots of things, like controlling spatial relationships, without getting into classifications, sequencing, and logic. That means that the planning system can quit worrying about the details, and simply solve conflict problems, perhaps without even having to know what the conflicting variables mean. One variable can't be in two states at once; that's all the planning system has to know.

Ray Jackson (920831) --

Welcome aboard, Ray! Be patient; the threads will intertwine and sooner or later the one you want will come along. Feel free to participate in anything, and don't hesitate to raise questions of interest to you.

Avery Andrews (920901) --

I'll wait until you get the post I just sent. Try again. I'm a little swamped.

Best to all, Bill P.

Date: Tue Sep 01, 1992 7:35 pm PST
Subject: Cooking; chapman; visual perception

[From Bill Powers (920901.1900)]

Penni Sibun (920901.1200) --

> ... maybe we can say that interactionists think that you can't
>describe or explain the individual w/o also the interaction.

Depends on what you want to describe about the individual. I come down more on the side toward the individual, because people live in a lot of different environments but they don't change their characteristics for every one. Of course a person can change his/her actions according to the circumstances, but it's the person who has the goals, like eating, not the environment. At least the nonsentient part of the environment.

>a primary interactionist field, ethnomethodology, came
>about explicitly in reaction to sociology. rather than asking ``how
>can we describe the[s]e here institutions,' it asks, ``are there
>institutions, and if so what, and if so how constituted.''

Sounds pretty much the same to me. In either case, the institutions are reified. Not that I don't believe in them. I just think they exist in people's heads, not in the environment.

> ``playing a role'' is a well-understood way of talking about these
>things. the advantage of talking about roles instead of perceptions,
>is that you don't have to ask perceived by whom.

Well-used, I'd say. Understanding how an object can play a role is something else. I think it makes a difference who is perceiving the object. If you objectify objects, you can't explain how the same object can have different roles depending on who's using it for what purpose.

>a machine is deterministic if there is only one way
>to get from one state to the next; it's nondeterministic otherwise.

That's what I was asking about. If it's nondeterministic in that sense, then what selects which of the possible paths will be taken? If there's some systematic selection method, then it's deterministic again. It's nondeterministic "all the way down" only if the choice is random. So the question is, when there are alternate paths, what determines the path actually taken?

>no, a history is what happens between one point in time and another
>point in time.

Do you mean that there are continuous processes taking place BETWEEN nodes? I had thought that an operation simply jumped you from one node to another.

>no, it says that the goal of breaking an egg is satisfied if an egg
>(any old egg) is broken. it doesn't have to be egg-47 or egg-13.

>> So the only egg that can
>> break is the one being operated on by the break operation and
>>there's no need to keep track of which egg is broken.]

that's true. why do you find this contradicts the above?

Because the way the proposition is stated, it doesn't matter whether the agent breaks the egg into the pan or whether someone else blunders through the room and breaks it on the floor, or if the agent finds it already broken in the carton. If it's broken the goal is satisfied, the way I read it.

>> This last summary is ever so much simpler than all the verbiage
>>that led up to it, because the attempt to present a general case has
>>been reduced to a specific case.

>in what way is it specific?

Now we simply have goal-setting, perception of the current state, comparison with the goal, and action that reduces the error. That I understand. What I don't get are all these nodes and arcs and that stuff, which sounds like the innards of a computer program, not something happening with someone doing it. chacun a son gout.

Avery Andrews (920902.0914)--

>By contrast, real planning, as carried out with some modicum of success
>by real people (we have 4 people, 1 car, 7 hard commitments & 3 wishes
>-- how are we going to get thru Saturday?) is carried out in domains
>where people have a lot of practical experience.

Let's see if I can tease out my idea on this subject, which I've tried to say before without much clarity.

What's essential about planning isn't that people, cars, commitments, and so on are involved. It's that you can't have two different sequences involving the same variable going on at the same time. Or that you can't have an object in two places at the same time. Or that A can't be both below B and above it at the same time. What kind of problem it is is the "practical experience" side of it. What practical experience shows is that when you try to do things like this simultaneously, the world won't cooperate and you end up in an internal conflict. There are some things you can't set as reference signals simultaneously.

So a true planning level has to be concerned with the conflict itself, not with what it's about. The conflict has to be discovered. It can be discovered

in ongoing experience, or it can be discovered when you try to run two parts of your imagined world at the same time. If your internal model is pretty good, you can discover conflicts without having them actually happen -- you can't even imagine a cup on a saucer and at the same time under it, although you could imagine it on the saucer and also upside down. You can imagine driving the car away, but you can't imagine your wife coming out of the house ten minute later and also driving it -- the the car you're driving -- away. That's the practical experience part.

But the planning part has to be something else. It has to be about how you resolve conflicts between variables of different types. One way is resequencing (you wait to go on your trip until your wife gets back). One way is changing spatial or temporal relationships (you drive your wife to her appointment and go on to yours and pick her up on your way back). At the planning level, all you work with are logical variables. At that level they don't mean specific experiences any more. If you can't do x and y at the same time, then you sequence them. If doing x rules out doing y, then you make a choice between x and y. And so on.

What I'm trying to do is get the details of examples sorted out from the specific functions that have to be performed. Some functions are of lower level than planning. What the planning level does is organize the way the lower-level functions are carried out.

Well, I can see I have only half an idea here. But it's worth throwing on the table. With the eggs.

RE: Chapman's book.

I now have a copy of Chapman's book on Sonja. Heavy going; he alludes to more than he explains. The part on visual perception isn't anywhere near as specific as I was lead to believe, ahem. He does have some sort of data on phenomena of visual perception, and that's good and useful. But the methods of achieving them that are proposed look suspiciously like computer-program methods that just happen to pop into mind -- bit blitting (filling out an area until a boundary is reached), search trees, and so on. I'd say that Chapman (or Uhlmann) is no closer to understanding form recognition than I am.

But something useful has already come out of reading this book. I looked at the way Chapman was handling the problems and thought, "No, that can't be how it works." And then I looked at my way and I thought "No, that can't be how it works." It may have been something Chapman said, or just the general atmosphere of his approach, but I suddenly realized that treating perceptual signals simply as "how much of a given perception is present" is just no good for the higher levels. Probably not for any level above sensations.

Maybe it was this. Chapman seems, although not clearly, to adopt the principle of "coded" perception, in which a system that receives a perceptual signal can tell what it's about from the information in the signal. In my model, of course, all perceptual signals are alike, so they can't report both the state of the perception and its identity. Mine is the pandemonium model, where it's up to higher-level systems to make sense out of the behavior of lower-level signals without knowing what they mean.

It was suddenly obvious to me that I was missing a bet with the pandemonium model. The meaning of a given perceptual signal's magnitude doesn't have to be just "how much." It can be "where on the scale of variation that is being sensed." I don't mean that this information is coded into the signal; I'm just talking about the significance of the signal magnitude that is created by the form of the perceptual function -- what it mean in terms of lower-level signals.

For example, there can be a signal that represents an arm angle in the horizontal plane. When the arm is swung all the way left, the signal is maximum and when it is centered, the signal is zero. Now you could say that this signal represents the amount of leftness, but we can also say that it represents the state of something as a position between the extremes of the range. In this case there's not much difference.

But then I thought of one of the control examples in Demo 1 (or in some demo or other), where you can alter the proportions of a diamond. By working the control stick you can make the diamond go smoothly from squat and wide to narrow and tall. You can hold it in any given shape against disturbances, or match it to another instance of a diamond. This is a change of shape, but it's taking place in some systematic space. So a single signal's magnitude, suitably derived from lower signals, could represent where on this continuum of shapes the diamond is right now. There isn't any "how muchness" to it. It's simply that the magnitude of the signal maps onto the continuum of shapes in a unique way. If you set the reference signal to a specific magnitude, you can work the stick until the perceptual signal has that magnitude, and then the diamond will be in a specific shape.

When you think of just one such continuum you aren't much further ahead. But now think of lots of different continua, different ways of mapping continuous changes in the visual field onto a single magnitude. For instance, one continuum might be the relative lengths of pairs of the sides of the figure (I won't say diamond now). Another might be the angular position of a line from the centroid to each vertex. Another might be the total size of the figure (the vertices moving radially in and out). If you have enough of these continua, sufficiently different and sufficiently independent, the "diamond" becomes a particular set of values of the perceptual signals indicating positions along each of these continua. A higher-level system could then come to recognize this particular combination as the meaning of "diamond." Perhaps even a weighted sum.

Maybe everybody else was already thinking about it this way, but I wasn't. Now I begin to get a little optimistic about actually modeling configuration perception. This is going to have to simmer for a while.

Best to all, Bill P.

Date: Tue Sep 01, 1992 8:21 pm PST
From: Martin Taylor
EMS: INTERNET / MCI ID: 376-5414
MBX: mmt@ben.dciem.dnd.ca

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Promo, Martin, Science

Thanks for the comments. They may mystify the wider audience, though. I appreciate your comment about memory. I do have a slide showing memory as a source of reference, but didn't think about the implications of that very much. Putting memory together with imagination is good. It's a different direction from where I have been concentrating. I'll add some more on those lines to the talk, and to the paper, if I ever make one out of the abstract.

Martin

Date: Tue Sep 01, 1992 9:02 pm PST
 Subject: Re: Delayed; misc subjects

[Martin Taylor 920902 0030] (Bill Powers 920901.1000)

>But doesn't this say that there "is" a CEV with its own definition Out There in
 > the environment, and that the percept simply represents it?
 >
 >I'm trying to make the percept the central concept, because that's what the
 > control system knows and what it controls. The CEV is found (by somebody who
 > knows more than we do) by working backward through the perceptual function
 and
 > finding what combinations of external variables would be sufficient to
 produce
 > that percept via that input function. In general there are many combinations
 of
 > external variables that will lead to the same percept -- that's the other
 side
 > of the concept "invariant."

More precision of language is clearly required. I'm almost sure we have the same ideas, but not quite sure. There is a problem distinguishing the "something" that might be "out there" and the percept of that something that is the actual object of control. I introduced the concept of the CEV some time ago with a vague notion that it was useful to separate the two concepts, and that the C*E*V (emphasizing the E) might be different from the percept to which it gives rise. (This is really difficult to put into precise words, partly because the idea is not precisely formed). Take an example--optical illusion. Suppose you are cutting the birthday cake so that "one cut, the other choose" is fair. But knowing the nature of illusion, you cut so that the bigger-looking piece is actually smaller. The other satisfies (apparently) the reference for maximizing cake, but actually chooses the smaller piece. Perceptually it was bigger, and only knowledge of a different kind allowed you, the cutter, to be satisfied that you actually got the bigger piece. I don't know whether that's a good example. Maybe it gets the idea across.

The base problem is that no ECS can know anything about what is "really" out there, so from one point of view, the idea of an external CEV is nonsense. But there are different ways of determining percepts that "should" be the same. Weighing the cake pieces "should" give the same result as looking at them, so far as telling which is bigger. But the two percepts that "should" be the

same are not. We (external to either) infer that there is something there that is being misperceived in one way or another, even though from the point of view of either ECS, there is no problem. You are quite right in calling it a third-party issue.

>I do want to cavil about control existing when one person happens to control
> something the way somebody else wants it controlled. I said this to Greg
> Williams but it's worth repeating. Control requires a closed loop. If one
> person is controlling a variable, it's highly unlikely that another can
control
> it at the same time without generating conflict.

What about four people carrying a bed by the four corners? At one level, they are all controlling the location of the bed in coordination with each other (Tom Bourbon demonstration, if I remember correctly). At another, they are controlling it to satisfy a reference that an employer should be happy, and the employer has provided input that each of the four adopts as a reference for the location of the bed.

> Exploratory behaviour, then, is NOT the control of perception, but the
> discovery of how to control perception.

>Tsk, tsk, Martin. Are you saying that the goal is to learn how to control,
but
> without any perception of whether you've learned to control? You forgot to
go
> up one more level.

I didn't forget. Maybe I wrote too much, and it got lost in the shuffle, but I explicitly limited the non-control aspect of exploration to one level, and more properly one ECS. An ECS above may well be in control, using, say, the gain of that ECS as something to be controlled (as you mentioned a month or so ago as a possibility). And certainly ECSs below would be in control during the exploration. But the point of the exploration is that the explorer cannot guess what percept will occur, and should not attempt to develop any particular percept. That would be "looking for", not "seeing what," which is the kind of exploration I'm considering. The explorer is, I still think, finding out what is likely to happen to its percept under different contextual circumstances when it generates an output.

>Exploratory behavior -- meaning observable actions -- is just like any
other
> kind of behavior. If you see someone sharpening a pencil, you can't tell
> whether he's sharpening it because he doesn't know what will happen and
wants
> to find out, or because he does know what will happen.

Yes, that's the permanent "Third Party Problem." You can't know, ever, though you may often make guesses that YOU think are pretty good. He may be sharpening the pencil just to get you wondering about that very question, with no concern either for getting the pencil sharp or for finding out what the sharpener does.

=====

It's still fun, but Paris beckons and I've got to get more sleep. Bye.

Martin

Date: Wed Sep 02, 1992 6:23 am PST
Subject: Dag's comments

[From Kent McClelland (920902)] Dag Forssell (920901-2)

Thanks Dag for your comments on my revised manuscript. Your points are well taken. The paper is into the second round of reviews at the American Journal of Sociology, and I have hopes that they'll take it without too much further revision. When my work calms down here slightly, I'll make this draft available on the net. Gary, what's the best way to handle that now? The paper, as you know, is not any shorter than the last time (unfortunately).

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: Wed Sep 02, 1992 7:40 am PST
Subject: Re: Exploration and Reorganization

[Martin Taylor 920902 11:15] (Rick Marken 920831.1930)

>Martin Taylor (920831 17:20) says:

>

>>"All behavior is the control of perception." But is it?

>

>Being the official sloganeer on the net, I would say:

>

>No. "All PURPOSEFUL behavior is the control of perception".

>

>This slogan captures the fact that there are many aspects of
>behavior that are neither the actions that are used to control
>nor the outcomes that result from control of perception.

It also comes perilously close to tautology, whether you are a PCT proponent or not. To have purposeful behaviour is to have a purpose or goal, with behaviour calculated to lead toward achievement of that goal. And how can that achievement or progress toward it be assessed other than by perception (in the broad sense in which you intend the term)? Even in PCT terms, the word "control" cannot be limited to smooth or monotonic results of monotonically directed actions--things break abruptly after smoothly increased forces are applied, for example. The breaking may be the achievement of the goal, or a step toward it, or an unwanted step away from it. So it is not clear to me what that slogan achieves, except to make people think about the importance of the feedback loop.

>Exploration and reorganization can "look" the same; but I think
>you could tell the difference using the test.

They aren't even in the same domain, so I don't see how they could "look" the same. Exploration involves some observable actions, reorganization is totally internal, even though it may strongly affect subsequently observed actions. I think of exploration as a means toward safe reorganization (performed at zero gain), not as a form of reorganization.

Martin

Date: Wed Sep 02, 1992 7:40 am PST
Subject: chapman; visual perception

[From Rick Marken (920902.0830)] Bill Powers (920901.1900) --

>RE: Chapman's book.

>The part on visual perception isn't anywhere
>near as specific as I was lead to believe, ahem.

Sarcasm on--

Which is EXACTLY what I guessed (and mentioned in a post to Avery). Proving once again that, while I am an asshole, I am also always right (which is appropriate since I am striving to be to Powers what Huxley was to Darwin; only problem is that I'm not quite as smart as Huxley; but there do seem to be plenty of volunteers for the part of Wilberforce [like just about every "scientific" psychologist] and they seem to be willing to be even sillier than the good Bishop).

Sarcasm off-- Best regards Bulldog Rick

Date: Wed Sep 02, 1992 8:25 am PST
Subject: Comments on Chapman's book

[From Bill Powers (920902.0800)]

RE: Vision, Instruction, and Action; David Chapman.

Chapter 2: The concrete-situated approach.

2.1 Routineness.

"The routineness idealization holds that we can and should study routine activity first, making only occasional reference to the novel elements that are introduced from time to time into the course of routine activity."

[By routine activity is meant such things as] "making breakfast, driving to work, reading the paper, typing forms, giving the kids a bath, and grocery shopping."

"Novel activity is possible only against a background of routine competence. In fact it typically proceeds by assembling routine pieces in novel patterns."

"The activity of animals up to the level of lower vertebrates is wholly routine. ... if there is distinct neural machinery responsible for novel activity, it must depend on the stable, previously evolved machinery for routine activity."

My remarks:

What is it that is repeated about a routine activity? Consider making breakfast. One certainly does not have the same thing for breakfast every morning, although there are certain little rituals like making the coffee, even after switching to decaf has made drinking coffee rather pointless. And unless you live alone, you never find things in exactly the same place. You run out of things. You find the milk has turned. There are only enough eggs for French Toast.

Nevertheless, at some level of perception, the "same thing" is repeated every day: we (some of us) have breakfast. I suggest that the level is categories. It isn't that in carrying out a routine one repeats the same actions in the same relationships to the world each time. It's that whatever actual relationships with the world are created each morning, we categorize them in the same way, as "making breakfast" or at least "having breakfast."

What is routine is that we select the same categories of activities each morning, categories known by their names more than by the details of what goes on. Category control requires only that whatever is going on, it be perceived as belonging to the selected reference-category. If someone offers you a chocolate sundae for breakfast, you turn it down not because you don't like such things, but because "that's not breakfast food." There is a category error, which you correct by refusing the food. In England you might well be offered fish for breakfast. Americans wouldn't normally consider a dish of fish to be breakfast and they'd pick something else, and try not to act revolted.

So routine activities are not fixed sequences or specific experiences; they are activities perceived and named from the viewpoint of the category level. Striking variations in the world can lead to category errors just as disturbances can produce errors at any level of organization. Control at the category level entails resetting reference signals for lower levels of control: control of relationship, event, transition, configuration, sensation, and intensity (make the coffee strong). Those reference signals are different every time we make breakfast.

I do not believe that the behavior of lower animals is strictly routine. I don't think that invertebrates are capable of varying their detailed actions so as to maintain the same category of relationship to the environment in the face of disturbances.

Back to Chapman.

2.2 Situatedness

" ... an agent's most important resource in computing what to do is its concrete situation. In driving, you are responsive to other cars, to road signs, and to the geometry of the road. Pouring milk over your breakfast

cereal, your hand constantly adjusts the flow to prevent milk from spashing off the irregular flakes and spilling out of the bowl. ... concrete action without perception is virtually impossible."

"Situations change continually. Algorithms which formally solve your problem are of no use if they terminate after the problem has changed or solved itself or turned into a disaster."

"Real situations are immensely messy, always more complicated than any representation of them can be."

"Real situations, in particular, make it hard to know exactly what the outcome of your actions will be."

"Perception is your only access to the situation around you. You can only see and hear things that happen nearby. Some things that matter to you will always be unknown and unknowable. Algorithms which require complete information are no use,"

"the concrete-situated approach addresses these four difficulties by suggesting that agents

- * Intimately interact with the changing situation.
- * Represent only relevant aspects of the situation.
- * Continually improvise and try out alternative ways of doing things.
- * Ground representation and action in perception."

My remarks:

This is straight control theory, with the exception of a few throwbacks to the stimulus-response view (as in responding to other cars and road signs). Chapman, however, does not know control theory. He has noticed that it is hard to know what the outcome of an action will be, but he has not noticed that we vary our actions to maintain the outcome the same anyway. He speaks of interaction with the environment, but doesn't see that this interaction is specifically one of control; it is not haphazard, but aimed.

He recognizes the messiness of the real world, but not that we achieve regular ends, in great detail, despite the messiness. He sees the importance of perception, but he appears to see perception as regulating behavior, rather than the other way around (grounding representation and action in perception). He uses "representation" in some specialized way that I don't understand, making it seem that perception is something other than representation. Perhaps this is his first glimmering of levels of perception. He does clearly understand that perception is the only access to the external world, which is fundamental to the understanding that behavior is the control of perception, not of the external world.

The "concrete-situated" approach, of course, has been taken for granted in PCT since its beginnings 39 years ago. This has made it hard to understand what Chapman, Agre, et. al. are going on about. I have been looking for some deep significance in these words, when all they are saying is that you have to consider how the environment and organism affect each other, and that the organism must act in the world as it is at any moment. To offer this as an

alternative to the traditional AI approach is really a devastating criticism of AI.

Back to Chapman.

2.3 Interactivity

" ... the organization of activity is emergent from interactions between an agent and its environment, rather than being a property of the agent (as in _mentalism_) or of the environment (as in _behaviorism_). Causality rapidly loops in and out of the agent, rather than looping around inside the agent's head and occasionally emerging to affect the world ..."

"An agent that accepts the world as an equal partner in organizing its activity does not continually try to force interactions to conform to a preconceived idea of how things should go. That's futile; other processes and agents in the world would constantly force you off track."

"Locating the determination of activity in interaction implies viewing life as a series of opportunities and contingencies; openings to participate in particular sorts of activities and events that arise as you do so. Life is not a series of problems, each handed to you as a unit and requiring a unitary solution. Life is constantly ongoing. The concrete-situated approach takes dancing and hanging out, rather than solving the eight [-queens?] puzzle or designing a circuit, as prototypical activities. These activities involve other people, who are likely to make it hard to impose your solutions. On the other hand, in these situations the other people will share the work of making interaction simple."

My remarks:

Noticing that interaction between a person and an environment is of central importance is certainly in line with the PCT view. Even noticing that "activities" depend for their details on vagaries of external forces and other people is consonant with PCT. The environment is full of disturbances.

But Chapman reveals his cultural background here. He is promoting a particular world-view, rather like the Tao and quite like Maturana's view, which says "go with the flow; the river flows by itself so don't try to push it; we're just ships drifting under the influence of wind and tide and where we end up is not up to us." There's a sense of groupiness when he talks about other people that sounds very young to me (but then what doesn't?). There's a sense of not wanting to face the world alone -- or to realize that this is what everyone has to do, by virtue of the relationship of the brain to what lies outside it. Penni echoed this sentiment when she asked rhetorically, "Why figure something out for yourself when you can ask someone else who's been there?"

As is obvious from my participation on the net, I have nothing against groupiness. It's nice. I seek it out. But one mustn't overlook the fact that seeking out the group and the niceness of being with and in it is a GOAL; if you don't decide that this is what you intend to do, and if you don't overcome disturbances and obstacles that stand in the way of doing it, you will not end up with the group. One of the obstacles is the group itself. Groups are not known for reaching out and trying to get people into them (evangelists aside). It's more the other way around; they tend more often to reject outsiders,

which is one of the things that feels so nice when you get inside. When people learn from those around them, which is the nice way of putting it, they also regress toward the mean and stop thinking for themselves, which is the nasty way of putting it.

The problem here is that not knowing PCT, Chapman doesn't realize that even the most trivial of behaviors is an OUTCOME that would not occur with any reliability unless ACTIONS varied, and unless they varied SYSTEMATICALLY to oppose disturbances and in relation to a GOAL. Organisms never just "hang out," if by that one means putting oneself into a situation, social or otherwise, and just being carried along by it. This simply doesn't happen. Even inside the coziest group, the people have goals for their relations to others and for how others treat them. Even just asking someone else a question implies that you want an answer, or at least a response. If someone tells you a fact or technique, that fact or technique will go in one ear and out the other unless you have prepared a place for it, and want to adopt it as yours, and do so for a reason, to suit a purpose of yours.

Chapman does realize that it's hard to impose your solutions on other people. In PCT we say that you can't control other people, at least in the long run (with reference to recent discussions of manipulation) without resorting to overwhelming physical force.

There is, in what he says, a kernel of truth. We do not accomplish our goals in the real world the hard way. Doing things the hard way not only reduces the quality of control, but violates other goals like avoiding physical exhaustion, avoiding offending people and arousing their opposition, and avoiding inner conflicts and physical contradictions. The optimum control organization is the one that maintains the perceived world as one wants to perceive it with the minimum necessary effort. And even an ordinary control process does not produce extra effort when doing so would actually prevent control. When a crosswind is blowing your car into a curve by just the right amount, you don't help it by turning the steering wheel the same way. This doesn't even require thought or principles: if you did act, the car would deviate from the path you want. The steering control system produces WHATEVER effort is required to keep the car in the turn, including none. Of course when you come out of the turn, you hope that the steering wheel is still connected.

Back to Chapman.

2.4 Dynamics

"A dynamic is a pattern in the interaction between an agent and its world... The concrete situated approach prescribes interleaving study of machinery and dynamics."

"Here is an example of dynamics... When you use a bowl, you put it back at the top of the stack. Over time, the bowls that aren't used often tend to sink to the bottom of the stack."

"* Dynamics are not causal agents..."

"* A dynamic operates only ceteris paribus, when an unbounded set of conditions holds. If you have a big dinner party, you may use all your bowls and they'll end up back in the cupboard scrambled..."

"* Dynamics are not processes in the agent's head. They are patterns of interaction which may be noticed by a theorist. Dynamics are typically not represented in [perceived by?] the agent...

"* A dynamic typically operates in many different domains. [Your records may end up sorted as the bowls do].

"* Beneficial dynamics often arise without their having been intended by the agents involved. This is not just an accident, but depends on subtle facts about the structure of activity ..."

"The concrete-situated approach shifts explanatory focus from things in the head to dynamics. We postulate new pieces of machinery only as a last resort... Machinery parsimony is simply good engineering."

" We have found that the deeper your dynamic understanding, the less machinery you need. Bits of machinery, when postulated, do not subserve particular capabilities; the entire agent is applied to every task..."

My remarks:

Chapman means by dynamics not differential equations as I would mean, but properties of the external world and the people in it, as one interacts with them. By "machinery" I take it that he means the design of the behaving system, the model. I strongly support his views that (a) no more machinery than necessary should be proposed, and (b) that the machinery should be general-purpose, not designed ad hoc to fit every special circumstance. One can't always avoid this; we say that the perceptual function of a configuration-control system produces a perceptual signal that stands for a configuration. But with the philosophy Chapman recommends, such ad-hoc chunks in the model are marked for replacement as soon as possible. As we do.

I don't know the whole story on Chapman and Agre's ideas about "beneficial dynamics," but the statement rings a bell. My concept of reorganization provides a way in which unperceived benefits of dynamics can influence behavior so that behavior comes to take advantage of them. The key to the idea is the word "benefits." A benefit must be an improvement in the state of the organism, somehow. It need have neither a logical nor a known relationship to anything the organism is controlling for in the CNS hierarchy. Through this kind of reorganization, as opposed to algorithmic reorganization, it is possible for behavior to reorganize at random until some quite remote deficiency in the organism is remedied. There never does have to be any knowledge in the CNS about WHY the change in behavior was beneficial. So it is possible for very subtle aspects of "dynamics" to result in behavioral organizations that take advantage of them, where doing so is to the benefit of the organism as a whole.

In comparing the PCT view with Chapman's, it's important to note that when Chapman speaks of what the agent "knows" he doesn't make any distinction between conscious and unconscious processes. The perceptions of which he speaks are almost always those of which we would normally be conscious. In PCT, of course, the term "perception" simply means the presence of a neural signal in a perceptual channel; nothing is implied about that perception being

within the field of awareness. So whether something is "perceived" or not tends to be judged by Chapman in terms of whether one would be aware of it or not -- which we see as only part of the story.

I'm not talking about "conscious" and "unconscious" in the Freudian sense. Awareness tends to dwell on certain levels of perception. But those perceptions, in the hierarchical model, could not exist if all the perceptual functions at all the lower levels, right down to spinal reflexes, were not generating their own perceptual signals as usual. This is because each level of perception does not start from scratch (as in Brooks' subsumption architecture) but derives its own type of perception by applying input functions to perceptions that already exist at lower levels. This is why the control hierarchy also acts not directly on the world, but by varying the reference signals for systems of lower levels, the same levels where the raw material for the higher level of perception comes from.

The result is that a lot of what Chapman talks about as the dynamics of "the world" would correspond, in PCT, to the dynamics of lower levels of perception and control. The world that is perceived at higher levels, where instruction and more complex interpretations take place, is still inside the organism and part of the activity of the machinery. To us, of course, it seems consciously to be "outside." But knowing what is "outside" is not as easy as that.

One last sample from Chapman:

2.5 Routines

"For us, the big question in designing an agent is, what sort of machinery will engage in the sorts of routines we want? The architecture we propose is not the only one that will engage in routines; in fact almost any system interacting with an environment will, because routines typically arise from physical determinism. If you put the same agent into the same situation twice, it will do the same thing. Even though real situations and real agents are never exactly the same twice, people and places change only slowly. Your desires, your daily routine, the arrangement of your office and kitchen, your route to work, and your relationship with your office mate, are all relatively constant. An agent's activity can be mainly routine because the world is mainly routine."

My remarks:

I already commented on routines at the beginning. Routines, when repeated, repeat only as categories; the details don't repeat. Not even in ordinary, simple, apparently fixed patterns. We simply perceive what happens on repeated occasions as "the same thing" because we are perceiving in terms of categories, not details.

When Chapman speaks of "doing the same thing" he isn't distinguishing, as we would, between an action and what the action accomplishes. This is strictly because of not knowing about PCT. It may be true that when you put the same agent into the same situation, the agent will ACCOMPLISH the same thing as last time. This would hold true even if you put the agent into a different

(but not WILDLY different) situation. But it is never true that when the agent is in the same situation as before, the agent's ACTIONS can be the same as before. That is simply impossible. The agent can't repeat its initial conditions, and neither can the situation. To a control system this makes no difference. To a system that assumes that commanding the same action as before will have the same result as before, it makes ALL the difference. Even a minute scarcely noticeable change in initial conditions will throw even fairly proximal results off to the point where nothing like the previous result is obtained.

This is the great glaring error of almost all the sciences of behavior -- the assumption that repeated outcomes simply follow from repeated actions. In a computer model, of course, in which we can predict exactly what will follow from prior conditions and known operations, and in which actions, too, are digitized and exactly repeatable, this failure of causality at the output is never encountered. So you can say that if Sonja aims at the monster, she will hit the monster. You can say that if Sonja aims, her missiles will go where she aims. You can say that if Sonja walks through an entrance and around a corner to where an amulet is, she will not run into the wall or grasp at a place 12 inches to the right of the amulet. All you have to do in the digital world is command an outcome, and it happens.

At one level of analysis this is OK. If control systems were used to accomplish all these subgoals, the reference-position would indeed be reached despite all the inherent variations and disturbances in the real world. The hand might start moving in the wrong direction because of a change in body position or a previous velocity, but by the time it was extended it would be in exactly the position required for picking up the amulet.

But at another level of analysis it is not OK. The system has to be designed so it can in fact produce consistent outcomes at EVERY level, despite realistic disturbances and uncertainties that always exist. The PRINCIPLE of control has to be incorporated at every level, or a real working model that has to function in reality instead of in a computer will simply fail.

I think that if Chapman were to learn the principles of control, he could incorporate them into his models in a way that would make them far more robust than they now can possibly be outside the world of bits and precise computations. In many ways he is already using them, but not in a systematic way, with understanding. What he is doing is fundamentally useful; he is exploring the higher levels of control organization where we have made little progress. He and Agre have abandoned the sterile approach of conventional AI for all the right reasons, and in doing so have gone a long way toward discovery of the principles of control. But to take Penni's side for a change, why go through that long painful process of reinventing the wheel, when you can ask someone who's already been there?

Best, Bill P.

Date: Wed Sep 02, 1992 8:42 am PST
Subject: Re: Exploration and Reorganization

[From Rick Marken (920902.0930)]

Martin Taylor (920902 11:15), in response to my slogan:

> "All PURPOSEFUL behavior is the control of perception".

says:

>It also comes perilously close to tautology

Well, slogans are slogans.

see Powers (920901.1000) where he says:

>Rick said it: all control behavior is control of perception.

etc

I said:

>Exploration and reorganization can "look" the same; but I think

>you could tell the difference using the test.

And Martin said:

>They aren't even in the same domain, so I don't see how they could "look" the
>same.

I was not clear. I think of "exploration" and "reorganization" as words describing potential mechanisms underlying an observed behavior; I think of them as descriptions of models. I interpret your comment above to mean that you think of exploration as a description of behavior and reorganization as a description of a model. In this case, there is indeed nothing to compare. And I would argue that much of the behavior that would be described as exploration is actually the result of a reorganization process; other behavior that would also be called exploration is actually systematic, control behavior. For example, when a rat is placed in a test chamber it "explores" -- it sniffs around, climbs the walls, pushes on stuff, etc. It produces a lot of changing perceptions for itself; I think many people would describe this as exploration. Now, what is causing this behavior? One explanation is that it is caused by systematic variation in the rats' references for various perceptions, this variation being caused by a higher order system trying to achieve its own perceptual goal (which might be verbally described as "a perception of high familiarity"). This explanation says that the exploration we see is a result of regular, systematic control processes -- and this could be tested using versions of the test for the controlled variable.

Another possible PCT explanation of the SAME behavior would be that it is the result of reorganization. This explanation would be particularly plausible if we knew that the rat had not eaten for a day before being placed in the test cage. If the "exploratory behavior" is the result of reorganization then we would expect to see the variability of the behavior gradually eliminated as control of some particular aspect of the environment had effects that reduced intrinsic error.

So that was the point of my comment; the SAME behavior (which could be verbally described as "exploratory") could have different internal causes.

This whole discussion may not be relevant to what you were proposing -- which seems like it may have been that "exploration" is an alternative to the reorganization model of learning. If so, then I would have to say "never mind". But, I would suggest that, along with your new model, you could tell us what observations it is meant to explain (that are not handled by the control hierarchy and reorganization model as it already exists) and how we might test your explanation vs the explanation provided by the existing model.

I have often carried on about the importance of the phenomenon of control -- and knowing what that is before getting to involved in the modelling. I would just like to add one more caveat; it is also important to understand the difference between phenomena and the models that are designed to explain them. I know that you understand that difference Martin. But I think that we all must be careful (when we are talking about control phenomena) to be as clear as possible (in our expression and our understanding) about the difference between what we experience (phenomena) and what we imagine (models that we hope will explain the phenomenon).

Best regards Rick

Date: Wed Sep 02, 1992 9:15 am PST
Subject: Scriptures and social control

[CHUCK TUCKER 920902 OR 9/2/92 12:23:30]

I was surprised that someone did not pick up on Greg William's suggestion (920827-2) to quote statements from B. F. Skinner. If you had you would have noted that Greg (quite unlike him) did not cite Skinner's entire paragraph. The sentence he did not cite was: "His behavior (i.e., "when a man controls himself") in so doing is a proper object of analysis, and eventually it must be accounted for with variables lying outside the individual himself." (pp.228-229)

This makes Greg's point (that Skinner used 'control' in a non-PCT sense) quite clear (to me at least).

Another quote from Skinner's chapter of "self-control" (his quotation marks) makes it also clear: "It appears, therefore, that society is responsible for the larger part of the behavior of self-control. If this is correct, little ultimate control remains with the individual. A man may spend a great deal of time designing his own life - he may choose the circumstances in which he is to live with great care, and he may manipulate his daily environment on an extensive scale. Such activity appears to exemplify a high order of self-determination. But it is also behavior, and we account for it in terms of other variables in the environment and history of the individual. It is these variables which provide the ultimate control." (p. 241)

What is so astounding to me is the lack of responsibility and (obviously) self-determination that this view (so-called theory) conveys. Why would anyone be convinced with a view that eliminates them from the process? I can believe it for sociologists who claim that "society made me do it" but for psychologists; where is their individualism??? Another thing that impresses me is the poor quality of writing and reasoning that I find in this work by

Skinner; no wonder he could never figure out the criticisms of him - he could understand them - he was just not very bright.

ON SOCIAL CONTROL

I will have more extensive comment on the posts since 920825 when the "paradox of control" and "manipulation" became a topic but it seems to me that y'all are returning to the question of social control. That's fine but again I would suggest that the conversation will be clear if you restrict the term 'control' to "self-control" that can only be performed by "a negative feedback purposeful process" which is either part of a particular organism or restricted to a single organism thus, by specification, the term 'control' cannot be used to describe any activity that goes on between "self-regulating systems". Words such as influence, manipulation, persuasion, communication, information or others can be used but NOT 'control'. I think this will clear up a number of issues immediately. Then I would suggest that illustrations provide some description of the recent history of the relationships that are used: are you talking about strangers meeting for the first time or are you talking about children that have lived with a person as a parent for X years? I think it makes a difference although "history" does not make people do anything either.

I have enjoyed the conversation and will get in on it next week.

Regards, Chuck

Date: Wed Sep 02, 1992 10:23 am PST
Subject: Re: chapman; visual perception

(ps 920902.1100)

[From Rick Marken (920902.0830)]

Bill Powers (920901.1900) --

>RE: Chapman's book.

>The part on visual perception isn't anywhere
>near as specific as I was lead to believe, ahem.

Sarcasm on--

Which is EXACTLY what I guessed (and mentioned in a post to Avery). Proving once again that, while I am an asshole,

it looks to me like all this proves is:

1) you're more inclined to believe bill than avery

and

2) you will happily believe things w/o going to the primary source to check for yourself

cheers.

--penni

Date: Wed Sep 02, 1992 11:00 am PST
Subject: Re: Comments on Chapman's book

(ps 920902.1200)

importance of perception, but he appears to see perception as regulating behavior, rather than the other way around (grounding representation and action in perception). He uses "representation" in some specialized way that I don't understand, making it seem that perception is something other than representation. Perhaps this is

right. this is part of the problem i've had talking w/ y'all. in an ai/ mentalist approach, perception is the input to representation. representation is datastructures that hang around in your head (whereas perception doesn't hang around in yr head, it's basically a process). chapman and i et al. argue against that view of perc./repr.

solving the eight [-queens?] puzzle or designing a circuit, as

no, you have a 3x3 square and 8 tiles you can slide around in it.

But Chapman reveals his cultural background here. He is promoting a particular world-view, rather like the Tao and quite like Maturana's view, which says "go with the flow; the river flows by itself so don't try to push it; we're just ships drifting under the influence of wind and tide and where we end up is not up to us." There's a sense of groupiness when he talks about other people that sounds very young to me (but then what doesn't?). There's a sense of not wanting to face the world alone -- or to realize that this is what everyone has to do, by virtue of the relationship of the brain to what lies outside it.

that's fascinating. if you talked to him in person (or were subjected to him on the net), you wouldn't have that impression at all.

Penni echoed this sentiment when she asked rhetorically, "Why figure something out for yourself when you can ask someone else who's been there?"

well, i can't really comment on the psychoanalysis of me! ;-} (but i don't mind it.)

As is obvious from my participation on the net, I have nothing against groupiness. It's nice. I seek it out. But one mustn't overlook the fact that seeking out the group and the niceness of being with and in it is a GOAL; if you don't decide that this is what you intend to do, and if you don't overcome disturbances and obstacles that stand in the way of doing it, you will not end up with the group. One of the obstacles is the group itself. Groups are not known for reaching out and trying to get people into them (evangelists aside). It's more the other way around; they tend more often to reject outsiders, which is one of the things that feels so nice when you get inside. When people learn from those around them, which is the nice way of putting it,

they also regress toward the mean and stop thinking for themselves, which is the nasty way of putting it.

you and agre would actually get along very well.

"* Dynamics are not processes in the agent's head. They are patterns of interaction which may be noticed by a theorist. Dynamics are typically not represented in [perceived by?] the agent...

see above.

This is because each level of perception does not start from scratch (as in Brooks' subsumption architecture) but derives its own type

i'm not sure this is right about brooks....

organization where we have made little progress. He and Agre have abandoned the sterile approach of conventional AI for all the right reasons, and in doing so have gone a long way toward discovery of the principles of control. But to take Penni's side for a change, why go through that long painful process of reinventing the wheel, when you can ask someone who's already been there?

well, yes, indeed. i do have sympathy for your position--even if i still don't agree w/ all of pct, i do think it's worth a look. but i don't have any bright ideas on *whom* you should get to look.

cheers. --penni

Date: Wed Sep 02, 1992 11:32 am PST
Subject: Re: Exploration and Reorganization

[Martin Taylor 920902 14:30] (Rick Marken 920902.0930)

>I was not clear. I think of "exploration" and "reorganization" as words
>describing potential mechanisms underlying an observed behavior; I think
>of them as descriptions of models.

Well, clarity, like everything else, is in the eye of the beholder, and in the case of writing, the writer beholding(?) sees differently from most others. I thought I was clear, too. But the rat example clearly fits one view of exploration, which was not what I had intended.

> For example, when a rat is placed in a test chamber
>it "explores" -- it sniffs around, climbs the walls, pushes on stuff, etc.
>It produces a lot of changing perceptions for itself; I think many people
>would describe this as exploration.

Yes, they would, but it wasn't really what I meant, unless neither of your proposed explanations apply.

>Now, what is causing this behavior?

>One explanation is that it is caused by systematic variation in the
>rats' references for various perceptions, this variation being caused by a

>higher order system trying to achieve its own perceptual goal (which
>might be verbally described as "a perception of high familiarity").
>This explanation says that the exploration we see is a result of regular,
>systematic control processes -- and this could be tested using versions
>of the test for the controlled variable.

I hadn't thought of this possibility. It comes close to an idea I had very many years ago, that organisms control to maintain a preferred level of variability in perception at all levels of abstraction. If things are too familiar to the rat, look for something new. If you are put in a sensory isolation chamber, turn up the gain in your sensory systems so that the noise can be interpreted as perceiving something. At a high level, don't allow yourself to get bored.

>Another possible PCT explanation of the SAME behavior would be that it
>is the result of reorganization. This explanation would be particularly
>plausible if we knew that the rat had not eaten for a day before being
>placed in the test cage. If the "exploratory behavior" is the result
>of reorganization then we would expect to see the variability of the
>behavior gradually eliminated as control of some particular aspect of
>the environment had effects that reduced intrinsic error.

This would also account for the actions observed by a third party.

My account ties in with the alerting function discussion. When there are "spare" degrees of freedom that could be used for control, and are not because all the references are reasonably well satisfied, then some are available for actions that find out "what is there." This is not to satisfy any reference for reducing hunger (or "thirst for knowledge"), but to discover the perceptual results of some action. That discovery permits the assignment of the correct sign to the output that was/led to the action, for a later occasion when the ECS in question has a reference that it must satisfy with its percept. Reorganization is a consequence of exploration, in this view, and it happens at zero gain.

I hope this is clear, but even if it is, the clarity may well be an illusion.

Martin

Date: Wed Sep 02, 1992 1:12 pm PST
Subject: Scriptures and social control;hearsay

[From Rick Marken (920802.1400)] CHUCK TUCKER (920902 12:23:30) --

>The sentence he did not cite was: "His behavior (i.e., "when a man
>controls himself") in so doing is a proper object of analysis, and eventually
>it must be accounted for with variables lying outside the individual
himself."
>(pp.228-229)

>This makes Greg's point (that Skinner used 'control' in a non-PCT sense)
quite
>clear (to me at least)

It doesn't make it clear to me. Do the "variables lying outside the individual" control behavior or influence it or what? All we have above is a new description of what these variables do ("account for") and a new description of the behavior that these variables account for ("controlling oneself"). Do these variables influence the behavior (is "controlling oneself" just what happened to result THIS TIME from the presence of these variables) or do they control the behavior (is "controlling oneself" the only result that could have occurred -- the variables would have somehow compensated for any deviation from that behavior).

Skinner just didn't really know the meaning of "control" so he used the word in whatever way made sense in context. If he knew what control was and tried to figure out how it could be done, all the verbal confusion would have disappeared -- he would have discovered PCT.

>I would suggest that the conversation will be clear if you restrict the term
>'control' to "self-control" that can only be performed by "a negative
feedback
>purposeful process" which is either part of a particular organism or
> restricted to a single organism

I don't like "self control" as a replacement for control. There may very well be a perceptual variable that could be called "self" that is controlled; but that is only a small aspect of control, and an unconfirmed one at that. I think the word control works just fine in the discussion of behavior control. Behavior refers to variable results of action. Control refers to variables that vary far less than would be expected given the variance of influences that are known to affect them. Thus, some behaviors are controlled and some are not. PCT says that controlled behaviors (controlled results of action) are controlled by systems inside the behaving system itself. If another system tries to control the results that are controlled by the behaving system (ie. if the other system is trying to control behavior), then the two systems will most likely be in conflict (if they have even slightly different references for the controlled results of action).

penni sibun (920902.1100) --

>it looks to me like all this proves is:

>1) you're more inclined to believe bill than avery

Not really. But Bill does have an awfully good track record. I've tested nearly all Bill's claims about PCT (with experiments and computer modelling) and he's right over 99% of the time (I have caught him a couple times, but only when he's speculating off the top of his head).

>2) you will happily believe things w/o going to the primary source to
>check for yourself

Not only do I go to the primary sources (when I deem it worth the effort based on what I've heard about them or based on excerpts that I have read from them -- as with the Chapman book), I also TEST claims. I don't just believe what I read (at least, when it is testable); I try to demonstrate it to myself and see if something is not what it seems. I believed Bill without going to the primary source (Chapman) because I've worked in the field of perception for 20

years and if someone were really doing something any better, in terms of modelling perception, than we've been doing then I would have been told about it already. If I had to read every piece of, er, stuff, that people thought of as a great breakthrough I wouldn't have time to read anything of real value -- like Jane Austen. So, indeed, I filter -- through the fine strainer of PCT.

And that stuff about being right every time -- it was sarcasm. I was kidding (sort of).

Regards Rick

Date: Wed Sep 02, 1992 2:16 pm PST
Subject: Re: Cooking; chapman; visual perception

(ps 920902.1400)

[From Bill Powers (920901.1900)]

>a primary interactionist field, ethnomethodology, came
>about explicitly in reaction to sociology. rather than asking ``how
>can we describe the[s]e here institutions,' it asks, ``are there
>institutions, and if so what, and if so how constituted.''

Sounds pretty much the same to me. In either case, the institutions are reified. Not that I don't believe in them. I just think they exist in people's heads, not in the environment.

well, of course, to interactionists, institutions are neither in the head nor in the environment, but.... this list, cs91, is an institution cause we are all participating in it. it's not my head, it's not in your head, and it's not out there in the environment somewhere just sitting around.

Well-used, I'd say. Understanding how an object can play a role is something else. I think it makes a difference who is perceiving the object.

yes, of course it does.

If you objectify objects, you can't explain how the same object can have different roles depending on who's using it for what purpose.

i don't see how this follows (unless ``objectify'' has a lot of connotations here).

>a machine is deterministic if there is only one way
>to get from one state to the next; it's nondeterministic otherwise.

That's what I was asking about. If it's nondeterministic in that sense, then what selects which of the possible paths will be taken?

that's not the machine's problem:

``...we shall...permit several possible `next states' for a given combination of current state and input symbol. the automaton [machine]...may choose at each step to go into any one of these legal next states; the choice is not determined by anything in our model, and is therefore said to be nondeterministic.'

--lewis, h. and c. papadimitriou, elements of the theory of computation, prentice-hall, 1981.

If

there's some systematic selection method, then it's deterministic again.

sure.

So the question is, when there are alternate paths, what determines the path actually taken?

as far as toast is concerned, i don't know if they actually say (i didn't see it). they may well have flipped a coin.

>no, a history is what happens between one point in time and another
>point in time.

Do you mean that there are continuous processes taking place BETWEEN nodes? I had thought that an operation simply jumped you from one node to another.

in section ``ontology of cooking tasks,' a&h say: ``a history is a function from natural numbers (representing `time') to world states.' roughly, a history is a record of what happens during a run of toast.

Because the way the proposition is stated, it doesn't matter whether the agent breaks the egg into the pan or whether someone else blunders through the room and breaks it on the floor, or if the agent finds it already broken in the carton. If it's broken the goal is satisfied, the way I read it.

you're probably right on that, and it's a good point.

Now we simply have goal-setting, perception of the current state, comparison with the goal, and action that reduces the error. That I understand. What I don't get are all these nodes and arcs and that stuff, which sounds like the innards of a computer program, not something happening with someone doing it. chacun a son gout.

oh, absolutely. can you see that that's exactly my reaction when y'all get into the details of reference signals etc. and all that math? in both cases, the theories and their models are disjoint, and it's an enormous act of faith (or science?!) to maintain the belief that they have anything to do w/ each other.

cheers.

--penni

Date: Wed Sep 02, 1992 3:33 pm PST
Subject: individualism, chapman

[Avery Andrews (920903.2000)]
(Bill Powers 920901.1900)

Re Penni on individualism, etc.:

>Sounds pretty much the same to me. In either case, the institutions
>are reified. Not that I don't believe in them. I just think they exist
>in people's heads, not in the environment.

>> ``playing a role'' is a well-understood way of talking about these
>>things. the advantage of talking about roles instead of perceptions,
>>is that you don't have to ask perceived by whom.

I'd put it like this. You can have any perceptual functions you like, but if you don't perceive the right things, you loose. And some perceptions are important because of things in the heads of other people (heresy in Medieval Spain, Political Incorrectness at some (many? most? all?) US universities). I think there are some very subtle questions lurking here, for which Tyler Burge's stuff on non-individualism is relevant reading.

My suspicion is that you can't understand the non-individualistic phenomena without understanding the relevant properties of the individuals first, but that you also can't understanding them without going beyond the individuals.

Re Chapman:

>The part on visual perception isn't anywhere
>near as specific as I was lead to believe, ahem.

Hmmm. It seemed to say relatively definite things about all sorts of topics that had puzzled me, via markers, addressing pyramids, etc. And what's wrong with bit-blitting? It seems like perfectly sensible thing for sheets of neurons to be able to do.

>Maybe it was this. Chapman seems, although not clearly, to adopt the
>principle of "coded" perception, in which a system that receives a
>perceptual signal can tell what it's about from the information in the
>signal. In my model, of course, all perceptual signals are alike, so

Well, that would depend on what kind of a system it was. E.g. if you have a whole bunch of input lines coming into a perceptual system, different patterns on the lines can signal different things being perceived, but the signals get their significance from where they come from and what they're going to, just like in PCT. As far as as I can see, Chapman & PCT have exactly the same attitude towards what kinds of computations are legit, namely ones that one can sort of see how to do with neurons (`essentially connectionist hardware', in Chapman's jargon).

Avery.Andrews@anu.edu.au

Date: Wed Sep 02, 1992 7:07 pm PST
Subject: make of this what you will (off sci.cog)

Xref: manuel sci.cognitive:338 comp.ai.philosophy:7142
Path: manuel!munnari.oz.au!uunet!trwacs!erwin
From: erwin@trwacs.fp.trw.com (Harry Erwin)
Newsgroups: sci.cognitive,comp.ai.philosophy
Subject: Mind Reading
Keywords: chaos, neurobiology
Message-ID: <713@trwacs.fp.trw.com>
Date: 2 Sep 92 12:47:59 GMT
Followup-To: sci.cognitive
Organization: TRW Systems Division, Fairfax VA
Lines: 43

I suppose we all use personal experience as an input to our research. I'm going to describe something here that has influenced my ideas.

I've learned over the years to read minds. This is not ESP, but rather a four-stage process as follows:

1. Instantiate a personality model,
2. Calibrate it based on observation,
3. Do a process switch into the calibrated model,
4. Observe the model.

The personality model appears to have some relationship to a full personality--it's often good enough to predict behavior in the short term. Observation includes defocused inputs. The process switch requires suppression of the sense of myself, and often requires a displacement of the point of view to the location of the other individual. Observation is difficult, since it involves deliberately "being" within the personality model, rather than in one's own comfortable personality. I find this whole process exhausting.

The results are often startling, and it is the simple existence of this sort of ability that leads me to speculate on distributed cognition in social groups.

A couple of insights have emerged from this:

1. Lou Pecota and Tom Stafford have shown that chaotic processes are very cheap to control. Freeman's work seems to indicate that a chaotic process is an efficient point to start from if you're doing pattern matching. Perhaps human personalities are based on chaotic processes (as Paul Raup's work seems to indicate) for efficiency of control.
2. The ease with which we innovate based on existing ideas may well reflect a chaotic process responding to small control signals. Since this "stream of consciousness" appears to be generated by a process akin to the process by which speech is generated, we may find that the evolutionary origin of this innovative process to be associated with the evolutionary origin of speech.

Cheers, Harry Erwin
Internet: erwin@trwacs.fp.trw.com

Date: Wed Sep 02, 1992 7:59 pm PST
Subject: Promoting PCT

[From Hank Folson (920902)]

(Dag Forssell 920901)

>>I won't embarrass Dag by listing all the different things he has
>>tried to get me interested in over the years, but there were many.

>Would it embarrass me?

Must have. You're responding, aren't you? -Just joking, folks. Dag's intention is to make a good point:

>Was I trying to interest you (focus on you), or just bubbling along
>about what interested me at the time (focus on me)? I got no error
>signal from your refusal to take my expert advice, did I?

If you had no error signal, my response was not the controlled variable. If so, I'm ticked off that you wasted my time going on about your selfish interests all those years....

As I typed this little witticism, it became embarrassingly clear to me that that this last sentence is a classic S-R view of life that has no place on the net, except perhaps as a bad example. Had I known about PCT then, I would have watched for your reaction to my courteous but non-committal response to your latest interest. Had I seen no hurt in your eyes (The Test), I would have understood that you were indeed just bubbling along about yourself. If my time was wasted, it was my doing, not yours, as I recall no violence in our past.

OR: My response was the controlled variable, but your expectation was framed in the negative (that I would NOT respond). This a nice defense system people have that creates a win-win situation. My negative response would create no error signal. A positive response creates an error signal, but a higher level is satisfied by my positive response which overrides the lower level error signal. Thanx to Ed Ford for introducing me to this in Freedom From Stress.

In any event, whatever your intentions, I, an independent control system, responded as I wanted to respond, which reinforces the difficulty in promoting PCT, which is where this all started.

In control, Hank

Date: Thu Sep 03, 1992 8:08 am PST
Subject: Skinner & Control; More cooking; variaboili

[From Bill Powers (920903.0800)] Chuck Tucker (920902.1223) --

It's hard to know what Skinner was thinking of when he said variables in the environment "control" behavior. He certainly couldn't have meant the same thing as when he said that an experimenter "controls" the behavior of a pigeon toward a preselected goal-state. He just wouldn't have spoken of the

environment controlling in that way. I don't think we can pin him down to a specific meaning, because I don't think he had one.

>... the term 'control' cannot be used to describe any activity that
>goes on between "self-regulating systems".

That puts my view succinctly.

-

Penni Sibun (920902.1200) --

>... in an ai/mentalist approach, perception is the input to
>representation. representation is datastructures that hang around in
>your head (whereas perception doesn't hang around in yr head, it's
>basically a process). chapman and i et al. argue against that view
>of perc./repr.

This is a sticky problem with words. Perception can mean a whole closed-loop process which includes the actions on the world required to bring a perception into some state. It can mean the activities of the perceptual part of the brain -- that is, a process of sensing and constructing higher sensory variables out of lower ones. Or it can mean the outcome of all these processes -- that is, something that is apprehended as a perception, or as Martin says, a percept. To compound that confusion, traditional uses of the term have separated sensation (registering of sensory input energies) from perception (interpretation, inference, and insight). I think that all these usages reflect confusion about the whole subject.

I decided to cut through the confusion (at least in my own head) by dropping all the overlapping and ambiguous usages and using perception to mean simply the presence of a neural signal in any afferent pathway at any level of complexity. I could do this because in the epistemology of control, what is controlled in the final analysis is the signal inside the brain that corresponds to some aspect of the external world. In fact, the "aspect of the world" becomes hypothetical, while the signals are all that are available to awareness. This puts the perceived world inside the head from the very beginning, so there is no question of "coding" or "representation." The world known to the brain, including all its dynamic phenomena, is a world of signals. That's the world that we see as being outside of us, because I assume that the interface between awareness and whatever else there is is always mediated by the brain. I haven't seen any other way of making sense of subjective experience in the same framework with observations of the behavior of others and the physical construction of a person -- and believe me, I have tried for a long time.

From your words, I get the impression that the AI concept of "representation" is something still different. We have the world, perception of the world, and representations of perception of the world. This could be a halting step toward the sort of hierarchy of perception that makes up HPCT. But it doesn't seem to be cast in terms of experienced aspects of the world at the highest level. It seems more like a conversion of higher aspects of experience into a mathematical system that is almost deliberately nonintuitive.

>> This is because each level of perception does not start from
>>scratch (as in Brooks' subsumption architecture) but derives its own

>>type

>i'm not sure this is right about brooks....

I think it is. "Higher" systems in Brooks' architecture get information from the same sensors that feed lower systems; when a higher system wants to act, it has to inhibit the lower system to keep it from interfering. Look at his overall diagram. He shows an input line that branches out to all the levels in the subsumption stack; input information gets to the higher levels without going through the lower systems.

In the HPCT model, the ONLY systems that can act directly on the environment are those in the lowest levels of control. Higher systems have to act by telling lower systems what states of the controlled variables at that level to maintain; they can't bypass the lower systems and operate the muscles directly. Also, many levels involve perception and control of signals that are already being perceived and controlled, individually, by lower systems. The higher-level perceptions are functions of already-existing lower-level perceptions that have been abstracted from sensory inputs by lower perceptual functions. So in general, the inputs to higher systems have already been processed by lower systems, many levels of them. This is very different from Brooks' arrangement.

--

Martin Taylor (920902.1430) -- RE: exploration as result of systematic control process.

>I hadn't thought of this possibility. It comes close to an idea I had
>very many years ago, that organisms control to maintain a preferred
>level of variability in perception at all levels of abstraction.

The problem with ideas like this is that putting them into a model, a block diagram that shows how they would work is very difficult. If a system is going to control for variability at all levels of abstraction, you have to show how it gets information about variability from all these levels, and how its actions affect the lower systems to change the variability in the necessary direction. I can't imagine how that could be done without interfering with the control systems at all those levels. And just what is it about an ECS that you would vary in order to change its "variability?"

The same problem shows up in your suggestions about a "alerting functions". What is it that they are affecting in the controlling systems? How do they know about the consequences of those effects? What, exactly, are these systems supposed to do? I can see some broadly-defined effects, but I don't see any models that would tell us what to expect when connections of this sort are actually set up and working.

We can't tell if conjectures like these are any good until they are cast in the form of a model from which we could predict the consequences of such an organization. To turn them into models, you have to make specific propositions about the details. Until that's done, there's nothing to test. You don't even know if the organization you suggest is feasible on its own merits, much less descriptive of behavioral organization.

--

Penni Sibun (920902.1400) --

>well, of course, to interactionists, institutions are neither in the
>head nor in the environment, but... this list, cs91, is an
institution >cause we are all participating in it. it's not my head,
it's not in >your head, and it's not out there in the environment
somewhere just >sitting around.

My problem with institutions is that they're nouns, whereas what goes on
between organisms are processes and interactions -- verbs. What something like
CSG-L "is" depends on who's looking at it. For me, it's an ongoing
conversation with a person who claims to be Penni Sibun, someone who uses the
name Rick Marken, a writer who expresses himself very much like Chuck Tucker,
and so forth. The only real person on this net is me. From someone else's
viewpoint, I'm imaginary. What the net "is" depends on each user's conception
of it, and that conception isn't out there in the world. It's in a head.

>> If you objectify objects, you can't explain how the same
>> object can have different roles depending on who's using it for what
>> purpose.

>i don't see how this follows (unless ``objectify'' has a lot of
>connotations here).

When the role is in the object, Out There, it's treated like a physical
property of the object. Everyone expects the role to be self-evident to
everyone else, as the role of a lawnmower is to cut grass. By objectivifying
the object, I mean projecting one's private experience of it into a world that
is assumed to be identical for all observers.

>> That's what I was asking about. If it's nondeterministic in that
>> sense, then what selects which of the possible paths will be taken?

>that's not the machine's problem:

>``...we shall...permit several possible `next states' for a given
>combination of current state and input symbol. the automaton
>[machine]...may choose at each step to go into any one of these legal
>next states; the choice is not determined by anything in our model,
>and is therefore said to be nondeterministic.'

Doesn't your quote say explicitly that it is the machine's problem? "The
automaton [machine] may choose at each step to go into any one of these legal
next states..."

I think that what the quote means is that it isn't the programmer's problem
-- that is, the simulation or model or whatever is making its own choices
based on current experience, rather than having those choices programmed in
from the start. But if you include the state of the environment and the
criteria for choosing, then the automaton is deterministic, even if the
programmer didn't determine its choices. The only way I can see to produce a
nondeterministic outcome is to make the choice using random numbers or the
output of a Geiger counter.

> >no, a history is what happens between one point in time and another

> >point in time.

> Do you mean that there are continuous processes taking place BETWEEN
> nodes? I had thought that an operation simply jumped you from one
> node to another.

>in section ``ontology of cooking tasks,' a&h say: ``a history is a
>function from natural numbers (representing `time') to world states.''
>roughly, a history is a record of what happens during a run of toast.

This still doesn't answer my question. If an operation like "put pan on burner" is carried out, I assume that the history would record the pan first in some other place, then after the operation, the pan on the burner. What I was asking was whether the position of the pan was recorded in this history over the whole trajectory between "on the counter" and "on the burner," or whether just the end positions were recorded. In other words, were the positions BETWEEN the nodes recorded, or just the positions AT the nodes?. My impression is that histories don't really use the REAL numbers to represent time, but only the CARDINAL numbers. Is "history" time really physical time? If you move a pot of coffee from the burner to the table too fast, will it slop over?

The reason I ask is that unless physical time is used, none of the physical properties of the elements of the task can be dealt with. Only abstract properties from categories up (in my scheme) can be handled.

>> What I don't get are all these nodes and arcs and that
>> stuff, which sounds like the innards of a computer program, not
>> something happening with someone doing it. chacun a son gout.

>oh, absolutely. can you see that that's exactly my reaction when
>y'all get into the details of reference signals etc. and all that
>math? in both cases, the theories and their models are disjoint, and
>it's an enormous act of faith (or science?!) to maintain the belief
>that they have anything to do w/ each other.

This is why we don't start with the math, but with the phenomenon of control. Once you have experienced and understood control as something happening, the PCT explanation becomes the only feasible one; it's clear that no other explanation can even come close to handling what you see happening. It would, in fact, take an enormous act of faith to believe that the control-theoretic explanation is wrong.

By this I don't mean that a specific model, assembled for purposes of analysis, is unique or self-evident. I mean that the CONCEPT of which the specific model is one embodiment seems unavoidable. If you're holding back from accepting PCT, it's only because you haven't seen the CONCEPT yet. This can't be done through mathematical analysis. You just have to look at your own experiences and actions, and realize that everything you know about them occurs first as a perception. And that all your actions are organized to make perceptions be a certain way, or not be a certain way. Until that gets through, none of the math will seem particularly interesting or compelling.

Best to all, Bill P.

Date: Thu Sep 03, 1992 8:43 am PST
Subject: Planning and Agre's Cooking Problem

Avery- I like the typology (relevant;relevant but routine;
relevant but uncontrollable). Try this on. I am working on
information systems which require different "granualities"
at different levels of use. The person at the top doesn't
wan to be deluged, so he/she wants aggregation, but not loss
of content. Those below want less aggregation, but not loss
of the leaderships "vision" from the top. I'm looking at a
model which puts everyone's information universe in a space
that, in two dimensions, is a pyramid. The base of the
shape is a time dimension (going right) and the hieght of
the shape is a measure of granularity.
People at the top have require less granularity, so the
height of their pyramid is greater than those below.(Showing
that their perception is more course - thier min scale is
greater.

The shape is trisected, so that the bottom part represents
those "inputs" that represent things which can be considered
stable during the relevant time span.(Relevant time span is,
by the way, a function of the time in which the perceiver
can respond to the input). The next portion up in the
pyramid represents data that is predicably variable; by
which I mean that it varies according to some pattern or
routine KNOWN TO THE PERCEIVER. The top part of the pyramid
represents information that the perceiver must treat as
either random, or varying in an unknown fashion; and
therefore is the part of the data spectrum faced by the
perciever which must recieve the greatest "attention."
Breakdowns in understanding of information among the layers
of the organization are related to the distribution of data
among the sections of the pyramids at each layer.
(lowerlayer having shorter time lines, and lower "ceilings",
because their perceptions are finer).
I'm still working on getting this in PCT language.

Date: Thu Sep 03, 1992 8:48 am PST
Subject: Deja vu all over again

[From Rick Marken (920803.0930)]

Bill Powers (920902.0800) quotes the following from Chapman:

>2.2 Situatedness

>"Situations change continually. Algorithms which formally solve your
>problem are of no use if they terminate after the problem has changed
>or solved itself or turned into a disaster."

>"Real situations, in particular, make it hard to know exactly what the
>outcome of your actions will be."

As Bill notes, these and some of the other quotes from Chapman sound about as close to PCT as you can be without being there. Reading them, I had a sort of deja vu -- and I realized that it was based on the fact that I had read nearly the same kind of stuff in the early 80s; it was written by the likes of Turvey, Kelso, Shaw, etc. They called themselves the "Bernstein" school then (after the Russian physiologist) who proposed the idea that we now know as "coordinative structures". The Bernstein school has now become the nonlinear dynamics group after learning some chaos theory.

The point is that the Bernstein school (like the interactionists or whatever Chapman calls himself) was as close to PCT as you can be; they understood that consistent behavioral results are produced in the context of variable disturbances (they called it context conditioned variability); they knew that it was impossible to compute the many outputs required to produce a particular result (they called this the degrees of freedom problem). Some had a glimmer that perception was important to the "invariance" of behavior (they were Gibsonians, after all). So they were exactly where Chapman is right now, then. But they steadfastly refused to consider PCT as a solution to the problems that they correctly identified -- in fact, they actively rejected PCT (in the Fowler-Turvey article). So now they are off, happily poking down trendy dead end after trendy dead end. I predict the same fate for Chapman et al. After the interactionist verbalizations get tiresome they'll move off to the next trendy thing; and another group of young turks will come along, rediscover control and not know what they are looking at, and proceed to repeat the cycle again.

And PCT will just keep slowly rollin' along, with solutions to all the problems that are being rediscovered again and again, and making slow progress because there are only two or three people doing the research and modelling work.

What is it about PCT? Why don't people like it? Why do they prefer trendy junk (Chapman, Turvey, Brooks, etc) to solid quality (Powers)? Is there something wrong with PCT? Or is it the influence of network TV?

Best regards Rick

Date: Thu Sep 03, 1992 9:17 am PST
Subject: The phenomenon of control

[From Dag Forssell (920903-1)]

Bill Powers (920903.0800) in response to Penni Sibun (920902.1400)

>This is why we don't start with the math, but with the phenomenon of
>control. Once you have experienced and understood control as something
>happening, the PCT explanation becomes the only feasible one; it's clear
>that no other explanation can even come close to handling what you see
>happening. It would, in fact, take an enormous act of faith to believe
>that the control-theoretic explanation is wrong.

Penni, if you give me your snail mail address, I will be pleased to send you my information packet which includes a do-it-yourself demonstration

of the phenomenon of control. The demo includes a narrative adapted from Bill's work and several illustrations. You perform the demo with another person (preferred) or by just reading and visualizing.

By getting some personal, physical experience with control and what it looks like in action, you may gain a fresh perspective on the discussions here.

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

Date: Thu Sep 03, 1992 9:17 am PST
Subject: Loose talk, or what?

From Greg Williams (920903)

>Bill Powers (920830.1545):

>Controlling "a person" means controlling the value of some variable
>associated with that person. Some of those variables (the amounts and kinds
>of actions) are controllable if the controlled variable is known. Others
>(perceptions, controlled variables) are not.

>Bill Powers (920903.0800):

([replying to] Chuck Tucker (920902.1223))

CT>>... the term 'control' cannot be used to describe any activity that
CT>>goes on between "self-regulating systems".

BP>That puts my view succinctly.

Say what???? Talk about major contradictions....

Please elucidate if possible.

Greg

Date: Thu Sep 03, 1992 9:47 am PST
Subject: Interactions; bit-blitting; Mind reading & chaos

[From Bill Powers (920903.1000)] Avery Andrews (920903.2000) --

A likely story -- it's 920903.1000 here. This time warp is sort of fascinating.

>I'd put it like this. You can have any perceptual functions you like,
>but if you don't perceive the right things, you lose.

Right. Absolutely. I'm probably reacting to something that nobody here is saying. The idea that needs dumping is that there are aspects of social organization that exist apart from the people in them, which is the old sociological view, still extant, and the current economic view.

People have properties and physical environments have properties. By interacting with them we learn what's important and possible to perceive about them (invariants) and then what's possible and important to control about them (as we perceive them). Each of us learns to do this. So the interaction between people is particularly complicated, in that as we learn perceptions of other people and learn to control the perceptions, we change our properties as other people experience us. They, in turn, have to alter how they perceive us and what they can control about us or their relation to us. When this process converges, we get a stable pattern of interactions in which each person is controlling for the things that matter to that person, while allowing others to do the same, within reason. The result can be seen, I suppose, as an "institution." Bruce Nevin sees language this way, and so do I.

I guess I'm pushing the idea that you have to understand the basic principles that make individuals work before you can even start to understand the consequences of interactions among individuals, or between individuals and the physical environment. As our model for individuals improves, so will our ability to run social models consisting of many individuals.

>My suspicion is that you can't understand the non-individualistic
>phenomena without understanding the relevant properties of the
>individuals first, but that you also can't understanding them without
>going beyond the individuals.

... so we agree completely. Going beyond the individuals means letting two more more models of individuals interact with each other and their common environment (which will be especially interesting when they don't perceive the common environment in the same way).

RE: Chapman.

>... what's wrong with bit-blitting? It seems like perfectly sensible
>thing for sheets of neurons to be able to do.

What I read was an idea that a wave of "activation" spreads from some test point until it is stopped by a closed boundary. Another test point far from the purported object is reached by the spreading wave (if the boundary isn't closed). If it's not reached, the original test point was "inside" the boundary. So by perceiving the state of the remote test point, we can perceive whether the initial point was or was not inside the boundary -- or inside SOME boundary.

This is how you'd have to do it on a computer screen using standard polygon-filling software and hardware. As all graphics programmers know, this is a slow process.

It seems to me that a much simpler way to determine inside or outside would be to compute the direction of the radius vector from a point that travels once around the boundary to the stationary point, which can be done in terms of x-y coordinates. If the direction makes one complete 360-degree rotation during a

complete scan, the point is inside. If the final direction angle is the same as the initial one, the point is outside.

This way feels much more like the way I do it. I run my attention (or my eyes) around the boundary, to see if my scan wraps all the way around the object or unwraps again before it's finished.

```

*****
*                               *
*                               *
*                               *
*           A           *   Wrap = 360 degrees
*                               *
*                               *
*****

*****
*                               *
*           *****
*           *
*           * A           *   Wrap = 0 degrees
*           *****
*                               *
*****

```

Note also that you can say that A is inside or outside even though the figure is "leaky" -- has holes between the points defining it.

My way uses far less computing activity than Uhlmann's.

```

-----
>if you
>have a whole bunch of input lines coming into a perceptual system,
>different patterns on the lines can signal different things being
>perceived, but the signals get their significance from where they
>come from and what they're going to, just like in PCT.

```

In the pandemonium model, these input lines form a bus, with many perceptual functions hanging off of it. Each perceptual function applies its own input process to its copy of the information, generating a single perceptual signal. So all the perceptual signals are alike, representing only magnitudes, the momentary values of the functions. The meaning of a given signal is determined by the function which derives it from some or all of the signals on the bus.

Of course ALL POSSIBLE FUNCTIONS of the bus signals are implicit in the bus signals. Implicit functions, however, accomplish nothing and produce no perceptions. Only those functions explicitly computed by neurons are perceptually real, have any existence for higher systems.

So to say that the different patterns on the lines signal different things states only part of the truth. The whole truth is that they signal infinitely many different things, and so signal nothing in particular. Even a single static set of signals on the bus could lead to an infinity of different perceptual signals, each a different function of those bus signals. If there are n signals on the bus, then at most n different functions of the signals could be independently controlled at a given instant. But over time, there is

no limit to the number of different sets of n functions that can be computed and controlled, starting with n signals on the bus.

I think that even a model in which the patterns of signals are supposed to be meaningful runs afoul of this same consideration. There is no "pattern" of signals until a receiver passes the signals through weighting or other processes. Pattern "detection" is really pattern CREATION. The receiving device contains a definition of a pattern, and its output simply reports the degree to which that definition is met by the current distribution of magnitudes in the set of input signals. One set of input signals can be the input for many simultaneous and quite different pattern-sensing devices. There is no one inherent pattern in the input signals. Pattern is in the eye of the beholder.

Harry Erwin (920902) --

Welcome to CSGnet, Harry!

>I've learned over the years to read minds. This is not ESP, but
>rather a four-stage process as follows:
>
>1. Instantiate a personality model,
>2. Calibrate it based on observation,
>3. Do a process switch into the calibrated model,
>4. Observe the model.

One of our PCT freaks who is a psychotherapist, David Goldstein, has been working recently with a "straight" therapist named Maher. David and I have been exploring a "method of levels" that uses a principle of awareness saying that awareness never registers the level of perception that it's working FROM, but only the information reaching that level from lower systems.

One of the tasks involved in conducting this method is for one person to listen carefully to what another is saying (on practically any subject), and try to get a feel for the attitudes, opinions, interpretations, and so forth that the other is using "behind the scenes" in talking about the surface subject matter. This involves a technique very much like the one you describe in your post. The conductor of this process basically has no thoughts of his own; the whole point is to understand the other's background thoughts, at least well enough to ask questions that bring them to attention. It's not true, of course, that the conductor's thoughts aren't his own, but they are a model of the other's thought processes, and exclude the sorts of thoughts and attitudes that the conductor would normally bring into a conversation.

The interesting thing about Maher is that he has developed something very close to this technique quite independently of PCT (as you appear to have done). Maher speaks of it as being the other person as nearly as you can. I'll let David Goldstein describe Maher's approach in more detail, as he is now intimately familiar with it.

The point I would like to make is that your mind-reading method has one shortcoming, which is the shortcoming of any model that is plausible but has not been tested. The method of levels, the PCT version of psychotherapy, does involve making a model of the other person -- I hadn't thought of it exactly that way until you said it. But it also involves continually TESTING that

model. The testing is quite simple: you propose to the other person that some background thought is going on, and the other person agrees, disagrees, or comes back with a modified version of your suggestion that's closer to right. This helps you modify your model, so the next suggestion or question is more likely to be fruitful.

I've used this method with a lot of people on a teaching/demonstration basis (I'm certainly not a psychotherapist), and I concur with you that it's an exhausting procedure. You should try it, however, with a friend, organized as an exploration as we do with the method of levels. When you're on the receiving end, you find that you can do "up a level" more than once, and that when you carry this procedure to its termination point (where, in effect, there is no more mind left to read at a higher level), the resulting state is most interesting.

Concerning Chaos, it may be that chaotic processes are cheap to control, but it's even cheaper to control non-chaotic processes and the result is much more precise control. Chaotic processes are hypersensitive to small changes in critical variables, which is what makes control cheap. But they are also subject to disturbance just as easily by extraneous variables, so they need a lot of controlling. Furthermore, using a controlled chaotic process to control something else, as in a hierarchy, is not likely to work very well, because the chaotic process can't be loaded much by contact with an environment without losing its character.

I think that "chaotic" is often used when it would be just as correct to say "random." Random processes can lead to systematic results when employed in the context of a control system -- our perennial example is the mode of steering of the intestinal bacterium *E. coli* (which can't steer, but can only swim in straight lines or tumble randomly to create a new direction of swimming unrelated to the previous one).

The proponents of chaos as a factor in perception are, no doubt, thinking of the way chaotic oscillators go from one organized state to another by way of chaotic intervening states. This suggests some natural way of settling into stable perceptual modes with a sort of random reorganization leading to the transitions. My own opinion, which isn't worth much in this field, is that there may be an analogy here, but the real mechanism is quite different. I don't see, for example, any signs that perception has an oscillatory component. Perhaps there is some other sort of physical phenomenon that would mimic the behavior of chaotic oscillators, but I don't think an oscillator is a very useful model of a perceptual function.

Greg Williams--

OK on the figure: we'll wait. I'll send for permissions to the editor too, just in case.

This is about getting the Little Man paper ready to ship out, folks. We're almost there.

Best to all, Bill P.

Date: Thu Sep 03, 1992 11:23 am PST

Subject: STELLA for PCT?

[from Gary Cziko 920903.1340]

To Macintosh programmers and users (e.g., Marken, Taylor)

In a recent campus newsletter on using computers for instruction, there is an article on one professor's use of the Macintosh program STELLA for dynamic modelling in the life sciences.

As someone who is quite computer literate but with only very basic (and in fact only BASIC) programming experience, I wonder if STELLA might be a way for me to do some PCT-type modelling without the pain of learning a traditional language. Can someone comment on the potential uses (advantages and disadvantages) of STELLA for PCT work? --Gary

Date: Thu Sep 03, 1992 11:34 am PST

Subject: Deja vu, PCT promotion

[From Dag Forssell (920903-2)]

Rick Marken (920803.0930)

>What is it about PCT? Why don't people like it? Why do they prefer
>trendy junk (Chapman, Turvey, Brooks, etc) to solid quality (Powers)?
>Is there something wrong with PCT? Or is it the influence of network TV?

There is nothing wrong with PCT, except that it lives up to the standards of a hard, explanatory science.

It compares to physics (Bill's background), chemistry, engineering, and such, which are explanatory and predictive with expectations of 100% correlation between science and results in the real world.

It is hard to comprehend for people thoroughly trained in the soft, descriptive sciences, which explain little and predict even less. You can tell the soft, descriptive sciences by the suffix science, as in science of psychology, life science, social science.

You are one of the very few psychologists who have demanded that the science of psychology should have something useful to say about the practice of psychology. You are one of the very few who likes to think for yourself based on the evidence, instead of conveniently accepting what your teachers told you and what the abundant literature of the prestigious association tells you.

You rant and rave against people who have another world view than you do. People who cannot perceive the inconsistencies you see and are not disturbed by the lack of clarity and application in what they "research." They have different systems concepts and are happy where they are.

You edited American Behavioral Scientist, from Sage Publications, Thousand Oaks, Ca, Volume 34 / Number 1, devoted to:

PURPOSEFUL BEHAVIOR: The Control Theory Approach.

I was reading it recently, and came across the following in Phil Runkel's article on: Research Method for Control Theory:Most psychological research, for example, does not assume the circular causation specified in Postulate 4; it assumes linear, input-output causation. I once sent a diagram of the input- output metatheory to 16 psychologists, asking whether they thought it was a fair representation of the theory underlying method in psychological research. Most of them gave me offhand answers such as "Looks OK to me." Four of the 16 wrote comments to the effect of "Method? I don't think about method; I just do research.".....

So when you say:

>What is it about PCT? Why don't people like it? Why do they prefer

you know the answer. People will like it, and do. You have to consider who you approach with this marvelous answer to the prayer of many people who are troubled and need a better understanding.

Address groups and professions where people are used to hard, explanatory, predictive thinking that have problems PCT can help them solve. Show them the theory in all its glory, but start with some aspect they can use right away, or they may not see the relevance and maintain interest.

Address groups where people can use more "common sense" common sense in their personal relations with Ed Ford's narrative approach and superb counseling book: FREEDOM FROM STRESS. Brandt publishing, Scottsdale Az (602) 991-4860

Stop worrying about the 1/10 of 1% who are already thoroughly trained in the science of psychology and whose prestige is wrapped up in defending what they already think they know, their school, teachers etc..

Approach the other 99.9% in the way PCT tells you to. Find and address people's error signals.

As you well know, there are three ways to solve a conflict.

- 1) Change what you do.
- 2) Change what you want.
- 3) Change your choice of perception.

You can change your want in terms of people you want to convert.
You can of course continue on course, if you enjoy and want to rant.

All the best, Dag

Date: Thu Sep 03, 1992 3:49 pm PST
Subject: Pyramids; Succinctness; Stella

[From Bill Powers (920903.1600)]

Thomas Baines (920903) --

You might consider looking at my Living Control Systems, the chapter called An Outline of Control Theory. In it are defined 11 levels of perception

ranging from intensities to system concepts. These are, of course, also levels of control. If the ordering is right, the lower levels operate the fastest; also, it is necessary to control variables of level n-1 in order to control variables of level n. I don't know how well these properties actually hold up in every instance, but you might find some suggestions in this hierarchical arrangement.

Use the term "model" carefully around here. We generally think of models as things you can set up as a program, turn on, and run -- at least in principle.

-

Greg Williams (920903) --

That's the trouble with succinctness. Somebody immediately comes up with a counterexample -- in the worst case, in your own words.

"Controlling a person" means controlling some variable associated with that person. This is possible without direct use of force only if the variable is an output variable. Otherwise [Tucker] "... the term 'control' cannot be used to describe any activity that goes on between "self-regulating systems"."

I was thinking of an activity "between" self-regulating systems, meaning one in which they are interacting. If I'm controlling your outputs, we aren't interacting; I'm just diddling around with side-effects of your control actions and you're ignoring me.

-

Gary Cziko (920903.1340) --

I've seen Stella in operation, and it would be a fine tool for simulating control systems on a Mac. I've heard that it's available for PCs too. But my impression is that it's VERY expensive. If Rick could get a copy I think he would freak out: it's really quite pretty. You just draw block diagrams using standard components like amplifiers and integrators, specify the parameters, connect the blocks by drawing lines, and run! You can get plots of various variables, read in data from files, and so on.

Greg and Pat Williams are going to be working on a general simulation program that will probably work similarly. The only problem: it will be written for the PC world. I presume it will be written in C, however, so a C-competent Mac user might be able to port it to the Mac world, in which case we'd all be able to run the same simulations. That would truly be a great step forward.

Best to all, Bill P.

Date: Thu Sep 03, 1992 3:59 pm PST
 Subject: Re: Skinner & Control; More cooking; variaboili

[Martin Taylor 920903 19:00] (Bill Powers 920903 0800)

(To Penni Sibun)

>By this I don't mean that a specific model, assembled for purposes of
 >analysis, is unique or self-evident. I mean that the CONCEPT of which

>the specific model is one embodiment seems unavoidable. If you're
>holding back from accepting PCT, it's only because you haven't seen
>the CONCEPT yet.

(To me)

>We can't tell if conjectures like these are any good until they are
>cast in the form of a model from which we could predict the
>consequences of such an organization. To turn them into models, you
>have to make specific propositions about the details. Until that's
>done, there's nothing to test. You don't even know if the organization
>you suggest is feasible on its own merits, much less descriptive of
>behavioral organization.

I don't have or expect to have a specific model of alerting systems. I can think of many. But the situation is like that of PCT itself. The basic proposition of PCT, that "all (purposeful) behaviour is the control of perception" is incontrovertible. There are many models that might instantiate it. So with the alerting systems. Given the basic proposition of PCT, their existence in an organism supplied with more sensory than motor degrees of freedom is incontrovertible. There are many models for how they might work.

As for "controlling for variability," I was more musing than proposing, but if you want a proposal, you mentioned a while back that if you had two one-way control systems back-to-back to make a two-way one, and if they had the right (square-law?) control function, then the pair could be controlled both for gain and for reference level. The normal (S-R?) view of adaptation is a change of gain in the perceptual input, put there's no intrinsic reason why that sensitivity shouldn't be in the output (also?).

I see here added complexity, but no implausibility, and I don't plan to model it (yet?). It doesn't seem to be a mainstream problem, as the alerting structure is.

=====

Unless I see something else for response in the next few minutes, this is likely to be my last posting until Sept 24. (But my addiction to CSG-L may make a liar of me until Saturday afternoon).

Martin

Date: Thu Sep 03, 1992 4:03 pm PST

Subject: Re: Interactions; bit-blinking; Mind reading & chaos

[Avery Andrews (920904.0950)]

(Bill Powers 920903.1000)

>This is how you'd have to do it on a computer screen using standard
>polygon-filling software and hardware. As all graphics programmers
>know, this is a slow process.

But presumably it would go quite quickly if each bit had its own processor looking after it, as is plausible the case for brains (c.f. sheets of neurons). But I do like your way better, on casual consideration (& it also seems quite consistent with the kinds of things that Chapman does with rays).

Re busses: fine. But as far as I can see, Chapman's architecture is consistent with this - as long as you have a finite system of categories sent along the bus, you can have a single wire coming off for each one.

Avery.Andrews@anu.edu.au

PS did Harry Erwin really send something to CSGNet? I posted something of his yesterday, but I also lost a whole swag of unread messages 2 days ago, so might well have missed a posting by him.

Date: Thu Sep 03, 1992 4:24 pm PST
Subject: A Strange Bedfellow?

From Greg Williams (920903 - 2)

Quoting B.F. Skinner:

The relation of organism to environment must be supposed to include the special case of the relation of scientist to subject matter.

-- THE BEHAVIOR OF ORGANISMS, 1938, p. 43

We often overlook the fact that human behavior is also a form of control. That an organism should act to control the world around it is as characteristic of life as breathing or reproduction. A person acts upon the environment, and what he [sic] achieves is essential to his survival and the survival of the species.... We cannot choose a way of life in which there is no control. We can only change the controlling conditions.... -- ABOUT BEHAVIORISM, 1974, pp. 189-190

Self-management raises the same question as self-knowledge: Who are the managing and managed selves? And again the answer is that they are repertoires of behavior.... The managed self is composed of what is significantly called selfish behavior -- the product of the biological reinforcers to which the species has been made sensitive through natural selection. The managing self, on the other hand, is set up mainly by the social environment, which has ITS selfish reasons for teaching a person to alter his [sic] behavior in such a way that it becomes less aversive and more reinforcing to others. -- ABOUT BEHAVIORISM, 1974, pp. 176-177

One person manages another in the sense in which he manages himself. He does not do so by changing feelings or states of mind. The Greek gods were said to change behavior by giving men and women mental states, such as pride, mental confusion, or courage, but no one has been successful in doing so since. One person changes the behavior of another by changing the world in which he lives. -- ABOUT BEHAVIORISM, 1974, pp. 180-181

In traditional terms, one person arranges positive or negative contingencies in order to create interests, provide encouragement, instill incentives or purposes, or raise consciousness in another person. In doing so, he brings him under control of various features of his environment. -- ABOUT BEHAVIORISM, 1974, p. 181

A person who has been exposed to the promise of heaven and the threat of hell may feel stronger bodily states than one whose behavior is merely approved or censured by his fellow man. But neither one acts BECAUSE he knows or feels that his behavior is right; he acts because of the contingencies which have shaped his behavior and created the conditions he feels. -- ABOUT BEHAVIORISM, 1974, p. 193

A physical world generates both physical action and the physical conditions within the body to which a person responds when a verbal community arranges the necessary contingencies. -- ABOUT BEHAVIORISM, 1974, p. 220

In self-management the controlling self is different from the controlled. But all selves are the products of genetic and environmental histories. Self-knowledge and self-management are of social origin, and the selves known and managed are the products of both contingencies of survival and contingencies of reinforcement. Nothing about the position taken in this book questions the uniqueness of each member of the human species, but the uniqueness is inherent in the sources. There is no place in the scientific position for a self as a true originator or initiator of action. -- ABOUT BEHAVIORISM, 1974, p. 225

It would be absurd for the behaviorist to contend that he is in any way exempt from his analysis. He cannot step out of the causal stream and observe behavior from some special point of vantage... -- ABOUT BEHAVIORISM, 1974, p. 234

Democracy is a version of countercontrol designed to solve the problem of manipulation.... Contingencies designed for explicit purposes can be called manipulative, though it does not follow that they are exploitative; unarranged contingencies must be recognized as having equal power, and also possibly unhappy consequences.... To say that all control is manipulative and hence wrong is to overlook important uses in education, psychotherapy, government, and elsewhere. -- ABOUT BEHAVIORISM, 1974, pp. 243-244

I have used technical terms in making a technical point. I have preferred a technical term elsewhere when it could be used at no great cost.... But elsewhere [in this book] I have freely used the lay vocabulary while accepting the responsibility of providing a technical translation upon demand.... Those who approach a behavioristic formulation for the first time may be surprised by the mention of self-control. Does this not suggest some kind of inner determination?... According to traditional definitions of self-control, ... [etc.] the behaviorist is indeed inconsistent, but according to his own definitions he is not... "If human behavior is as fully determined as the behaviorist says it is, why does he bother to write a book? Does he believe that anything matters?" To answer that question we should have to go into the history of the behaviorist. Nothing he says about human behavior seriously changes the effect of that history. His research has not altered his concern for his fellow men... -- ABOUT BEHAVIORISM, 1974, pp. 247-248

When two different response rates occur in the presence of different stimuli, the response is under STIMULUS CONTROL. ... The girl whose facial expressions make her look "approachable" instead of "aloof" is more likely to be asked for a date. She may assume an "approachable" expression to exert stimulus CONTROL over a young man's behavior. ... When a distant doorbell rings, you may "make a mistake" and go to the phone. The doorbell exerts some CONTROL over going to the phone. This phenomenon is called GENERALIZATION.

... A man may CONTROL the behavior of another man by arranging relevant conditions. Also, he may control his own behavior by arranging the same kinds of conditions. A mother may put candy out of sight to DECREASE the probability that her child will ask for it. She may do the same thing to reduce the PROBABILITY that she will eat the candy. In analyzing cases in which one response controls another, we distinguish between the controlling response and the controlled response. Putting candy out of sight is the CONTROLLING response; eating candy is the CONTROLLED response. ... Putting candy out of sight to keep from eating it CONTROLS behavior by REMOVING an SD [discriminative stimulus]. -- THE ANALYSIS OF BEHAVIOR [a programmed text; "correct" answers in BOLDFACE]

When you act to control yourself and other people, you are doing 'what God allows you to do rather than forces you to do.' And you do it by constructing a SITUATION that controls you or the others. Frazier does not control the members of Walden Two. The world which he designed and which they maintain is the controller. -- NOTEBOOKS, 1980, p. 112

In short, for Skinner, the results of one person's responses to his/her environmental stimuli (which CONTROL him/her) can be another person's CONTROLLING stimuli.

And on statistics:

There are at the present time two quite different modes of approaching the behavior of organisms which are hard to distinguish theoretically but which are clearly different in practice. The statistical approach is characterized by relatively unrefined methods of measurement and a general neglect of the problem of direct description. The non-statistical approach confines itself to specific instances of behavior and to the development of methods of direct measurement and analysis. The statistical approach compensates for its lack of rigor at the stage of measurement by having recourse to statistical analysis, which the non-statistical approach in general avoids. The resulting formulations of behavior are as diverse as the methods through which they are achieved. The concepts established in the first case become a part of scientific knowledge only by virtue of statistical procedures, and their reference to the behavior of an individual is indirect. In the second case there is a simpler relation between a concept and its referent and a more immediate bearing upon the individual. It may be that the differences between the two approaches are transitory and that eventually a combination of the two will give us our best methods, but at the present time they are characterized by different and almost incompatible conceptions of a science of behavior. It is obvious that the kind of science here proposed naturally belongs on the non-statistical side of this argument. In placing itself in that position it gains the advantage of a kind of prediction concerning the individual that is necessarily lacking in a statistical science. The physician who is trying to determine whether his patient will die before morning can make little use of actuarial tables, nor can the student of behavior predict what a single organism will do if his laws apply only to groups. Individual prediction is of tremendous importance, so long as the organism is to be treated scientifically as a lawful system. Until we are spared the necessity of choosing between the two approaches, we must cast our lot with a non-statistical investigation of the individual and achieve whatever degree of reliability or reproducibility we may through the development of techniques of measurement and control. -- THE BEHAVIOR OF ORGANISMS, 1938, pp. 443-444

It's a shame Skinner never got together with Phil Runkel!

Greg

Date: Thu Sep 03, 1992 5:07 pm PST
Subject: Strange Bedfellows

Subj: RE: A Strange Bedfellow?

[FROM: Dennis Delprato (920903)] (Greg Williams 920903)

Greg, I don't get the point of all the quotations from Fred Skinner on control. You sum up: "In short, for Skinner, the results of one person's responses to his/her environmental stimuli (which CONTROL him/her) can be another person's CONTROLLING stimuli." Do you mean he is a strange bedfellow for himself? Seriously, what is the message I am missing?

Skinner's thinking is more advanced than that of the typical psychologist, hopelessly behind that of others, and he is extremely careful in his selection of words and their organization. The effect of the latter is that it is virtually impossible to pin him down--he's slick.

As far as individual-subject research, it was Skinner who offered that "no one goes to the circus to see the average dog jump through a hoop significantly oftener than untrained dogs raised under the same circumstances." A snappy line for which I give him credit. Too bad Runkle cited not one (as I recall) Skinner-inspired paper on actuarial vs. single-subject research. Don't know where Phil has been since 1938.

Date: Thu Sep 03, 1992 6:26 pm PST
Subject: Control/Manipulation; Skinner in Bed

From Greg Williams (920903 - 3)

>Bill Powers (920903.1600)

>"Controlling a person" means controlling some variable associated with
>that person. This is possible without direct use of force only if the
>variable is an output variable.

Two questions:

1. Is controlling a person possible with THREAT of force (which might be included in what you mean by "direct use of force")?

2. Suppose I yell "Fire!" when we're both in a theater. You (being prudent) join the crowd heading for the exit. Were you, BOTH before and after I yelled, ACTIVELY controlling for not getting burned (or some such)? If you weren't before, but were after, then it could be said that I had altered WHICH reference signal you were ACTIVELY trying to match, and perhaps that I had controlled, without direct force (or threat of same), a non-output variable

(namely, which reference signal is currently being ACTIVELY matched) associated with you.

I suspect that you will say that you were ACTIVELY controlling for not getting burned both before and after I yelled, and that my yelling was a disturbance which resulted in your acting in the manner (attempting to exit) necessary to counter the disturbance. (Note, however, that it might be an interesting conundrum if you claimed that, in another instance, somebody was controlling for drinking a Coca-Cola BEFORE they ever knew about Coca-Cola! Also note that we presumably are quite busy ACTIVELY controlling all of those potential perceptions which might come in handy to control at any moment!) If so, then the situation is a (successful) "manipulation" as I described it in the Prolegomenon, assuming that I yelled so as to control for my perception, say, of you leaving the theater.

My problem is with your downplaying the interactivity and, especially, significance of manipulation. You say that the manipulee "ignores" the manipulator. But this is simply not true in general, even though it might be true sometimes. Complex interactions are the rule, not the exception, between con men and marks, for example. There is a complex "dance" going on between influencer and influencee so as to maintain the controlling relationship. In many such "dances," control can even go BOTH ways.

Furthermore, your claim that manipulating is "just diddling" seems purely ideological, in the sense that it ignores the potential effects of some examples of manipulation (for ill OR good) on the manipulee. It seems to me that you are going to extremes to defend an extra-PCT, imported notion of extreme autonomy and anti-environmentalism. PCT itself says that one's actions (output variables) are conjointly determined by the person's reference signals AND disturbances. It does NOT say that either the reference signals or the disturbances are to be construed as in some way more important or significant in producing the actions. You are sounding like Skinner turned upsidedown -- he (wrongly, I think) overemphasized environmental determination, while you are overemphasizing organismic determination. True, one's INANIMATE environment doesn't control one's variables (of ANY sort), but one's ANIMATE environment can control some of one's variables IN WAYS WHICH RESULT IN IMPORTANT CONSEQUENCES TO THE CONTROLLEE.

Don't you see that, much (but, I'll grant, not all) of the time, people want to control others' ACTIONS, NOT others' DESIRES? That means that the slogan "you can't be controlled without overwhelming physical force" is, in practice, not such a Big Deal. We are being controlled and are controlling others on a daily basis without direct physical force -- with, I claim, hugely significant effects. To mention just one example: The physician who advises his patient wrongly (we have a God's-eye-view here) because he doesn't think it will matter much that he doesn't know what he is supposed to know about some esoteric matter, and besides, he needs to see another patient NOW; the physician is not employing overwhelming physical force -- he is manipulating the patient (who doesn't even suspect that the doctor might be wrong), who DIES. That's a trivial result? What's trivial is that "you can be controlled by overwhelming physical force." Of course you can! But the notion that your outputs can be controlled WITHOUT direct force is subtle, practically important, not at all obvious with regard to its potentials and limitations, and so deserving of exploration far beyond rubber-banding, in my opinion.

The pervasiveness of manipulation argues for a theoretical account, based on PCT, which isn't biased by pre-judging the enormity of its effects so as to support a "but (if you don't beat me up) you really can't hurt/help me -- only I can do that" mentality.

>Dennis Delprato (920903)

>Greg, I don't get the point of all the quotations from Fred Skinner
>on control. You sum up: "In short, for Skinner, the results of one
>person's responses to his/her environmental stimuli (which CONTROL
>him/her) can be another person's CONTROLLING stimuli." Do you
>mean he is a strange bedfellow for himself? Seriously, what is
>the message I am missing?

The "strange" referred to some of Skinner's ideas (in particular, those on "control") being so different from some PCT ideas as to be, I suspect, mutually incomprehensible to parties on either side who aren't extraordinarily broad-minded. The "bedfellows" referred to some of Skinner's ideas (in particular, those on "statistics") being similar to some PCT ideas.

>Skinner's thinking is more advanced than that of the typical
>psychologist, hopelessly behind that of others, and he is
>extremely careful in his selection of words and their organization.
>The effect of the latter is that it is virtually impossible to pin
>him down--he's slick.

Another point I was trying to make is that he wasn't dumb. (Or, to belabor a point, inconsistent: once again, because he said that one's RESPONSES to controlling stimuli can provide others' controlling stimuli. This ploy makes Rick Marken's quest for a "paradox of Skinnerian 'control'" hopeless.

>Too bad Runkle cited not one (as I recall) Skinner-inspired paper on
actuarial
>vs. single-subject research. Don't know where Phil has been since 1938.

You mean Phil RUNKEL. I, in turn, don't know why it took so long for anybody to follow up on Skinner's ideas. Guilt by association, perhaps?

Greg

Date: Thu Sep 03, 1992 9:03 pm PST
Subject: Re: Skinner & Control; More cooking; variaboili

(ps 920903.2200)

[From Bill Powers (920903.0800)]

--

Penni Sibun (920902.1400) --

>well, of course, to interactionists, institutions are neither in the
>head nor in the environment, but.... this list, csgl, is an
institution >cause we are all participating in it. it's not my head,
it's not in >your head, and it's not out there in the environment

somewhere just >sitting around.

My problem with institutions is that they're nouns, whereas what goes on between organisms are processes and interactions -- verbs.

yes, i agree. that's suggested by ``participation.''

What

something like CSG-L "is" depends on who's looking at it. For me, it's an ongoing conversation with a person to claims to be Penni Sibun, someone who uses the name Rick Marken, a writer who expresses himself very much like Chuck Tucker, and so forth. The only real person on this net is me. From someone else's viewpoint, I'm imaginary. What the net "is" depends on each user's conception of it, and that conception isn't out there in the world. It's in a head.

well, i don't suppose solipsism is very useful. unless you want to assume that everything is imaginary, then this list is something real. its existence depends on our actively maintaining it.

```
>``...we shall...permit several possible `next states' for a given
>combination of current state and input symbol. the automaton
>[machine]...may choose at each step to go into any one of these legal
>next states; the choice is not determined by anything in our model,
>and is therefore said to be _nondeterministic_.'
```

Doesn't your quote say explicitly that it is the machine's problem?
"The automaton [machine] may choose at each step to go into any one of these legal next states..."

I think that what the quote means is that it isn't the programmer's problem -- that is, the simulation or model or whatever is making its own choices based on current experience, rather than having those choices programmed in from the start. But if you include the state of the environment and the criteria for choosing, then the automaton is deterministic, even if the programmer didn't determine its choices. The only way I can see to produce a nondeterministic outcome is to make the choice using random numbers or the output of a Geiger counter.

well, i probably confused you by saying ``machine' '--let's stick to ``automaton.' at any rate, neither is a program: an automaton is a description, a theoretical abstraction. one can perfectly rigorously say whether an automaton is deterministic or not; i gave the def. above. determinism does not describe what the automaton does, it describes how it is built.

i recommend the book i cited for learning about computational theory. interestingly, chomsky can be blamed for a lot of it.

```
>in section ``ontology of cooking tasks,' a&h say: ``a _history_ is a
>function from natural numbers (representing `time') to world states.'
>roughly, a history is a record of what happens during a run of toast.
```

This still doesn't answer my question. If an operation like "put pan

on burner" is carried out, I assume that the history would record the pan first in some other place, then after the operation, the pan on the burner. What I was asking was whether the position of the pan was recorded in this history over the whole trajectory between "on the counter" and "on the burner," or whether just the end positions were

the natural numbers are the positive integers starting from 1, so a function from natural numbers to world states implies discreet rather than continuous time.

penni

Date: Fri Sep 04, 1992 8:33 am PST
Subject: The phenomenon of control

Dag- could you send me the info & demo too? I need all the insight I can get.
Thomas B. Baines
ARGONNE NATIONAL LABORATORY
EID900
9700 South Cass Avenue
Argonne, IL 60439

Date: Fri Sep 04, 1992 8:54 am PST
Subject: Deja vu, PCT promotion

[From Rick Marken (920904.0930)]

Dennis Delprato -- Thanks. I received the programs and accompanying materials. I have not looked them over yet, in detail, but I will this weekend. I have the feeling, however, that the problems you had with the program are most likely caused by bugs in the compiler (rascal) than in the very simple code I saw (rascal is nearly the same as pascal).

Dag Forssell (920903-2)--

Thanks. I feel much better now.

>You are one of the very few who likes to think for yourself based on
>the evidence, instead of conveniently accepting what your teachers told you
>and what the abundant literature of the prestigious association tells you.

And you are a very astute judge of character.

>So when you say:

>>What is it about PCT? Why don't people like it? Why do they prefer

>you know the answer.

Yes, I do. Of course. I know that people are controlling for system concepts that are often fundamentally inconsistent with PCT. So PCT is just a disturbance.

But I am a control system too. And I am controlling for (among other things) being part of a happy community of scientists exploring the nature of control in humans and other systems. I am part of that community already -- but it's rather small (relative to my reference for what its size should be, of course). I want a larger group because I need help from people who appear to have great skill and energy (people like Chapman, Agre, Brooks, Beer, etc). There are, in fact, people with quite formidable skills who are "doing the work" of understanding life from a PCT perspective. But there could be more -- and it hurts to see people (like Chapman, say) squandering their talent on ill-considered nonsense. I know that I cannot control a variable whose value depends on the settings of reference signals in other people; that it, I know I can't MAKE people understand and love PCT. But, being a control system myself, it's hard to avoid reorganizing (isn't that what ranting is?) in an attempt to do SOMETHING that might end up with me perceiving myself in my dream community of scholars. But I must admit that, since I first got into PCT (in about 1979) the community of interested (and CAPABLE) PCT scholars has grown -- SLOWLY, but steadily. And it's nice to have you as a part of that community.

Best regards Rick

Date: Fri Sep 04, 1992 9:18 am PST
Subject: control, manipulation, influence, etc.

[From Bill Powers (920904.0700)]

Greg Williams (920903 - 2) --

The problem with reading all those quotes from Skinner (for which labor many thanks!) is that one mistakenly tends to put meanings into his words that make common sense (and in our case that jibe with PCT). Skinner arranged his choices of words this way deliberately. Common sense generally sides with control theory. But it is possible to talk what is heard as good control theory in words that are defined by the speaker in absolute contradiction with control theory. This makes it seem that the speaker was ahead of his time, when in fact he was crusading against the purposive overtones of common-sense descriptions in part by trying to pre-empt the language.

He said:

>I have used technical terms in making a technical point. I have preferred a
>technical term elsewhere when it could be used at no great cost.... But
>elsewhere [in this book] I have freely used the lay vocabulary while
>accepting the responsibility of providing a technical translation upon
demand..

An example of this:

>In traditional terms, one person arranges positive or negative contingencies
>in order to create interests, provide encouragement, instill incentives or
>purposes, or raise consciousness in another person. In doing so, he brings
him
>under control of various features of his environment.

You would think from this that people can have interests, feel encouragement, seek incentives and have purposes, and be conscious, wouldn't you? But if you asked Skinner what he meant by each of these terms, he would describe them "technically" in terms of the environmental causes of the behaviors from which we infer these inner states. In other words, he did not believe that these inner states were causal in any sense. At best they were epiphenomena. In the above quote, he links his technical terms, positive and negative contingencies, with common-sense terms, making it seem that the one flows naturally into the other. This can seduce the reader into believing that interests and purposes are the sort of thing that contingencies produce. And the swift passage across "positive or negative" simply asserts a characteristic of contingencies and hurries on to the main point, having left these immensely important little adjectives to be taken for granted.

Also, although Skinner often uses the term "control" in contexts where we would read it as "negative feedback goal-directed control of perception," the fact that he always, in the end, referred the control back to the environment shows that he could not have meant the term in that way. The environment is simply a collection of stimuli that have reinforcing properties and which have the capacity to control behavior -- which we know is not possible for anything but a purposive environment.

We must also recognize that Skinner consciously bypassed the problem of variable behavior producing consistent results. He did this by EXPLICITLY defining behavior -- responses -- as classes of outcomes, not outputs. This means that when he spoke of external agents or things controlling _behavior_, he had to mean controlling outcomes, not actions. We know from control theory that this is precisely what external agencies can't control, except by winning conflicts. While Skinner couldn't see the contradiction here, we can.

Here's another problem:

> A physical world generates both physical action and the physical
>conditions within the body to which a person responds when a verbal
>community arranges the necessary contingencies.

While he never says so out loud, he is implying that the physical world generates _specific_ physical actions and conditions, those particular effects that we observe as systematic behavior. This strange directedness of physical effects is handled in the same way he handles the problem of variable actions producing consistent effects. He defines a stimulus not as a specific physical situation or energy input, but as a class of situations defined by their effect on behavior! So the actual physical situation can vary widely, but if the same behavior persists, by Skinner's definition the "same stimulus" exists.

Yet Skinner always spoke as if he were simply applying basic principles of "science" to an analysis of behavior. He blithely invoked mysterious nonphysical properties of the environment to make this view sound plausible. All of the critical questions were begged. But he was so clever at the use of language that he could hide these gaps in his argument and make it seem that he was simply applying logic in an impersonal and rigorous way.

Another thing that Skinner did very well was to slip in non-sequiturs that easily passed without notice:

>In self-management the controlling self is different from the controlled.
>But all selves are the products of genetic and environmental histories.

Wait a minute: even if all selves are products of genetic and environmental histories (which applies, of course, to everything in the universe, not just selves), what does this have to do with controlling selves and controlled selves? Control systems are products of genetic and/or environmental histories, but they still control. This is like dismissing behavioral theories by saying that they all come down to physics and chemistry in the end, so why bother? It makes a hell of a difference WHAT the physics is and WHAT the chemistry is. In fact that makes ALL the difference.

If Skinner had not dismissed a controlling self a priori, he might have come across the concept of hierarchies of organization. But he did not. Instead he made it seem as if an irrelevant statement was a refutation of the concept of a controlling self, by prefacing it with "But". He got away with this a lot.

Skinner was also dishonest:

>It would be absurd for the behaviorist to contend that he is in any way
>exempt from his analysis. He cannot step out of the causal stream and
>observe behavior from some special point of vantage...

But Skinner constantly did this. He knew something about the rest of us that we did not know. There was no question of his view of behavior being simply a product of random environmental influences; if he had really believed that, he could not have argued against anyone else's view because that would be like saying the other person should have been exposed to different contingencies.

Skinner's view admits of no concept of truth or correctness of an idea. Yet Skinner always spoke from conviction, and he did not hesitate to criticize other views. Trying to sound nice, he advocated manipulating contingencies so that the world would be a better place, but he never admitted that non-scientists could pick their own definitions of a better place, and it never occurred to him that his idea of a better place might seem like Hell to someone else. There was always a layer of attitudes in Skinner that remained a secret, mainly from himself. Perhaps it is unfair to call him dishonest. But if he was not dishonest, he was singularly unaware of his own assumptions.

Much of what Skinner criticized about other approaches to behavior theory was worth criticizing. Trait psychology and statistical explanation were correctly assessed by him, in my view. But Skinner often used the obvious wrongness of other approaches as a way of making his seem right by contrast. It just doesn't follow that if you see what's wrong with the other fellow's idea, your ideas are automatically right. But that's how Skinner used these arguments.

Greg Williams (920903 - 3) --

>1. Is controlling a person possible with THREAT of force (which might
>be included in what you mean by "direct use of force")?

Yes, I think so. It doesn't even have to be an offered threat; it can simply be a perceived threat. A protester might see cops in the vicinity, and think "If I do what I was planning to do, they'll throw me in jail. So I guess I'll do something a little less spectacular." So you could say (imprecisely) that the cops, just by being visible, have controlled the protester's choice of goals (higher-level behavior). You could say this even if the cops were really on their way to some other location and knew nothing of this protester. It's the perceived threat that matters. If the cops threaten to arrest the protester, but the protester doesn't believe they actually will, the threat has no effect. Until it's carried out, of course.

>2. Suppose I yell "Fire!" when we're both in a theater. You (being prudent) >join the crowd heading for the exit. Were you, BOTH before and after I yelled, >ACTIVELY controlling for not getting burned (or some such)?

What does "actively" controlling mean when there's no disturbance? If you ask whether my reference level for being burned was zero before and after you yelled, the answer is yes. If you ask whether the pathways in the associated control systems existed, again the answer is yes with some reservations because of the possibility of switching connections. If you ask whether I was actively countering a perceived or imagined disturbance before you yelled, the answer is "No."

Did your yelling "FIRE" affect my behavior? Certainly; I took it for an honest report (this time). I was avoiding getting burned, or so I thought. Did your yelling "FIRE" control my behavior? Well, to know that we'd have to have some more details. Suppose I was one of those who didn't start toward the exit. How would you alter your yelling of "FIRE" to counteract that error between what I did and what you wanted me to do? If you did try something else, and it worked, then you would have been controlling my behavior. But even then, you couldn't control the _outcome_ I was controlling for by varying my position. That is, you couldn't make me want or not want to get burned by yelling "FIRE".

I still think you're confusing "control" with "affect" in many of your examples. You can't tell whether control is present by looking at a snapshot. You can't distinguish control from mere influence by looking at apparent cause-effect pairs. You can't identify control by looking at a single action. Control is a process that takes time and involves relationships that persist through time.

>My problem is with your downplaying the interactivity and, especially, >significance of manipulation. You say that the manipulee "ignores" the >manipulator. But this is simply not true in general, even though it might be >>true sometimes. Complex interactions are the rule, not the exception, between >con men and marks, for example. There is a complex "dance" going on between >influencer and influencee so as to maintain the controlling relationship. In >many such "dances," control can even go BOTH ways.

While all control of one person by another involves interactions, not all interactions between people (however complex) involve control. Two people dancing in the opening of a narrow door, trying to get through in opposite directions, are not controlling each other; they're simply trying to get through the door. Their control systems, taken as a single system, are

oscillating. If one of them is deliberately and systematically stepping in front of the other to prevent passage, the other will realize that after a time and take action to keep from being controlled: a vigorous punch in the snoot, for example, would clear the way if the other guy is littler than you.

>Furthermore, your claim that manipulating is "just diddling" seems purely >ideological, in the sense that it ignores the potential effects of some >examples of manipulation (for ill OR good) on the manipulee.

I don't ignore that at all. If the effect of manipulating output begins to matter to the manipulee -- if it disturbs something that IS under control -- the manipulee will push back, and the situation will be converted from manipulation into conflict.

>It seems to me that you are going to extremes to defend an extra-PCT, >imported notion of extreme autonomy and anti-environmentalism.

Hey, I like spotted owls a whole lot better than I like people who go around cutting down beautiful old forests.

And anyway, what's "extreme" about autonomy? When you track down the causes of behavior, you don't end up back in the environment, but at the highest levels, and if you try to backtrack farther than that, you end up with the whole species' struggles to stay alive and in control, over the last 3 billion years. That sounds perfectly reasonable to me. Conservative, even.

>PCT itself says that one's actions (output variables) are conjointly determined >by the person's reference signals AND disturbances. It does NOT say that either >the reference signals or the disturbances are to be construed as in some way >more important or significant in producing the actions.

Oh, but the reference signals are far more important than all the disturbances put together, because they define what constitutes a disturbance. It doesn't work the other way around. Anyway, what matters to the organism is not its action, but the outcome, and when there's no immovable obstacle or irresistible force, the reference signal determines the outcome quite independently of the environment. The environment has very little to say about the outcomes that matter to the organism. That's what the actions are for: to make sure it doesn't.

>You are sounding like Skinner turned upsidedown -- he (wrongly, I >think) overemphasized environmental determination, while you are >overemphasizing organismic determination.

Skinner thought that environmental influence was determination; I'm trying to avoid that mistake. Even the effect of disturbances on actions is not deterministic; the reference signal can change, too, which means that in the long run nobody can predict the effect of the disturbance. Particularly not if there are other disturbances too, as there usually are. The only deterministic aspect of the organism-environment interaction is the effect of the reference signal in determining what the input variable's state will be within narrow limits, and that holds true only when external interference isn't too radical.

>True, one's INANIMATE environment doesn't control one's variables (of
>ANY sort), but one's ANIMATE environment can control some of one's
>variables IN WAYS WHICH RESULT IN IMPORTANT CONSEQUENCES TO THE
>CONTROLEE.

But so can inanimate variables have important consequences. If you're caught in an avalanche, that will have important consequences. The severity of the effect isn't what makes the difference. It's whether the effect is intended, and whether changes in the effect give rise to opposing changes in the manipulation that is producing it. Only that is control. People bump into each other all the time, probably more often than they bump into the environment. In these encounters there are many mutual effects. Sometimes there are unforeseen effects that turn out to be extremely inconvenient, like lethal. People adapt their own modes of control to avoid the worst of these effects -- not instantly, but over time. Sometimes they don't succeed. Usually they do. And they make these adaptations to suit their own higher intentions, not someone else's.

>Don't you see that, much (but, I'll grant, not all) of the time, people want
>to control others' ACTIONS, NOT others' DESIRES? That means that the slogan
>"you can't be controlled without overwhelming physical force" is, in
>practice, not such a Big Deal.

I'd say that most of the time, people are controlling for their own perceptions, and treat the actions of others simply as disturbances to be counteracted or avoided. I'd say that most people don't even think about controlling other people's actions -- that is, they don't systematically vary their own actions in order to keep someone else producing a given action. The other's action is usually just an incidental part of what one is doing or wants to happen. When I buy toothpaste, I hand the clerk my money and expect to get my change, and the toothpaste, back. I don't care how the clerk does that, or whether another clerk does it instead. My reference level is for the outcome that matters to me. It's not for the clerk's behavior.

>We are being controlled and are controlling others on a daily basis
>without direct physical force -- with, I claim, hugely significant effects.

We are AFFECTING AND BEING AFFECTED on a daily basis. That's not control. When General Motors lays off 5000 people, it doesn't do this in order to put any person out of work. It is laying them off so it doesn't have to pay them salaries, so it can become leaner and meaner, with the consequence of course that it is removing buying power from the hands of its own customers, but businesses in America never look that far ahead. These are hugely significant effects, but they aren't control.

There is control involved in such actions, at another level. It's backed up by a credible threat of overwhelming physical force, called The Law. A laid-off worker can't show up for work the next day, and at the end of the week demand a paycheck. There is no way that worker can continue to get a paycheck from that company -- no way, that is, that will not end up with physical force being applied to prevent it. The worker might counter that threat by bringing a gun, but that would be hopeless, too, in the long run. The force available to The Law is greater than what any individual can resist or overcome. There are too many people on the side of the law.

It isn't that I don't believe that control exists in our society. I do believe it. I believe it is part of the fabric of our society, and all the others I know of in this last part of the 20th Century. It is true control. It's not diddling around manipulating people's actions. It rests solidly on the ability to seize a person and throw him or her in a cell with bars, and to kill the person in the process if necessary, in total disregard of that person's attempts to remain free. It's the kind of control you see on a football field, or in a drive-by shooting, or in torture of political dissidents, or in shelling a city or loosing a firestorm of bombs over an army for two weeks.

This is the only kind of social control that I consider serious enough to worry about. It shows a complete ignorance of human nature. It is just as if people can't distinguish living systems from rocks.

It never crossed Skinner's mind that his control over his trained animals depended entirely on his size and strength relative to theirs. Skinner could reduce a rat's weight to 80 percent of the free-feeding weight just by seizing the rat, putting it into a cage from which it could not escape in search of food, and not feeding it. He could force the rat to press a bar by lifting it up and putting it in another cage in which the only way to get food was by pressing the bar. The "control" of the rat's behavior by operant conditioning was an effete and polite little exercise in comparison with the brute force that made the experimental procedure possible in the first place. This is why you hear about operant conditioning being used so often in prisons, mental hospitals, and "detention centers" in which the contingencies -- the rules -- can be set up without fear that the inmates can refuse to live by them. Play the game or it's off to maximum security conditions, walking or being dragged. This underlying threat of physical force, made credible by frequent application, makes a mockery of any attempts to "influence" behavior by more theoretical means. The only reason any method works in prison is that the inmates know that if it doesn't work, they will suffer the consequences.

So I am not naive, Greg, about social control. I just want the conversation to stay focused on the real problem, which comes from not understanding how people actually work. Our social system needs an overhaul. We need to stop training people from birth to treat everyone else as an enemy to be overcome by force. We need to develop a little respect for the will of others. We need to figure out what people want and try to give it to them, instead of withholding it to gain an advantage over them. We need a conception of life together that goes beyond competition, swaggering, and bullying. The message of control theory is that yes, you can control other people, but no, you can't make them hold still for it, and no, you're not going to get what you really want from life in that way.

I believe passionately in autonomy, not as a recommendation but as a fact. It is the nature of an organism to be autonomous. And it is the nature of an organism to control. When the one is invalidated by the other, we must consider this to be a malfunction, whether it is one person or a whole society that is destroying its own means of survival.

Best, Bill P.

Date: Fri Sep 04, 1992 9:18 am PST
Subject: AI stuff

[From Bill Powers (920904.1100)]

Penni Sibun (920902.1400) --

>>What the net "is" depends on each user's conception of it, and that
>>conception isn't out there in the world. It's in a head.

> well, i don't suppose solipsism is very useful.

That's too easy an answer. Control theory is not about solipsism. It just recognizes that we don't all experience the same environment, so whatever you say about the environment is probably not true for everyone else. I don't doubt that there's an environment there. What I do doubt is that our perceptual representations of it are isomorphic to it.

>well, i probably confused you by saying ``machine''--let's stick to
>``automaton.'' at any rate, neither is a program: an automaton is a
>description, a theoretical abstraction. one can perfectly rigorously
>say whether an automaton is deterministic or not; i gave the def.
>above. determinism does not describe what the automaton does, it
>describes how it is built.

OK, I think that's what I said. The automaton itself could behave in unpredicted ways while still being deterministic. I suppose your definition would hold even if the automaton is deterministic, but one of its computing elements computes a random action when it's operating.

Automata that are theoretical abstractions aren't very interesting, are they? I prefer mine to be "concrete-situated."

>the natural numbers are the positive integers starting from 1, so a
>function from natural numbers to world states implies discreet rather
>than continuous time.

OK. The problem with discrete time is that it rules out real physical phenomena unless it's handled in a way I haven't seen in any AI models. Physical phenomena take place continuously, and their properties determine how they behave through time. If you apply a force to an object at time 1, how far will it have moved by time 2? If there's no link between discrete time-points and the underlying continuum, there's no way to tell. How far it moves will depend on the clock speed of your computer. You have to define a dt in real seconds, so you can derive distance moved from the double integral of force/mass TIMES DT. If you make one computing cycle equal to dt seconds, like 0.01 sec, you'll get one final position. If you make the computing cycle equal to a longer dt, like 0.1 sec, the object will move farther in the same number of computing cycles.

This problem is present any time that the computer has to deal with processes that actually are continuous. When it's ignored, as it usually is, the results of a simulation don't really mean anything. Not if you're trying to model a real system, that is.

I've written a persuasive letter to Chapman, by the way. It's worth another try.

Best, Bill P.

Date: Fri Sep 04, 1992 9:35 am PST
Subject: Concepts & models; Busses

[From Bill Powers (920904.1130)]

Martin Taylor (920903.1900)]

RE: concepts and models

>I don't have or expect to have a specific model of alerting systems. I can
>think of many. But the situation is like that of PCT itself. The basic
>proposition of PCT, that "all (purposeful) behaviour is the control of
>perception" is incontrovertible. There are many models that might
instantiate
>it. So with the alerting systems. Given the basic proposition of PCT, their
>existence in an organism supplied with more sensory than motor degrees of
>freedom is incontrovertible. There are many models for how they might work.

It's not self-evident that left-over perceptions or control systems must have the function of "alerting." I don't dispute that there will be perceptions and systems not in use, because of the df problem.

The "concept" of control that I was talking about was the concept that behavior controls perception. This concept is seen as fitting experience: that is, any behavior you happen to notice turns out to be a perception, on closer consideration. Cut off the relevant perception and you can't do the behavior, particularly in the presence of disturbances.

The model is the CT model. It shows the kind of organization that's required to create the effect observed. With due regard to alternative but equivalent forms of the same model, it's the only known kind of system that can produce this sort of phenomenon. There's no alternative that I know of.

In the CT diagram, every block and every arrow is assigned a specific function. By instantiating these functions it's possible to construct a working model of a control system. Even in cases where we can't yet instantiate the functions, we can see that if we could, or if we could fake it, we'd have a runnable model.

This is what I'm after regarding your "alerting systems." You show them as control systems not connected to lower systems, but sending arrows over to the active control systems. What is the meaning of these connections? What KIND of effect are they supposed to be having on the target system, and what is there about their origins and destinations that makes this effect plausible? Are you including enough machinery in this model, even in block form, to enable it to do the things you say it should do?

>As for "controlling for variability," I was more musing than proposing,

>but if you want a proposal, you mentioned a while back that if you had
>two one-way control systems back-to-back to make a two-way one, and if
>they had the right (square-law?) control function, then the pair could
>be controlled both for gain and for reference level.

I can see that the variability of the perceptual signal relative to the reference signal would be changed by changing the loop gain, loosening or tightening control. But that is AFFECTING variability, not CONTROLLING FOR variability. In order to control for variability, the system that is varying the common-mode reference signal must perceive something about the lower system that constitutes its variability. What would that be?

-

Avery Andrews (920904.0950) --

>Re busses: fine. But as far as I can see, Chapman's architecture is
>consistent with this - as long as you have a finite system of
>categories sent along the bus, you can have a single wire coming off
>for each one.

But then for each category wire, you need a category-perceiver at the lower level. This just moves the problem down a level. Each category-perceiver has to receive all the elements that could be categorized, and respond to the one category it is designed to detect. There could still be an infinity of categories, with only those actually perceived having any effect in the system.

>PS did Harry Erwin really send something to CSGNet? I posted something
>of his yesterday, but I also lost a whole swag of unread messages 2
>days ago, so might well have missed a posting by him.

I thought the post was from him. Sorry.

Best, Bill P.

Date: Fri Sep 04, 1992 10:18 am PST
From: Penni Sibun
MBX: sibun@parc.xerox.com
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: The phenomenon of control

thanks.

penni sibun
xerox parc
3333 coyote hill road
palo alto, ca 94304

Date: Fri Sep 04, 1992 11:00 am PST
Subject: Re: The phenomenon of control

Dag,

You probably saw my 'intro' on the net a few days back. Anything along the lines of PCT in industry that you could throw my way would be much appreciated. By the way, I found your performance review proposal informative and helpful.

I've gotten much material from Ed Ford and Tom Hancock -- in fact, Tom gave me copies of your 2/4/90 & 2/11/90 drafts concerning Deming and control systems. I'm really struggling to make PCT a real mindset to some people (including myself), and eventually, to an entire organization here.

I would appreciate it if you could put me on your 'snail mail' list; please let me know what postage and handling fees would be appropriate.

Regards,
Ray

Ray L. Jackson
3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson
602-963-6474

Date: Fri Sep 04, 1992 11:07 am PST
Subject: Re: AI stuff

(ps 920904.1200)
[From Bill Powers (920904.1100)]

Penni Sibun (920902.1400) --

> well, i don't suppose solipsism is very useful.

That's too easy an answer. Control theory is not about solipsism. It just recognizes that we don't all experience the same environment, so whatever you say about the environment is probably not true for everyone else. I don't doubt that there's an environment there. What I do doubt is that our perceptual representations of it are isomorphic to it.

i didn't say control theory was about solipsism. you said everyone ``out there'' is imaginary. i don't think that has explanatory power. and i certainly don't say that perc. reps are isomorphic to the env.--shoot me as a cognitivist if i do! and remember, i'm not trying to talk about the enviroment. a message ago, we were agreeing that instituions were processes, not objects. you don't hold that processes are out there in the environment, do you?

OK, I think that's what I said. The automaton itself could behave in unpredicted ways while still being deterministic. I suppose your definition would hold even if the automaton is deterministic, but one of its computing elements computes a random action when it's operating.

automata are just math. when you say ``4 + 7 = 11'' it doesn't matter 4 what or 7 what or what it looks like when you've got the 11 together. that's just not part of the mathematical description.

Automata that are theoretical abstractions aren't very interesting, are they? I prefer mine to be "concrete-situated."

no. but yr original question had to do w/ terminology, and i tried to explain that terminology. i do think that, as w/ any other field, to really understand what the practitioners are talking about, you need to have some idea of the foundations. theory of computation is one such for ai.

>the natural numbers are the positive integers starting from 1, so a
>function from natural numbers to world states implies discreet rather
>than continuous time.

OK. The problem with discrete time is that it rules out real physical phenomena unless it's handled in a way I haven't seen in any AI models. Physical phenomena take place continuously, and their

right. ai models are like this, basically cause the run (or are supposed to) on digital computers. i don't think you'll find anyone that would argue that there's something missing by not being able to model continuous time. the issue is more how close and/or useful your approximation is.

I've written a persuasive letter to Chapman, by the way. It's worth another try.

what did you try to persuade him of? just curious.

cheers. --penni

Date: Fri Sep 04, 1992 12:36 pm PST
Subject: Phenomenon - Demo

[From Dag Forssell (20904-1)]

Tom Baines, Penni Sibun, Ray Jackson:

Mailing information packet this weekend. Glad you asked.

Dag

Date: Fri Sep 04, 1992 12:42 pm PST
Subject: Re: Concepts & models; Busses

[Martin Taylor 920904 15:40] (Bill Powers 920904.1130)

Very quick, or it will be 3 weeks till I answer...

>It's not self-evident that left-over perceptions or control systems
>must have the function of "alerting." I don't dispute that there will
>be perceptions and systems not in use, because of the df problem.

Self-evident is a funny term. It's self-evident to me, in that I can find very little that is more "atomic" than the simple statement of the df situation, from which the existence of alerting systems follows directly, given the basic axiom of PCT (not Greg's 6 axioms, but the axiom that behaviour is the control of perception).

Perhaps the problem is that we have a different sense of "alerting." The way you use it, it sounds as if you make some connection with consciousness. All I want to say is that the presently uncontrolled percepts are monitored, and the consequences of this monitoring are that when something relevant happens, control moves to percepts that are marked as possibly having become important to control. That means either that there were already "spare" degrees of freedom that were not being used actively for control (I suspect this is the normal case) or that some degrees of freedom that were being controlled are dropped in favour of the ones that were "alerted." I take "alerting" to describe the marking of df as being probably useful to control, and don't put any more on it than that.

In the diagram to which you refer. the arrows were meant to indicate this kind of non-specific pointer function, not to specify a mechanism. On other occasions I have proposed gain functions with dead zones as being a way that ECSs can perform the alerting function on themselves. In that case, there is no overt mechanism, but the df shift occurs as a consequence of conflict between the "alerted" ECS and previously active ones with correlated perceptual functions. There are probably dozens of other possibilities that come within my concept of the necessity of "alerting."

As a related topic, it seems to me probable that there are usually far fewer output degrees of freedom used in active control than are available, because the use of all available ones seems rather like a very high workload condition. If this is so, then the activation of an output degree of freedom as a consequence of an alerting situation need not take a degree of freedom away from what was already being controlled. It also implies that output degrees of freedom may, in a way, "drift" in and out of use by the actively controlling ECSs. There's a difference between a cat standing guard by a mousehole and the same cat in a similar attitude, resting. The guarding cat is actively controlling a lot, without moving, whereas the resting cat is not.
 =====

>In the CT diagram, every block and every arrow is assigned a specific
 >function. By instantiating these functions it's possible to construct
 >a working model of a control system. Even in cases where we can't yet
 >instantiate the functions, we can see that if we could, or if we could
 >fake it, we'd have a runnable model.

This is the construct I called "symbolic" PCT, and that I see as the reason you had problems with your reorganization simulation in flipping the signs of output-to-reference links in an active controller. Symbolic PCT is surely easier to program than is distributed PCT, but I doubt it is as robust or as true to life.
 =====

>I can see that the variability of the perceptual signal relative to
 >the reference signal would be changed by changing the loop gain,

>loosening or tightening control. But that is AFFECTING variability,
 >not CONTROLLING FOR variability. In order to control for variability,
 >the system that is varying the common-mode reference signal must
 >perceive something about the lower system that constitutes its
 >variability. What would that be?

Again musing rather than proposing...Imagine a perceptual function consisting of a leaky integration of a rectified time-difference signal. If I remember my rusty electronics--

---->|--C1-----

```

      | |
      R C2
      | |
  
```

where $C2 \gg C1$ (I think I put the rectifier in the wrong place)

One can imagine other perceptual input functions that might do the job. They don't have to be complicated, and they can combine the inputs from many lower-level ECSs, either before this circuit or after it. depending on whether they perceive something like the variation of the average or the average variation of a lower-level ECS.

Martin

Date: Fri Sep 04, 1992 1:03 pm PST
 Subject: Sciences and Ideologies - I

From Greg Williams (920904)

>Bill Powers (920904.0700)

>The problem with reading all those quotes from Skinner (for which >labor many thanks!) is that one mistakenly tends to put meanings into >his words that make common sense (and in our case that jibe with PCT). >Skinner arranged his choices of words this way deliberately. Common >sense generally sides with control theory. But it is possible to talk >what is heard as good control theory in words that are defined by the >speaker in absolute contradiction with control theory. This makes it >seem that the speaker was ahead of his time, when in fact he was >crusading against the purposive overtones of common-sense descriptions >in part by trying to pre-empt the language.

I hope that netters will not misunderstand my purpose in posting the Skinner quotes. I did not want to make it seem as if Skinner was a control theorist. Rather, as I said to Dennis Delprato in my previous post, I wanted to show that Skinner's ideas are essentially -- radically -- in opposition to those of (at least some) PCTers. (I also wanted to counter the claims that Skinner was (a) stupid/incoherent and/or (b) inconsistent.) To be more precise about my claims about the radical opposition (I'm beginning to think of it as a complementarity, wherein each pole is self-consistent and, to its supporters, seems complete and in a sense "finished," yet each is only part of a larger, more complete story) between Skinner's ideas on "control" and -- to make it concrete -- Bill's ideas:

I think Skinner's IDEOLOGY of radical environmentalism wrongly (because he can't get as far as he could otherwise in explaining behavior) downplays what is happening in the present moment. I also think Bill's IDEOLOGY of radical autonomy wrongly (because he can't get as far as he could otherwise in explaining behavior) downplays what has happened up to now. These two ideologies are complementary -- they are actually two mutually incompatible world-views which are each self-consistent but, I think, fundamentally inadequate to the task of explaining behavior.

So what is adequate to the task? Skinner's SCIENCE, as opposed to his IDEOLOGY, is certainly NOT adequate. I believe he allowed his ideology to sway his notions of an adequate science of behavior to the degree that he was self-defeating.

PCT SCIENCE -- again, not IDEOLOGY -- IS, I think, adequate. Skinner was model-less; PCT is based on models. And those models, aside from confirming folk notions of how organisms "work," are constructed in such a way as to make the two mutually incompatible "complementary" world-views described above fit together in a plausible way. I think the science could heal that ideological rift, and lead to a new sort of ideology noting the importance of BOTH history and the present situation. PCT SCIENCE is capable of doing this because it is a sufficiently DETAILED, MECHANICAL (by which I simply mean LAWFUL) model for explaining behavior. Yet there is the tendency of Bill (also Rick, and others) to see PCT science as supporting the ideology of radical autonomy and denying an ideology of radical environmentalism. I have little hope of changing their opinion on this; but others without ideological axes to grind might see some merit in my arguments. I would not like to see the potential of PCT science sacrificed on the altar of ideology, as Skinner sacrificed his "science of behavior."

So much for the outcome of my argument. It will take some time for my outputs to lead (I hope) there.

>>Greg Williams (920903 - 3) --

GW>>1. Is controlling a person possible with THREAT of force (which might GW>>be included in what you mean by "direct use of force")?

BP>Yes, I think so. It doesn't even have to be an offered threat; it can BP>simply be a perceived threat.

All right, then, I THINK we arrive at the following more complete formulation of your definition: "Controlling another person" means altering variables associated with that person so as to maintain certain of your perceptions as you desire. Another person's output variables (actions) can be so controlled (I have been calling this "manipulation") if you have appropriate knowledge (or adequate prediction) of the person's control structure, without the threat of or direct use of overwhelming physical force by you. Another person's control structure (reference signals) cannot be so controlled regardless of the threat (or perceived threat) of or direct use of overwhelming physical force by you. Nor can another person's CONTROLLED (by him/her) perceptions be so controlled without the threat (or perceived threat) of or direct use of overwhelming physical force by you.

Question: Do you have any problems with the above paragraph? I need to know whether it is OK before proceeding further. Thanks!

To be continued,

Greg

Date: Fri Sep 04, 1992 1:09 pm PST
Subject: Manipulation and models

[From Rick Marken (920904.1330)]

Bill Powers (920904.0700) --

>I believe passionately in autonomy, not as a recommendation but as a
>fact. It is the nature of an organism to be autonomous. And it is the
>nature of an organism to control. When the one is invalidated by the
>other, we must consider this to be a malfunction, whether it is one
>person or a whole society that is destroying its own means of survival.

Well, there's the frontpiece to my book.

That was a wonderful post, Bill . I particularly liked your succinct description of the "real life" lessons of PCT:

>We need to develop a little
>respect for the will of others. We need to figure out what people want
>and try to give it to them, instead of withholding it to gain an
>advantage over them. We need a conception of life together that goes
>beyond competition, swaggering, and bullying. The message of control
>theory is that yes, you can control other people, but no, you can't
>make them hold still for it, and no, you're not going to get what you
>really want from life in that way.

Each of these suggestions will be a disturbance to one system concept or another. I bet some people just can't accept the idea that a society can work without competition. Others will claim that non-contingent giving is the path to sloth. Others will say that a society without controllers will be chaos. Similarly, each of these suggestions will be consistent with one system concept or another as well.

What PCT offers (that goes beyond ideology) is a MODEL.

Capitalists can argue for the importance of incentives, competition and lack of control while communists argue for the importance of non-contingent giving, cooperation and strong controls. But control theorists can take out their model of hierarchical control systems and show what will happen IF people give without contingency compete without restraint, cooperate, provide incentives, and try to control. The point of doing PCT is to develop the model person (based on research) so that this modeling effort can be done with a high degree of coherence and fidelity. Then, instead of a fight between ideologies, it will just come down to whether people want to believe the model or not. People can hold out against the models (there are still "flat earthers" running around); but most people do seem to eventually come around, at least when the model is very good.

I would like to suggest that the debate about control of behavior could be settled rather quickly if we could sit down in front of a model hierarchical control system (I volunteer my spreadsheet model) and see what really happens when we try to control it. Eventually, we should be able to link our terminology to operations on the model. Then our arguments about whether or not control systems can be controlled could be aimed at the model (which can readily be changed, if necessary) instead of at one another (who cannot be so readily changed).

Best regards Rick

Date: Fri Sep 04, 1992 1:57 pm PST
Subject: Sciences and Ideologies - I

[From Rick Marken (920904.1430)]

Greg Williams (920904) --

> Skinner was
>model-less; PCT is based on models. And those models, aside from confirming
>folk notions of how organisms "work," are constructed in such a way as to
make
>the two mutually incompatible "complementary" world-views described above fit
>together in a plausible way. I think the science could heal that ideological
>rift, and lead to a new sort of ideology noting the importance of BOTH
history
>and the present situation.

Talk about synchronicity. I had just posted my last post on using models to heal our apparent rift on "control of behavior" and then I read this. Sounds good to me, Greg.

> Yet there is the tendency of Bill (also Rick, and others)
>to see PCT science as supporting the ideology of radical autonomy and denying
>an ideology of radical environmentalism. I have little hope of changing their
>opinion on this; but others without ideological axes to grind might see some
>merit in my arguments. I would not like to see the potential of PCT science
>sacrificed on the altar of ideology, as Skinner sacrificed his "science of
>behavior."

I don't want to see PCT sacrificed on the alter of ideology either! Help. Stop me before I sacrifice again! Tell me what I am doing that's wrong. How can I stop?

I swear. I am (at least, I think I am) only trying to describe (verbally) what I think is implied by the PCT model of behavior (at least, as far as control of behavior goes). I agree that you can control behavior -- via deception, disturbance to controlled variables (including changes in the feedback function -- ie operant conditioning), asking, etc. It's just that you can't do it arbitrarily (you might pick a variable to control that the other person is also controlling) or for long (the controllee might eventually catch on). That's all. My only ideology (that I know of) is that I believe (based on my understanding of the nature of control) that controlling is not a very good

way for controllers to interact, my own behavior notwithstanding. (As I always say to my kids, "do as I say, not as I do; better yet, do as your mother says; even better yet, do what you want").

Best regards

Rick

Date: Fri Sep 04, 1992 2:48 pm PST
Subject: RE: Sciences and Ideologies - I

Curt McNamara (920904.1730 CDT)

> [From Rick Marken (920904.1430)]
>
> My only ideology
> (that I know of) is that I believe (based on my understanding of the nature
of
> control) that controlling is not a very good way for controllers to
interact,
> my own behavior notwithstanding. (As I always say to my kids, "do as I say,
> not as I do; better yet, do as your mother says; even better yet, do what
> you want").
>

Really? Is this how you raised your kids? No time-outs,
loss of privelieges, etc.? Or was there something more?

Seriously, I would very much appreciate those of you who have
raised/are raising children would pass on your tips on "aligning thier
control structures" with those of civilization.

BTW it appears to me that some CSG-L folks don't believe in
manipulation because they think everyone's control structures are
perfectly aligned with their highest principles. In fact, I would
guess that none of us are actually organized this well. The
non-alignment is what allows manipulation to occur.

Curt

```
*****
Curt McNamara (mcnamara@mgi.com) | "Mistakes are part of the dues
Mgmt. Graphics, Inc. | one pays for a full life."
1401 E. 79th St. | Sophia Loren
Mpls., MN 55425 |
*****
```

Date: Fri Sep 04, 1992 3:00 pm PST
Subject: Robot Olympics--opportunity?

[Martin Taylor 920904 19:00]

Here's a forwarded posting from a local Toronto newsgroup. Some handy CSG-er
might want to follow it up.

=====

An update on:

The Second BEAM Robot Olympics and Micromouse Competition:

Ontario Science Centre,
Toronto, Ontario, Canada.
Apr 22 - 25, 1993

Sponsored by:

IEEE Canada, VSPANS Robotics, and the University of Waterloo.

The BEAM Robot Olympics is not so much a series of technological competitions as a chance for robot enthusiasts to present their designs to each other, the press, and the public. It is also an open forum for anyone who wants to get started in the field to compete and compare. Any and every robot will be considered so long as it does not come exclusively from a kit or store. Robots of similar ability will be pitted against each other in simple competitions, but generally robots will be judged on sophistication of behavior, novelty of design, efficiency of power source, and quality of hardware innovation.

Basically, if you built it, we'd like to see it.

The BEAM Olympics main events are (but not limited to) the following:

SOLAROLLER: A self-starting robot dragster race.
PHOTOVORE: Robots face a closed "world" and each other.
HIGH JUMP: Robot creature leaps, lands on feet.
ROPE CLIMBING: First up, first down, self-starting.
LEGGED RACE: Walking creatures run for the money.
INNOVATION MACHINES: Electronic chopsticks, for example.
ROBOART/MODIFICATION: Mechanical/electronic aesthetics that move.
ROBOT SUMO: Push/Bash an opponent out of a ring.
LIMBO RACE: How low can you build?
NANOMOUSE: A smaller and simpler form of the...
MICROMOUSE: Where metal mice race for aluminum cheese.
AEROBOT: 3D robot flying challenges.

The Robot Olympics thus features eleven formal competitions which range in difficulty from simple to complex. To this end, there is a guide available which contains competition rules, "get-started" instructions, discussion into the new science of Artificial Life (Alife), and full information for anybody who might want to run their own BEAM Robotic Games.

Fresh copies of the Guide are available now at a cost of \$10.00 Cdn, \$15 American [US funds], and \$20 international [US funds] for cost of copying and postage. Please make cheques or

money-orders payable to BEAM Robotics Inc. Competitors are asked to register and fill out a "behavior sheet" for their robot(s) and more competitions will be run based on the variety of robots who do show up. Those in a class by themselves are still eligible for major awards, and everybody will be included in the subsequent BEAM Olympic portfolio, documentaries, and videos.

All venues are open to the interested, young or old, so grab your soldering iron, raid the junk pile, and we'll see you there.

For more information on the BEAM Robot Olympics, other robotic competitions, and weird robotics in general, contact:

BEAM Robot Olympics
 c/o: Mark W. Tilden
 MFCE, University of Waterloo
 Ontario, Canada
 N2L-3G1
 (519)885-1211 x2454

E-mail: mwtilden@watmath.uwaterloo.ca

The rest of this article concentrates on general competitor guidelines and the entrance form. If you want to be put on the real-mail list, please complete and send in the entrance form by real-mail. If possible, include a picture of your competitor. Updates and further details will be posted to this and other relevant newsgroups as they become available.

Rules and General Guidelines:

These are the rules and guidelines for the Second BEAM Robot Olympic Games. Although the spirit of these Olympics stress few formal restrictions, these are some guidelines which every competitor should follow. If a design does not fit these parameters or may but you're not sure, please contact the organizers for a ruling. Consideration will be given to very innovative or imaginative designs. The BEAM Olympics is meant as an informal meeting of mind and robo-critter with the public and the media. It is hoped that everyone will abide by this spirit.

All those wishing to compete should fill out and send in a registration form for all robotic competitor(s) before Monday, April 5, 1993 to be included in the show dossier. Entries will be allowed after that, but they will not make that years dossier. A copy of the registration form is included at the back of this post. All who register will be put on a mailing list and kept informed of upcoming events.

GENERAL GUIDELINES FOR ROBOTIC COMPETITORS:

- All robotic devices must be either entirely custom built or a heavily modified toy. No commercial, store bought or kit robots will be allowed without heavy physical modification (modified or improved software is not sufficient modification). "Heavy" defined as permanent structural/electronic additions which extend, replace or enhance a functional aspect of the device (ie: replace batteries with solar engine, add functional arms, interface a unique touch/vision system, etc.). Any devices made from commercial construction kits (ie: Lego, Mecanno, etc.) must also feature obviously non-commercial, custom elements to indicate sufficient intent of innovation.

- All robotic entries should be self-contained or have an option where they can execute behavior without human intervention. Tele-operated mechanisms are allowable only if it is obvious they have been designed around some BEAM competition guidelines, or are for some autonomous task-oriented purpose (ie: tele-operated moon rover with retrieval claw). Commercial radio controlled kit models will not be allowed unless they have undergone severe technical modifications. Devices which do not feature some autonomous ability are allowed, but will lose critical style points.

- Any robotic entry may be disqualified if too large (ie: bigger than a standard upright refrigerator). The sole reason for this is that display space, access doors, and power are limited. Exceptions may be granted for exceptional, self-powered entries. Please contact the organizers for a ruling.

- Any robotic competitor which is obviously of mass-produced, commercial manufacture and/or performs an obviously commercial task shall be disqualified as a competitor. We don't want this to be a blatant advertising opportunity for major corporations, that's what trade shows are for. Corporate research and design prototypes are the exception but they must be represented by their designer at the competition, not by the corporation. However, robotics companies are allowed to hand out cards, flyers, posters, and device specifications.

- If a robotic entry must use wall current, it can use only one plug at a nominal amperage (120 VAC, 3 Amps maximum drain). Any robot which uses a combustion process will be discouraged for safety reasons, and will not be allowed to run inside the competition auditorium.

- For the most part, no "violent" robotic competitors will be allowed; that is, no competitor may have a functionally destructive capacity (ie: drills, cutters, soldering iron, flame-thrower, chainsaw, etc.) although decorative or whimsical elements along this line are allowed (ie: waterpistols). Likewise, any robotic competitor which may damage the competition courses, other robot competitors, organizers or audience will be disqualified. Robotic competitors may interfere with each other during the course of simultaneous runs (where rules per-

mit) so long as they do not violate the "no damage" rule. Anybody who justly feels his/her device could be damaged by another entrant will not suffer penalties for not competing in that trail. Likewise, entrants who accept that their designs could suffer major damage will also be respected (ie: No-holds-barred Robosumo). BEAM Robotics and affiliates take no responsibility for damage incurred by or afflicted on robotic devices, persons, or reputations during the course of competition.

- No robotic competitors can employ biological components (rats, chickens, lemon/potato batteries, jello, etc.) except by special permission from the organizers. Even then, no devices will be allowed which harm biologics in any way.

- Multitalented robotic competitors are encouraged, however such capabilities must be declared ahead of time on the registration form as part of the robot's behavior description.

- Partially finished robotic competitors are allowed (even encouraged) to compete so long as there is some function they can exhibit to show what they may eventually be capable of.

Again, if you are in doubt as to the validity of your entry, please contact the organizers directly at the University of Waterloo. Exceptions will be made for particularly imaginative entries.

Is all.

```
--
Mark Tilden: _-_-_-_-_-_-_- / (glitch!) M.F.C.F Hardware Design Lab.
_-_-_-_-_- | \ / \ / U of Waterloo. Ont. Can, N2L-3G1
|_-_-_-_-| \ / (519)885-1211 ext. 2454
"MY OPINIONS, YOU HEAR!? MINE! MINE! MINE! MINE! MINE! AH
HAHAHAHAHAHAHAHAHA!!"
```

Date: Fri Sep 04, 1992 8:58 pm PST
Subject: Controlling Kids

[From Rick Marken (920904.2030)]

Posting from home -- here is some leader.

Curt McNamara (920904.1730 CDT) in reply to:

```
>> [From Rick Marken (920904.1430)]
>>controlling is not a very good way for controllers to interact,
>>my own behavior notwithstanding. (As I always say to my kids, "do as I say,
>>not as I do; better yet, do as your mother says; even better yet, do what
>>you want").
```

says:

> Really? Is this how you raised your kids? No time-outs,
>loss of privelieges, etc.? Or was there something more?

Kids, like people, often do things that disturb variables we are controlling (like the loudness level in the room, the number of pieces of china on the floor, etc). We are usually stronger and smarter than kids so we can exert some arbitrary control. I did it with my kids occasionally, of course. What I didn't do was think that I was doing it "for their own good"; I was doing it for the reason I try to control anything at all -- for MY own good (to control the variables I wanted to control). I used control procedures (like those you mention) very rarely. In fact, the kids were generally wonderful -- my wife and I served mainly as helpers and tutors to the extent that we could understand (and tolerate) what the kids wanted to achieve.

I think parents have remarkably little to do with the way their kids turn out. I say this for both scientific and "religious" reasons. The scientific reason is PCT which shows that you cannot control behavior in any serious way, even when you have access to the system from birth. The religious reason is this: I can't believe god, even if he is the blood thirsty bozo portrayed in the bible, would actually let the behavioral development of EVERY child depend on the child rearing skills of the millions of idiots running around having them. Children, when they become people, are not nearly as idiotic as they might be expected to be. Just look at my incredible children; even Skinner had great kids.

Having said that, I should add a caveat. Although they cannot control their children (make them become doctors, lawyers, merchant men or thieves), parents can screw them up pretty good. The best way to do this is to try your best to control them -- ie. get them to behave in the way you want them to behave. When you continually try to get your kid to behave a particular way (assuming that the kid does not want to behave that way for whatever reason) you will get a kid who is constantly reorganizing in an effort to defeat the goals of your controlling. I watched this happen over a 2 decade period with my beloved brother; 20 years of efforts by well-intentioned parents to turn my non school loving brother into an intellectual managed to produce a man who can barely read.

It's a VERY ugly thing to watch.

The movie "Dead Poet's Society" may have been corny, but it HAPPENS. After seeing that movie, my mother told my wife and I that the father (of the kid who committed suicide) should have been a lot tougher on the kid -- not letting him stray from the father's intended "medical school" track. I've seen control up close and personal, Curt, and it ain't pretty when you're on the wrong end.

> Seriously, I would very much appreciate those of you who have
>raised/are raising children would pass on your tips on "aligning thier
>control structures" with those of civilization.

They will align them fine on their own. Cut 'em some slack.

>BTW it appears to me that some CSG-L folks don't believe in
>manipulation because they think everyone's control structures are

>perfectly aligned with their highest principles.

Not this CSG-L folk. I don't believe in manipulation of other control systems because I know that it is most likely to lead to CONFLICT. I would be happy to control people if the results were as consistent as those that I get when I control my computer or my piano. To the extent that you can manipulate people without creating conflict, go for it. But I think you will be disappointed in the long run. It's just not a good way to interact with people (or kids) because it is virtually certain to produce CONFLICT -- that's the yelling and screaming that you hear when you try to get the kid to do X(1) but they seem to want to do X(2).

You do need to exert control over kids sometimes; but making it a child-rearing POLICY is cruel and, almost certainly, catastrophic.

Best regards Rick

Date: Sat Sep 05, 1992 4:15 am PST

Subject: No (even) small sacrifice

From Greg Williams (920905)

>Rick Marken (920904.1430)

>I don't want to see PCT sacrificed on the alter of ideology either! Help.

>Stop me before I sacrifice again! Tell me what I am doing that's

>wrong. How can I stop?

Just keep thinking about what you are saying. Eventually it will dawn on you that the PCT model explicitly says that one's history (genetic and environmental) plays an important role in determining what one's current control structure is. (PCT doesn't say that one's past environment SOLELY resulted in one's current reference levels, but that it CONJOINTLY -- with one's control structure history -- determines the current reference levels.)

I think you agree with this in practice, but perhaps not in principle:

>Rick Marken (920904.2030)

>When you continually try to get your kid to behave a particular way (assuming

>that the kid does not want to behave that way for whatever reason) you will

>get a kid who is constantly reorganizing in an effort to defeat the

>goals of your controlling. I watched this happen over a 2 decade period

>with my beloved brother; 20 years of efforts by well-intentioned

>parents to turn my non school loving brother into an intellectual

>managed to produce a man who can barely read.

Thaaaaaat's historical (co-)determination of current control structures, folks!!! Stay tuned for a commercial from Behavior Mod, Inc., the folks with the motto: Give us your dumb, your weak, your tired... they're EASY to control! [PCT says why.]

>They will align them fine on their own. Cut 'em some slack.

Controlling for not altering variables associated with others is a form of control, isn't it? Such deliberate non-controlling-others control, I think, can also result in screwed-up persons, sometimes. Not because of conflicts within the non-controlled persons, but because their reference signals (unmodified by history of interactions with non-controllers) don't mesh with the desires of people around them (not just the non-controllers). Example: rich spoiled brats. Perhaps -- accepting the PCT idea that reorganization stops (with a new control structure) when the problem which started the reorganization is solved (meaning that one's environment is influential in "choosing" among possible new control structures proposed "randomly" (which might mean "truly" randomly or pseudorandomly) by the mechanisms within an individual -- there is an optimal level of "attempting to control others" with respect to their meshing with others' control structures? Perhaps; I suppose it depends strongly on the particulars of a given situation through time.

So keep it up, Rick, and together we can forge a broader ideology based on PCT science. History AND physiology....

Greg

Date: Sat Sep 05, 1992 9:49 am PST
Subject: Imagination, social control

[From Bill Powers (920905.0930)]

Penni Sibun (920204.1200) --

>i didn't say control theory was about solipsism. you said everyone
>`out there' is imaginary. i don't think that has explanatory power.

Imagination, in PCT, is a technical term and has explanatory power. A control system of level n is in the "imagination mode" when its output, instead of serving as a reference signal for systems at level n-1, loops back into its own perceptual function. Normally, sending the output downward as a reference signal results in the lower system matching its perceptual signals to that reference signal by acting on still lower systems. Since the higher system constructs its perceptions from the lower perceptual signals, its perception then matches its own reference signal.

So the effect of the imagination connection is as though the lower systems had acted instantly and perfectly. The level-n system, however, is supplying its own perceptual signal. It is "imagining" the perception instead of receiving it from sensory inputs and achieving it through action on lower systems and the world.

When I say I am imagining you and Rick and Chuck, this does not mean that THEY are imaginary. It means that I am supplying my own picture of them to myself. In your case, as I have never seen you, there is no way to tell if I am imagining correctly. For example, is your beard really curly and blond?

Imagination is confusing only when we believe that our perceptions are really outside us. If what we experience were really outside, it would be hard to explain how we can experience things that aren't there, or instead of what is there. Normal perception contains a large component of imagination at all

times. We imagine that there is a back on the terminal we're looking at. We imagine that the table extends continuously underneath the keyboard, and is of the same color there. In thousands of tiny ways we fill in what is actually missing from the picture of reality we get through our sensors, treating incomplete perceptions as if they were complete.

Note that the imagination connection explains both perception and imagination (and remembering, dreaming, planning, creating, and so on) by means of one simple switched connection, and does not require a separate perceptual system for things imagined and things "really seen." All experience, real and imagined, results in signals in the same perceptual channels. Imagination can take place at different levels; at low levels it is vivid and realistic because it includes imagined sensations. When the short-circuit connection occurs at higher levels, we experience only from that level up, without the realism. We call that imagining "in the abstract" or "imagining THAT" xyz is occurring without, in fact, imagining x,y, or z.

I've always been amused when scientists say things like "You're just imagining that" as if that dismisses the subject. Imagination is a phenomenon like any other, and requires explanation in any model of the brain.

Martin Taylor (probably too late)(920904) --

>Perhaps the problem is that we have a different sense of "alerting." The way you use it, it sounds as if you make some connection with consciousness.

No, I wasn't thinking of consciousness. I was just wondering what effect those signals were supposed to have on the other control systems.

>I take "alerting" to describe the marking of df as being probably useful to control, and don't put any more on it than that.

Then in your diagram you have to include the perceptual functions that perceive "possible usefulness" and say what it is that they are perceiving that tells them about usefulness, and you have to say how an output of the alerting system "puts a marker on" something, and what effect that marker has, and what knows about that effect, and so on. You're really proposing an arrangement with a lot of hidden complexities even at the block diagram level, not to mention the implementation level. It seems to me that you're proposing dormitive principles, to use Bateson's image [for those who don't know it, a sleeping draught makes you sleepy because it contains a dormitive principle].

>On other occasions I have proposed gain functions with dead zones as being a way that ECSs can perform the alerting function on themselves. In that case, there is no overt mechanism, but the df shift occurs as a consequence of conflict between the "alerted" ECS and previously active ones with correlated perceptual functions.

This is NOT a dormitive principle, but a perfectly good model. However, I don't see the necessity of classifying it as an "alerting" function. When you've said it has a dead zone and that a large enough error can wake it up and cause conflict, you've completely described the relation of this system to the others. To say that it's an example of an alerting function implies that this classification adds something to what is happening. And I fear that by creating such a class, we would be tempted to include other quite dissimilar

processes in it, and leave the impression that there is something common to all the processes in this class. Classification is arbitrary. The way we classify things isn't what makes them work.

>As a related topic, it seems to me probable that there are usually far fewer >output degrees of freedom used in active control than are available, because >the use of all available ones seems rather like a very high workload condition.

I think that all the 600 plus or minus 200 muscles are always part of active control systems at least while we're awake. Muscles are always in a state of nonzero tonus. The tendon and stretch receptors are always connected and ready to respond if not already responding. If you're standing, you can't raise an arm without producing an adjustment in the tension of a toe muscle.

"Active control," as I said to Rick Marken, is a somewhat ambiguous concept. A balanced pair of control systems receiving a zero error signal and not being disturbed is not varying its outputs, but it is still connected and "live" even though it's not "doing" anything (i.e., no derivatives).

In considering output degrees of freedom, don't we have the same problem in counting them that we have with input degrees of freedom? As soon as you introduce the time dimension, the number of degrees of freedom multiples greatly. A pianist can play music on an 88-key keyboard with only 10 fingers and a foot-pedal. And if absolutely necessary, a nose.

>Symbolic PCT is surely easier to program than is distributed PCT, but I >doubt it is as robust or as true to life.

Show me.

RE: your rectifier model.

Very clever. This would indeed pick up AC variations, which would qualify as a perception of variability. You just need a resistor to ground on the rectifier side of the series capacitor.

Neural signals are natural rectifiers, as they can represent only one sign of frequency. If you subtract the average value of a signal from the signal, and use a long-time-constant neuron as the receiver, you'll also get RC-smoothed rectification of variations.

Speaking compulsively, this would be a transition-level perceptual function, wouldn't it? I wonder if perception and control of variability would naturally occur at that level and not higher ones. (????).

Greg Williams (920904) --

>I also think Bill's IDEOLOGY of radical autonomy wrongly (because he >can't get as far as he could otherwise in explaining behavior) >downplays what has happened up to now.

I don't really want to be an ideologue or downplay anything that should be up-played. I'm not sure what you mean by "what has happened up to now." The model surely contains mechanisms for change, for the organism to alter itself

as well as the environment, so at least some aspects of "what has happened up to now" are included in the overall picture. What is missing?

Where I resist is where you seem to be pushing for "control without control" which makes no sense to me. I've acknowledged that there are ways to affect people's ACTIONS (their outputs) without specific consent of the system being affected, and that a manipulator can produce such effects intentionally. If the effects are produced accidentally I just call this a disturbance, with the controlled variable being corrected as usual by the change in the action. If they're produced intentionally (which would be hard to do in any more than a trivial way, but admittedly can be done by a person who devotes enough attention and time to observing an individual), there is still no problem for the person affected unless some internal goal is disturbed, hierarchical or intrinsic. I think it's unlikely that an external agent could intentionally violate another person's inner goals without creating direct resistance. That, of course, takes us into a different universe of discourse.

Another aspect of this debate that seems to have been overlooked is what I have said about the time-scale involved. Fool me once, shame on you; fool me twice, shame on me. Anything is possible in a brief encounter with a con-man. If you follow the con-man around, you'll get the impression that people are gullible fools who are easy to control through deception and so forth. But if you follow the person who was conned, you'll find someone who is far less likely to be conned again. And if you follow the much smaller population that has been conned twice, you'll find even fewer whose gullibility has not been replaced by a deep suspicion even of their friends. The minute few who go through life being conned again and again can't blame the con-men any more: they're malfunctioning.

It is possible that large numbers of people are malfunctioning. One kind of malfunction is to adopt and try to live under a conception of human nature that is incompatible with the way people are actually organized.

>PCT SCIENCE -- again, not IDEOLOGY -- IS, I think, adequate. Skinner
>was model-less; PCT is based on models. And those models, aside from
>confirming folk notions of how organisms "work," are constructed in
>such a way as to make the two mutually incompatible "complementary"
>world-views described above fit together in a plausible way. I think
>the science could heal that ideological rift, and lead to a new sort
of >ideology noting the importance of BOTH history and the present
>situation.

Perhaps what's bothering you is that PCT in itself has nothing to say about social interactions. It has seemed to me that you are looking for ways in which the society affects individuals, including their processes of reorganization, in a way that is purposive or at least directed. So the properties of a society in one sense DETERMINE how an individual must become organized in order to live in it. This, of course, creates an unavoidable image of the individual against The System, with the system inevitably winning by direct force or subtle and untiring influences.

What is wrong with this image (which may be a straw man, but never mind) is that there is no System. There are only the interactions among individuals. These individuals are ALL adapting themselves to the worlds in which they find

themselves, in service of their own deepest needs, their intrinsic reference levels that only heredity can influence.

This means that the System is not an entity with goals and purposes; it is only the consequence of ALL people continually adjusting to the presence and actions of ALL other people. The shape of the social system is in constant flux; it isn't going anywhere by itself; it's just changing form and drifting as the individuals each look for a state of minimum internal error. Furthermore, there isn't just one System; there are thousands of overlapping systems of different sizes, even inside what we think of as a single culture. There are small groups of people within which each person has found a local minimum of internal error; take a person out of the local group and transplant him or her into a different one, and the pushing and jostling will recommence as EVERYONE has to seek a new equilibrium. There is a large overall system only in a few regards, a few threads that run through all the smaller ones.

I react with alarm to the prevalence of social control through the use of force. But I also see this not as a concerted plot, a conspiracy to control people, but as a malfunction. When you grow up in a society that has not progressed beyond competition, you are growing up amid people who see no alternative but to push back as hard as they are pushed. That's how they treat you; and by pushing back you learn how to become a pusher of others. The company that locks out the workers and posts armed guards at the gate is run by people who are afraid, with justification, that without the armed guards the strikers would break down the gates and destroy the livelihoods of the decision-makers. Each person selects and justifies his or her own actions on the basis of what the other people are perceived as doing and wanting. The compassion of the rich for the poor is greatly circumscribed by their fear of the poor, and by anything that threatens to make them one of the poor. And the poor do not behave in any way calculated to reassure the rich because they are convinced, quite correctly, that the rich need poor people to support the life-style of the rich at sufficiently low cost. There are simply not enough resources, material or physiological, for everyone to live on the tail of the curve.

I don't see the solutions to these problems as coming out of continuing in the same vein. If you tell some people to stop controlling other people by force, and use force to make them do it, you threaten to remove their only means of defense against force. You will be ignored or removed. Our whole society is in a state of minimum internal error, but that minimum is not very low. It entails fear, misery, and pain. It's just that a move in any so-far-imagined direction away from this state of equilibrium entails even more fear, misery, and pain. People say "Yeah, but what's the alternative?" All the alternatives they know about would be even worse, and have been.

The only cure I can see is to fix the malfunction. Teach people how they work. Make it comprehensible to people that when they push on someone else, the natural and inevitable result is a push back. This result is under the control of the pusher. It is not perversity or ambition or hostility that makes others push back: it is you, trying to make people work in a way contrary to their natures. Make sure that each person grows up understanding exactly why others react as they do, and who is responsible for what.

This understanding will make it possible to find a new equilibrium in which all the errors are lower. Maybe not zero; maybe still not what we would want.

But surely lower than they are now. This is a direction not tried; in this direction lies something better. So there is nothing to prevent our moving that way.

The problem I see, Greg, is that people ARE autonomous, but they keep trying to behave as if all the other people are not. That is the malfunction.

Your paragraph seems fine to me. No error.

My only proviso is that we agree on this:

You say:

> Another person's output variables (actions) can be so controlled (I
> have been calling this "manipulation") if you have appropriate
> knowledge (or adequate prediction) of the person's control structure,
> without the threat of or direct use of overwhelming physical force by you.

This is true, but I maintain that if such appropriate knowledge actually existed, the malfunction of which I speak would have been cured, and the desire to control other people would all but vanish. It is not likely that anyone now classed as a "manipulator" has anything approaching the kind of knowledge of even one other person to accomplish what you describe. The kinds of non-coercive manipulation we see around us now could easily be avoided by anyone with a modicum of self-understanding, if avoiding them were necessary to maintain one's inner integrity.

Curt McNamara (920904.1730) --

Well, this discussion seem to be tempting some of our silent listeners. Nice to hear from you, Curt.

> BTW it appears to me that some CSG-L folks don't believe in
> manipulation because they think everyone's control structures are
> perfectly aligned with their highest principles. In fact, I would
> guess that none of us are actually organized this well. The
> non-alignment is what allows manipulation to occur.

A person who is in internal conflict can be manipulated in the sense that control of some variable has been lost or crippled, so someone else can control it without much effective resistance. But there is a defensive mechanism, in that if the controlled variable gets pushed far enough to one side, the inner effort that was helping in that direction will relax, leaving the other side to push back. So the controller can control only in the region where the person's inner conflict prevents any net action at all.

I do agree that when internal conflict is removed, people become MUCH harder to control. Paradoxically, however, they see fewer situations as threatening their control, and so become less unwilling to be used for another's purposes if there's no serious inconvenience involved. I think that people who are afraid (with justification) that they have few defenses against control learn to avoid getting into situations where others can take advantage of their handicap.

Bill

Date: Sat Sep 05, 1992 11:18 am PST
Subject: Don't know much about history

[From Rick Marken (920905.1200)]

Posting from home again. So here's some more header.

Greg Williams (920905) says:

> Eventually it will dawn on you
>that the PCT model explicitly says that one's history (genetic
>and environmental) plays an important role in determining
>what one's current control structure is.

I don't understand how I might have given the impression that I don't think that this is the case. Genetic history provides intrinsic reference levels; environmental history provides the constraints within which reorganization must discover control organizations that end up producing inputs that satisfy these intrinsic references.

>(PCT doesn't say that one's past environment SOLELY resulted in one's
>current reference levels, but that it CONJOINTLY -- with one's control
>structure history -- determines the current reference levels.)

I think you might have misspoken here. It's not one's current REFERENCE LEVELS that are the conjoint result of environmental and control structure history. What is the result of this history is one's CONTROL ORGANIZATION: the perceptual functions that determines the variables controlled, which higher order system(s) will use these systems by providing them reference inputs, the output function that determines the environmental effects that will be used to control the perceptual variables and the connection from error to output (in terms of polarity and gain) that determines the dynamic characteristics of the systems). So one conjoint result of environmental and control structure history might be a control system that controls the perception of the DISTANCE BETWEEN yourself and other people. But this history cannot determine the appropriate reference level for this perception; in fact, the reference for this perception must vary in order to control higher order variables. All that environmental and control structure history can do is find that the "distance between yourself and other people" is a perception that should be controlled and it can find out HOW to control it. It is likely that people with other environmental and control histories will not end up developing control systems that control the distance between themselves and others. In this sense, environmental and control history does influence the nature of the control hierarchy one ends up with. But I don't think that it is correct to say, then, that one's environmental and control structure history CONTROLS one's current behavior -- or even one's current behavioral potential (defined as the control hierarchy). This is

because it is impossible to be sure what "solution" reorganization will come up with (in terms of resulting control organizations) to deal with a particular environment. I don't think you can "hem in" the development of control organizations so that only certain organizations are created. Something like this kind of "hemming in" IS done in the classic operant situation; the rat MUST learn to control the position of the lever or it doesn't eat. But even here, it is not clear that each rat is building the same control organization. Some rat's may be building control systems that control the visual position of the lever. Others may develop control systems that control their own position in such a way that the lever always ends up going down. You can usually produce the same physical result (such as closing the circuit with the lever) by controlling different perceptual variables.

>Controlling for not altering variables associated with others is a >form of control, isn't it?

Possibly.

>Such deliberate non-controlling-others control, I think, can also >result in screwed-up persons, sometimes.

I don't buy it. Your statement implies that at least SOME controlling is good FOR THE CONTROLLEE (the kid in this case). I've already admitted that controlling another person (like a kid) can be good for the CONTROLLER. Your example of "spoiling" is not what I would call spoiling. You describe a situation where a kid doesn't have to develop certain control systems in order to produce certain perceptions that it wants to control. I want to perceive myself living in a house but I never learned to build a house because my houses have always been built for me -- am I spoiled?. My way of producing the perception of me living in a house is to plop down a big down payment and qualify for a loan. I could learn to build a house though -- just like the rich kid could learn to go out and get a job in order to buy a car if, for some reason, asking Dad for it didn't work any more. I think of real "spoiling" as a kind of controlling -- where parents give their kids stuff in order to keep the kids "out of their hair". The parent is controlling the kid, keeping it "away". If the kid wants to be close to the parent (and most kids apparently do) there is conflict between parent and kid (they are both controlling for the same variable -- "closeness"-- but relative to different reference levels) and the parent is likely to have the upper hand so the kid will be constantly reorganizing -- so kid will be a pain in the ass -- ie. "spoiled".

I think what you are saying is that it is possible to raise a kid in such a way that it develops none of the control systems that will be necessary to survive in the "real world". I agree that this is possible -- but I would argue that the way to solve this problem is not to control the kid but to EDUCATE IT. I think that there is a big difference.

>So keep it up, Rick, and together we can forge a broader >ideology based on PCT science. History AND physiology....

What is this "ideology based" PCT thing? I'm describing how I think the PCT model works. If I'm wrong about something just let me know. I'm willing to be corrected if I'm making a factual mistake. I think of ideology as as belief that rules out such correction. I'm really willing to be corrected if I have made a mistake about the PCT model. What's my ideology???

Best regards Rick

Date: Sat Sep 05, 1992 11:23 am PST
Subject: Re: Misc as listed:congratulations

[From Dick Robertson - 9/05/92]

Hi Dag, congratulations to you and Christine on your first customer. I just read your post of 0830. It is great stuff. I want to use it in a History of psychology class I'm teaching, in which I start with Kuhn and end with HPCT. Also, I'll soon be sending you the Plooiij paper that I promised.
best, Dick

Date: Sat Sep 05, 1992 12:21 pm PST
Subject: controlling reference signals; non-controlling control

[from Bill Powers (920905.1400)] Rick Marken (920905) --

A great great post, Rick. You've picked up a problem that keeps flitting into my mind and flitting right out again. Of course reference signals aren't set by external events OR genetics. They have to be flexible if the hierarchy is to work. They are the MEANS of control of higher systems. Only the top level of reference signal seems to be a problem, and even that problem goes away when you realize it has to be adjustable by reorganization. The only reference signals that are DETERMINED are the intrinsic ones -- and the "history" that affects them is so slow that human history doesn't even count.

You are right about Greg's non-controlling control, but you missed a good opportunity for a needle. Actually, I'm really a powerful non- controlling controller, as are most people. I'm not doing anything to control about 4.999 (999...) billion people. Just think: all of those people under my power, and I don't have to lift a finger.

Although educators are often controllers (with truant officers to supply the force), I agree with you, Rick, that Education is not control.

And anyway, who's smart enough at this stage of the game to know how to raise kids without EVER controlling them?

Best, Bill P.

Date: Sat Sep 05, 1992 12:44 pm PST
Subject: Our Flat Earth Brothers & Sisters

[From Hank Folson (920905)]

To All:

A couple of times now disparaging remarks have been made on the net about the the Flat Earth Society. Twenty years ago I watched an interview on TV with a member of the Society. As I understand now, the interviewer was controlling to show how smart he was and how out of touch with reality the Flat Earther was. But the quiet very in control elderly gentleman talked more about not taking our view of everything for granted. In spite of the interviewer's pushing for details about how the earth could be flat, he continued to make his main point: the obvious, what we take as truth, may not be true. Some time later, his real message finally soaked in.

If this group had formed before 1492, the name would then have been the Round Earth Society. Everyone would have made lots of fun of them, until of course, 1492.

The Flat Earth Society has much in common with the position CSGNet and PCT are in now. They are our brothers. They too are doing some very useful, if unappreciated, work.

It's all perception...

Hank

Hank Folson, Henry James Bicycles, Inc.
704 Elvira Avenue, Redondo Beach, CA 90277
310-540-1552 (Day & Evening) MCI MAIL: 509-6370
Internet: 5096370@MCIMAIL.COM
Date: Sat Sep 05, 1992 9:04 pm PST
Subject: Sciences & Ideologies II

From Greg Williams (920905)

I said the following to Bill Powers in a post of mine dated 920904:

>>All right, then, I THINK we arrive at the following more complete
>>formulation of your definition: "Controlling another person" means altering
>>variables associated with that person so as to maintain certain of your
>>perceptions as you desire. Another person's output variables (actions) can
>>be so controlled (I have been calling this "manipulation") if you have
>>appropriate knowledge (or adequate prediction) of the person's control
>>structure, without the threat of or direct use of overwhelming physical
>>force by you. Another person's control structure (reference signals) cannot
>>be so controlled regardless of the threat (or perceived threat) of or direct
>>use of overwhelming physical force by you. Nor can another person's
>>CONTROLLED (by him/her) perceptions be so controlled without the threat (or
>>perceived threat) of or direct use of overwhelming physical force by you.

>>Question: Do you have any problems with the above paragraph? I need to know
>>whether it is OK before proceeding further. Thanks!

In reply, Bill said (920905.0930):

>Your paragraph seems fine to me. No error.

>My only proviso is that we agree on this:

>You say:

>>Another person's output variables (actions) can be so controlled (I
>>have been calling this "manipulation") if you have appropriate
>>knowledge (or adequate prediction) of the person's control structure,
>>without the threat of or direct use of overwhelming physical force by
>>you.

>This is true, but I maintain that if such appropriate knowledge
>actually existed, the malfunction of which I speak would have been
>cured, and the desire to control other people would all but vanish. It
>is not likely that anyone now classed as a "manipulator" has anything
>approaching the kind of knowledge of even one other person to
>accomplish what you describe. The kinds of non-coercive manipulation
>we see around us now could easily be avoided by anyone with a modicum
>of self-understanding, if avoiding them were necessary to maintain
>one's inner integrity.

Thanks again for replying on this, Bill. I still think it is important that we be able define our terms, and that, of course, requires your cooperation. I'll attempt to deal with your caveat later; right now, I want to call attention to a different sentence in the paragraph which "seems fine" to you: "Another person's control structure (reference signals) cannot be so controlled regardless of the threat (or perceived threat) of or direct use of overwhelming physical force by you." (Slightly) unpacking this with the aid of the first sentence in the paragraph, we get the following: "Another person's control structure (reference signals) cannot be altered by you so as to maintain certain of your perceptions as you desire, regardless of the threat (or perceived threat) of or direct use of overwhelming physical force by you." So far, I suppose, so good -- just substituting in what you define "controlling another person" to be. I dispute this contention.

I dispute it in two ways: (1) with regard to what might be termed "short-term" control of another's control structure; (2) with regard to "long-term" control.

(1) Consider again my yelling-"fire"-in-a-crowded-theater scenario which I set forth earlier this week. I yell "Fire!" and you (being prudent -- that is, in order to satisfy a reference signal in you having to do with not getting burned) move toward the exit. I am controlling for my perception of seeing you move toward the reference signal, so if you didn't move thusly, I would alter my actions so as to (I hope) see you move thusly. The situation as construed would appear to fit my definition of "manipulation" -- I am controlling your actions, as per the first sentence in the paragraph above which you OKed. But I claim that I am also altering your control structure: undeniably, it is operating differently than it would have had I not yelled "Fire!" We can mince words about whether your reference signals per se have been altered, but it is clear that the operations of your control structure are different than they would have been otherwise. I must quickly dispel the notion that, because I cannot control to perceive your control structure (operating in ANY way -- I simply don't have perceptual access to your control structure "innards"), then I am not controlling your control structure. The definition of "controlling another person," above, which you agreed to, says only that my control of

certain variables associated with you involves control of certain of my perceptions, and is neutral regarding the connection between your variables I am controlling and my perceptions I am controlling in order to control you. The connection might be very indirect. The point is that I have altered which of your lower-level reference signals you are (let us say) "needing" to use, given your high-level reference signal which (as we have discussed before) can be said to be continuously operating throughout (with no disturbance prior to my yelling "Fire!").

Perhaps you will object that "alter" in the definition of "controlling another person" means "being the sole agent involved in causing a change in the variable." But you didn't object in the "manipulation" scenario, where actions are being controlled. Assuming good control, an action is determined conjointly by the (relevant) current reference signal and the (relevant) current disturbance. So alter must mean "being a contributory agent involved in causing a change in the variable." That is, to alter a variable associated with another person is to be necessary, but not necessarily sufficient, for a change in the variable.

Perhaps you will object that my (assumed successful) manipulation in some sense REALLY didn't alter your control structure, only the way it was operating. That your reference signals weren't altered, only which were disturbed. OK, then if we take "which reference signals are disturbed" as a variable associated with you, I was controlling that variable, wasn't I? If so, we have another kind of variable associated with another person which can be controlled by another: which reference signals are disturbed, and hence, I shall say, "active."

However, all of the above having been said, I really am not very interested in convincing you of the possibility of "short-term" control of another's control structure. I've gone through the motions mainly to set the stage for a discussion of the possibility of "long-term" control. A "dry run," as it were.

(2) I admit freely that the "short-term"/"long-term" nomenclature isn't optimal. I mean to suggest a distinction between the controller's altering the potential controlee's control structure without/with reorganization of the latter's control structure being involved. "Long-term" alterations unquestionably involve changes in the controlee's reference signals, since such changes are the result of reorganization. And in "long-term" control of another's control structure, the controller acts to control certain of his/her perceptions which will be controlled if the controlee's control structure reorganizes so as to "solve a problem" posed by the controller and which the controllee wants to solve (shorthand for saying his/her higher-level reference signal(s) require solution of the problem to be achieved). As it happens, the controllee cannot solve the problem posed (that he/she wants to solve) using his/her existing control structure. Assuming that the controllee does not decide to give up wanting to solve the problem, then his/her control structure reorganizes in an attempt to solve the problem. Reorganization will cease (and a new control structure -- with different reference signals -- will result) when (and if) the problem is solved. If the problem is indeed solved thusly, then the controller (who controls for his/her own perception that the controlee solves the problem) can be said to have altered the controllee's reference levels. The control could involve direct force or threat, or not.

So, such control is possible. Its likelihood is another matter; its success requires that the controlee not give up wanting to solve the "insoluble without reorganization" problem and that he/she reorganize until a solution is found. Of course, the solution might be impossible altogether, regardless of how much reorganizing occurs (remember Liddell's "experimentally neurotic" animals who were given impossible discrimination tasks?), but since the controller wants to see the controllee actually solve the problem, presumably the controller wouldn't pose an impossible problem except by mistake.

Again, you might object that the new (reorganized) controllee's control structure is not determined solely by the controller. That is manifestly true. But will you object that the new structure is determined solely by the controllee? As I understand your theory of reorganization, a series of random (pseudo- or not, and perhaps biased in certain ways -- that possibility doesn't matter to my argument) alterations in the control structure occur, and each is tested to see whether it solves the problem (or "conflict") which gave rise to the reorganization process in the first place. When the problem is solved ("well enough," in some reasonable sense), reorganization stops with the control structure which was capable of solving the problem. So the new control structure depends on both the random "walk" in control-structure space, which I take to be "immediately" the result of the controllee's control-structure (and possibly "lawless" to some extent, as you suggested not long ago, as an ultimate refuge of non-traditional-free-will autonomy) and, conjointly, the problem set by the controller. Had a different problem been set (and had the controllee reorganized successfully to solve it), then the resulting control structure would have been different. In short, the controller's control is necessary, but not sufficient, for the particular changes in the controllee's control structure following reorganization. That is, the controller can alter the controllee's control structure (reference signals), contra the sentence I've been analyzing.

Finally (for now, since I want to get back to writing this month's HortIdeas), if you object that what you really meant by not being able to control another person's reference signals is that you can't alter them reliably in a precise way which you have pre-specified, I agree. But you can control them as per the definition above of "controlling another person."

We are still distant from sciences and ideologies -- and Skinner. I beg patience while I attend to other affairs. I promise to return. A hint of the future: I take education and indoctrination and friendship and organized cooperation and so forth and so on to rely on control of others' control structures (sometimes with threats and use of physical force, but often not).

And to nip one potential problem in the bud, I do NOT think PCT supports notions of trans-organismic control. Control is always the matching of a particular organism's perceptions with that organism's reference signals. Colloquially, I say, for example, "I control the path of the joystick." That is shorthand for "I control my perception of the path of the (perceived by me) joystick."

Again, to be continued,

Greg

Date: Sat Sep 05, 1992 10:38 pm PST

Subject: Raising children

from Ed Ford (920905:2335)

>Rick Marken (9209040) Seriously, I would very much appreciate those of
>you who have raised/are raising children would pass on your tips on
>"aligning their control structures" with those of civilization.

Rick, I think the first and without doubt the most important key to raising children is how you spend time with them. If you look at the hierarchy of control, you'll find that we create our concept of the world we live in including those with whom we live from lower order experiences. Your children's perception of you is created by them from their experiences in which you are or are not involved. With my experience in working and consulting in schools, corrections, mental health and residential treatment centers over many years, plus raising my own eight children, I believe nothing is more important than our individual alone time with each of our children. Children create their perception of you based on the time you have spent with them and the quality of that time. As all of us have learned, we tend to listen to and respect those whom we perceive respect us, who care about us, and have expressed a belief (both verbal and in action) that we can make it.

The second aspect in raising children has to do with creating an atmosphere in which they have chance to develop into responsible children. To me, responsibility is the willingness (commitment or reference level) and the ability (a skill where parents teach their children how to think responsibly ((notice I didn't say behave, but think - which comes from my perception of others as living control systems))) - the ability to obey rules and standards and ultimately to set their own, without infringing on the rights of others to do the same. All this means the parent must set standards and rules which reflect the parent's own values and beliefs (systems concepts level) and follow through with the natural consequences of not following the rules. I'm not speaking of hurting children (punishment) nor do I believe in being permissive. All children eventually must learn to respect the rules of the culture in which they live or they will be in conflict with that culture. Again, my experience with my own children has taught me that children tend to adopt the standards and values of their parents IF they've establish a close, loving relationship with them and that the values and standards in the home have brought them a satisfying life.

Ultimately, if children perceive you as caring about them, as believing in them, if they recognize the existence of reasonable standards within the home, they are more likely to work cooperatively to find a way to get along. (I might suggest to those who perceive the Bible in a negative way to look at the more positive aspects, especially St. John's gospel and his emphasis on love).


As I said in an earlier post, I think the way to deal with (read talk with) children is something I am beginning to do quite extensively in schools systems throughout the country - using PCT as the basis. Since you can't control another LCS (living control system), I begin my training by teaching the participants to just ask questions. The participants are teachers and selected parents who will be working with district school parents who want help in working with their children in a method compatible with what the school district teachers and administration is learning. I call this section

Teaching Responsible Thinking (again, notice I've gotten away from the word behavior). I give them a little card and I've rewritten the section on teaching responsible thinking here. The first section has to do with spending quality time as a vehicle to helping the child create a loving perception of the parent. The second section has to do with setting rules, standards and consequences. (I'll be glad to send anyone on the net one or more copies of this card, but be sure to post your complete snail mail address.)

C. TEACHING RESPONSIBLE THINKING

ASK THEM WHAT THEY THINK - KEEP QUESTIONING THEM - don't tell them what you think; ignore excuses & don't ask why; be non-judgemental; be specific & stay focused.

1. Exploration: Basic questions to ask are...
 - What is it that you want? What are the rules or standards?
 - What were you doing to get what you want?
2. Evaluation: Getting them to think responsibly...
 - Is what you're doing getting you what you want?
 - Is what you've doing against the rules or standards?
3. Commitment: Getting them to choose responsibly...
 - Are you willing to work at resolving your problems?
4. Teach them how to create a plan...
 - a. establish a specific are of needed improvement.
 - (keep plan small to assure success)
 - b. set a measurable goal for needed feedback.
 - c. have them think through then explain in detail how they're going to achieve their measurable goal.
 - d. set up a chart which shows progress in time increments
 - (hourly, daily, or weekly, etc.)



(the above chart example might have to do with a child's plan to do 40 minutes of homework a night)

I train my students (in small groups of three) to develop the skill of just asking questions. Assuming children are willing to talk with you, the only way you can access them is through asking questions if they let you (that's the reason for establishing strong relationships). It is only through questions that they will examine their own world. This is done in #1, exploration. You find out what their various reference levels, how they perceive various things, including the rules, etc. and what actions their are

taking to achieve their goals. In #2, evaluation, you get them to compare their reference level with their perceptual variable. (cause that's how we're designed - see PCT) This is the first step toward taking responsibility. They are making a comparison, and, in the act of making that comparison, they are evaluating whether or not it is working and whether it is against the rules. They might also evaluate whether they are willing to live with the rules, change them, or move, or what ever.

If they recognize and are willing to admit that what they are doing is against the rules and/or isn't getting them what they want, then you can go on to the third step, #3, commitment. This step tests the strength of the reference signal, especially if it is a revised signal, such as their willingness to work at getting along with their parents. Do you want to work at resolving your problem?

Finally, #4 is teaching them how to make a plan. And boy does PCT help here! Once you get them to establish a specific area of needed improvement, you get them to set a measurable goal. This is so lacking in any kind of treatment facility or school. You cannot attain a goal unless you establish a specific reference signal and a specific feedback that reflects that goal, otherwise the system just will not work as efficiently. Imagine driving down the street and observing a traffic sign that says "drive at a safe speed." What speed would you drive? And even if the sign said "speed limit 55" but your speedometer wasn't working or just wiggled up and down, you still couldn't achieve the goal of driving 55 MPH. Both the reference signal AND the feedback or perceptual signal would have to be specific enough for any LCS to operate efficiently. If you connect the x's in the above graph, you'll see the perceptual variable. Most children are never taught to make effective plans because teachers have no idea how the brain works. In the above chart, the perceptual error is the difference between the feedback and the goal.

Once you understand how living control systems work, it opens up a whole world of understanding of how to help others operate more efficiently and on their own. It also teaches you where your limitations are in dealing with others. To quote Clint Eastwood, "A man has got to know his limitations".

Once students begin to sense the practical stuff I teach, they begin to respect the theoretical basis from which I create my ideas. That, to me, is where PCT will begin to achieve any kind of large scale recognition. When one's home life, personal life, or business life begins to benefit and the connection to PCT is perceived, I think our chances of getting PCT some national recognition will improve. People like Dag, David, Brent Dennis, Dick Robertson, myself, Jim Soldani, and others who work daily with people with problems see the results. There's no doubt that PCT works, both to us and the people we help. Rick, you and Bill and the rest have a much tougher road. Except for our comparatively small group, when someone you've taught finally says Aha!, that means their intellectual life might have moved light years ahead, but they are now doomed to many years of the frustrations you've been experiencing. One thing we have in common, along with our family-like relationships, is that we're working with ideas that are bound to make this life of ours a little better. And that's a lot.

Finally, I spent last weekend with my son, John, and his family in Villa Park, California. John and I and his three children went to the beach Saturday morning while Hester and John's wife Bridget ran an garage sale. John's

youngest child, a boy named Breen, will be two in December. A month earlier, John mentioned how Breen was walking in the ocean with the water barely at ankle level. As the entire surface of water move back toward the ocean, John noticed Breen lean to one side as if he were trying to compensate for the water's movement. Apparently, at his age, Breen was using the visual feedback of the land as part of what helped him to stand and walk. Since his visual feedback of what probably appeared to him to be what he was walking on was moving, he tried to compensate by leaning toward the water moving away from him and he fell several times. Now, he has learned (cause it doesn't happen any more) to negotiate the water along the shore line without any difficulty. Apparently he has learned to use the sense of feel of his feet on the sand as feedback and not so much visual. I think I have this right? At any rate, isn't it nice to understand PCT?

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU 10209 N. 56th St., Scottsdale, Arizona 85253
Ph.602 991-4861

Date: Sun Sep 06, 1992 4:41 am PST
Subject: Raising Kids/Education = ?

From Greg Williams (920906)

>Ed Ford on raising children (920905:2335)

Both Pat and I think that Ed's comments are wise. His techniques are important, in our opinion.

As outlined by Ed, his suggested interactions with children include instantiations of purposefully influencing them -- what I have been calling "manipulation." Recall: influencing another is altering some variables associated with the other, and manipulating another is influencing certain of the other's "don't-care" percepts so as to control certain of one's own percepts.

The trouble is, while Pat and I and certainly Ed believe that such manipulation can have tremendously important consequences, Bill Powers has characterized the consequences of manipulation as "unimportant" to the manipulee. Bill, are the consequences of manipulation largely or even solely important only to the manipulator?

Also, Bill and Rick: You both said in posts of yesterday that "education is not control." Please explain what education is, as related to PCT notions. Can the concept of education even be approached via PCT ideas?

Thanks for clarifying,

Greg

Date: Sun Sep 06, 1992 4:42 am PST
Subject: Imagination

[From Bill Powers (920906.0600)]

not just to mean "have some effect on" but in the sense of "have a particular preselected effect on," which is not influence, but control (that is, it can't be done by any means but control).

>I want to call attention to a different sentence in the paragraph which
>"seems fine" to you: "Another person's control structure (reference
>signals) cannot be so controlled regardless of the threat (or perceived
>threat) of or direct use of overwhelming physical force by you."

I had missed the significance of the parenthesized (reference signals), which as Rick Marken pointed out is not the same as a control structure. A control structure is (to me) a physical structure, not the signals being passed around in it. But that gets ambiguous, so let's let that go, too.

I think you're missing the point of "overwhelming physical force." If I resort to that, then it doesn't matter what you do or don't do. I may try a threat first. If that doesn't work, I can try hurting you. If that doesn't work, I can try seizing you and forcing your limbs to move as I want them to move. As long as I have superior resources, you WILL do what I want, even if I have to do it myself using you. As a final resort, I will just kill you and go looking for someone easier to control (why work up a sweat?). To a controller of people, people are interchangeable. You're thinking of the use of force as being constrained by some kind of civilized limitations. Unless one gives up the goal of control, there are no such constraints. You do whatever it takes, to quote a prominent humanitarian.

>(1) Consider again my yelling-"fire"-in-a-crowded-theater scenario
>which I set forth earlier this week. I yell "Fire!" and you (being
>prudent -- that is, in order to satisfy a reference signal in you
>having to do with not getting burned) move toward the exit. I am
>controlling for my perception of seeing you move toward the reference
>signal, so if you didn't move thusly, I would alter my actions so as to
>(I hope) see you move thusly. The situation as construed would appear
>to fit my definition of "manipulation" -- I am controlling your
>actions, as per the first sentence in the paragraph above which you
>OKed. But I claim that I am also altering your control structure:
>undeniably, it is operating differently than it would have had I not
>yelled "Fire!" We can mince words about whether your reference signals
>per se have been altered, but it is clear that the operations of your
>control structure are different than they would have been otherwise.

It is undeniable that my control systems are operating differently in the presence of a disturbance from the way they would be operating without the disturbance. There is nothing different about their organization, of course, but only in the signals that are moving around in the system. The way you're arguing, that would mean that ALL disturbances are "controlling my control structure."

I don't need to mince words about whether reference signals have been changed. They have. I can't lift a finger without changing reference signals. I can't resist gravity without resetting reference signals when I change position. For ALL systems in the hierarchy, acting means changing the reference signals of lower systems. There are no FIXED reference signals in the whole (healthy) system, except perhaps at the highest level -- and even those are subject to change, I think. I hope.

When you yell "FIRE" you insert a disturbance into my system at some level. You do this by altering an uncontrolled perception that is a component of a higher-level controlled perception. To combat this disturbance, I naturally have to shift reference signals at all lower levels. That is how action works. I alter the controlled lower-level perceptions (by means of altering reference signals) so that when they are put together with the uncontrolled one you have changed, the net result is no disturbance at the higher level.

The simplest example of this that I know of is also the first one I fell for. Someone says "Hey, your sister just fell off the swing!" I spin around and look, and see nothing of the kind. The other kid screams out "Made-ya-look, made-ya-look!" and collapses laughing.

Once this has happened, it creates a conflict. What should I do the next time someone tells me of some dire or interesting happening? If I don't look I may fail to do something necessary, or miss something interesting (it would have been interesting to see my sister fall off a swing). The other has presented me with a problem (as you say). He has not, however, presented me with a solution. It's up to me to figure out that it doesn't really matter to me if someone else makes me look by giving me false reports, not in comparison with seeing something interesting or helping with a problem (I really would have helped my sister). So the next time I'll fall for it cheerfully, knowing that most of the time I can believe what people tell me (except that kid, but even when he's the reporter I'll take the chance because he might be telling the truth once in a while). From then on I may look, but it's no problem. Not looking would cost too much. Let him control the direction I'm looking for a few seconds -- what's the big deal?

That kid, by the way, stopped doing that when someone else beat him up for doing it.

>The point is that I have altered which of your lower-level reference signals
>you are (let us say) "needing" to use, given your high-level reference signal
>which (as we have discussed before) can be said to be continuously operating
>throughout (with no disturbance prior to my yelling "Fire!").

And my point is, "That's life." Any disturbance can do that. It doesn't matter who is doing it or what is doing it. It's still not control of anything that matters to me. In the final analysis, all that I will NOT sacrifice in order to counteract a disturbance is my system concepts.

>Perhaps you will object that "alter" in the definition of "controlling
>another person" means "being the sole agent involved in causing a
>change in the variable."

Not at all; that's not what controlling means. I may tense my muscles to keep from falling down when gravity is the sole agent involved in causing a change in my posture. Gravity isn't controlling me; gravity can't control anything. If, in addition, someone is pushing on me, I will just shift the forces accordingly (by altering the appropriate reference signals).

>So alter must mean "being a contributory agent involved in causing a
>change in the variable." That is, to alter a variable associated with
>another person is to be necessary, but not necessarily sufficient, for

>a change in the variable.

This is getting positively Talmudic. Yes, I sort of agree with this. Except that even if you do nothing, other causes may change the variable. So "necessary" isn't in it. When you change your effect on the variable, and allow that other influences may also be present, the variable might do anything, including changing oppositely to your influence.

>Perhaps you will object that my (assumed successful) manipulation in some
>sense REALLY didn't alter your control structure, only the way it was
>operating. That your reference signals weren't altered, only which were
>disturbed. OK, then if we take "which reference signals are disturbed" as a
>variable associated with you, I was controlling that variable, wasn't I?

No manipulation alters my control structure. It only alters signals. All manipulations alter reference signals -- at the levels lower than the place where the manipulation has its direct effect. The manipulation is always resisted, successfully, at the level where its direct effect occurs. The changes in reference signals that accomplish the resistance take place at lower levels. That causes the manipulee no problems unless those lower-level changes alter some other higher-level perception that is under control. Then direct conflict will result (if you insist on having an effect). If you yell "FIRE" and can't show me where it is, I will put you under citizen's arrest for public endangerment.

>(2) I admit freely that the "short-term"/"long-term" nomenclature isn't
>optimal. I mean to suggest a distinction between the controller's altering
>the potential controlee's control structure without/with reorganization of
>the latter's control structure being involved. "Long-term" alterations
>unquestionably involve changes in the controlee's reference signals, since
>such changes are the result of reorganization.

You're using a very general term -- structure -- that isn't warranted by the discussion, which so far has concerned only signals handled by a structure of constant properties. All disturbances call forth changes in (lower-level) reference signals. So do reorganizations, but for a different reason. So with respect to whether reference levels change, there is no substantial difference between long-term and short-term. The main difference is in WHY they change.

>And in "long-term" control of another's control structure, the controller
>acts to control certain of his/her perceptions which will be controlled if
>the controlee's control structure reorganizes so as to "solve a problem"
>posed by the controller and which the controllee wants to solve (shorthand
>for saying his/her higher-level reference signal(s) require solution of the
>problem to be achieved).

Reorganization is not affected by higher-order reference signals, not the way I model it. It's affected by intrinsic error, regardless of its source.

If you attempt to control some aspect of my behavior by creating circumstances that cause intrinsic error in me, you will certainly be "influencing" my behavior. But you will not, with equal certainty, be controlling my behavior. Even I can't predict what will result when I reorganize -- how could you? I might end up solving the problem you have posed, or I might end up solving a different one that makes your puzzle irrelevant. Or I might just say "The hell

with it -- you're just making life difficult." There would be nothing you could do to stop me -- without the use of force. But my point is that, as I said above, while you can present the problem, you can't force what you think of as the solution on me. I will find my own solution; in fact, nobody else can. Only I can decide what constitutes a "solution."

>If the problem is indeed solved thusly, then the controller (who controls >for his/her own perception that the controlee solves the problem) can be >said to have altered the controllee's reference levels. The control could >involve direct force or threat, or not.

How many times do I have to scream "TO AFFECT IS NOT TO CONTROL?" It is only to present a situation. All the hopeful controller can do is present a situation by altering some variables. What the behaving system will do about that, if anything, is completely out of the controller's hands. All that the would-be controller can do is disturb, usually without any concept of what is actually being disturbed in the supposed controllee. A con man may get his money even if the mark sees through the scam, perhaps because the mark feels sorry for the con-man and doesn't mind handing over a few bucks to give this pitiful creature a momentary sense of success. Haven't you ever done that for a panhandler with a long sad unlikely story? The con-man then thinks he is controlling one thing; in fact he is affecting something else in the controllee and his sense of control is a fantasy. [Notice, by the way, how it's OK to use sexist language when you're saying something bad about someone. Nobody insists on "con-person"].

>So, such control is possible.

You're satisfied with your argument, but I'm not.

>Its likelihood is another matter; its success requires that the controlee >not give up wanting to solve the "insoluble without reorganization" problem >and that he/she reorganize until a solution is found. Of course, the solution >might be impossible altogether, regardless of how much reorganizing occurs >(remember Liddell's "experimentally neurotic" animals who were given >impossible discrimination tasks?)

Now you're destroying your own argument. What is done about the presented situation is out of the control of the putative controller. So the putative controller is not a controller, but only an affecter, only a disturbance. With respect to Liddell's rats, wasn't it lucky for Liddell that so many rats volunteered to be in his experiment?

>... but since the controller wants to see the controllee actually >solve the problem, presumably the controller wouldn't pose an >impossible problem except by mistake.

The controller can't make the controllee solve the problem, or predict what solution will be found, or even predict that the controllee will really try to solve the problem rather than just get rid of the pesky controller. The controller isn't a controller; you can't make him one just by calling him one. The controller is the problem, not the solution.

>Again, you might object that the new (reorganized) controlee's control

>structure is not determined solely by the controller.

It's not determined at all by the controller. Either you have sole determination or you don't have determination at all. "Determined by" means "unconditionally dependent on." If other influences can have an effect on the outcome, actually altering it, then there is no control or determination by anybody. If B is determined by A, then given event A nothing but B can happen. "Joint determination," a phrase I vow never to use again, is an oxymoron.

>But will you object that the new structure is determined solely by the
>controlee?

No. Determination is not control, although control may entail determination (as in, the reference signal determines the reference level of the input signal). Control is varying an action so as to maintain an input near a reference state despite disturbances. You're shifting around among three different usages of the word "control," only one of which is consistent with PCT.

>As I understand your theory of reorganization, a series of random (pseudo-
>or not, and perhaps biased in certain ways -- that possibility doesn't matter
>to my argument) alterations in the control structure occur, and each is
>tested to see whether it solves the problem (or "conflict") which gave
>rise to the reorganization process in the first place.

Well, yes, but you have to be more specific. If there's a conflict in the hierarchy, the conflict itself may be sufficient to start reorganization because of the resulting error signals, but the real reason may be that failure to control something (because of the conflict) has led to a different intrinsic error -- being in conflict about eating may cause your blood sugar level to fall. The criterion for ending reorganization is not that the conflict be resolved, because the reorganizing system knows nothing about conflict. It's that the blood sugar, or the hierarchical error signal, be restored to its proper state. It doesn't matter why it is restored. It could be restored by accident, and reorganization would still cease. The "test" for a proper solution has nothing to do with the solution. It has to do with the intrinsic effect of finding any solution that corrects the intrinsic error -- not a specific solution.

>Finally ... if you object that what you really meant by not being able to
>control another person's reference signals is that you can't alter them
>reliably in a precise way which you have pre-specified, I agree. But you
>can control them as per the definition above of "controlling another person."

No, that's not my objection, although it's a good objection in that it says you can't control unless you actually control. My objection is that influences and effects are not control.

Perhaps we could save some time here if you would state the conclusions toward which you are shaping all these arguments. I find your uses of words very slippery, but they all seem to be aimed at some final conclusion. Just what is this merging of the environmental and control views that you want to accomplish? If it takes this elaborate and convoluted an argument to support the idea that the environment (or other people) actually does have some

controlling effect on organisms, I can't help wondering which came first: the conclusion or the reasoning.

I'm with Rick on this subject: as far as I can tell, I'm just making straightforward deductions from the principles of PCT. Any "ideology" that's involved follows from these deductions. I don't see anything straightforward about your deductions; you keep switching around among "affect," "influence," "determine", and "control" as if these were equivalent terms. You can't prove control by demonstrating influence or any of the other terms. Physical environments can't control anything, lacking as they do the required organization. Other people can't control a person's behavior except in a trivial sense; they can control (non-coercively) only variables of lower level than the variable they disturb. As far as I can tell, these results follow strictly from the theory, and need no ideology to sustain them.

So what's your ideology?

Best, Bill P.

Date: Sun Sep 06, 1992 7:59 am PST
Subject: What's the Rule?

[From Bill Powers (920906.0900)] Ed Ford (920905.2335) --

Ed, we all appreciate your wise words about raising children. Where were you 35 years ago when I needed you to explain to me what control theory means in the real world?

RE: responsible thinking.

Let me register a complaint. I think that the "asking questions" approach can easily be overdone. It would be very strange, I should think, when everything an adult says to you ends on a rising note. "Would you like to eat your cereal? Do you think it's nap-time now? Am I going to get your milk and cookies now? Do you think it's all right to stick your finger in Baby's eye?" It shouldn't take much of this before the child will start mentally erasing the question mark and treating the communication like any sentence telling you what to do and not do: "Eat your cereal. Take a nap. Get the hell away from Baby."

I think you should ask questions when you really want to know the answer, not as a way of disguising an order. Don't you?

I'm also disturbed by the content of some of the questions. One is "What's the rule?"

Children, according to my uninformed but cheerfully confident thinking, start out mastering low levels of control and progress to higher levels. At any given time, there's a level where they're experimenting with reproducing perceptions, but have no higher-level reason for doing so. The adult who is alert to these levels can easily guide the child at the newly-forming level just by engaging in the pertinent activities with the child, and when communication is established, suggesting various things to try at that level. There's no conflict, because you're not controlling anything that the child is

also controlling for. So if you can stay just ahead of the child's progress through the levels, you can provide many more possibilities than the child could discover unaided. Of course you can't tell which ones will be taken up, but generally kids are willing to try anything for a while. If it happens to be at the level where they're still exploring.

When you get to the rule level -- logic, level 9 it says here -- children at the right stage will simply accept rules because you say that is what the rule is. They're still trying to learn to perceive and control in terms of rules. It's fun; they slavishly learn the rules of countless games, and talk to each other like lawyers when some little proviso is violated or interpreted differently by another kid. They'll rush to the rule-book to settle arguments. As far as they're concerned, rules are a natural feature of the universe. So when you tell them about the family's rules or society's rules, they just say "Oh. That's how it is." And then they conscientiously go around saying please and thank you and hello sir or madame as the case may be, until they get bored with that and go on to another game. Rules are just the way things work.

So you can remind children of a certain age of the rules (by questioning, perhaps, or simply by reminding them), and they won't object.

There comes that age, however, when the next level begins to develop, whatever it is. You sing out as you've done for a few years "Nine o'clock, time for bed," and the kid says "Who says so? Why can't I stay up until 10?" A little later, you say "You're grounded for staying out past 11," and the kid says "Like hell I am." Then you start spouting principles: "I'm trying to do what's best for you; I know you'd rather stay out but don't you think I would let you if I thought it was safe?" And to your surprise, the kid says, "Well, OK."

But if you say "Listen, smart-ass, the rule is that you get in by 11, and if you don't, you don't go out at all if you want to live in this family." Now you can express this any way you like -- you can ask "What's the rule?" if that's the way you frame statements? But however you say it, you're doing the same thing: you're entering into a conflict at the level where the kid has begun choosing his own reference levels for higher reasons, and you're referring to your greater coercive power to make it stick. In effect you're saying "Obey the rule or get out," knowing full well that the kid would be terrified and grief-stricken at being thrown out on his own.

So the point I'm raising is that "What's the rule?" is probably the wrong question, or statement, when the person has passed beyond the stage where rules are taken as natural phenomena, and have realized that somebody, some mere person, makes the rules and enforces them. That's when people start wanting to make their own rules, and start questioning the rules that others want to impose on them.

That's the stage when kids start turning into philosophers. They ask, in the rhetorical way, "A person ought to be loyal to her friends, shouldn't they?" That's not a rule; it's a principle. This principle may work out that you should defend your friend by driving by the house of someone who looked crosseyed at the friend and shooting through the windows. But it's still a principle, and if you want the kid to reconsider how to behave, you can't do it by saying "Isn't the rule No Shooting Into People's Houses?" The kid now knows that that's not HIS rule. It's somebody else's.

The only place where you can argue with this kid and have a hope of an effect is now at the principle level. You can even ask questions: "Is this how you've like to be living when you're 40? Would you like your own kids to be shot at? Can you imagine what this neighborhood would be like if the adults believed the same way you do?" And so on.

I don't have any particularly brilliant ideas about how you would actually do this: play it by ear out of long experience, I suppose, if you have the experience. But I think it's kind of a sham to ask "What's the Rule?" of a teenaged hellion in a detention center, when all you're doing is reminding the person of who has the power to slap them in solitary and take away their privileges. You can control people that way, but it won't stick. If you don't care about recidivism, you can keep everyone in line by asking what's the rule, but you won't have addressed their real problem, which is utter confusion about principles and probably system concepts too.

So if you want to ask questions of people who already pick their own rules, you should ask "What's the Principle," at the very least. You've left rules behind. No matter what you wish, the person is going to do the rule-picking from now on, unless credibly threatened with physical force. If you want to persuade, you have to back up to the principle level and talk philosophy, generalizations, morals, will- this-work-if-everyone-does-it? Basically you ask the same things you always do: what are you doing? Is it working? But now you're talking about principles, like sticking up for your friends. Or you're talking about system concepts, like being a cool dude or whatever it is now. You're not going to alter any attitudes toward rules by this time; you can only hope to persuade the person to try some new principles or even selves.

I don't know if any of this can actually be done with an 18-year-old offender. Maybe it's too late. On the other hand, you seem to suggest that it's not too late in as many cases as people are willing to think.

And by the way, Greg, as you can see I do believe that one can strongly affect growing children. But not for long, and you can't make them choose from your smorgasbord just the items you consider most tasty.

-

Best? Bill?

Date: Sun Sep 06, 1992 10:22 am PST
Subject: Misc topics

[From Rick Marken (920906.1030)] Good morning CSGNet!

Bill Powers (920905.1400) says:

>Rick Marken (920905) --

>A great great post, Rick.

I'm kvelling. Thank you so much, Bill. What a wonderful weekend this is turning out to be.

Hank Folsom (920905) says:

>The Flat Earth Society has much in common with the position
>CSGNet and PCT are in now. They are our brothers.

Not from my perspective. It's great to be skeptical but I think that the Flat Earthers are just being silly. PCTers don't reject linear causal models (SR, cognitive, etc) because we don't like them. It's because there is NO EVIDENCE for them. There is no data that supports these models -- period. The data presented in the journals is meaningless; I was just reading an article last night (in the Human Factors journal) where an r square of .36 was called "strong". The PCT model is based on data where the relationships between variables is ALWAYS in the range of r squared = .98. The data in favor of PCT is overwhelming (since there is virtually no data in favor of the other views). I would say that, based on the data, PCT is as solidly the model of choice for living systems as is "spherical" the model of choice for the shape of the earth -- maybe more so. The Flat Earthers seem to me much more like the SR and cognitive psychology types -- trying desperately to make a false model fit the available data.

Ed Ford (920905:2335) says --

>Rick Marken (9209040) Seriously, I would very much appreciate those of
>you who have raised/are raising children would pass on your tips on
>"aligning their control structures" with those of civilization.

Actually, I didn't ask -- it was Curt McNamara. My kids are raised (at least, from their point of view).

>All this means the parent must set standards and rules which reflect
>the parent's own values and beliefs (systems concepts level)

The parent has no choice, right? You don't have to tell a control system (the parent in this case) to control.

>and follow through with the natural consequences of not following the rules.

They are "natural consequences" only if the kid's behavior is part of what the parent is controlling; the parent can't help but "naturally" take action (ANY action) in an effort to move the kid's behavior back to the parent's reference for it.

>I'm not speaking of hurting children (punishment) nor do I believe in being
>permissive. All children eventually must learn to respect the rules of the
>culture in which they live or they will be in conflict with that culture.

Skinner also was big on non-punitive control. How do you "non-punitively" get a kid to produce a perception for themselves (and, incidentally, for you) that they have no intention of producing?

This amazing new approach to child rearing based on PCT sounds pretty much like the old version; set limits, teach children to follow the rules; show that their actions have "natural" consequences (a cop out that tries to make us believe that the environment is trying to control us so we have to learn to deal with it), etc. It sounds to me like you are saying what Greg appears to be saying -- a little control is good for the kid. I claim that there is no

reason (based on the PCT model) to think this is the case. The kind of control you are suggesting, Ed, is fine for the PARENT, but it is not necessarily great for the KID.

Confirming the above hypothesis, Greg Williams (920906) says:

>Both Pat and I think that Ed's comments are wise. His
>techniques are important, in our opinion.

>As outlined by Ed, his suggested interactions with children include
>instantiations of purposefully influencing them -- what I have been calling
>"manipulation."

And to the extent that that is the case I think that Ed's recommendations, if carried out seriously (REALLY enforce your references for how the kid should behave -- ie manipulate the kid's behavior) are prescriptions for conflict and pain (for the kid, at least). In practice, I believe that Ed does not treat kids as he suggests; he might be verbally firm and he surely tries to influence people (by talking to people, getting them to ask questions, explaining possible consequences, etc). But I'm sure that Ed would never pursue the inevitable violence that would result from attempting to get a person to "behave according to my rules and standards, OR ELSE"

I think there is an idea lurking in the background here. The idea is that those who recommend against "manipulation" based approaches to child rearing are therefore in favor of allowing kids "complete license" to do whatever the hell they want. But that is a false implication. When kids do things that disturb the variables you care about then you do something to stop it. The alternative to manipulative child rearing is not to become a doormat. It's to become a TEACHER. The goal is to help the kid learn to control mode skillfully -- so that it's activities don't have unfortunate side effects and so that it doesn't run into conflicts with other control systems. But you don't get a kid to be skillful by manipulating it into this state -- the control model implies that this just cannot be done. Educate -- don't control. We have to learn to educate better, not to control better. And don't worry, unless parents adopt the "let the kids walk all over you" reference levels of "summerhill" fame, they can be trusted to take care of themselves just fine. You don't have to tell a control system to control; you do have to tell a control system what to expect when it's dealing with other control systems. I'm just trying to do the latter; it sounds to me like Greg and Ed are trying to do the former.

>Also, Bill and Rick: You both said in posts of yesterday that "education is
>not control." Please explain what education is, as related to PCT notions.
Can
>the concept of education even be approached via PCT ideas?

Education takes advantage of aspects of the control hierarchy that we don't discuss that much on the net -- imagination, memory -- and also aspects that nobody understands too well -- how we turn word perceptions into the imagined perceptions that those words represent. Education involves teaching a control system to control. Words can be used to try to describe the kinds of perceptions that can be controlled in order to control a higher order perception. Like in teaching tennis -- the "teachee" usually comes in with a good idea of the higher level perception to be produced (a great tennis game) but doesn't know which sensations, configurations, transitions etc to control

in order to get this to happen. The coach tries to communicate what s/he thinks are good perceptions to control and, possibly, good reference level to get these perceptions too. The teacher (a good one) doesn't MAKE the teachee exhibit a particular behavior; the teacher tries to help the person learn to control perceptions that will allow the person to control. Education is not control exerted by the teacher -- it is a process where the teacher uses the teachee's ability to understand, imagine, remember, and reorganize in order to help the teachee be able to control.

Bill Powers (920906.0600) on the imagination connection:

>Anyway, this ought to be a lot easier to put into that spreadsheet model, Rick.

I think this is the way I had it set up in the first place (error goes directly back into perceptual signal) but I changed based on the discussion we had some time ago with Martin where he proposed that the imagination connection go through the perceptual input. We liked his idea better at the time. I'll go take a look.

Best regards Rick

Date: Sun Sep 06, 1992 10:54 am PST

Subject: Dealing With Children

from Ed Ford (920906:1150)

>Bill Powers (920906.0900) Where were you 35 years ago when I needed
>you to explain to me what control theory means in the real world?

The same place you were, struggling to figure out what to do when you didn't know what to do with a bunch of crazy, mixed-up kids who were going through the trauma of the 60's.....

>I think that the "asking questions" approach can easily be overdone.

Absolutely. Especially with one's spouse. In that case, it is rarely done, and even then, very carefully.....

>I'm also disturbed by the content of some of the questions. One is
>"What's the rule?"

I hope my brief remarks on the subject were understood in the context intended, namely, as a brief outline only. There are a lot of techniques involved in this approach that take many sessions to teach. I covered the basic content, enough to give people a sense of what I was trying to do and a sense of the direction I find most helpful and efficient. Once a person begins to experience working with a child in this way, hopefully they'll do a little creative thinking on their own and expand the approach TO FIT THE SITUATION.

When working with the more difficult juvenile, I tend to have them establish in their minds the highest priority systems concept they want at the time, which is generally (in a lock up facility) to "get out of here." Then I'd ask,

"Is punching that kid going to get you out of here?" The juvenile might respond, "He gets in my face." I'd then repeat my question, quietly and persistently, "But is punching that kid going to get you out of here or is it going to delay your stay?" They then might respond "But what else can I do?" Then I'd say, "Do you want to work at a better way of dealing with kids who get in your face and possibly get out sooner, or do you want me to leave you alone? What do you want?" In the above short dialogue, the rules and standards are implied but not directly asked. This is what I mean by creatively moving about an upset living control system, trying in as subtle a way as possible to help him/her deal with life.

>So the point I'm raising is that "What's the rule?" is probably the >wrong question, or statement, when the person has passed beyond the >stage where rules are taken as a natural phenomena.....

As you can see from my above remarks, I'm not saying that my little card should be used exactly without using any creative thought on the part of the parent/counselor/teacher or who ever. You have to look around, fuss with ideas as they come to you, try different ways of asking, as you suggested when you said..

>"Is this how you'd like to be living when you're 40?"

I think the real key for middle and high school students, especially juveniles in lock up or residential treatment centers, is to look for their reference levels with the highest priorities, such as "Do you want to get out of here?" or "What do you really want?" Often you have to get beyond the immediate problem, which is really symptomatic of a higher prioritized but not yet articulated problem.

>I don't know if any of this can actually be done with an 18-year-old offender.

My experience is - it depends whether they want help or not. Many don't want to work with you, others do. I've had them say "Hey, Man, I make more in a week than you do in a year and I don't pay taxes." He's right, of course. But then my chances of getting shot, killed, sent to prison or just beaten up are less. My theory is work with the control systems who want to work with you, they're tough enough.

Ed Ford

Date: Sun Sep 06, 1992 4:27 pm PST
Subject: Just a word

[From Bill Powers (920906.1830)]

Rick Marken (920906) --

Rick, I think Hank was trying to say that the Flat Earth Society is a deliberate put-on; that they support "wrong" ideas as a matter of principle. Is that right, Hank?

Ed Ford (920906) --

I guess I knew that you wouldn't do anything like an automaton.

Best, Bill P.

Date: Sun Sep 06, 1992 4:29 pm PST
Subject: Sciences and Ideologies - III

From Greg Williams (920906 - 2)

>Bill Powers (920906.0700)

>Greg, you keep wanting to make me say that people can't influence each other, when what I'm saying is that they can't control each other >(except with respect to output).

No, I am not. I realize that you know that people CAN influence each other, and I realize that we are in agreement on this. But you also say that such influence is "not important"/"doesn't matter" to the influencee. I agree with you that it is not important in the sense that it doesn't cause conflicts for the influencee or prevent the influencee from continuing to control. But influence often has other results which I think are important for the influencee, including, possibly, the demise of the influencee, as we've discussed before. At this point, twice recently, you have interjected that influences due to one's inanimate environment can also have nasty results like killing the influencee. I agree with that, of course, but the point I was trying to make is that people can PURPOSEFULLY influence others. You told me in no uncertain terms a few posts back that such purposeful influencing (if possible at all, I'll add here in an attempt to be diplomatic) is "garden variety control." And I immediately agreed that it is. It is control simply because the influencer is matching certain of his/her perceptions with certain of his/her reference signals. Then I tried to codify the explanations I had gotten from you regarding "controlling other people" into a definition/specification purportedly (my purport) setting forth your ideas. I asked you whether you had any problems with the definition, and you said no (with a caveat which I haven't addressed so far, but plan to eventually).

After your last post, I think we need to reconsider the definition of what you mean by "controlling other people." Why all of this homing-in on what you mean? Because I think it is important to make explicit what PCT says about how people can interact. I think it is important because I think PCT has a LOT to say about the nature and limits of such interactions, and I (and, I believe, a lot of other folks -- some of them applied psychologists, some clinicians, some sociologists, some political scientists, some laypersons, etc.) would appreciate having a solid theoretical foundation for thinking about human interactions. That's part of my ideology, since you asked. Furthermore, my ideology includes getting more folks interested in and working with PCT, because I think more work on PCT probably will lead more quickly to the BIG problem which I'd like to see solved in enough detail to satisfy me: how complex perceptual signals are built up. This leads to mentioning another part of my ideology, which is my belief that more work on PCT can be fostered by emphasizing how it connects to other ideas, rather than how everybody else is wrong. No, I don't think PCTers should NOT tell folks when they are wrong (witness our highly critical paper being submitted to SCIENCE in my name as

well as yours); but I think it will help to tell folks when they are right or close to right or have something that could tie in to PCT. I think it is not very helpful to emphasize the isolation of PCT ideas from "all" or nearly all previous ideas in the biological and social sciences. I don't want to have a "revolution" that nobody attends. Getting back to particulars in my ideology with respect to PCT, I think that the tone of interactions with nonPCTers needs to be less "us" against "them" and more "give-and-take." This is most obvious, in my opinion, in the ideology of some PCTers which denies any "importance" to a person of the influences (current and historical) of other people. This ideology, which I think is not supported by PCT (and have been trying to show why I think that in recent posts), makes a mockery of common sense, as does Skinner's opposing ideology of environmental determinism. I think each view is a partial truth, or perhaps even a whole truth from a single point of view -- but not the whole truth from all points of view. I see PCT as saying that the organism is in control (usually) in the here-and-now, but "the here-and-now organism" is not solely the product of itself, nor solely the product of its ancestors and itself. PCT is a CYBERNETIC theory that shows how -- precisely in what ways -- no organism/lineage is an island. That that theoretical image "fits" with common sense is a testament to the truth in control theory. To try to dismiss the effects of either "outside" or "inside" the organism-environment line will be seen as folly by clinician, researcher, and laymen alike. As Skinner's environmentalism was rejected, so will the organismism of some PCTers. I can't say that I am sad that Skinner's science was also rejected in part because of his ideology, but I would be very sad if PCT science is rejected because of the ideology of some PCTers. So much for my ideology.

>The confusion is coming partly from your using influence not just to mean
>"have some effect on" but in the sense of "have a particular preselected
>effect on," which is not influence, but control (that is, it can't be done by
>any means but control).

I defined "influence a person" very strictly as what happens when don't-care perceptions of the person are altered. They could be altered by inanimate portions of the environment or by animate portions of the environment. Skipping microbes, higher plants, and non-human animals, if the perceptions are being altered by another person, the person might or might not be controlling certain of his/her own perceptions which are related in some way to altering the first person's don't-care perceptions. Here, I think, is the breakdown in our communication. You say above that control is "having a particular preselected effect on." In the definition of "controlling another person," the controller has a preselected effect on his/her OWN perceptions, and so is controlling; the controller has some effect on a variable associated with another person, but is certainly (I thought I made this clear in the Prolegomenon) NOT controlling a variable associated with the other person. Is THIS why we're not communicating very well?

Back to square (nearly) zero. A person controls his/her own perceptions, not ANYTHING else. To belabor the point a bit, a person CANNOT control ANY variables (actions; perceptions, controlled or not; reference signals; control structure) associated with another person in the sense that the first person CANNOT control ANYTHING except his/her own perceptions. What happens in manipulation, I have said and maintain, is that one person alters (doesn't control, since they aren't the first person's perceptions) another's don't-care perceptions PURPOSEFULLY so as to achieve control of certain of his own

(the first person's) perceptions. Manipulation is the subclass of influence where influence is PURPOSEFUL, and can be done only by control systems. Now I can put forth the suggestion that successful manipulation requires some sort of connection between the first person's controlled perceptions and the second-person's influenced (don't-care) perceptions. Otherwise, the first-person's supposedly controlled perceptions would be independent of the second-person's influenced perceptions, and the feedback loop would be broken.

Notice that I am forced by the basic PCT definition of control to surrender any claims that I can control ANY variables associated with you, not just the ones we have been disputing about. All I can control (I promise I won't belabor it again after this) is my own perceptions. But I CAN control those perceptions by altering at least some of the variables associated with you. (At least that is my claim, and has been all along.) Given this refinement in language, perhaps we can make some progress toward resolving our dispute.

Now, I am controlling certain of my perceptions. The question, for me, is whether I can do that by altering some of your don't-care perceptions, by altering some of your reference signals, by altering some of your controlled perceptions, or by altering something else about your controlling. I think the answer is yes in all cases. There is also a further question as to whether I must threaten to or use force to do any of this.

You have agreed that, if I have a good enough model of your controlling, I can alter some of your don't-care perceptions so as to control some of my own perceptions. (Again, this is my "manipulation," and it could be divided into "facilitation" and "exploitation," based on whether, after a manipulation had been successfully completed, the manipulee judged the manipulation as "good" or "bad" for him/her.) I further claim that manipulation can have consequences for the manipulee judged by him/her to be important to him/her. And I claim that "facilitation" happens a lot in education/training, and con games are frequently occurring examples of "exploitation." Propagandizing, advertising, and patriotic urging-on also include many instances of manipulation, which tend to be harder for the manipulees to judge as "good" or "bad" for them. Finally, I claim that physical force or threats of physical force need not be present for manipulation to be successful. You have agreed.

Despite your arguments, and partly because I don't think your model of reorganization has matured sufficiently beyond the preliminary verbal stages to the simulation testing level, much less the empirical testing level, I'll be recalcitrant for now and continue to claim that it is possible, at least under some PCT models of reorganization (yours doesn't have to be the only one, does it?), to sometimes alter another's reference signals and/or something else about their controlling besides reference signals, actions, or controlled perceptions so as to control certain of one's own perceptions. For now, we can agree to disagree. Basically, I'm satisfied with your acknowledging that manipulation can exist, even though you seem to think it unproblematic/uninfluential in everyday life, in contrast with human interactions leavened by physical force or threats of physical force. I see this as an empirical question which won't be decided by your examples and my counterexamples. Be that as it is, I claim that physical force or threats of physical force needn't be present here, either.

The bottom line here is that I believe one can often (i.e., in education/training, therapy, cooperative work, etc., etc.) guess correctly enough about another's controlling structure and dynamics to allow reliable control of some

of one's perceptions (related, of course, through the other's actions, to their altered reference signals, and/or "something else" about their control structure). NOTE THAT THE ALTERED VARIABLES ASSOCIATED WITH YOU WHICH I ALTER ARE NOT CONTROLLED, AND DO NOT NEED TO BE ALTERED IN PARTICULAR WAYS EXCEPT SO I CAN CONTROL MY CHOSEN PERCEPTIONS. I NEED KNOW NOTHING ABOUT HOW YOUR VARIABLES ARE ALTERED (translation: I need not know anything about the source of disturbances to my control), I NEED ONLY CONTROL. An aside: Altering internal variables of another control system is not guaranteed to result in changes in the control system's actions which will allow control of your perceptions. Knowledge of the control system is important for the would-be controller here, too. I leave the mathematics of possible control-by-alteration for the future.

Last (thankfully), there is the possibility of altering some of another's controlled perceptions so as to control some of your perceptions. A private post by Martin Taylor has swayed me to think that even this might be possible without threat or force (though probably not likely). As Martin says, control isn't perfect. If your gain is about equal to mine and we each want a lever to be at a different place, the lever will end up between the two places, and both of our (I suppose I would say) "trying-to-be-controlled" perceptions of the lever will be altered. But perhaps you count this as an incident involving physical force. Is it "overwhelming"? I don't know. Maybe we need to define "overwhelming physical force" better (ugh!).

And so, we continue (mutually) to push back -- and I, for one, am learning a lot.

Greg

P.S. Are you going to send me the figure captions for 300-dpi output? I think you could go ahead and use TrueType fonts on your new dot-matrix if you want to save time -- SCIENCE doesn't EXPLICITLY say that the captions MUST be laser-printed, and I don't think they'll disqualify on that basis. But I'd be happy to do it here if you think it best.

Date: Sun Sep 06, 1992 6:22 pm PST
Subject: What's a teacher to do?

From Greg Williams (920906)

>Rick Marken (920906.1030)

>[Replying to Ed Ford:] It sounds to me like you are saying what Greg appears
>to be saying -- a little control is good for the kid. I claim that there is
>no reason (based on the PCT model) to think this is the case.

I say that SOMETIMES manipulation can be judged (afterwards) by a kid as "good" for him/her. The PCT model does not imply that this is impossible or even improbable. Of course, the outcome of manipulation depends on the kid's control system, the model which the would-be manipulator has of the kid's control system, and the skill of the would-be manipulator -- it takes a wise counselor like Ed Ford to FLEXIBLY size up the chances for manipulation ending up being "good" for the kid (as judged by the kid).

But there is certainly nothing in the PCT model implying that a parent or teacher controlling for perceptions of his/her own OF A PARTICULAR KIND by altering variables associated with a kid will INVARIABLY or ALMOST INVARIABLY result in the kid judging such manipulations "good" for him. Unfortunately, you apparently believe that PCT DOES imply that a PARTICULAR KIND of manipulation (NOT control) is best for teachers to practice:

>Educate -- don't control. We have to learn to educate better, not to control >better....
 >Education involves teaching a control system to control. Words can be used to >try to describe the kinds of perceptions that can be controlled in order to >control a higher order perception. Like in teaching tennis -- the "teachee" >usually comes in with a good idea of the higher level perception to be >produced (a great tennis game) but doesn't know which sensations, >configurations, transitions etc to control in order to get this to happen. The
 >coach tries to communicate what s/he thinks are good perceptions to control >and, possibly, good reference level to get these perceptions too. The teacher
 >(a good one) doesn't MAKE the teachee exhibit a particular behavior; the >teacher tries to help the person learn to control perceptions that will allow >the person to control.

You say that the teacher should "teach... control... by [attempting to communicate what s/he thinks are good perceptions to control... The teacher (a good one) doesn't MAKE the teachee exhibit a particular behavior; the teacher tries to help the person learn to control perceptions that will allow the person to control." So the "good" teacher should [note the normative, which I don't think you can derive purely from PCT -- you must import some of your own ethical presumptions, which are exterior to the science; but I'll let that pass] control for his/her perception of actions by the teachee which can be interpreted by the teacher as showing that the teachee is controlling well, eh? How do you get from the PCT model to the claim that this kind of manipulation is always best for the teachee? I just see too many potential ifs, ands, and buts, even disregarding the is-to-ought problem. Why should the teacher have the audacity to presume that he/she should suggest what are "good perceptions to control" for this kid. Haven't you ever heard anyone tell their sad story of holding a tennis racket "the way the coach showed me -- he said I was controlling my serves EXCELLENTLY" and, years later, "accidentally" discovered that holding the racket completely differently resulted in much better scores? Sorry, but I can't buy your "teaching control" as always best, and I certainly can't buy that your recommendation is implied by PCT.

>Education is not control exerted by the teacher -- it is a process where the >teacher uses the teachee's ability to understand, imagine, remember, and >reorganize in order to help the teachee be able to control.

Here you are agreeing with me that education makes great use of manipulation. My major problem is with your recommendation that one certain kind of manipulation is best for manipulees in general. I think your recommended kind of manipulation ("teach control") is probably OFTEN a good recommendation -- but not always. That's based on my own ideology (which comes out of ALL of my experience, at least as I recall it), not -- because it can't be -- solely on the PCT model.

>Ed Ford (920906:1150)

Pat and I still think you are wise and your techniques are important. And we also appreciate your hands-on understanding that manipulation (for the last time today, NOT control) of others can in some cases be beneficial to those who are manipulated -- I'm starting to term this "facilitation." You show courage in not being swayed by theoreticians from using methods which you've SEEN work. And I bet you are able to construct incredibly (to non-clinicians) accurate models of parts of clients' control systems after interacting with them for a short time -- you'd have to be able to do that to make quick but reliable predictions of who will be willing to be facilitated and who won't! Finally, Pat and I appreciate your appreciation of the importance of flexibility in dealing with others.

Keep up your Good Work, Ed,

Greg

Date: Sun Sep 06, 1992 7:50 pm PST
Subject: Quality time

[From Dag Forssell (920906)]

Dedicated to Ed Ford. Article in the View section of today's paper.
Is it not interesting how this kind of message is more noticeable when it comes from a man.

Los Angeles Times, Sunday, Sept 6, 1992

(FIRST PERSON/RHOLAN WONG)

FULL-TIME FATHERHOOD: HARDEST JOB OF ALL

Like an increasing number of people, I found myself unexpectedly without a job a few months ago. No, that's not exactly right. Rather, while my wife, Debbie, spent her days as a social worker, I suddenly had two jobs, both at home part-time free-lance writer/public relations executive and full-time househusband and parent for Derek, our 4-year-old son.

I had little preparation for the latter career, but intensive on-the-job training has taught me some of the basics of being a homemaker. I offer the following lessons I've learned for the consideration of other men who may be spending more time at home.

LESSON NO. 1

Being a househusband is the toughest job I've ever had. Whenever Debbie or another woman used to mention how hard it is to be a housewife, I would always nod my head in sympathy and mouth politically correct statements--"Are you now working inside or outside the home?" But I was really thinking, "Aw, c'mon. You should have a real job like mine, with constant, unreasonable deadlines, repetitive tasks and tyrannical and immature bosses."

Now I know that no deadline is as inflexible as a child's cry for breakfast, lunch or dinner, that the cooking--to say nothing of cleaning and washing--must be done over and over and over, and that a child can be the most demanding tyrant of all.

I never realized how much work it takes to keep a house running. The first few weeks at home, I was amazed to find that shopping, cooking and cleaning up for three meals could take the whole day.

But the household chores have been the easy part. The heavier burden--by far--has been entertaining, educating and disciplining my son.

Part of that burden is that I can never let down. I always have to keep up with a boy who has unlocked the mystery of perpetual motion. Every minute he may be about to eat a spider or throw a book through a window.

Even harder is the emotional toll. I never stop feeling for very long that I'm not doing enough for Derek. Yes, I may be giving him the right things to eat, but am I giving him the educational stimulation he needs? Is he watching too much television? Am I making enough opportunities for him to play outside? And do those playtimes foster an active imagination and develop motor skills? Am I giving him enough time with friends so that he's learning social skills? And will those friends be a good or bad influence?

Oh, and is he happy?

Looking out for Derek also means always having to look at myself. I have to examine everything I do around him, for he is studying me and will surely mirror my slightest immaturity. I was a television sports addict until I heard Derek say that he couldn't talk to me because he would miss one of his cartoons.

Dealing with all these physical and emotional burdens takes a lot of time, which brings me to

LESSON NO. 2

"Quality time" is a myth- If I want a relationship with Derek, I have to put in the hours. If I don't, I miss out on my son's life. These past months have given me not just a chance to catch up on Derek's intellectual growth, but to establish a deeper, richer relationship with my son. Recently, Derek intentionally misbehaved at a church picnic, and I punished him. In the past, discipline meant I got a few hours, perhaps as much as half a day, of the cold shoulder treatment from him. This time, after about five minutes, he walked up to me, plopped down on my lap and laid his head on my chest.

I thought about what had opened Derek's heart to me. I hadn't been giving him more presents, as our reduced income has eliminated most trips to the toy stores, video arcades and amusement parks. I had changed my behavior in only one way I spent more time with him at home.

Looking back, I know I wanted so much to believe I could have it all, that I had accepted the peculiar notion that the quality of my time with Derek could make up for the lack of quantity.

But, I made no such assumption in any other area of my life. When I went out to dinner with Debbie, I made sure we could linger at the restaurant as long as we wanted. When I wanted to become a better jazz pianist, I practiced more. And I'd never told one of my bosses that I could go home early because the morning had been quality time.

Of course, spending more time with Derek leaves less time for other parts of my life, but I don't mind at all, because of:

LESSON NO. 3

Being a househusband is worth it.

One of the most demanding jobs in the world is also one of the most rewarding, with times where my joy is as strong as any I've ever felt. I've always known I feel something special for my son, but in these past few months I've grown hopelessly, uncontrollably, head-over-heels in love with him.

Since I've been home, I've gotten to play catch with him, and hear him squeal with delight when I toss a fluttering knuckleball--that he then asks me how to throw; take him to the doctor when he was sick, and hold and comfort him when a nurse had to take blood and give him shots; watch him imitate me as I exercise and hear him say, "Hey, Dad, we're doing the same thing and we're twins, right?"; watch him as he sleeps at the end of the day and feel amazed at the intensity of my love for him.

I got a job. My time at home ended. I accept the reality that, as Robert Bly says, the Industrial Revolution has assigned men jobs outside of the home. Plus, we need the extra money and security that a steady job will provide, and I genuinely enjoy my public relations career. But for as long as this period of unemployment lasted, I'm thankful.

-----Rholan Wong lives in West Los Angeles.

Date: Sun Sep 06, 1992 8:03 pm PST
Subject: control/influence discussion

[From Bill Powers (920906.1900)]

Greg Williams (920906 - 2) --

If we can manage to boil down these long posts into some simple issues, we may finally get somewhere. I'll try to be brief.

1. I have said that control by disturbance is not important, doesn't matter, to the controllee. This is because this mode of control DOES NOT DISTURB ANYTHING THAT MATTERS TO THE CONTROLLEE, one way or the other (harming or helping). The effect of the disturbance on the controlled variable is canceled by the action that the controller perceives and controls. If taking these actions (changes in lower-order reference signals) produces side-effects in other higher-order control systems in the controllee, those side-effects are also cancelled. So there is no important effect on any controlled variable in the controllee. IF THESE CONDITIONS ARE NOT MET THEN THE CONTROLLEE WILL PUSH BACK AND THERE WILL BE CONFLICT.

2. Now, "don't-care" variables and perceptions. If a perception is truly a don't-care perception, then changing it causes no error of any kind in the controllee: the change neither violates a goal nor satisfies one. The controllee considers the result neither good nor bad, judgements which always take place relative to a reference level. If, however, the perception is an uncontrolled component of a successfully-controlled perception (the position of a mouse when a cat is chasing and catching it), then changing it (which is easy to do because it is not directly controlled by the other) will produce a disturbance of the controlled perception, and the effect of the disturbance (the mouse's relative position) will be canceled before it becomes significant, through altering the other components of the perception (the position of the cat, which the cat can control). This will entail changes of reference signals at lower levels, but not at the level where the disturbance is actively opposed.

To sum up so far:

A controller is free to control the actions of another person without conflict if and only if neither the controlling influence nor the elicited actions cause uncorrectable errors in the controllee.

A controller is free to control [environmental variables corresponding to] uncontrolled perceptions in the controllee without conflict only if the perceptions in the controllee are truly don't-care cases -- that is, not components of controlled perceptions and therefore not related to any goal of the controllee.

To go on.

3. Purposeful influence on another that does not involve control of the influencee (or violence in some form) can be purposeful only in the sense that the influencer purposefully controls environmental variables on which perceptions in the influencee are thought to depend. As long as these environmental variables are not involved in the influencee's own control actions, this can be done without conflict.

4. There are only two ways in which external influences (purposefully manipulated or naturally-occurring) can have an effect on the influencee: through altering sensory information, or through affecting the physiology of the influencee directly. Effects on physiology fall in two classes: sensed effects, and intrinsic effects. These classes may have a nonzero intersection.

5. If the effects are sensory, the results are covered in numbered paragraphs 1 and 2 above.

6. If the effects are intrinsic, then the influencer is actually applying coercive force to the influencee, inducing malaise, pain, injury, hunger, thirst, suffocation, or other signs of malfunction. This almost certainly will require overriding the efforts of the influencee to prevent these changes. This violates the condition that we are discussing non-coercive influence of another's behavioral system.

The argument so far thus leads to the conclusion that while a putative controller can control certain actions of the controllee, and can control

environmental variables not related to any goal of the controllee, in no case can the effects disturb any variable under control by the controllee without inducing conflict between controller and controllee.

This sums up what HPCT has to say about the possible interactions of people with respect to their controlling each other without coercion.

I wrote a lot more, but this is enough to talk about. I don't think there is any ideology in the above: it is, as far as I could make it, a straight deduction from the principles of HPCT as they stand today.

There is one passage in your post to note before quitting for now:

>... the ideology of some PCTers which denies any "importance" to a
>person of the influences (current and historical) of other people.

I have never said any such thing. The influence of other people on me, current and historical, has been of tremendous importance to me. But assigning and defining that influence and its importance was my act, never theirs. Just as you decide for yourself what importance to give my offerings concerning HPCT, and what influence to allow them to have on you.

Best, Bill P.

Date: Sun Sep 06, 1992 10:33 pm PST
From: Uwe.Schnepf
MBX: Uwe.Schnepf@gmd.de
Subject: information packet

Dear Dag,

I would be very interested in the information packet as well. I am working in behaviour-based robotics, and I would like to see what else is around in the field of control.

many thanks

Uwe Schnepf
AI Research Division
German National Research Center
for Computer Science (GMD)
Schloss Birlinghoven
P.O. Box 1316
5205 Sankt Augustin 1
Germany

e-mail: usc@gmdzi.uucp
usc@gmdzi.gmd.de
phone: +49-2241-142704
fax: +49-2241-142618

Date: Mon Sep 07, 1992 3:37 am PST
Subject: I agree (S & I IV)

From Greg Williams (920907)

>Bill Powers (920906.1900)

>If we can manage to boil down these long posts into some simple
>issues, we may finally get somewhere. I'll try to be brief.

Sounds reasonable to me.

>1. I have said that control by disturbance is not important, doesn't
>matter, to the controllee. This is because this mode of control DOES
>NOT DISTURB ANYTHING THAT MATTERS TO THE CONTROLLEE, one way or the
>other (harming or helping)... So there is no important effect on any
>controlled variable in the controllee. IF THESE CONDITIONS ARE NOT MET THEN
>THE CONTROLLEE WILL PUSH BACK AND THERE WILL BE CONFLICT.

I agree. Your "control by disturbance" is my "manipulation." Manipulation works as you say, requiring "cooperation" of the manipulee. I must say, however, that "control by disturbance" is important in another way to BOTH the controllee AND the controller, since they ARE both controlling in ways contributory to the outcome. And, in this sense, I think what we control for IS important to us.

>2. Now, "don't-care" variables and perceptions....

I agree. Again, I say that this is how "manipulation" works.

>A controller is free to control the actions of another person without
>conflict if and only if neither the controlling influence nor the elicited
>actions cause uncorrectable errors in the controllee.

I agree.

>A controller is free to control [environmental variables corresponding to]
>uncontrolled perceptions in the controllee without conflict only if the
>perceptions in the controllee are truly don't-care cases -- that is, not
>components of controlled perceptions and therefore not related to any goal of
>the controllee.

I agree.

>3. Purposeful influence on another that does not involve control of
>the influencee (or violence in some form) can be purposeful only in
>the sense that the influencer purposefully controls environmental
>variables on which perceptions in the influencee are thought to
>depend. As long as these environmental variables are not involved in
>the influencee's own control actions, this can be done without
>conflict.

I agree.

>4. There are only two ways in which external influences (purposefully
>manipulated or naturally-occurring) can have an effect on the
>influencee: through altering sensory information, or through affecting
>the physiology of the influencee directly. Effects on physiology fall
>in two classes: sensed effects, and intrinsic effects. These classes
>may have a nonzero intersection.

Sounds OK to me. I'm not waffling because of ideology, just would need to have some time to think about the technicalities. I agree for now, and I don't think it would alter our debate if later I noticed a minor "but if."

>5. If the effects are sensory, the results are covered in numbered paragraphs 1 and 2 above.

OK

>6. If the effects are intrinsic, then the influencer is actually
>applying coercive force to the influencee, inducing malaise, pain,
>injury, hunger, thirst, suffocation, or other signs of malfunction.
>This almost certainly will require overriding the efforts of the
>influencee to prevent these changes. This violates the condition that
>we are discussing non-coercive influence of another's behavioral
>system.

I agree.

>The argument so far thus leads to the conclusion that while a putative
>controller can control certain actions of the controllee, and can
>control environmental variables not related to any goal of the
>controllee, in no case can the effects disturb any variable under
>control by the controllee without inducing conflict between controller
>and controllee.

Yes. This is why I originally was FORCED (by PCT) to define manipulation as I did. It works by manipulee "cooperation."

>This sums up what HPCT has to say about the possible interactions of
>people with respect to their controlling each other without coercion.

I don't think so. I believe that PCT also says that how an individual is controlling at any time depends significantly on past interactions with others (and the environment in general) as well as on his/her previous controlling. Resolving this difference between us will involve examining and evaluating candidate models of reorganization, some of which, I believe, imply that "new" controlling (after a reorganization) does not SOLELY depend on previous controlling (and perhaps an "unlawful" internal switching of some sort), but also on the environment.

>I wrote a lot more, but this is enough to talk about. I don't think
>there is any ideology in the above: it is, as far as I could make it,
>a straight deduction from the principles of HPCT as they stand today.

I agree, except for the last statement.

GW>>... the ideology of some PCTers which denies any "importance" to a
GW>>person of the influences (current and historical) of other people.

>I have never said any such thing. The influence of other people on me,
>current and historical, has been of tremendous importance to me.

That is why I think many people will welcome a theory of influence based on PCT. We are ALL tremendously affected by each other.

>But assigning and defining that influence and its importance was my act,
>never theirs. Just as you decide for yourself what importance to give my
>offerings concerning HPCT, and what influence to allow them to have on you.

Here is where ideology might be overriding deduction (from PCT). Your statement is in line with the (only) statement above with which I disagree. It appears to me that you have left out "half" (figuratively) of the full implication from PCT, the half about your (current) deciding for yourself what importance to give something being the conjoint product of BOTH your past controlling AND environmental effects on you. To leave out the second half implies its unimportance, and that is simply autonomistic ideology, not a deduction from PCT. PCT says that no control system is an island -- or, rather, that "a control system" does NOT include only the ME, but the ME's environment also. The loop closes THROUGH the environment (often containing other MEs -- I suppose you would prefer THEYS).

Thanks for being so cogent. I do think we are making progress.

Best, Greg

Date: Mon Sep 07, 1992 5:57 am PST
Subject: replies and a demo

[From Hans Blom (920907)]

Some replies and remarks. Sorry for the delay. My modem broke down and I had to wait for a new one.

Avery Andrews 920901.0957

>The anesthesia monitoring systems sounds like an instructive integration
>of thought and action - is there a comprehensive writeup anywhere?

Expert control of the arterial blood pressure during surgery. International Journal of Clinical Monitoring and Computing 8: 25-34, 1991.

Rick Marken (920831.1100)

>Right. "The test for the controlled variable" is how you determine whether
>or not control is involved in any observed behavior.

I agree. But... This isn't always feasible. Please run my coli simulation (included below). Now 'The Test', in my opinion, would mean to pick up coli, put him down somewhere else, observe its behavior and check whether the change in position is resisted or undone. I do not see resistance. I do not see behavior that is different from coli's normal behavior. Or is it there, but is the problem that I cannot see it?

As I see it, coli is too primitive to be able to control. At most it _tries to_, but doesn't quite succeed. Maybe, in its environment, this is optimal?

>> Two: what do we mean by 'control'?

>Production of consistent results in the face of disturbance.

I do not recognize consistency of results when there is so little regularity in behavior in the first place.

>By the way. Your description of your e. coli model was somewhat puzzling.

>When, for example, does your e. coli model change to a new direction?

At every step, i.e. every new run of the simulation.

>What is the criterion for change?

There is no criterion for change. That would require a memory, which I do not think coli has. Every decision is new, based on the current observation only.

>In our models, a "counter" is set after each tumble. The rate at which it

>counts down depends on the currently experienced gradient of attractant.

In my model, coli has neither a counter (memory) nor a sensor to sense gradients. I assume it is too primitive for either.

```
program e_coli;
```

```
    { simulation of E. coli behavior }
{ written in Turbo Pascal for IBM-PC or compatible

simulation ends      1. when a key is pressed
                    2. when coli escapes into infinity
                      (which so far it always did, eventually)}
```

```
uses
```

```
  crt, graph;
```

```
const
```

```
  scale = 1.4;      {graphics scaling; adjust to zoom in/out}
  rad0 = 50.0;     {initial radius; coli is let go here}
  radx = 100.0;    {influence radius; outside this radius sensor no good}
  radw = 10.0;     {optimum feeding radius}
  rmax = 500.0;    {beyond this radius coli is considered lost}
```

```
var
```

```
  x, y, r:         real;      {current coordinates and radius}
  xm, ym, xxm, yym: integer;  {screen pixel sizes}
  xs, ys:         real;      {scaling constants}
  c:              char;      {stop simulation?}
```

```
procedure grinit; {initialize graphics stuff}
```

```
var
```

```
  driver, mode: integer;
```

```
begin
```

```
  detectgraph (driver, mode);
```

```

initgraph (driver, mode, '');
if graphresult <> grok then
  begin writeln ('cannot find graphics driver'); halt end;
xm := getmaxx; ym := getmaxy;
xxm := xm div 2; yym := ym div 2;
xs := scale * xm / 200.0; ys := scale * ym / 150.0;
{ plot circles of radii 10, 50 and 100 }
ellipse (xxm, yym, 0, 360, round ( 10.0 * xs), round ( 10.0 * ys));
ellipse (xxm, yym, 0, 360, round ( 50.0 * xs), round ( 50.0 * ys));
ellipse (xxm, yym, 0, 360, round (100.0 * xs), round (100.0 * ys))
end;

procedure move; {take one step}
var
  s, u, dx, dy : real;
  t: string [12];
begin
  {calculate response strength}
  s := r; if r > radx then s := radx;           {s = 0 .. radx}
  s := 0.2 * abs (s - radw);                   {step size, but ...}
  if s < 0.1 then s := 0.1;                   { ... minimum step size}

  {calculate random direction unit vector}
  dx := 2 * random - 1; dy := 2 * random - 1; {select random direction}
  u := sqrt (dx * dx + dy * dy);
  dx := dx / u; dy := dy / u;                 {random unit vector}

  {do one move of length s}
  x := x + s * dx; y := y + s * dy;           {update position}
  r := sqrt (x * x + y * y);                  {compute new radius}

  {plot position as a pixel}
  putpixel (xxm + round (xs * x), yym + round (ys * y), white);

  {print radius in upper left corner}
  setviewport (1, 10, 60, 20, true); clearviewport;
  str (r:6:0, t); outtext (t);
  setviewport (0, 0, xm, ym, true)
end;

begin
  randomize;                                  {initialize random generator}
  grinit;                                     {initialize graphics}
  x := sqrt (0.5 * rad0 * rad0); y := x;      {initialize position}
  r := sqrt (x * x + y * y);
  repeat
    move                                       {take one step }
  until (abs (r - radw) > rmax) or keypressed; {till coli lost or you tired}
  c := readkey;
  closegraph
end.

```

Date: Mon Sep 07, 1992 7:41 am PST
Subject: Influence; e. coli model

[From Bill Powers (920907.0900)]

Greg Williams (920907) --

Let's stay with the basic ideas for a while and make sure we understand each other.

>Your "control by disturbance" is my "manipulation." Manipulation works as
>you say, requiring "cooperation" of the manipulee. I must say, however,
>that "control by disturbance" is important in another way to BOTH the
>controllee AND the controller, since they ARE both controlling in ways
>contributory to the outcome. And, in this sense, I think what we control
>for IS important to us.

I'm controlling variable v by means of action a . You apply a disturbance d to v . This causes v to depart from v^* , my reference level for it, by approximately $d/\text{myloopgain}$; my action causes a to vary enough to produce this result, so that a is nearly $-d$. So by varying d , you can cause a to vary. This is your control of my action. You can have a reference level a^* for a . My action a will come to a state of about $a^*/\text{yourloopgain}$.

By this way, this takes no cooperation from me -- I don't have to pay any attention to my own action. All I have to do is continue doing what I was doing before: keeping v near to v^* . From my point of view, your variations of d are just another disturbance. It doesn't matter to me whether those variations are systematically aimed at a goal of yours or are random.

By our postulates (which forbid conflict), I must not have any goal that forbids my action a to be controlled by you. So my action can't be "important" to me -- that is, I can't have any goal for it, but must be willing to let it vary as you choose. I can't have any preference for a particular state of my action if you are to be able to control it without conflict with me. This procedure works best if I (higher level) pay no conscious attention to my actions at all, but just let the (lower-level) control system operate them.

It has just occurred to me that we are probably using the term "important" in different ways. I suspect that you're using it to mean objectively important -- that is, important in ways unknown to me but known to the manipulator (in the manipulator's opinion). The manipulator may believe that if he can get me to bring my action into the state a^* , something good for me will result, or something bad for me will be avoided. The manipulator may, for example, have the opinion that if I could be maneuvered into learning the moves that result in taking square roots correctly, this action by me will be of future benefit to me even though I'm indifferent to it now and am not doing it "on purpose."

Is this what you're talking about when you speak of importance?

Hans Blom (920907) --

I puzzled mightily over your e. coli Pascal model until I went back and read your description: "In my model, coli has neither a counter (memory) nor a sensor to sense gradients. I assume it is too primitive for either."

What I was puzzled about was the lack of sensing of the gradient -- there can be no control system without sensing. I was saying to myself, "Why, this is just a random-walk generator -- how could it ever go up or down a gradient?" The answer is, of course, that it can't, except by accident. This is why it always escapes off the screen eventually.

To make this model behave like e. coli it's necessary to make the step size, or the delay between tumbles, depend somehow on the sensed rate of change of concentration. Your model is much too simple; the real e. coli (the intestinal bacterium) does have such rate-of-change chemical sensors, and does vary the interval between tumbles according to the time rate of change of concentration.

What you have shown is useful, in that it shows that a bacterium that is too simple -- that is not a control system -- can't do what e. coli actually does.

Rick, why don't you post the code for our e. coli model? I'm sure that Hans can translate the Apple-specific parts into his version of Pascal.

By the way, Hans, could you set your right margin to about 70 or 72? Your lines are exceeding the screen width here, and I end up with lots of single words on a new line. I suspect this is true for most others too.

Best, Bill P.

Date: Mon Sep 07, 1992 9:55 am PST
Subject: Sciences & Ideologies V

From Greg Williams (920907 - 2)

>Bill Powers (920907.0900)

>Let's stay with the basic ideas for a while and make sure we
>understand each other.

Sounds good to me.

>I'm controlling variable v by means of action a. You apply a
>disturbance d to v. This causes v to depart from v*, my reference
>level for it, by approximately d/myloopgain; my action causes a to vary
>enough to produce this result, so that a is nearly -d. So by varying
>d, you can cause a to vary.

So far, so good.

>This is your control of my action.

>You can have a reference level a* for a.

I can? By PCT, I can ONLY control for MY PERCEPTIONS. (I could say HOW MANY TIMES DO I HAVE TO SHOUT THIS?, but I won't.) Now, I could have a reference level corresponding to my perception that you are performing action a* -- and my perception could be flatly wrong, for some reason or another; or I might have a reference signal for some other perception which is (possibly complexly) RELATED to your action which I am manipulating.

>My action a will come to a state of about $a^*/\text{yourloopgain}$.

IF I am controlling for my perception that you are performing a^* and I am not fooled in some way.

>By this way, this takes no cooperation from me -- I don't have to pay
>any attention to my own action. All I have to do is continue doing
>what I was doing before: keeping v near to v^* . From my point of view,
>your variations of d are just another disturbance. It doesn't matter
>to me whether those variations are systematically aimed at a goal of
>yours or are random.

I'm sorry I led you to this confusion. What I meant by saying in my previous post that "cooperation" is necessary for manipulation to work is simply that the manipulee's control must NOT be conflicted. The manipulator intends (and how well he/she succeeds depends on how good a model he/she has of the manipulee's controlling) to "set up" the situation so that the manipulee will (given that situation) "naturally" control so as to allow the manipulator to control certain of his/her own perceptions. So, the conned mark must be "cooperatively" greedy for the game to "work" (with respect to what the con man wants to happen).

>By our postulates (which forbid conflict), I must not have any goal
>that forbids my action a to be controlled by you.

Yes, that's what I was just trying to say.

>So my action can't be "important" to me -- that is, I can't have any goal for
>it, but must be willing to let it vary as you choose.

You've ALMOST got it. I can set up a situation so that you will act in the way you want to act (given the situation) AND I can control the perceptions I want to control. What I CANNOT set up (for manipulation to work) is a situation so that you will want to NOT act that way. So the action in question can be EITHER "unimportant" (your sense) to you or "important" to you -- it CANNOT be "important" to you that you DON'T do the action I want you to do. Your controlling (especially: HOW YOU PERCEIVE THE SITUATION -- and you might MISperceive it or perceive it INCOMPLETELY, because an exploiter has set it up thusly, using his/her model of your controlling, which could be wrong) must be such that you will "go along" (I early said "cooperate," but that was confusing) with the manipulation.

>I can't have any preference for a particular state of my action if you are to
>be able to control it without conflict with me.

No. You can prefer to do it; you cannot prefer to not do it (if I am able to control your action without conflict). If you are greedy and the con man shows you (you believe at the time) how to get a lot of money by "going along" with him, you might get to the point where you are even suggesting how the (unbeknownst-to-you) scam should be played out. You can positively LOVE to be performing the actions which you THINK will satisfy your wish for moola. Your gain for those actions can be sky-high (and that makes the manipulator's job very easy). At any rate, I can see (sometimes) a fairly loose coupling between the manipulator's controlled perceptions and the manipulee's actions which

result in control by the manipulator of his/her perceptions. The short answer: by "fostering" (another confusing word?) your "natural" control of particular actions which you might have VERY STRONG PREFERENCE FOR, I can manipulate you. Manipulations of truly "don't-care" actions are also possible. We are indeed progressing in interesting directions!

>This procedure works best if I (higher level) pay no conscious attention to
>my actions at all, but just let the (lower-level) control system operate
them.

Not necessarily, as per my comments immediately above.

>It has just occurred to me that we are probably using the term "important" in
>different ways. I suspect that you're using it to mean objectively
important
>-- that is, important in ways unknown to me but known to the manipulator (in
>the manipulator's opinion).

We needn't get THAT objective. How about just "judged by you AFTER the
manipulation of you (or after you know the full "facts" about an attempted
manipulation) to be important to you? But I see that you're catching on.

>The manipulator may believe that if he can get me to bring my action into the
>state a*, something good for me will result, or something bad for me will be
>avoided.

Yes, and I would term that an attempt at "facilitation" -- it would BE
"facilitation" if the manipulee, after the successful manipulation
[facilitation is a KIND of manipulation], judges the result to be "good" for
him/her. I'll go out on a limb and say that I think our major difference
overall is that I think successful manipulations (both facilitations and
exploitations) are extremely common, and you think they are extremely rare. Am
I right about that? It is an EXTRA-PCT question -- an empirical issue which I
think depends on the particulars of individual manipulations. PCT says only
that manipulation CAN fail, and that, in general, the Test probably can be
used to increase the chances of a manipulation succeeding. Most importantly,
PCT does NOT say that manipulation is necessarily bad for the manipulee. And
it does not say that ONE kind of manipulation is always best -- that's what I
chided Rick about "prescribing."

>The manipulator may, for example, have the opinion that if I could be
>maneuvered into learning the moves that result in taking square roots
>correctly, this action by me will be of future benefit to me even though I'm
>indifferent to it now and am not doing it "on purpose."

Well, I didn't think of that -- luckily, I didn't say HOW long after the
completion of the successful manipulation the manipulee should judge its worth
to him/her. At any rate, I suppose that in many cases, the manipulee might not
be able to say whether the manipulation was beneficial or not -- perhaps NEVER
could. A very interesting point (remember I had said that the manipulee might
find it difficult to judge the worth of propaganda and advertising, in at
least some cases) -- that's something that makes life complex and interesting,
I think!

>Is this what you're talking about when you speak of importance?

Basically, yes -- as laid out above.

Best, Greg

P.S. Shouldn't you tell Hans B. about how E. coli actually controls, based on Koshland's work? He seems to think it is conjectural, not empirically based.

Date: Mon Sep 07, 1992 1:19 pm PST
Subject: control and children

from Ed Ford (920906:1415) Rick Marken (920906)

>The parent has no choice, right?

I have no idea what you mean by this statement.

>You don't have to tell a control system (the parent in this case) to
>control.

Like any group of people organized to live or work together such that they can cooperatively get along, many parents seek help and instruction (I was one of them) on the best way to set standards in such a way that everyone would cooperatively get along. Control systems naturally control, but in any community or between any two individuals, part of growing and maturing is to learn to control in such a way that you do no harm to yourself and that you learn to respect the rights of those in your environment. That is not something we do naturally, and thus, in a family setting, I believe it is the job of the parent to teach their children in as patient and loving way as possible how to use their control systems to their greatest advantage and with the least amount of violence to anyone's system.

>They are "natural consequences" only if the kid's behavior is part of
>what the parent is controlling; the parent can't help but "naturally"
>take action (ANY action) in an effort to move the kid's behavior back
>to the parent's reference on it.

There is a big difference between trying to "move the kid's behavior back to the parent's reference on it" and trying to teach children in conflict to look for (reorganize safely) a better way of solving their internal conflict which is reflected in their actions that violate the rights of others. For example, hitting, screaming, yelling, cursing by children violates the rights of others in their community. I believe their social contact ought to be reduced to a point where they harm no one (sent to a smile chair or their room) until they make the choice to work things out within the community in which they live. That would be the natural consequence of such action. Beating the kid, spitting in his face, chaining him up, all that is punitive. Just like the natural consequence of speeding is for a policeman to give you a ticket. Pulling you out of the car, shooting you through the foot that operates the accelerator, that's punitive. Once children have a reference for working things out, then they PERCEIVE THE PARENT AS A TEACHER who teaches them how to make an effective plan, namely, establish an alternative way of work at solving the problem they had in the first place in a satisfactory manner.

>How do you "non-punitively" get a kid to produce a perception for
>themselves (and, incidentally, for you) that they have no intention of
>producing?

You don't. You can only work with those who are willing to work with you, or who are willing to set a reference for your help. What I do is first try to find a reference within children that they are willing to deal with which will lead toward the solving of a problem. I don't call this manipulation. In working with the juveniles in lock up, which aren't exactly your nicest people to work with, I first ask (if they're upset) if they want to talk or should I come back later. If they say they want to talk and I then ask them what they want. If they express confusion or they're upset, and they say they don't know what they want, I might ask if they are satisfied with the way things are. If they are ticked at someone, perhaps a guard or case manager, I'll ask them if they'd like things to be better. I just keep going till I can achieve something with the juvenile, some kind of a plan, something. I don't call that punitive, or permissive. And I'm certainly not manipulating what you call his behavior. As long as the juvenile has a reference or willingness to work with me, I can deal with him/her, but I'm not controlling his behavior. I'm teaching him (providing he is willing) how to deal with his internal world in such a way that his/her life becomes better. You'll never help people unless they perceive you as someone from whom they want help and secondly, they perceive you as someone that can help you, they perceive you as a teacher.

I think when you use the word manipulation you have to consider how each person perceives the other person. If a juvenile perceives me as someone who can help them, someone who can teach them a better way of dealing with their life, and they submit to that person's help, I don't see that as manipulation any more than I see Bill Powers manipulating me when I seek his counsel or help. I think the same is true of our children. If they perceive me as a loving parent who cares and believes in them and believes in their ability to succeed in life, and I ask them if they want my help and they say yes, I don't see that as manipulative. With children who really refuse help or don't want to deal with me at the time of the problem, sending them to their room (or to a smile chair) is trying to teach them at a time they may not want to be taught. But later on, assuming a close relationship, they generally are appreciative of my actions.

You keep talking about controlling behavior as if it can be controlled. What I am talking about is something entirely different. I'm talking about TEACHING RESPONSIBLE THINKING. It is how we think that determines our actions. It is how we focus on our various references, how we perceive a variety of things, only a small part of which are the actions we take. I really don't think I can control how a person thinks through a problem. I can ask questions in such a way that help them (if they have a reference for it) or teach them how to focus in areas about which they may want to reflect on, to evaluate, etc. But first they have to want to think through a problem before they are willing to deal with me. Forcing an upset child to a room doesn't force their thinking, it is only the non-punitive consequence of the violation of another's rights. Their thinking is what is going to ultimately determine their actions and their future happiness.

>How do you non-punitively" get a kid to produce a perception for
>themselves (and, incidentally, for you) that they have no intention of

>producing?

You don't. I don't think you'll ever get someone to produce a specific perception. First place, you haven't any control over getting anyone to produce a perception or a reference. Secondly, you wouldn't know what they produced because they would be explaining it to you through their words and you would be perceiving their words, reflecting on what you'd perceive to be their thoughts and what you'd create in your own world would be different from their world. As I said above, all you can do is help them deal with their own internal world, the only world they'll ever know.

>The alternative to manipulative child rearing is not to become a doormat.
>It's to become a TEACHER. The goal is to help the kid learn to control more
>skillfully --so that it's activities don't have unfortunate side
effects...

That's exactly what I've been talking about. Perhaps you've perceived what I've been suggesting as manipulative, but it isn't. Or, what is teaching, if it isn't teaching a person how to recognize, think through, use, and evaluate the areas within their system where thinking is done, namely the various references, at different levels, what and how they perceive various areas within their perceptual system, compare, etc. I'm not making anyone do anything. If my children or an employee or a juvenile in a detention center doesn't want to work with me, he'll have to live with the consequences of that choice. But he does have that choice. And I can't force him to allow me to teach him.

As you've said before, it's all perception. If a person perceives you as someone with whom they want to interact, concerning no matter what, they'll set a reference to do just that. If they perceive you as someone they dislike, don't trust, or whatever, they'll probably reject your offer. All I see myself as is a teacher, which is someone who cares and who is willing to help others make a better world for themselves. But I can't do a thing unless they set a reference for me to help. That to me is what it is all about.

Ed Ford

Date: Mon Sep 07, 1992 6:50 pm PST
Subject: Flat Earth Society

[From Hank Folson (920907)]

>[From Bill Powers (920906.1830)]
>Rick Marken (920906) --
>Rick, I think Hank was trying to say that the Flat Earth Society
>(FES) is a deliberate put-on; that they support "wrong" ideas as
>a matter of principle. Is that right, Hank?

It is, but I messed up my communication. I was surprised to get any response, as I did not intend to hit a controlled variable. A more practiced practitioner would have stopped after:
>>If this group had formed before 1492, the name would then have
>>been the Round Earth Society.

But I added in an attempt at humor:

>> ...They are our brothers...<< Had I not linked PCT and the FES the error signals would not have been created. Mea Culpa. This has been a good PCT lesson for me. I'm learning.

I will not be active on the net for the next 3 weeks. I just learned that I can share a booth at the major bicycle trade show and must prepare. Part of my preparation will be rewriting a 50 page booklet promoting custom bicycle frames to reflect the principles of PCT. This will be used in an ad campaign aimed at potential buyers of custom frames. A couple of conversations with people on the net have already led to significant changes.

Au Revoir, Hank

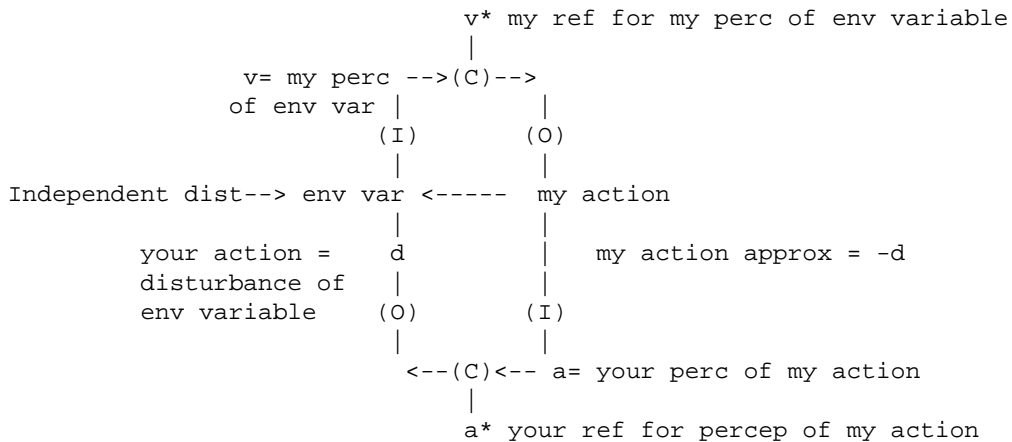
Date: Mon Sep 07, 1992 9:02 pm PST
 Subject: Influence; Ed's advice; e. coli apology

[From Bill Powers (920907.2030)]

Greg Williams (920907-2)

We seem to have a technical disagreement.

Here is a diagram of you controlling your perception of my action while I control my perception of the environmental variable you are disturbing:



This is a non-conflict situation. The two control systems can act independently. You can make your perception a equal your reference signal a* quite independently of my making my perception v equal my perception v*. I added the independent disturbance later -- we can still both control our

perceptions independently even with an independent disturbance acting on the environmental variable.

However, I MAY NOT have a second control system that prefers any particular value of my action. My action must depend, in amount and direction, strictly on the error, $v^* - v$. If some other system in me wants a particular value of the same action, then I will be in internal conflict.

>>You can have a reference level a^* for a .

>I can? By PCT, I can ONLY control for MY PERCEPTIONS. (I could say HOW MANY >TIMES DO I HAVE TO SHOUT THIS?, but I won't.) Now, I could have a reference >level corresponding to my perception that you are performing action a^* -- >and my perception could be flatly wrong, for some reason or another; or I >might have a reference signal for some other perception which is (possibly >complexly) RELATED to your action which I am manipulating.

You're talking about perceiving THAT I'm performing an action a^* , which is very different from what the above diagram shows. The above diagram applies only when you're perceiving an action that I am using to control a perception, and when you're using a disturbing variable to affect the variable I'm controlling. You can't have any systematic effect on your own perception if you've guessed wrong about either the variable I'm controlling or the action I'm using to control it. You don't have to be EXACTLY right, but you have to be reasonable close. If you're far off, then your action will have only a statistical effect on my action as you perceive it, or something related; then you'll have classical S-R theory in its full glory, statistics and all.

And of course you can't pick a reference level for the amount and direction of my action until you're perceiving the correct action.

You can have reliable control over my action only when your perception of it covaries closely with the state of my output effects, and when your action directly disturbs my perception. Also, this reliable control depends entirely on my success in keeping the environmental variable, or my perception of it, in a close match to v^* . If I lose control, you do too.

>>I can't have any preference for a particular state of my action if you >>are to be able to control it without conflict with me.

>No. You can prefer to do it; you cannot prefer to not do it (if I am >able to control your action without conflict).

The amount and direction of my action must be totally determined by my error signal. I can't also prefer that it be in any particular state. So I cannot prefer to produce a large amount of the effort, or no effort, or an effort in the opposite direction. Since that is the range of your choices for my action (if you're perceiving the right one), I can neither hinder nor help you. That's what I mean by saying that you can control my action only if I have no goal for it: only if its state has no importance to me.

If you try to preserve your argument by going up a level, the same problem will arise. If my action consists of selecting from a range of (lower-level) actions, then your perception will be which action in that range I am using. A disturbance of my controlled variable would then require shifting from one

KIND of action to another. Now the control system must be free to select any action within the range of kinds, depending on the disturbance. There cannot simultaneously be another system that prefers for a given action to be used, or not to be used -- if inner conflict is to be avoided. Again, you cannot control for my using a particular kind of action unless I have no preference at all, for or against any of the actions. That's because I can't control if I have such a preference. Your success at controlling my action (now with regard to kind) depends on my success in varying the kind of action to control my controlled variable, whatever it is. And my success depends on leaving the choice of kind of action completely free.

If it's important to you that I somehow prefer the action that you create by disturbing me, then you're going to have to reconsider whatever deductions you have made from this assumption. Whatever phenomena you're trying to explain in this way will have to be explained in a different way.

I'll come back to "importance" later. I don't think you have yet seen the impact of this subject on the general argument.

>P.S. Shouldn't you tell Hans B. about how E. coli actually controls,
>based on Koshland's work? He seems to think it is conjectural, not
>empirically based.

I thought I did -- didn't I send that post? When you eliminate objective importance, you're back to one control system interacting with another one. When you talk about objective importances, you've taken off your theoretician's hat and are just one control system among many, trying to work things out.

--
Ed Ford (920907) --

Those were good, solid, well-directed answers to Rick's queries. I think you have the proper alternative to manipulation -- to "making" people do things, whether for their own good or not. You express the basic idea that I'm trying to get across: that all impetus for change has to come from within the person. And you show how to get in contact with that inner desire to change.

Each person's continuing problem is how to set up an inner structure of goals that will work in the physical and social environment that exists. The actual setting up has to be done from inside. Even if you know EXACTLY what the person needs to do to get along better (however unlikely that is), you still have to get the idea across to the other, who has to understand it and translate it into the private terms in which all behavior is carried out -- and put it into effect. The limit of your influence on the other person is reached when you present your case.

When you write your next book, Ed, get that post into it somewhere.

-
Hans Blom (920906) --

I got your program to compile and run with no trouble. I see that it's not just a random-walk generator. The perceptual function is hidden in the condition for choosing long or short steps, so that the steps get longer as

the distance from the circle increases. In effect, you've given the bacterium the ability to sense distance (which could be done in a fixed gradient by sensing concentration), and you've given it a reference level for sensed distance -- the radius of the inner circle. So I apologize for jumping to the conclusion that this isn't a control system just from reading the code without running it.

From the way the model behaves, there seems to be some positive feedback in it -- the bigger the error, the larger the steps. So when the error gets above some amount, there's a runaway effect. I don't see what's causing it.

Did you see my post in which I described Koshland's findings about how e. coli is actually organized?

Best to all, Bill P.

Date: Tue Sep 08, 1992 2:36 am PST

From Greg Williams (920910)

>Bill Powers, recent posts

No time to look at most recent posts now -- headed for Lexington to take the Sept. HI to the printers. But I woke up with a clear vision of how what I have been terming "manipulation" goes beyond rubber-banding.

Throwing caution to the winds (i.e., being unpicky about absolute precision of wording):

Rubber-banding-type control of another involves disturbing a variable which you know to be controlled by another so they will "correct" for your disturbance by performing actions which you want to see.

Manipulation involves changing another's uncontrolled perceptions so that they will control in a way so as to perform actions you want to see. They will control in that way (assuming your model of the way they control is adequate) because, given the situation you have led them to perceive, that is the way they "naturally" control.

In rubber-banding, control works best when the controlee doesn't pay attention to the actions you want to see. Manipulation can be aided by the manipulee having a high gain for, and paying close attention to, the actions you want to see.

Does this clarify? See you tonight or tomorrow morning.

Best, Greg

Date: Tue Sep 08, 1992 7:44 am PST

Subject: Clarifying clarification

[From Bill Powers (920908.0730)]

Hey, Olson, don't take this personally.

In the con game this happens when the wallet (v1) containing the valid bond or Treasury note or whatever (v2) is switched for the similarly- wrapped one containing newspaper. The Mark has control only of the wallet (v1), but is shown by the con man that the wallet contains something valuable (v2). The Mark doesn't want to perceive the wallet alone (v1) but the wallet with something valuable in it, which is the perception $p = f(v1,v2)$.

The action that gets the wallet and v2 into possession of the Mark is to hand over cash to the con man. This does get the wallet into the possession of the Mark. The Mark perceives this, however, as getting a wallet with something valuable in it. The con man sees to it that the Mark gets the wallet, but without the something valuable in it.

This is the same con game played by advertisers; when you hand over your money, you get the flashy car, but the nice pretty girl isn't in it any more and never will be. What you get is a dumb promiscuous broad who can be conned into thinking that a flashy car contains a terrific guy, and when she hands over the price she will give you a baby you didn't particularly want, or AIDS. And when she looks in her package all she will find in it is plain old you. Great design for a relationship.

I agree with you that lots of this kind of deception goes on. What's good about it?

The only way for the Mark to stop being conned is to make sure it's his own perception of v2 that's in the package. This goes whether he is buying wallets from strangers, cars from television advertisers, religions from books, or educations from experts. Or theories from anybody.

Best, Bill P.

Date: Tue Sep 08, 1992 9:04 am PST
Subject: stella; the first naive question

i've been working with stella on the mac for a while now, and the recent set of inquiries with regard to its utility in pct modeling motivated me to play with it. a couple of disconnected comments, and a question, now arise...

1) the last time i talked to high performance systems, they did not mention anything about porting stella to the pc (at least i think someone here asked that question). at the same time, neither is dynamo going to go to the mac. unfortunate, for while it is not nearly as slick, it is a lot more flexible in some ways.

2) it occurred to me that the classic system dynamics diagram, as developed by forrester, takes into account one of the central notions of pct, about which mr powers was trying to educate me earlier... equations which represent an entity's response to a change in a system level always take as input, not the actual state of that level, but a transformation of it, often labeled "perception of...".

3) i put the following plot into stella the other night:

```

-----level<-----valve----->source/sink
--->perceived level          /\

                               \/
reference----->error----->rate of change
level          /\

                               disturbance

```

where the only material flow is the increase/decrease of the level, and all other connections are information flows. the first question is: does this at all capture the basic pct structure (without the hierarchy)? it looks to me like the basic cybernetic structure, and follows the essentials of forrester's method. now, if 'the entity' is that part of the structure which receive the perception, contains the reference, calculates the error, and acts upon the rate of change, then i think i have some idea of the distinction between the phrases 'controlling' and 'controlling for'. but i might also have it all wrong. it seems to me that the entity is controlling the environment (acting upon the rate of change) in order to obtain (is controlling for) a certain perception. yes? no? next question: is the disturbance a disturbance of the environment, as i have shown it, or a disturbance of the perceptual transformation? or both?

hmmm... i have also just thought of some other fundamental confusions on my part, not the least of which is the unclear distinction between entity and environment implied by the dependence of rate of change on error. but let's start simple, shall we?

```

-----< Cognitive Dissonance is a 20th Century Art Form >-----
Eric Harnden (Ronin)
<HARNDEN@AUVM.BITNET> or <HARNDEN@AMERICAN.EDU>
The American University Physics Dept.
4400 Mass. Ave. NW, Washington, DC, 20016-8058
(202) 885-2748
-----< Join the Cognitive Dissidents >-----

```

Date: Tue Sep 08, 1992 9:17 am PST
Subject: e. coli model, again

[Hans Blom (920908)] Bill Powers (920907.0900)

You say

```

>I puzzled mightily over your e. coli Pascal model until I went back and read
>your description: "In my model, coli has neither a counter (memory) nor a
>sensor to sense gradients. I assume it is too primitive for either." What
>I was puzzled about was the lack of sensing of the gradient -- there can be
>no control system without sensing.

```

Agreed. There can be no control system without sensing. But a control system can operate without sensing a gradient. My simulation shows just that. You can improve the quality of control by setting the minimum step size to zero.

Try it and see coli stay at its optimum spot for what I think is an indefinite period (that is, if it gets there at all).

I was saying to myself, "Why, this is just a random-walk generator -- how could it ever go up or down a gradient?" The answer is, of course, that it can't, except by accident. This is why it always escapes off the screen eventually.

No. The demonstration shows that it can, strange enough, walk down a gradient even if the gradient cannot be sensed. Change the minimum step size to zero and coli does not escape off the screen eventually but stays at the optimum radius once it gets there. Initially actions are almost random, eventually control is perfect.

>Your model is much too simple; the real e. coli (the intestinal bacterium)
>does have such rate-of-change chemical sensors, and does vary the interval
>between tumbles according to the time rate of change of concentration.

I understood that it had a chemical sensor, and I modeled that. I did not understand that it had stereo-sensing, so no modeling of that. In my model, the interval between tumbles (= step size) varies with the sensed concentration minus the optimal (reference) concentration. My implementation may be different (varying step sizes vs. a counter that determines the step size; i.e. my time scale is different), I think it does the same thing. A time rate of change, i.e. a memory, is not required in my simulations. Remember that my attempt in these simulations was to show the 'origin' of control: how simple can a system be and yet control (well enough).

Greg Williams (920907 - 2)

P.S. Shouldn't you tell Hans B. about how E. coli actually controls, based on Koshland's work? He seems to think it is conjectural, not empirically based.

Thanks for the reminder. I immediately abstracted away the real thing. I shouldn't call it a simulation of E. coli. It is not. It is an attempt to construct the simplest controller thinkable: one sensor, one goal, one comparator, one actuator, and the simplest possible modulation of actuator action by the error.

Bill Powers (920907.2030)

From the way the model behaves, there seems to be some positive feedback in it -- the bigger the error, the larger the steps. So when the error gets above some amount, there's a runaway effect. I don't see what's causing it.

No, it's the other way around: the large steps are normal. Steps get smaller when coli gets closer to its goal. At the goal, they have zero length (or are very small) so that coli stays there (for some time). It is as if - very loosely speaking - my coli has discovered how to freeze Brownian motion into stasis at places where it feels at home. (Oh, how I antropomorphize!)

A question: is, in my simulation, the feedback positive or negative? Would it matter if coli consistently stepped into the direction exactly opposite to the direction it goes now?

Did you see my post in which I described Koshland's findings about how e. coli is actually organized?

Yes, and thank you very much for that inspiration. As I said, I abstracted away the real thing and started to theorize, wondering how such a tiny thing, not much larger than a big protein molecule, actively moves around its world. Maybe its 'natural laws' aren't much more complicated than the simple, random movement that we know as Brownian motion. How simple can you get? An animal is usually defined as a living entity that can move by itself. What is the simplest animal (with respect to movement)? And from that: what is the simplest type of control? Sometimes people 'flow along' with the river of life and let things happen. Sometimes people are in control and make things happen. Are these two mutually exclusive or are there intermediate states? Now it seems that there are, at least in my simulation.

Date: Tue Sep 08, 1992 10:23 am PST
Subject: Stelli models'

[From Bill Powers (920908.1030)] Hans Blom (920908) --

I understand now.

Here is some more grist for your mill. This program really isn't a model of e. coli -- it's a demonstration in which YOU play the part of e. coli. This gives lots of insights into "control by random output."

When you try it, you'll think you're using lots of sophisticated perceptions, like which way the dot is moving in space, how close to the target it's projected direction will come, and so on. Most of these perceptions are superfluous.

You can prove this by changing the display so that the spot moves only in X, and X represents the radial distance of the spot from the target (modulo a screen wrap). You can make the spot approach the target just as well even though you've lost all 2-D information and have only a one-dimensional perception. This is E. coli's situation.

You can try putting in some slow random disturbances of direction. You will get to the target anyway. This is a control system!

```
program ecoldemo;
uses dos,crt,graph;

var x,y,x0,y0,dx,dy,theta: real;
    ch: char;
    i,j,hsize,vsize,vcenter,hcenter: integer;

{ This is lifted from your program }

procedure grinit; {initialize graphics stuff}
var
    driver, mode: integer;
begin
    detectgraph (driver, mode);
    initgraph (driver, mode, '');
    if graphresult <> grok then
```

```

begin writeln ('cannot find graphics driver'); halt end;
end;

begin
clrscr; gotoxy(1,12);
writeln(
'Experiment in "directed randomness." Use space bar to change
direction');
writeln(
'of moving dot, and steer it into the circle. "r" to reset, "q" to
quit. ');
write('
                                Press space to start');
ch := readkey;
grinit;
i := 20; j := 0;
hsize := getmaxx; vsize := getmaxy;
hcenter := hsize div 2; vcenter := vsize div 2;
x0 := hsize - 75; y0 := vcenter;
theta := 0.75; randomize;
repeat
clearviewport;
outtextxy(0,vsize - 20,
'TAP SPACE TO CHANGE DIRECTION; RESET=R  QUIT=Q');
ch := chr(0); circle(hsize - 75,vcenter,10);
i := random(vsize); j := random(hsize div 4);
x := 0.0; y := 0.0;
setcolor(white);
repeat
theta := 2 * pi * random;
while (not keypressed) and (abs(x - x0) + abs(y - y0) > 10) do
begin
x := x + 1.0 * cos(theta); y := y + 0.7 * sin(theta);
if x > hsize then x := 0 else if x < 0 then x := hsize;
if y > vsize then y := 0 else if y < 0 then y := vsize;
i := round(x); j := round(y);
putpixel(i,j,white);
delay(30); { CHANGE THIS DELAY TO MAKE SPOT MOVE AT GOOD SPEED }
end;
ch := chr(ord(readkey) and $df);
until (ch = 'Q') or (ch = 'R');
until ch = 'Q';
closegraph;
end.
-----
-
Eric Harnden (920908) --

```

Bully for you, Eric. Another modeler, by golly. This is really welcome news. Can Stella setups be transmitted as Ascii code? It would be great if you could send models to Gary Cziko, to help him learn Stella and work out problems with you. Do you know if there's a PC version of Stella, and what it costs? Rick Marken, would Aerospace pop for Stella for your Mac? I'll bet somebody there already has it.

>3) i put the following plot into stella the other night:

```

>
>      |-----level<-----valve----->source/sink
>      |--->perceived level                /\
>      |                                   |
>      \/\                                |
>reference----->error----->rate of change
>level
>                                   /\
>                                   |
>                                   disturbance
>

```

>where the only material flow is the increase/decrease of the level, and
>all other connections are information flows. the first question is: does
>this at all capture the basic pct structure (without the hierarchy)?

It sure does. As shown, it's a one-way control system: by setting the reference signal ("level") you can make the level rise to the reference level, but the system can't make the level fall if the reference level decreases. Put a leak in the tank and it will be able to go both ways. Or is your system set up so the valve can cause losses as well as gains of level? I guess that's what "source/sink" means, now that I think of it.

>it seems to me that the entity is controlling the environment (acting upon
>the rate of change) in order to obtain (is controlling for) a certain
>perception. yes? no?

Yes! I would not use the word "control" for the effect on the rate of change, however, because that implies that there is some preferred rate of change that is brought about by the system. I would just say that the system VARIES the rate of change of valve position (my interpretation). I know that your usage of "control" is common in servo circles, but I'd rather do without it. It's too confusing when you say that the control system controls by controlling.

To show that valve position or its rate of change isn't controlled, put a variable leak in the tank, so the amount escaping varies randomly (or as a sine wave or whatever's handy -- but make it slow). Now you'll find that the valve position varies in more or less the same random way, while the controlled variable -- the perceived level -- remains near the reference level. So the system clearly isn't trying to make the valve assume any particular position or rate of change. I think it's less confusing if we reserve the word "control" for the whole process of making the perceptual signal match the reference signal.

Your disturbance of the rate of change is perfectly OK. Disturbances can happen anywhere in the environment, and in more than one place at the same time. That won't bother the control system. Try the variable leak while leaving the disturbance you already have in place. It will still work -- but you won't get the impression that the rate of change is a controlled variable!

Can you plot the behavior of variables against time? I seem to remember that Stella could do this. If so, try superimposing plots of the disturbance and the rate of change of valve position.

If everyone out there had Stella it would be SO EASY to illustrate critical points about PCT! Static diagrams just don't capture the relationships.

--

Best to all, Bill P.

Date: Tue Sep 08, 1992 10:48 am PST
Subject: RE: Sciences and Ideologies - I

Thomas Baines :

My wife & I have three children. Girl, boy, boy. All are out of college & are "responsible adults"...as far as we know. All three are doing well in professions & seem to be sane. We know of no serious anti-social behavior.

What I am about to say will probably mortify some folks, but I believe it to be accurate. Our kids were raised much the same way we raise our dogs.

First they are housebroken & taught the rudiments of personal hygiene & responsibility for personal actions. With the kids its a little easier, because the language barrier is not as difficult. The dogs learn faster up to a point, however.

The next thing is learning how to learn. This requires figuring out what the kid (or dog) pays attention to, then using that as the channel through which to impart knowledge & stimulate curiosity.

Next, necessary limits on curiosity and behavior must be imparted. Both kids & dogs must be socialized lest they get killed crossing the road or someone dangerous.

Also, they must be socialized so that they do not become unable to learn how to recognize when their behavior is likely to result in their not being able to obtain adequate food, clothing, shelter & affection.

In order to make any progress on these latter points, one must, again, establish reference signals that are personally meaningful to the child/dog. Any insistence on my part to force the kid or pup to see things my way have been pretty fruitless. I have to see things their way, then establish why their viewpoint is in error.

By the way. I am an ex-Army Special Operations Officer. I don't think I'm a patsy. BUT!! I have NEVER had to resort to "Because I said so." Both kids & dogs will come to respect your judgement about things if you have established a way of expressing that judgement in reference to things they value, & the result of which is an experience that they can rely on for making their own judgement in the future.

Inconsistent messages or behavior on my part always causes more problems than the kind of working through the situation so that both of us are satisfied.

The latest dog, by the way, fetches the paper (in the snow), closes the door behind herself, brings her leash when its time to go for a walk, leaves the kitchen or dining room when the humans sit down to eat, & is generally a great pleasure to share time with. Whatever we're doing, it works.

Date: Tue Sep 08, 1992 10:51 am PST
Subject: control and children, connections

[From Rick Marken (920908.1030)]

Ed Ford (920906:1415) says:

Rick Marken (920906)

>The parent has no choice, right?

I have no idea what you mean by this statement.

What I was responding to was the following statement by you:

> All this means the parent must set standards and rules which reflect the
>parent's own values and beliefs (systems concepts level)

It sounded to me like you were suggesting that parents set references for standards and rules FOR THEMSELVES -- since we know that nobody can set references for other people (kids in this case). So the parent can only control his or her own system concepts by varying (setting) the references for standards (principles) and rules (programs) as necessary to keep their perception of system concepts matching their reference. Thus, it SOUNDED like you were saying that it was important to tell parents to control their system concepts. Thus, I was puzzled and said "the parent has no choice, right" meaning there is no choice but to control system concepts by setting references for lower level perceptions (assuming the parent can perceive and has a reference for a particular level of a system concept -- not always a reasonable assumption).

Your answer to my suggestion that you don't need to teach control systems to control was absolutely beautiful; worth a second quote: Like any group of people organized to live or work together such that they can cooperatively get along, many parents seek help and instruction (I was one of them) on the best way to set standards in such a way that everyone would cooperatively get along. Control systems naturally control, but in any community or between any two individuals, part of growing and maturing is to learn to control in such a way that you do no harm to yourself and that you learn to respect the rights of those in your environment. That is not something we do naturally, and thus, in a family setting, I believe it is the job of the parent to teach their children in as patient and loving way as possible how to use their control systems to their greatest advantage and with the least amount of violence to anyone's system.

This is an EXCELLENT paragraph. It captures what, to me, is the big problem that is made palpable (and understandable) by PCT, namely; how do you teach (NOT MANIPULATE OR CONTROL) control systems to get along with other control systems, knowing full well that you cannot arbitrarily control these systems nor stop them from controlling.

I guess I get hung up on one little point. In your statement above it is where you say "the best way to set standards". The way you have stated it here it sounds OK to me because I can hear it as a recommendation to vary ones own references at all levels so as to avoid conflict and act cooperatively. This can often be achieved only by reducing the gain of certain control systems or

seeing the situation from a higher level so that lower level possibilities for reference variations become apparent.

Maybe it is a communication difficulty, but I sometimes get the impression from your writing, Ed, that you think that there is some RIGHT setting for "standards" (which I take to mean references for principles and/or programs) that will produce good results (which I take to mean lack of intra or interpersonal conflict). The paragraph above suggests that I might just be laboring under a misconception -- so that I think we can agree that ANY setting for a standard that eliminates intra and interpersonal conflict is "good". I think, then, that we can also agree that it is impossible for an outside observer to know what the setting of any standard should be that would accomplish this result. The way you discuss standards in "Freedom from stress" I get the uncomfortable impression that you imagine that there is some particular set of standards (values, beliefs, whatever you want to call them) that is "best" in a general sense (meaning, for all people, not just for Ed Ford). The paragraph above suggests that maybe my interpretation of your position is wrong and that we can agree that references for standards can only be set situationally (due to varying disturbances) and contextually (context provided by the continuously changing values of other references in the system).

I think that it is more consistent with the PCT model to say that "the best way to set standards" means: always be able to VARY your references for the perception of "standards" in such a way that no malfunctions (conflicts) are created in one's own or other people's control systems.

Greg Williams (920906 - 2) says --

> another part of my ideology is my belief that more work on PCT can be
>fostered by emphasizing how it connects to other ideas, rather than how
>everybody else is wrong.

I think my "Blind men" paper fills that bill. It says that psychologists are right in the sense that control DOES look like response to stimulation, adaptation to constraint (reinforcement) and output generation. But these appearances, taken at face value, give a misleading impression of how behavior works. At no point in the paper do I say that anyone is "wrong" -- just that they have missed one little thing: the fact that all these appearances are aspects of the phenomenon of control. The paper also connects current approaches to psychology with the PCT approach. Is this consistent with your ideology? I hope so.

Regards Rick

Date: Tue Sep 08, 1992 11:52 am PST
Subject: Dealing With Children
Thomas Baines:

ED- My little brother (now 50 yrs old, 6'1", 245) is a Federal Probations Officer. He has had some success in reducing recidivism in his case load by starting with the questions "Do you like what has happened to you? Would you like to avoid having it happen again?" Like I said elsewhere about dogs & kids, start with trying to see what they see & what they think counts.

Date: Tue Sep 08, 1992 1:52 pm PST
Subject: Teaching, control, and what matters

[From Rick Marken (920908.1400)] Greg Williams (920906) says:

>I say that SOMETIMES manipulation can be judged (afterwards) by a kid as
>"good" for him/her. The PCT model does not imply that this is impossible or
>even improbable.

I agree. Sometimes, hauling the kid out of the street IS the right thing to do from your point of view and (eventually) from the kid's too. But that's just your call in that situation -- and we are assuming that it was a good call that time. All I'm saying is that, if one takes the position that this kind of "wise" controlling is a good thing to do they will be creating problems for themselves and the controllee AS A MATTER OF PRINCIPLE. I agree that it often is "right" to do some emergency controlling of a kid when you are a parent. After all, the kid can't control perceptions at the level you can, yet. I just object to the idea that control of the kid is good as a matter of principle.

Now that I think of it, Greg (and Pat too, I presume) maybe we can find some common ground on child rearing in terms of levels of control. I think I could be comfortable with the recommendation that parents CONSIDER controlling a kid's behavior when the kid needs to be able to control a variable that is at a higher level than those that the kid is currently able to perceive and control. For example, a 2 year old can probably control the perception of it's relationship to a rolling ball by adjusting it's running speed, etc. But sometimes, while controlling this variable, the kids behavior becomes part of a higher order variable that the kid cannot perceive/control but the parent can -- for example, a program variable like "IF kid follows ball into street WHILE car approaches at high rate of speed THEN possible squashed kid". The parent can, I think, guess pretty well whether or not the kid can perceive and control this higher order variable and will CONTROL THE KID by forcibly changing a result it intends to produce (getting close to the ball) in order to control a higher level variable the way the parent would have controlled him/ herself if the parent had been chasing the ball into the street. F. Plooi found that chimp mothers do something like this intuitively; they help the kid by controlling variables at levels that the kid cannot control.

An important corollary of this concession is that parents should stop controlling higher level variables for the kid once the kid can control those variables. Thus, once the kid can perceive and control programs like "running after the ball into the street" then it's time to stop controlling program level variables for the kid. Besides, by the time the kid starts to understand program level perceptions, the kind of physical force control that worked on the three year old becomes a bit more difficult -- and probably a lot less justifiable.

>you apparently believe that PCT DOES imply that a PARTICULAR KIND of
>manipulation (NOT control) is best for teachers to practice:

No. There is no manipulation or control involved in what I described as education.

I said:

>Education is not control exerted by the teacher -- it is a process where the
>teacher uses the teachee's ability to understand, imagine, remember, and
>reorganize in order to help the teachee be able to control.

>Here you are agreeing with me that education makes great use of manipulation.

I am imagining a person who wants to produce some results and doesn't know how. The teacher tries to communicate POSSIBLE perceptions to control and possible ways to control them so that the person can produce the desired end. I imagine no manipulation on the part of the teacher. The teacher says, in effect, "try to create this perception". The teachee does not "have to" produce it; and might come up with a better way. My idea of a good teacher does not try to MAKE the teachee do things the RIGHT way. The teacher is a suggestor, not a manipulator.

Thomas Baines says:

> What I am about to say will probably mortify some folks, but
> I believe it to be accurate. Our kids were raised much the
> same way we raise our dogs.

Thank you, Thomas. This helped put things back in perspective. We have been doing TONS of talking about PCT and child rearing. But when push comes to shove it turns out that how we talk about child rearing is probably only incidentally related to how we actually interact with our kids -- and those interaction styles probably vary HUGEYLY. And, yet, most kids come out just fine. Skinner had great kids and so did Dr. Spock, as far as I know. Thomas says he got great results from raising his kids like dogs; I got great results raising mine like humans. I really think that all kids (unless their is brain damage or some other physiological problem) will turn into fine, functional control systems regardless of who is raising them; it just doesn't matter that much unless the kids are subjected to the systematic violence (literally) of control that you usually only read about in national enquirer. How many people lock their kids in closets in restraints and pass food under the door? REAL control, the kind that produces REAL resistance and, hence, real violence, is what I care about. This is the control that keeps people out of jobs, schools, parks, neighborhoods, etc because of their race, religion or accent. It is the control that forces some people to adapt to the violence of poverty because they are seen as incompetant or lazy. It is the control done by all of us who have (and want to make sure we keep having) at the expense of those who don't. Bill Powers went through the litany already. It is THAT kind of control that matters -- not whether you say to your kid "the one who has the gold, rules". This latter kind of control is annoying but basically harmless, unless you really, REALLY do what is done to MANY people in this society -- let them have the gold ONLY if they are white, smart, clean, "hard working", and do only PRECISELY what YOU want. Only armed groups of people are able to exert THAT kind of control over control systems. And it creates all the problems it is designed to solve. PCT explains why.

Best regards Rick

Date: Tue Sep 08, 1992 6:56 pm PST

Subject: Child-rearing magic

[From Bill Powers (920908.2000)]

Thomas Baines (920908) --

You have some pretty attentive ears out here. Your description of organism-raising could provide some answers to the general problem of how independent autonomous systems can get along together, and how to "raise" children of any species by some means that doesn't end up teaching them violence and crippling their self-starters.

On the other hand, a skeptic could need some convincing that this is a method that others might want to use. First off, I've never heard anyone describe a method of child-rearing they worked out and used and conclude by telling us that the kids turned out terrible as a result of the method. The problem here is that the parent naturally uses his or her own reference levels in judging the result. If the children behave in a way consistent with those reference levels, the parent will point to them as proof that the method works. I expect that the Ayatollah Khoumeni's Dad was quite proud of the job of child-rearing he did. So to judge the worth of any method of raising children (or pets), one has to get outside of the framework of the parent doing the raising. That's where a theoretical framework can help.

Your description leaves itself open to several interpretations, depending on what the reader thinks is right and proper. Many people think they know what children and dogs "pay attention to:" pain, hunger, thirst, parental (or masterly) rejection, and loss of freedom. So they might interpret this phrase a little differently from the way you meant it (I presume):

>This requires figuring out what the kid (or dog) pays attention to,
>then using that as the channel through which to impart knowledge &
>stimulate curiosity.

Some people would mean by that, "Find out what they value, and then take it away from them until they cave in." Considering your interest in PCT, I rather doubt that this is what you meant, but you see the problem.

I think everyone would like to hear the details, for comparison with PCT principles and probably for more practical reasons than that. How, for example, do you do the housebreaking? How about some examples of "rudiments of ... responsibility for personal actions" and how you actually teach them?

The one that really gets me is

>Any insistence on my part to force the kid or pup to see things my way
>have been pretty fruitless. I have to see things their way, then
>establish why their viewpoint is in error.

As a father of teenagers I was terrible at trying to figure out (a) what mistake they were making, and (b) expressing it in a coherent way that they might understand. How the hell did you do it, aside from using your natural genius? I'll bet a lot of people would like to know.

Best, Bill P.

Date: Tue Sep 08, 1992 7:02 pm PST
Subject: Re: Child-rearing magic

The only really magical trick that Cindy & I discovered for child-rearing was one to stop the chorus of 'me-firsts' whenever goodies are distributed. The rule was that whoever said 'me-first' first was last. This seemed to work very effectively, without provoking any resentment on the part of the control-/manipul-ees. Why, I don't know.

Avery.Andrews@anu.edu.au\

Date: Tue Sep 08, 1992 8:44 pm PST
From: RSCH287E MBX: RSCH287E@cl.uh.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: Industrial/Organizational Psychology in a PCT construct

Mr. Dag Forssell:

I am extremely interested in discovering more about Industrial/Organizational Psychology in a PCT construct. Would you happen to know of any texts that may apply to this topic? If not, would you know of any journal articles?

Any insight you may be able to provide would be greatly appreciated. Thank you very much for your time and cooperation.

Owen Daniel Primeaux rsch287e@cl.uh.edu

Date: Wed Sep 09, 1992 2:36 am PST
Subject: Blind Men and Scientific Progress

[from Gary Cziko 920908]

Rick Marken (920908.1030) in response to Greg Williams (920906 - 2) says:

>I think my "Blind men" paper fills that bill. It says that psychologists are
>right in the sense that control DOES look like response to stimulation,
adapt-
>ation to constraint (reinforcement) and output generation. But these appear-
>ances, taken at face value, give a misleading impression of how behavior
works.
>At no point in the paper do I say that anyone is "wrong" -- just that they
have
>missed one little thing: the fact that all these appearances are aspects of
>the phenomenon of control. The paper also connects current approaches to
>psychology with the PCT approach. Is this consistent with your ideology?
>I hope so.

I started thinking about the Blind Men paper again while reading the philosopher of science Imre Lakatos. He makes the case that science progresses when new theories account for everything that the old theories account for plus account for some things they don't account for. Of course,

this is not a new idea (although perhaps some of the nuances Lakatos slips in are new).

So, Rick, if you still have problems getting this wonderful paper published, maybe you should add this angle to it as well, i.e., a little history of science like how Newton showed that Galileo and Kepler were partly blind and how Einstein showed that Newton was partly blind, etc. In each case, a more encompassing theory was proposed that pulled together what first appeared to be separate, unrelated phenomena (e.g., Galileo's terrestrial mechanics and Kepler's planetary mechanics). This is exactly what your Blind Men paper does as well.

What is particularly intriguing to think about, however, is how in 100 years or so PCTers will be considered partly blind! I wonder what the next step will be after PCT to pull what appear to be disparate, unrelated phenomena together.--Gary

Date: Wed Sep 09, 1992 4:03 am PST
Subject: You got it... but there's more to get

From Greg Williams (920909 [really!])

Chuck Tucker asked for the full citation for the following; maybe other netters also will be interested in this book (especially, I think, the chapter titled "Intentionality in Animal Conditioning"; because our older son probably would be called "dyslexic" by the (some) psychologists, Pat and I were intrigued by the chapter by a university-trained mathematician who essentially cannot read):

L. Weiskrantz, ed., THOUGHT WITHOUT LANGUAGE, Clarendon Press, Oxford, 1988, xvi + 533 pp. ["A Fyssen Foundation Symposium"]

>Bill Powers (920908.0730)

>>Greg Williams (92090[8]) --

Too early for me to get the date right... sorry.

>>Manipulation can be aided by the manipulee having a high gain for, and
>>paying close attention to, the actions you want to see.

>Disagree. Control is never done by paying attention to actions, unless
>it's specifically the perception of the action that you're controlling
>(as in dancing).

Sorry. I meant "can be aided by the manipulee having a high gain for, and paying close attention to, his/her perceptions which result in the actions you [the manipulator] want to see."

I also want to make it clearer that manipulation, as I want to define it, can be of the "rubber-banding" (purposive disturbance of controlled perceptions) OR of the "setting-the-context-for-'natural'-control" types.

>If the manipulator can get the Mark to accept the manipulator's description

>of v2, instead of the Mark perceiving v2 directly, the manipulator can
>misrepresent the state of the environment and thus gain control of the Mark's
>action.

Right!

>In the con game this happens when the wallet (v1) containing the valid
>bond or Treasury note or whatever (v2) is switched for the similarly-
>wrapped one containing newspaper. The Mark has control only of the
>wallet (v1), but is shown by the con man that the wallet contains
>something valuable (v2). The Mark doesn't want to perceive the wallet
>alone (v1) but the wallet with something valuable in it, which is the
>perception $p = f(v1,v2)$.

>The action that gets the wallet and v2 into possession of the Mark is
>to hand over cash to the con man. This does get the wallet into the
>possession of the Mark. The Mark perceives this, however, as getting a
>wallet with something valuable in it. The con man sees to it that the
>Mark gets the wallet, but without the something valuable in it.

>This is the same con game played by advertisers; when you hand over
>your money, you get the flashy car, but the nice pretty girl isn't in
>it any more and never will be. What you get is a dumb promiscuous
>broad who can be conned into thinking that a flashy car contains a
>terrific guy, and when she hands over the price she will give you a
>baby you didn't particularly want, or AIDS. And when she looks in her
>package all she will find in it is plain old you. Great design for a
>relationship.

You got it, in spades!

>I agree with you that lots of this kind of deception goes on.

YOU GOT IT!!! But there's still more to get....

>What's good about it?

First, there is usually nothing good -- for the mark -- about exploitation, which is "negative" manipulation. The mark ends up sorry he/she was fooled. But there is also (I claim, plenty of) manipulation which ends up being good for the manipulee (as judged by the manipulee, whether in fact or hypothetically judged). Remember from several days back that I am trying to develop a precise way of talking about (right now, I'll generalize it to) human interactions, based on PCT ideas, and my most general category is influence (altering another's variables, purposively or not). One kind of influence is manipulation (so termed because it requires some skill to do -- according to PCT; NOT so termed because it is necessarily negative for the manipulee OR because it necessarily involves deception) -- Rick used the term "suggestion" in a recent post, and perhaps that has a less-negative "folk" connotation than "manipulation," but I'll continue to use "manipulation" for now. There are four kinds of manipulation (I claim): those which, if successful (from the manipulator's point of view) result in the manipulee's (after completion of the manipulation) judging them (to him/her) harmful ["exploitation"], neutral ["?"], good ["facilitation"], or mixed/who knows/can't say ["?"].

I claim that all four types are extremely common. I further claim that much of what parents, teachers, therapists, and counselors (such as Ed Ford) do is facilitation. In saying that, I want to make it even clearer (especially for Rick) that my own judgement about whether facilitation (remember, defined as "good" from the facilitatee's view) is a "good" thing to do is an IDEOLOGICAL judgement NOT derived from PCT. Facilitation need NOT involve deception: Ed tells the juvenile offender the full tale about how Ed sees the relationship between what the kid might do and what the results will therefore be -- but the kid might otherwise have not seen the relationship quite so clearly or might not have thought through what he/she wants MOST if Ed hadn't been there. It CAN involve deception: the mother puts the medicine in applesauce and says, "Here's some yummy applesauce," NOT "Here's some yummy applesauce with medicine."

Once again, PCT says that for manipulation to work (even of the rubber-banding kind, since the manipulator needs to have confidence that the perception under control won't change unpredictably), using The Test is important. I'm not saying that many teachers/therapists/counselors know explicitly about The Test, but they use it without knowing its name or theoretical basis because they NEED to. However, I suspect that part of why Ed Ford is such an EFFECTIVE facilitator is that he understands the PCT-basis for what he needs to do.

Clearly, facilitation can be viewed as patronizing. My own view on patronizing depends on the particulars of the situation, and can extend far beyond "non-coercive" (non-physical/non-threatening) facilitation to tying up an attempted suicide. That's more ideology, not PCT. The bottom line is that patronizers can sometimes be VERY wrong in their predictions about how the facilitatee will ultimately judge the outcome.

He who would do good to another must do it in Minute Particulars:
 General Good is the plea of the scoundrel, hypocrite & flatterer,
 For Art & Science cannot exist but in minutely organized Particulars
 And not in generalizing Demonstrations of the Rational Power.

-- William Blake, JERUSALEM

Bill, I think we've reached an important consensus about my claims. Now, we can go back and nit-pick some more, or work together toward exploring the nature and limits of manipulation from a PCT point-of-view, or press on toward considering what I have termed "long-term" influence (involving reorganization) and my critique of "autonomism," or... what? Your choice... influenced, of course, by the history of our recent interactions.

>Rick Marken (920908.1030)

GW>> another part

GW>>of my ideology is my belief that more work on PCT can be fostered by
 GW>>emphasizing how it connects to other ideas, rather than how everybody else
 GW>>is wrong.

>I think my "Blind men" paper fills that bill. It says that psychologists are

>right in the sense that control DOES look like response to stimulation, adapt-
>ation to constraint (reinforcement) and output generation. But these appear-
>ances, taken at face value, give a misleading impression of how behavior works.
>At no point in the paper do I say that anyone is "wrong" -- just that they have
>missed one little thing: the fact that all these appearances are aspects of
>the phenomenon of control. The paper also connects current approaches to
>psychology with the PCT approach. Is this consistent with your ideology?
>I hope so.

Absolutely! Right on!!

>Rick Marken (920908.1400)

Greg Williams (920906) says:

>Now that I think of it, Greg (and Pat too, I presume) maybe we can find some
>common ground on child rearing in terms of levels of control....

Yep. No problem. Our ideologies mesh in large measure. But that fact doesn't follow from PCT.

GW>>you apparently believe that PCT DOES imply that a PARTICULAR KIND of
GW>>manipulation (NOT control) is best for teachers to practice:

>No.

Good.

>There is no manipulation or control involved in what I described as education.

Uh oh. See the above for one (last?) attempt at defining "manipulation."

GW>>Here you are agreeing with me that education makes great use of
GW>>manipulation.

>I am imagining a person who wants to produce some results and doesn't
>know how. The teacher tries to communicate POSSIBLE perceptions to control
>and possible ways to control them so that the person can produce the desired
>end.

Thaaaaat's manipulation, folks!!! (Is there an echo in here?)

Best wishes, Greg

Date: Wed Sep 09, 1992 4:31 am PST
Subject: MANIPULATION AND FORCE

CHUCK TUCKER [920992]

RE: On the conversation about controlling another (manipulation)
 [A getting B to control B as a direct effect of A's behavior]
 WTP 920828.1800; 920829.2130; 920830.1545; 920901.1000 + MT 920902
 0030; 920902 11:15 + GW 920829; 920830; 920830-2; 920903-2; 920903-
 3 {last post}

Let's return to to Bill's (my WTP) statement of the difference
 between his and Greg's (my GW) position in WTP 920829.2130):

"Greg's thesis is that there are ways of controlling the behavior
 of other people against their will or without their awareness other
 than by the use of overwhelming physical force.

My thesis (WTP) is that such apparent control (a) is an illusion,
 or (b) is of not importance to the "victim", or (c) is carried out
 only with the aid of the "victim", or (d) rests on a threat of
 force in the background. I contend that a properly functioning
 control system cannot be controlled from the outside, so that any
 apparent success by a manipulator results from some defect of
 organization in the control system."

Later, Greg (920830) stated the he was mainly concerned with (c)
 yet he had a few examples where some of the others may be involved
 but let us consider all (a-d) of these ways that another apparently
 is controlling a person (A is attempting to control B).

ILLUSION

If A asks B to do X and B is in the early phases of doing X and
 completes the act then it will appear that A's request is
 responsible for B's act; it will appear like A is "controlling" B.
 If A is driving down the street behind B and B's auto has its
 right-turn signal flashing and B slows down and turns right, A
 might claim responsibility for B's right turn and/or an observer
 might judge that A caused B's turning act. There are many acts
 where one person will ask or figure out what another will be doing
 next and either do it with them or ask him/her to do some within
 that act. Most "elementary collective forms" (i.e., queues, rings,
 arcs) appear as if A has "controlled" B (and even C, D, . . .) but
 that is not the case; this can be clearly illustrated with the
 simulation program "Gatherings" (formerly "CrowdV2"). In fact,
 driving on the highways and byways may be one of the best
 illustrations of persons controlling themselves while appearing to
 be "controlled" by another. [I recall an article written years ago
 in the Pacific Sociological Review with a title like "Taking the
 Role of the other in Driving" which made a similar point.]

NO IMPORTANCE or AID OF OTHER ["Victim"]

This apparent control technique does assume that B is controlling a
 perception but rather than being of "no importance" I see it as
 controlling for "cooperation," "participation," "involvement" or
 other names for "joint acts". It could look like Bill suggests
 that B is controlling for a perception contrary to the request of A
 but it is so trivial as compared with A's request that is "a

difference that makes no difference". [When someone tells me that it "makes no difference" I take them at their word; on some occasions I have found that they either did not "mean it that way" or "changed their mind later" but I had a heck of a time telling the difference] Thus, only the TEST would help determine whether B was controlling for "no importance" or for "cooperation" or the like. But in either case it seems to me that that B is controlling for a perception.

All of the examples and illustration stated by Greg and those that I infer from statements cited by me and Greg from Skinner assume that B has a perception for an act similar to that of A's request or a perception at a higher level that would be accomplished by complying with A's request. What Skinner does (which he does not recognize because he wants to be a SCIENTIST) is that he takes what is a purpose of B and puts it outside of B in either the natural or social environment as a stimulus or makes what was an external stimulus into the "history" of B. This is, in part, what Bill pointed out in "Skinner's Mistake".

The "Con man" (A in this example) can't get to "first base" without B (the "mark") having a purpose that A will use to exploit B. My Father ran a "boiler room" back after WWII where he and his friends "sold" ads for a party for the Veterans and appealed to patriotism. They did have a party for the Veterans (everything donated) and printed enough booklets so each advertiser could see their ad and they pocketed the rest of the money (profit on the scam was about 1000%). Jerry Lewis just "conned" people out of 43 million dollars this last weekend based on the notion that they were helping "victims" of a disease. We know about Jim and Tammy Bakker and those of their ilk who swindled people out of their life savings on the promise that they would get into Heaven. In all cases of this type of "manipulation" (and all others) those people could not be conned if they did not have a purpose (one of the most widespread purposes used for a con today is "getting something for nothing"- send your letter to Ed McMann). All "con jobs" are founded on a person having a purpose (some of my students believe that a degree will lead to the "good life") and "con man" (like myself) exploit that purpose [actually I tell my students that it is not true but few believe me]. I would limit "manipulation" to these instances where deception is involved (like 95% of the experiments in the so-called social and behavioral sciences, e.g., Milgram) rather than MAKE IT AS BROAD AS THOSE SUGGESTED BY GREG (I DON'T THINK PAT IS DECEIVING THEIR CHILDREN; DO YOU GREG?) [GOFFMAN HAS A WONDERFUL piece on this process called "Cooling the Mark Out".

FORCE OR THREAT OF PHYSICAL (EVEN PSYCHOLOGICAL) FORCE

Force or threat of physical force will not work unless a person has a concern for staying alive, healthy or avoiding physical harm. A will have only limited success in getting B to do much of any importance by using direct force since such acts make it difficult for B to perform. This is perhaps the reason that Southern plantations had so many slaves so some could be healing while others were working. Forced labor and torture requires many people

to work and it is not a very efficient system. It is much more efficient to beat a few people as examples of what could happen if you don't obey while at the showing pity and concern while impressing upon the people how important staying alive is for them. If a person loses all hope of staying alive and that life is not worth living then the possibilities of using force or threat of force for compliance is severely reduced if not eliminated.

All of the above can be stated for the notion of "psychological" health and well-being; this is why Watson could make more money in business than in academia. Modern techniques of manipulation focus on "mind control". But even here, it won't work if a person does not believe his/her psychological well-being is a perception worth controlling.

CONCLUSION

If a control system is functioning properly by controlling its own conduct with regard to maintaining low error between reference signals and perceptual signals thus being "self-governing" or "self-regulating" then no one can DIRECTLY determine any of its signals or the processes that are involved in such controlling. The human organism (which I call 'self' or 'selves') is self-programing, self-governing and autonomus.

CODA: THE DARK SIDE OF PCT

Since we can't directly determine through our action how another will control him/herself then we must use indirect means to get another to behave consistent with our intentions. In so far as people use theories of human behavior based on the notions that "stimuli cause response or behavior" then they will be unsuccessful (or, at the best, randomly and unwittingly successful) at getting others to comply with their requests and do what they intend for them to do. Their success will increase or improve to the extent that they understand PCT and how to use the TEST to determine with great precision "what and how a person is controlling him/herself". When one learns that and then finds out what might be the "occasions" for reorganization then he/she is prepared to devise ways to get another to better comply with requests and way to devise arrangements which serve the purpose of those requests. This, of course, could be used by those who want people to harm themselves or others; for purposes that we might consider "inhumane" or "destructive" to humankind. We cannot directly determine that others won't use these ideas for other purposes. We can only attempt to stress that when people do not treat each other as living control systems which are self-governing then conflict will arise and perhaps subvert the accomplishing of everyone purposes. But, of course, we can only do this indirectly.

QUESTION AND COMMENTS

What are the references to the work(s) of Chapman and/or Agre? I just missed it.

There is an article that will be in a future issue of the AMERICAN PSYCHOLOGIST by James O. Prochaska, Carlo C. DiClemente and John C. Norcross "In Search of How People Change: Applications to Addictive Behaviors" which is quite relevant to PCT. I would like someone's comment on it when it is published.

Regards, Chuck

Date: Wed Sep 09, 1992 9:34 am PST
Subject: Fileserver News

[from Gary Cziko 920909.0545]

Bill Silvert has reorganized the CSG directory of his fileserver. The fileserver is on biome.bio.ns.ca and the directory in which these files are located is pub/csg/.

There are two main types of files--program and document. Document files can be obtained by anonymous FTP or by sending a mail message to server@biome.bio.ns.ca with the command "get file-id" as the first line of the message. For example, to obtain the latest PCT bibliography compiled by Greg Williams send the following command to server@biome.bio.ns.ca:

```
get csg/documents/biblio.pct
```

and you will receive the bibliography via e-mail (longer documents may be sent as a set of several smaller ones). To obtain program and other binary files, anonymous FTP (file transfer protocol) should be used since I have found that converting binary files to ASCII for mailing and then back to binary is a real pain and there is no standard way of doing this. If you are at an educational institution or high-tech commercial establishment, someone there can probably show you how to use FTP. You probably cannot use FTP if you use CompuServe, MCIMail, ATMail or other commercial service for e-mail.

The two most important programs (I think going through each of these is essential for anyone wanting to participate intelligently on CSGnet) are Bill Powers's Demol and Demo2. Their file-id's are pub/csg/programs/msdos/demla.exe and pub/csg/programs/msdos/dem2a.exe. To get Demol using anonymous FTP, use the following commands:

```
cd pub/csg/programs/msdos binary get demla.exe
```

The fileserver can also be accessed using the Gopher program. To do this, telnet to biome.bio.ns.ca as gopher. This is an experimental gopher server, but it may not always be functioning.--Gary

P.S. Thanks to Bill Silvert for setting this all up for us and for checking the accuracy of this message.

Gary A. Cziko

Date: Wed Sep 09, 1992 10:43 am PST
Subject: stella models

as i mentioned before, the last time i talked to high performance systems, i got the impression that they were not considering a port of stella the the pc. although i believe that there is some talk at pugh-roberts of putting a graphic front end on dynamo for windows. as i remember, however, the cost was going to be exorbitant.

as for moving stella stuff around, it can't actually be transported in ascii, but like any mac file there's no reason it can't be binhexed for transport on the wires.

another alternative would simply be to use the dynamo equivalents of the stella diagrams. these diagrams are, after all, just exactly the kind that dynamo modelers draw before committing their models to code. they can be sketched from code just as well. for example, the dynamo code for the drawing i posted is as follows: (with exhaustive comments... please bear with me.)

```
l level1.k=level1.j+(dt)(rate1.jk)
  this is a level equation, stating that a level variable named level1
  will at time k be equal to its value at time j (the previous timestep,
  k-1), plus some difference defined by a rate times the integration
  interval.
n level1=1000
  an initialization statement, setting the value of level1 at the beginning
  of the run.
r rate1.k1=-10+error.k+noise()
  the rate of change of level1 is defined as a constant depletion of -10
  per timestep (there's your hole in the tank) plus the current error value,
  plus the output of a noise function (whose parameters i have forgotten...
  i think they are mean and sd. one could also use a suite of function
  generators.)
a error.k=reflevel-perclev.k
  an auxiliary equation defining the error as the difference between a
  reference and the current perceived level.
c reflevel=500
  constant definition for reference level.
a perclev.k=delay1(level1,5)
  one possible transformation of actual level1 to pereceived level. the
  perception of level1 lags by 5 time intervals (not integration intervals).
  other possible transformations include table lookups, where the perceived
  level is a function of real level, which can even be drawn on a graph.
```

and that's basically it. without all the comments it's only 5 lines, leaving out a couple of lines of program run specs. the language is fairly straightforward, and is, for me at least, a comfortable crossover from mathematical to textual representation. the equations also clearly define the drawings from which they derive. and you don't have to worry about order of calculation, either. the dynamo compiler manages that for you.

all of which is to say that i think we've got workable alternatives to ascii graphics for the effective communication of model structures.

Eric Harnden (Ronin)

Date: Wed Sep 09, 1992 10:50 am PST
 From: mcnamara
 Subject: RE: You got it... but there's more to get

> From Greg Williams (920909 [really!])
 >
 > There are four kinds of manipulation (I claim): those which, if
 > successful (from the manipulator's point of view) result in the manipulee's
 > (after completion of the manipulation) judging them (to him/her) harmful
 > ["exploitation"], neutral ["?"], good ["facilitation"], or mixed/who
 > knows/can't say ["?"].
 >

This reminds me of a "General Theory of Stupidity" that I saw in Whole Earth Review some time ago. The basic concept was that intelligence had nothing to do with IQ (whatever that is), but rather was a function of an individuals' interactions with other humans. There were four types of exchanges, classified on whether or not they resulted in a gain for the participants. The highest level was where both parties benefitted (genius), the next where one benefitted to the extent that the other lost (thief), the third where neither benefitted (neutral), and the fourth where both lost (stupid). The idea was that the first three interactions don't increase the entropy of the universe, but the fourth one does. It was couched in terms of information gain or loss, but could be recast in control terms. I'll provide the reference in a few days.

Curt

Date: Wed Sep 09, 1992 10:58 am PST
 Subject: Rejoinder to PCT Critique

[from Gary Cziko 920909.1335]

I have written a response to a critique by Amundson et al. to my article "Purposeful Behavior as the Control as Perception: Implications for Educational Research" which will appear in Educational Researcher which I have appended below.

Amundson et al. don't do much to actually critique PCT, but rather attack me on my non-PCT-related arguments. Since I see the PCT part much more important, I concentrate on this in my rejoinder.

I realize that this might not make a whole lot of sense without seeing Amundson et al.'s paper, but I have sent it anyway in case there are some obvious things I can do to improve my reply.

Also, the emphasized words don't show up as such in this ASCII version of the my paper.

I will need to send this off in a few days, so if anyone has comments, please send them as soon as possible.--Gary

=====
 On The Threats to Educational Research

Not Faced by Amundson, Serlin, and Lehrer

Gary A. Cziko
University of Illinois at Urbana-Champaign

Amundson et al. (1992) charge that I misunderstand or misrepresent the views of the authors I cite in support of my claims as well as the views of Lehrer et al. (1990). They also find that there is nothing in my discussion of statistical methods and perceptual control theory which justifies my criticism of the methods of mainstream quantitative educational and psychological research. I will consider each of these charges in turn.

Concerning my "misunderstanding or misrepresentation" of authors such as Cronbach and Snow on educational research, Lakatos and Einstein on the philosophy and methods of science, and Gould on biological evolution, I can easily provide additional evidence to support my portrayals. For example, while Amundson et al. chide me for using "the narrowest of readings and most selective of quotations" . . . "to depict Lakatos as an advocate of prediction and control" (p. 5), it did not take me long to find an even stronger Lakatosian endorsement of the role of prediction in the very same essay quoted repeatedly by Amundson et al.:

"For the sophisticated falsificationist, learning about a theory is primarily learning which new facts it anticipated: indeed, for the sort of Popperian empiricism I advocate, the only relevant evidence is the evidence anticipated by a theory, and empiricalness (or scientific character) and theoretical progress are inseparably connected." (Lakatos, 1970, p. 123, first two emphases added)

I could provide similar support for my understanding of Einstein (that he ascribed great importance to predictive tests of his theory) and Gould (that he considers unpredictable contingency to play a major role in the evolution of species and that he is not anti-selectionist). But since these issues are secondary to the purpose of my paper, I will not discuss them further here.

On the other hand, it seems that I did unwittingly (and not intentionally) misrepresent the views of Lehrer et al. on educational research, particularly in stating that " . . . Lehrer et al. appear convinced that progress in educational research is a matter of adding more and more variables and sophisticated statistical analyses to a one-way, input-output, external-causation model of human behavior" (p. 28) and they protest that "in fact, quite the opposite is true" (p. 16). But what does "quite the opposite" mean? Does this mean they reject a one-way causal model of human behavior? Or does it mean they maintain hope that we will eventually hit upon a small number of the "right" variables which will account for human behavior in educational settings? Unfortunately, Amundson et al. do not make their views very clear in this respect.

The primary purpose of my paper was to introduce perceptual control theory (PCT) to the educational research community and discuss its implications for educational research. My very brief discussion of group-based statistics was included to show that the typical results achieved using these methods are not very impressive by any reasonable standards, particularly when one attempts to make predictions concerning individuals. My discussion of PCT was then intended to provide an explanation of why this is the case. The main point is

this: If PCT is correct in positing a closed-loop, negative-feedback relationship between individuals and their environment, then the basic model (and the assumptions) underlying group-based statistical analyses as used in educational and psychological research is simply inadequate to the task of providing explanations for purposeful behavior. This is the major point of my article and a point which Amundson et al. do not address.

In an attempt to make this argument clearer, let us consider the basic building block of statistical analysis as used in educational and psychological research. This can be expressed in the simple formula $y = bx + a$, where x refers to an independent variable (usually some environmental condition or stimulus), y to some dependent variable (usually some response or behavioral outcome), and b and a are the regression weight and y -intercept (predicted value of y when x is zero), respectively. Correlations, t -tests, analysis of variance (univariate and multivariate, with or without covariates) multiple regression, canonical correlation, discriminant analyses, and path analysis all built upon this basic equation. For all these analyses it is assumed that y depends on x , i.e., that there is a one-way direction of causality from x to y . And in this important sense, despite Amundson et al.'s protest on page 6, the great majority of the group-statistical methods used in educational and psychological research are monolithic.²

But research on PCT has shown that such a model is incomplete and therefore inadequate for even the most mundane of purposeful behaviors. In a simple tracking task (in which a subject seated before a computer screen manipulates a mouse or joystick in order to keep a cursor positioned at a certain goal location in spite of constant and unpredictable disturbances applied to the position of the cursor), the subject's response ("dependent" variable y) does indeed depend on the position of the cursor on the screen ("independent" variable x), but x also depends on y at the same time. Thus, there are no observable "independent" or "dependent" variables in the usual sense of these words. In fact, the only truly independent variable is the internal goal or reference level of the subject, a perceptual standard which is not determined by any environmental stimulus. While this reference level is not directly observable by outsiders, it can nonetheless be inferred by observing that the subject maintains the cursor's position at or near a certain location in spite of the disturbances influencing the cursor. That is, we infer purposive behavior and perceptual control when there is a low correlation between what are typically considered "stimulus" and "response" since while the subject's actions must continually vary in order to compensate for disturbances, the position of the cursor changes relatively little since it (actually its perception) is being controlled by the subject. I argue that if a one-way model of causality cannot account for something as simple as this tracking behavior, it is likewise inadequate to account for more complex purposeful behaviors such as those involved in teaching history or learning calculus. I read the inconsistent and noncumulative nature of the existing quantitative research base in education as supportive of this argument.

The notion of a controlled perceptual variable is so foreign to mainstream quantitative researchers in education and psychology that I cannot pretend to be surprised if Amundson et al. provide no evidence of making much sense of my description of it and its implications for research. But fortunately, they and other readers need not rely on my brief introduction. Runkel's 1990 book *Casting Nets and Testing Specimens* provides a clear and (I find) convincing argument for how group-statistical analyses make unwarranted one-way causal

assumptions concerning human behavior. Runkel also explains how such analyses are actually analyses of relative frequencies and as such cannot provide us with the explanations we seek concerning the psychological functioning and behavior of living organisms as autonomous systems (my two quotes from Taylor made the same important point, a point also not approached by Amundson et al.). And Powers has developed two quite remarkable computer programs which demonstrate, statistically analyze, and simulate the phenomenon of perceptual control.³

Amundson et al. state that "there is nothing in Cziko's presentation of perceptual control theory that precludes the utility of 'clinical trials' (Schrag, 1992)." This is not the case if they mean to endorse the utility of educational trials in which different treatment conditions ("independent" variables) are compared for their effects on educational outcomes ("dependent" variables). Again, such a research methodology assumes a one-way model of causation which cannot account for either teaching or learning as purposeful behaviors. As absurd as it may initially seem (although probably no more absurd than the initial proposals that the earth is round and that species evolve), PCT explains why it is that external environmental factors are not the cause of behavior and as a consequence educational treatments do not cause educational outcomes. Indeed, without the understanding of purposeful behavior that PCT offers, Schrag is correct in stating that "the positivist paradigm is hard to avoid" (p. 7). But once it is realized that Schrag is also apparently committed to a one-way view of causation for which there is now a more adequate alternative, his positivist paradigm becomes quite avoidable indeed.

Amundson et al. conclude their critique with the prediction: "Important results will come from actual research, not from a priori armchair declarations." I couldn't agree more. What Amundson et al. (and the rest of the mainstream educational and psychological research community) do not realize, however, is that research results already exist which convincingly demonstrate that the one-way causation model which underlies virtually all group-statistics-based educational and psychological research is inadequate to account for even the simplest of purposeful behaviors. Examples of such research can be found in Powers (1989, 1992), Marken (1992) and in many of the papers included in the September/October 1990 issue of *American Behavioral Scientist* (Vol. 34, No. 1) and interested readers can easily replicate the phenomenon and findings of perceptual control by using the computer programs designed for this purpose by Powers. In addition to the empirical and predictive adequacy of PCT in those areas in which it has been tested, it also receives high marks for Amundson et al.'s criteria of "simplicity of foundation" and "logical coherence."

There are undoubtedly many educational researchers who are either content with the current methods of mainstream quantitative educational research or believe that if problems exist it is because these existing methods are not being properly used. For these individuals, my article and the work of Powers and other PCT researchers constitutes disturbances to be resisted, ignored, or ridiculed. This is exactly what PCT would predict. But there are undoubtedly many others who are dissatisfied with the results that quantitative educational research methods have delivered to date. I can only hope that some individuals in this latter group will find my article of interest and the PCT ideas it introduces worth pursuing. And I look forward to hearing from them (and hopefully Amundson et al. as well) on CSGnet4 where we can continue this conversation.

References

Marken, R. S. (1992). Mind readings: Experimental studies of purpose. Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1989). Living control systems: Selected papers of William T. Powers. Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1992). Living control systems II: Selected papers of William T. Powers. Gravel Switch, KY: Control Systems Group.

Runkel, P. J. (1990). Casting nets and testing specimens: Two grand methods of psychology. New York: Praeger.

Shrag, F. (1992). In defense of positivist research paradigms. Educational Researcher, 21(5), 5-8.

Notes

1This is not technically true for simple bivariate relationships derived from correlational (non-experimental) studies where it is recognized that the direction of causality is unknown and might be from y to x instead of from x to y, or that some other variable z might be the independent variable determining both x and y. Nonetheless, in all of these cases it is still assumed that there is a one-way direction of causality, whichever way and from whatever source it may be.

2Amundson et al. ask "who are these people" (p. 7), namely, the people who use the research methods I criticize. As a introduction to this community of researchers, I suggest they begin by looking through the last four issues of American Educational Research Journal (fall 1991 through summer 1992) where in the section on "Teaching, Learning, and Human Development" two-thirds (16) of the 24 published studies employed statistical analysis of group-data in an attempt to find relationships between one or more independent variables and one or more dependent variables.

3Powers's computer programs can be run on an IBM-compatible personal computer with a mouse or joystick. These programs are available via anonymous FTP from biome.bio.ns.ca as pub/csg/programs/msdos/dem1a.exe and pub/csg/programs/msdos/dem2a.exe.

4CSGnet is an electronic forum on PCT with the address CSG-L@VMD.CSO.UIUC.EDU. It is also accessible as a Usenet (NetNews) group as bit.listserv.csg-l. Individuals can subscribe to CSGnet via LISTSERV@VMD.CSO.UIUC.EDU or can contact the author (G-CZIKO@UIUC.EDU) for more information.

Date: Wed Sep 09, 1992 12:01 pm PST
Subject: Imagination objectives

[From Chris Love (920909.15:30)]
[To: Anyone concerned with the control system imagination issue]

Hello again,

I've been *off the air* although I've been heavily involved in the *Little baby's* development. I read Bill's post on imagination [Bill Powers (920906)] and felt that something was either left un- explained or missing. This dealt specifically with the objective(s) of imagination. To me, an objective is to assist ECS' that are having difficulty arranging their weights to percieve the desired control variable (I think you guys call it CEV now). Now I like Bill's proposal since it does away with having to *pair up* each and every percept (input) with the corresponding *output* of an ECS in the imagination loop. This is the way I originally percieved this to occur. By Bill's proposal on imagination, I mean that he uses a single signal, that signal being the processed error signal of the ECS, which he feeds back into the comparator (is that right Bill - this is how your diagram looked).

So what's my problem? The objective of imagination, in my view is to better organize the control system. I mean if you didn't do something to improve the operation of the system then why do it? Would it be simply to know *if* the control system can control the vble. or more specifically if an ECS (or group) can control their percepts?

So what's my proposal? Although I really like the idea of working *only* with that 1-D error signal, it appears that imagination requires the merging of each percept with the corresponding (matched) output. The end result is the same; that result being that their is a closed path through the ECS.

What's the difference (between Bill's proposal and this one)??? Well, first and most important, as I see things, is that an imrpovement results from using imagination, as it does for us in most cases. For instance, if I imagine (or in my own words ***picture***) how it would be to throw a dart into the centre of a dart board, I do certain things like feel the weight of the dart, use past experience about throwing, test out the inertia of my arm, etc, and use all this information to make the perfect toss. This preparation process is providing me with references that I feel will bring about an improvement to my game.

In terms of the control system, these *improvements* occur when the imagination loop is running and adjusting the weights (percept/output). Now, what I'm not proposing yet, but will soon be investigating is *HOW* this should be done. That's my next step. Sometimes, (this is more a feeling) I see the effect of imagiانتion as being similar (maybe only in results) to the big discussion on random reorg.

This finishes my opinion on imagination.

Where am I in the Little baby? Well, I have just finished developing the beautiful recurrsvive lowpass filter (2nd order) that will be used in the control system. Based on the Nyquist theorem, etc., each ECS beginning at the lowest level, will have an ever decreasing bandwidth, wich will always be half of its lower level neighbor. This is also nice from another perspective. This is to say that if you desire to *skip* a layer, you will still be safe in terms of bandwidth/aliasing since the highest layer into which the percept enters will always have an adequate passband necessary for anti-aliasing AND their is a nice *somewhat* linear phase delay in the filter too!!

So now the question arises as to how I will propose to test out these theories? I now want to make a pseudo control system, where most everything that can be preprocessed is, excluding a single ECS layer. Here I can isolate the effects of the various learning/reorg. algorithms that the group here at DCIEM (Martin Taylor's group) will develop. I hope this will provide a nice testbed for any proposals that come up on the net for percept learning and/or output learning. I hope this to be able to provide a fast and easy way of verifying their worth rather than disputing them in theory. I think Bill will agree here as well as others. AS they (who are they anyways???) say - "There's no substitute for the REAL THING", huh!

Back to the books/Prograph (software).

Best,

Chris Love. (DCIEM)

Date: Wed Sep 09, 1992 1:51 pm PST
Subject: Nits; Stella etc.; Gary's paper; raising organisms

[From bill Powers (920909.1430)] Chuck Tucker (92-09-09)

Every now and then you decide to put all the arguments into some kind of order. Very nice. If you keep doing this, we may yet figure out what we're talking about.

-

Greg Williams (920909!) --

>Bill, I think we've reached an important consensus about my claims.
>Now, we can go back and nit-pick some more, or ...

I want to nitpick some more.

>>I agree with you that lots of this kind of deception goes on.

>>YOU GOT IT!!! But there's still more to get....

I think that we agree that deception is a form of manipulation. We also seem to agree that rubber-banding is a form of manipulation (if anyone doesn't understand what rubber-banding means, ask). But one of these forms is control, while the other is not.

When you manipulate action by applying a disturbance to a known controlled variable, the action will correlate with the disturbance somewhere close to -1.00, as long as the other person continues to control that variable relative to the same reference level. Accurate predictions can be made of individual actions over considerable periods if the reference level stays the same for a while, as it will usually do if you request it, or if the level is high enough.

When you manipulate action by deception, you are betting that your deception will create a perception in the other person that will lead to the action you want. This may work and it may not. Furthermore, it's basically an open-loop process, in that deceptions typically follow a formula which is applied in the expectation that the predicted action will occur. Deceptions clearly do not

work on everyone; those who employ them rely on population statistics to find people who do in fact behave as the manipulator wants. Advertisers pay to modify the perceptions of tens of millions of people, but they're satisfied if perhaps one person out of a hundred is deceived, and one out of ten of those produces the intended action as a result. In an audience of 10,000,000 that's 100,000 sales.

So manipulation by deception resembles an S-R phenomenon more than a PCT phenomenon. A theory of manipulation by deception can't predict actions of an individual. Its predictions would be statistical, and therefore would apply only to a population. And proving that manipulation by deception actually works would be extremely difficult, because you have to take into account the people who would perform the wanted act (buy the car, drink the beer, hand over the money) even without the deception. A theory of manipulation by deception exists on the same plane as theories of diet and psychoanalysis.

Manipulation, therefore, seems to mean at least two things that are so different from each other than they don't intersect anywhere. The rubber-banding kind is reliable, precise, repeatable, and predictable from basic principles. The deception kind is statistical, unreliable, and dependent on peculiarities of individual experience. It is possible to predict the outcome of the rubber-banding kind of manipulation for an individual. It is not possible to predict (in the same sense) the outcome of a deceptive manipulation of an individual, either from experience or from basic principles.

Are there any forms of manipulation that have the qualities of a control-type manipulation? Or are the others all statistical and relevant only to populations?

--

Eric Harnden (920909) --

The Dynamo approach get too close to programming for most people to learn it easily. Anyway, it's expensive, too, isn't it?

I think we should start leaning on Pat and Greg Williams to get their simulation toolkit running. They write in C, so someone with a Mac could adapt it for the Mac world. And maybe there are some mainframe types out there who would do the same for other machines. Since this will be a language for talking about PCT, it won't need all the industrial whistles and bells.

> ... all of which is to say that i think we've got workable alternatives
>to ascii graphics for the effective communication of model structures.

That's what we need, all right. But cost and skill at programming are still barriers for lots of us.

-

Gary Cziko (920909) --

The paper would be easier to read if you would do a search on "and" and convert most of them to periods and a new sentence.

I hope it works -- sounds like you mashed down pretty hard on some toes, there. Your appeal to those who AREN'T satisfied is probably the most realistic way to go. The majority of people who publish in any field are going to defend what they're doing as perfectly adequate. Proving that it's not will not deter them for a moment. They don't even bother to understand your arguments; they're too busy thinking up how to demolish your conclusions.

The real test will come with publication.

--

Thomas Baines (920909) --

>I'm not very far along with PCT so you'll have to wait 'til I catch up.
>Try this, however...

I can't even understand it. How can I try it?

>This reception/emission sets up patterns which reenforce one
>another in some places & cancel one another elsewhere.

It's much too soon to start offering alternatives to PCT. Most people take about 2 years to get comfortable with control theory, so they can reproduce it from scratch from basic principles. I have a feeling it would take me at least that long to be able to reproduce your theoretical presentation from basic principles, particularly as you haven't said what they are.

The kids are fine. Roxanne (the dog) has become fond of chasing rabbits. I want to alter that behavior. I think she is controlling for the hunt & chase. I don't think it is for food. I am trying to find a way to get her to substitute tennis balls & frisbees. Stay tuned.

I'll bet you aren't going to do this by setting up reception/emission patterns that reinforce in some places and cancel in others. I'll bet you're going to do it pretty much the way PCT would describe it. That is, if we on this net could ever agree on how PCT would describe it. In the mean time, let us know what you do and how it works. That, at least, is DATA.

With small pups, start with two boxes which fit together to make a covered sleeping space just slightly larger than the pup. They will quickly be ready to wait for the yard or the paper.

I said you were a genius.

-

Best to all, Bill P.

Date: Thu Sep 10, 1992 3:34 am PST
Subject: I disagree

From Greg Williams (920910) >Chuck Tucker (92-09-09)

Chuck, your post makes several good points. Our remaining disagreements were addressed by my last post (yesterday) -- I think. Let me know of particular questions you still have, if any.

Bill Powers (920909.1430)

>I want to nitpick some more.

The manipulation (facilitation, I hope!) worked!!

>I think that we agree that deception is a form of manipulation.

Back to precision (Talmudicity?): Some kinds of manipulation, including examples of ALL FOUR types I enumerated in my last post, can SOMETIMES MAKE USE OF deception. They are NOT deception, just MAKE USE OF IT sometimes.

>We also seem to agree that rubber-banding is a form of manipulation...

Yes. It is a kind of purposeful influence (= manipulation) in my definition.

>But one of these forms [amended to be called disturbing actions associated
>with a controlled perception of the manipulee (rubber-banding) and altering
>certain of the manipulee's perceptions (suggestion? I don't really have a
>name for it, but it DOESN'T always involve deception!) -- GW (920910)] is
>control, while the other is not.

Both forms involve, we assume from the definition, control by the manipulator of certain of his/her own perceptions, and so are "control" from that point of view. But perhaps one or both aren't "control" from a different point of view. At any rate, we need to be clear ABOUT the meaning of "control" here.

>When you manipulate action by applying a disturbance to a known
>controlled variable, the action will correlate with the disturbance
>somewhere close to -1.00, as long as the other person continues to
>control that variable relative to the same reference level.

Yes. Note the "as long as"!

>Accurate predictions can be made of individual actions over considerable
>periods if the reference level stays the same for a while, as it will usually
>do if you request it, or if the level is high enough.

SOME manipulees will "do as you request." Not all. Go up to a homeless person on the street and ask him/her to stand on one leg for two minutes. Lots of luck....

>When you manipulate action by deception, you are betting that your
>deception will create a perception in the other person that will lead
>to the action you want. This may work and it may not.

How true, how true.

>Furthermore, it's basically an open-loop process, in that deceptions
>typically follow a formula which is applied in the expectation that the
>predicted action will occur.

Sometimes it is "open-loop," but it is often CLOSED-LOOP: the "dance" of manipulative interactions of which I have spoken in earlier posts. And the same holds for instantiations of "suggestion" manipulations which DON'T use deception. As noted above, there ARE such instantiations, particularly prevalent, I might add, when the manipulations are "good" for the manipulee. Ed's interactions with the offender "willing to work with him" alter the offender's notions about being able to control outcomes of importance to him/her -- and Ed doesn't DECEIVE the offender about the relationship between (certain of) his/her actions and the (important to him/her) outcomes. And if Ed finds that the offender is becoming, let's say, "hostile," then Ed will modify his altering in ways which he thinks are appropriate.

>Deceptions clearly do not work on everyone;

Right again.

>those who employ them rely on population statistics to find people who do in
>fact behave as the manipulator wants.

This is true in some cases, i.e., when advertisers want to control for profits generated by selling to a population en masse. But the encyclopedia salesperson (Pat used to be one -- a POOR one, because she really thought it wasn't huckstering) manipulates INDIVIDUALS. So do con men. Statistics would be unhelpful, to say the least, for these manipulators, and many more -- but not for ALL deceptive manipulators, of course.

>So manipulation by deception resembles an S-R phenomenon more than a PCT
>phenomenon.

No, per the above comments.

>A theory of manipulation by deception can't predict actions of an individual.

No, the theory says that prediction of actions of an individual requires exploring (by the Test, in essence) the controlling of the individual. The same holds for predicting that the rubber-banding manipulee will "cooperate" sufficiently. Here, PCT is the theoretical basis for saying how such predictions can be improved.

>Its predictions would be statistical, and therefore would apply only to a
>population.

They COULD be, but don't ONLY have to be. Here, PCT would NOT be the theoretical basis for notions of how best to predict; rather, statistics math would be.

>And proving that manipulation by deception actually works would be
>extremely difficult, because you have to take into account the people who
>would perform the wanted act (buy the car, drink the beer, hand over the
>money) even without the deception. A theory of manipulation by deception
>exists on the same plane as theories of diet and psychoanalysis.

Only the STATISTICAL theory of prediction of manipulation by deception working or not. That is, it would be valid for a population of cases but not for individual cases.

>Manipulation, therefore, seems to mean at least two things that are so
>different from each other than they don't intersect anywhere.

I think it does include two basically different kinds of "approach" by the manipulator. (On the other hand, both "approaches" could be intertwined, and I suppose that Ed, say, uses both at times.)

>The rubber-banding kind is reliable, precise, repeatable, and predictable
>from basic principles.

No more so than the other kind (construed more broadly than you have been construing it: as NOT NECESSARILY making use of deception). To work, both kinds require adequate predictions about another's controlling, which can be made only (unless "luckily") by employing The Test. Many times The Test is implicit, as when one predicts (almost always successfully) that a student -- THIS student, sitting HERE -- will "cooperate" in rubber-banding with ME, the teacher. The student might suddenly go berserk. I've never seen it happen yet, of course. The "of course" (of course) is related to my faith in continuing uniformities of certain sorts in the world, as (to an extent) appropriately criticized by Hume. If you want to rubber-band with a stranger and the outcome (for you, the manipulator) is VERY IMPORTANT, you better talk with the would-be manipulee and decide on the basis of what he/she says and how he/she acts whether to TRUST the "truth" of his/her saying, "Yes, I will go along with this."

>The deception kind is statistical, unreliable, and dependent on peculiarities
>of individual experience.

I disagree.

>It is possible to predict the outcome of the rubber-banding kind of
>manipulation for an individual.

Yes, often, based on using PCT ideas to guide the predicting.

>It is not possible to predict (in the same sense) the outcome of a deceptive
>manipulation of an individual, either from experience or from basic
>principles.

No. PCT ideas can be used to guide the predicting here, as well. Why do you think a con man has to spend so much time interacting with a mark before deciding whether the mark is suitable or not. It isn't just because it takes a while to find out how big a bankroll the mark has. The con man must (in essence) perform The Test with respect to whether the mark is greedy enough (and has other properties of control) to give the exploitation a good chance of success.

>Are there any forms of manipulation that have the qualities of a
>control-type manipulation? Or are the others all statistical and
>relevant only to populations?

I suspect that all forms of manipulation can involve predictions at EITHER the individual level (based on PCT) OR at the population level (based on statistics). Which is appropriate depends on whether you are manipulating an individual or a population. The government loves epidemiology. I'd rather see my family doctor.

>I think we should start leaning on Pat and Greg Williams to get their
>simulation toolkit running.

You mean manipulating???? OK, we're thinking about it. In the interim, some folks might be interested in TutSIM, which costs about \$130 and is for IBM-compatibles only. Non-graphical, somewhat "clunky," and slow (especially without a math coprocessor), but very capable. It is a "block diagram" language which requires less training in math than an "equation" language. I can provide more info if wanted. The Cadillac of IBM (286 or better) simulation programs is Granino Korn's DESIRE. I use it regularly. Blazing speed!!! (Requires a math coprocessor.) The NEUNET version allows easy construction of neural networks (PDP-type) which can be hooked into up to 1000 1st-order differential equations. (!) Unfortunately for novices, it is an equation language. Unfortunately for just about everyone, it costs an arm and a leg (ca. \$500).

>They write in C, so someone with a Mac could adapt it for the Mac world.

Any MACees out there who want to get in on this from the start? Maybe it is even possible to work toward a "stone soup" version (a la FRACTINT, the world's best fractal plotter, done by cooperative efforts and offered as "freeware")?

Best, Greg

Date: Thu Sep 10, 1992 6:13 am PST
Subject: therapy

To: Harry Erwin, Bill Powers, interested others on CSG-L.
From: David Goldstein
Subject: Mahrer and method of levels
Date: 09/04/92

Bill Powers(09/02/03) has pointed out the similarity between the therapy methods of Alvin Mahrer, the HPCT technique of the method of levels, and the ideas which Harry Erwin expressed (09/02/92). He asked me to describe the Mahrer approach so that others could see the similarity.

Alvin Mahrer's brand of therapy is called Experiential Psychotherapy. He has written a brief manual describing it entitled: How to do Experiential Psychotherapy. It was published by the University of Ottawa Press in 1989. Unlike standard practice, two hours are set aside for a session. During each session, he takes a patient through the same four steps:

(a) Step 1--The patient is asked to let awareness go to whatever is the top thing on the patient's mind. The therapist helps the person achieve a very

intense level of whatever experience is selected. There are very specific ways which Mahrer spells out how to do this. The purpose of intensifying the experience is that it becomes easier for the therapist to experience what the patient is experiencing and also, Alvin Mahrer says that these are the optimal circumstances which he has found to encourage significant changes of experience within the patient. Failure to achieve a strong level of experience may result in the decision to end the therapy session.

During this process, the therapist and patient have their eyes closed and are sitting in easy chairs which are placed side by side. The patient describes the experience and the therapist lets the words and the way they are spoken create an experience inside the therapist. As much as possible, the therapist is suspending his/her own experiences during the process and plugging in the described experiences of the patient. Alvin Mahrer says he is not thinking when he is doing this. Unlike most other therapies, and like HPCT Therapy, Mahrer does not give interpretations.

(b) Step 2--Alvin Mahrer has found that when the level of experience becomes very strong, patients undergo a change in experience. The new experience which occurs is described as being at a deeper level. The therapist helps the patient identify, welcome and accept the deeper experience no matter what it is. The deeper experience becomes the center of attention for the remainder of the session. Usually, the deeper experience is one which the patient does not usually or only rarely allows himself/herself to experience.

In HPCT terms, I think what has happened is that the patient has "gone up a level." The deeper potential, in other words, is a higher level reference experience. It becomes the wanted experience. The next two steps are designed to allow the person to experience the deeper potential within the session.

(c) Step 3--The patient is asked about past situations in which the deeper experience was or could have been present to some degree. The patient is asked to relive these past experiences. This time the patient is encouraged to try to achieve the deeper experience as much as possible. This may require them to be/ behave very differently than they did when the past situation actually happened.

(d) Step 4--The patient is asked about present and future situations in which the deeper experience occurs or could occur. The patient is asked to imagine and describe how they will be/ behave in these present and future experiences in order to achieve the deeper experience. They rehearse and make a commitment. Mahrer again spells out very specific ways in which the therapist can achieve this step.

I hope from this brief description you can see why Bill Powers sees a similarity between Experiential Psychotherapy, HPCT Therapy and the ideas of Harry Erwin. Steps 1 and 2 are basically the method of levels done one time. Steps 3 and 4 are ways to help a person control the new experience which was identified in the first two steps. Harry Erwin's procedure sounds very much like Mahrer's therapy and the HPCT method of levels.

Date: Thu Sep 10, 1992 9:38 am PST
Subject: Re: MANIPULATION AND FORCE (long - 370 lines)

[Ray Allis 920910.0845]

The following six posts are from an e-mail list concerned (primarily) with the use of computers in education. Their struggle for clarification seemed to me to sort of parallel some recent discussions here.

Ray Allis

----- Begin Included Message ----- From: KDF1@PSUVM.BITNET

I am glad that Joseph (Parsons) decided to clarify the distinction between the concept of "negative reinforcement" and that of "punishment." I am in complete agreement with the substance of his clarification, but would seek to re-state it since its present form is rather brief, and very technical. Those who are not familiar with psychology as a scientific discipline may welcome a less technical statement of the concepts.

Quite simply, as defined by American behaviorism, punishment is the application of unpleasant (i.e. "aversive") stimuli following performance of a behavior, in order to decrease the frequency with which that behavior occurs. Organisms "escape" the unpleasant stimuli by ceasing to perform the behavior which results in it. Punishment suppresses a given behavior without necessarily promoting another (except as that other behavior incidentally falls into the category of an "escape"). This may be significant for learning in general, since learning most usefully refers to a specific content learned, and not merely to the suppression of some previous learning.

On the other hand, negative reinforcement promotes the specific behavior (albeit also an "escape") which has led to the removal of aversive stimuli. The difference is one of "proscription" versus "prescription." Or, re-stated yet again, we may say that generally speaking, "reinforcement" promotes a given behavior by providing a positive reward for performing it. In the case of negative reinforcement, the reward lay not in the presence of something, but in its removal (i.e. in the absence of aversive stimuli). One can see, then, that the absence of something, cannot simultaneously mean its presence, and negative reinforcement cannot be the same as punishment.

All of this may seem excessively subtle, and I would not want to be seen as a militant advocate of psychological behaviorism (in fact, I am a phenomenological existential and humanistic psychologist). But nearly everyone will admit that human behavior is for the most part neither simple nor transparent. For the sake of methodological clarity and scientific progress, behaviorists have been obliged to refine their operational definition of concepts. More often than not these efforts have been made in good faith.

Kipling D. Forbes
Department of Philosophy
The Pennsylvania State University, U.S.A.

----- End Included Message ----- ----- Begin Included Message -----

I would like to clarify a part of my previous clarification of the concept of "negative reinforcement."

I wrote: <we may say that generally speaking, "reinforcement" promotes a given behavior by providing a positive reward for performing it. In the case of negative reinforcement, the reward lay not in the presence of something, but in its removal (i.e. in the absence of aversive stimuli). One can see, then, that the absence of something, cannot simultaneously mean its presence, and negative reinforcement cannot be the same as punishment.>

The clarified statement should read: we may say that generally speaking, "reinforcement" promotes a given behavior by providing a positive reward for performing it. In the case of negative reinforcement, the reward lay not in the presence of something, BUT IN THE ABSENCE OF SOMETHING ELSE (i.e. in the absence of aversive stimuli). One can see, then, that the absence of something (aversive stimuli), cannot mean the simultaneous presence of the same thing (aversive stimuli), and that negative reinforcement cannot be the same as punishment.

Frankly, I think most people would have accurately interpreted the first version, in so far as it was capable of being accurately interpreted. However, the "clarified" version probably exemplifies a shift from a rhetorical mode of communication to a more strictly logical one. I need hardly add that formal logic will not always suffice for all forms of human communication ^ but feel constrained to note also that the limits of a dialectical logic (if that be not too much of a redundancy) has not yet been fully explored.

Kipling D. Forbes
Department of Philosophy
The Pennsylvania State University, U.S.A.

----- End Included Message ----- ----- Begin Included Message -----
From: Karen O'Quin <OQUINK@SNYBUFVA.BITNET>

Examples of punishment and negative reinforcement:

For negative reinforcement, probably the easiest example is that of putting up your umbrella in order to keep the rain off. The behavior is increased (which makes it reinforcement of some type), in order to keep away an aversive consequent (rain on your new suit). If the behavior increases in frequency, it is by definition reinforcement. (With positive reinforcement, the consequent is positive).

In fact, the whole concept of negative reinforcement is illustrated by the old joke about therapy: The client tells the therapist that clapping his hands keeps the elephants away. The therapist says "But there are no elephants around here!" And the client says, "See, it works!"

Punishment can be identified because the behavior decreases in frequency because of a negative consequent. Touch a hot stove, get burned, and your stove-touching behavior will decrease in frequency.

(To complete all the possibilities, the last category is called extinction, in which behavior decreases because nothing positive results. Tell a joke, no one laughs, and the joke-telling behavior will decrease.)

If behavior increases in frequency, it's reinforcement (either positive or negative). If behavior decreases in frequency, it's either punishment or extinction. By the way, to increase consistency, some textbooks are calling the latter two positive and negative punishment. Somehow, I just can't get used to those terms!

So that's my PSY 101 lecture on learning for the day!

Karen O'Quin, Ph.D.	Bitnet: OQUINK@SNYBUFVA
Department of Psychology	Internet: OQUINK@SNYBUFVA.CS.SNYBUF.EDU
SUNY College at Buffalo	Phone: 716-878-4139
1300 Elmwood Avenue	716-878-6215 (msg)
Buffalo, NY 14222	

----- End Included Message ----- Begin Included Message -----

From: KDF1@PSUVM.BITNET

Barbara Marantz' recent formulation of the distinction between punishment and negative reinforcement is basically correct. She wrote:

>As I understand it, the essential difference between punishment >and negative *reinforcement* is that punishment follows a >response, is "painful" and therefore *decreases* the frequency of >the response it follows. On the other hand, negative >*reinforcement* also follows a response. However, it is >"pleasant" (in that it removes something painful), and, therefore, >*increases* the frequency of the response it follows. That's why >it's called "reinforcement."

I would only caution that behaviorists prefer to speak of "aversive stimuli" rather than "pain" and "punishment." This is not just an attempt to gloss over the plain truth. When dealing with young children or the mentally disabled, stimuli such as a puff of air blown in the face, or cold water from a water gun (again directed to the face) can be very effective aversive stimuli; they are unpleasant, but not necessarily *painful*.

I would also add that in the clinic, or in the treatment environment in general (this may be a custodial residential setting) aversive stimuli are seldom used, primarily because of the dangers of abuse, but also because of what Joseph Parsons mentioned earlier, namely that extinction and negative reinforcement are better ways of shaping behavior, and also because of what I previously said, namely that punishment (i.e. use of an aversive stimulus) does not necessarily offer a behavioral alternative, but merely suppresses previous learning.

The speedy suppression of behavior is sometimes desirable, and may surpass any other consideration. There are cases, for instance, of autistic children who like to break glass (all kinds -- windows, mirrors, etc. using only their hands). Obviously, such behaviors can be dangerous to themselves and to others, and may have to be suppressed at almost any cost.

Generally, however, one needs approval from a committee comprising THE WHOLE WORLD in order to use such behavioral programs, and so for this reason, as well as the more intrinsic ones mentioned earlier, workers in the field avoid

... stuff deleted...

> After all, somebody "applies," somebody "provides" it, and
> if I'm a person I'm likely to be extremely aware of that
> (responding to social situations is people's stock in trade).
> It seems to me that the real difference between "negative
> reinforcement" and "punishment" is that to call it "punishment"
> acknowledges that someone is doing it to someone else. But it
> also seems to me that the normal stance in behaviorist theory
> doesn't have any room for that distinction, because the social
> transaction is bracketed out. That's ok for rats and pigeons,
> who don't engage in social transactions, and don't know the
> difference between a universe which rains on them and a person
> who pours water on them. But people? I don't resent the
> universe; I do resent the experimenter.

Actually, your question sounds more like, "What is the role of the individual's understanding of the social/environmental transaction?"

It is an interesting question, but one that behavioral theory does not try to address. The reinforcement/response paradigm is a simplistic view of behavioral responses to environmental stimulus. What the individual believes or rationalizes as to the motives of the "environment" would fall into the behaviorist's black box (i.e., the mind), an issue that behaviorists believe cannot be observed and therefore should not be subject to scientific inquiry.

Cognitive psychology/science has arisen in this "black hole". They and cognitively oriented social psychologist would be intrigued by your question. Your right, behaviorist have bracketed this issue out.

A comment to everyone in this discussion, remember that your personal definitions of reinforcement and punishment and the definitions used by behavioral theory DO NOT HAVE TO BE RELATED. Just like everything else in science, life and science try to be related, but don't bet on it! :)

ciao, danny

Danny Choriki, Environmental Psychology,
The Graduate School of The City University of New York
33 West 42nd St., Room 1707, New York, NY 10036-8099
Sol III, Milky Way, Local Group

internet: dac@aquila.gc.cuny.edu cis: 71470,3060
cern: dac@cunyvms1 voice: 212.642.2714

----- End Included Message -----

Date: Thu Sep 10, 1992 9:39 am PST
Subject: Portable Anti S-R Demo

[from Gary Cziko 920910.1000]

Rick Marken should already know that my all-time favorite experimental report in PCT is his "The Cause of Control Movements in a Tracking Task" (which is included in his book Mind Readings available for only \$18 from CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328).

This is such a neat experiment because it yields results which make absolutely no sense without PCT since it clearly shows how you can get the same "responses" when the "stimuli" are very different. So I was trying to figure out how this could be done with rubber bands. Here's as far as I got.

Use the classic setup of two rubber bands looped together and thus joined by a knot. The disturber inserts a piece of chalk in his or her end (of the rubber-band, that is!) and the controller does likewise. They start out so that the knot is over the reference spot and the controller is asked to keep the knot there. Now, the disturber slowly draws a pattern or letter or writes a short word while the controller compensates (controls the knot). Then they move to a different spot on the board and do it again (same disturbance pattern used).

Now, the purpose of this is to show that while the disturbance and response patterns are (essentially) the same in the two runs, THE MOVEMENT OF THE KNOT (the "stimulus") IS NOT THE SAME. This is the magic of control. But while it is easy to get a record of the behavior of the controller and the disturber, I don't see an easy way to get a record of the KNOT'S movement.

Hm, maybe if I use TWO sets of rubber bands and use long pieces of chalk (attached to the rubber bands at the bottom and top) I can join the rubber bands on another piece of chalk and have it leave a record of the knots movement. I'll have to try this out. Meanwhile, I would appreciate any other suggestions for this portable demo.--Gary

P.S. I promised at Durango that I would revise my portable demo list and make it available on CSGnet. I haven't forgotten, but I am still working out a few more demos to include and describe (as above).

Gary A. Cziko

Date: Thu Sep 10, 1992 9:39 am PST
Subject: Gary's reply, Neural imaging, Turing

[From Rick Marken (920910.0900)] Gary Cziko (920909.1335) --

>I have written a response to a critique by Amundson et al. to my article
>"Purposeful Behavior as the Control as Perception: Implications for
>Educational Research" which will appear in Educational Researcher which I
>have appended below.

And an excellent response it is, indeed. It is very well-written; clear, concise and it makes the right points. Of course, there is no chance that Amundson et al will get it -- but you slyly mentioned that in the paper. It seems to me that you are now trying to publish papers for the same reason that I do -- so that your grandchildren can hold up your articles and smile smugly at Amundson's grandchildren and say (in a very kind, non- controlling way, of course) "See, Grandpa Gary was right. Boy was he smart".

Here is part of an ad for a conference that was posted to the sci.psychology newsgroup:

> FUNCTIONAL NEUROIMAGING: Looking at the Mind

>Functional neuroimaging is the application of instruments to view the
>changes in physiological state which accompany the work of the brain.

>Topics include:

>__Human Sensation and Motor Control__

>Characterizing the response of the brain to sensation, and tracing the
>activity backwards to look at the physiological states which presage our
>movements and our speech.

And I bet they'll find those physiological states that "presage" behavior. I don't suppose anyone wants to tap them on the shoulder and suggest that all that expensive equipment might be more effectively used study how those "physiological" states "presage" what they really control -- perception. I'd say the SR view is as entrenched in functional neuroanatomy as it is in cognitive psychology. But it's like racism: you only hear it from them when you overhear it.

There was also an article in comp.robotics that dealt with the Turing Test. There is apparently an article on the subject by S. Harnad (my favorite philosopher): Here is the reference: Harnad, S. (1992) The Turing test is not a trick: Turing indistinguishability is a scientific criterion. SIGART Bulletin 3(4) (October 1992) pp. 9 - 10. Appears preceded by an Editorial on the Turing Test by Lewis Johnson, pp. 7 - 9, and followed by another commentary by Stuart Shapiro, p. 10]

I don't know if I'll waste my time reading it (NB-- Penni Sibun; I judge that I will learn nothing from it because I have read Harnad before. Might be a mistake but I'll take the risk). But I do think that the Turing Test is an EXCELLENT example of the behavioristic basis of AI (and cognitive science) etc. Harnad's contribution makes it even clearer -- he suggests a Total Turing Test meaning its not enough to get a simulation to answer questions like a real person -- you must also get it to behave in all ways like a real person -- ie -- brush teeth, play soccer, build model airplanes, etc. Of course, what is interesting is that behavior is defined completely superficially -- it is what you SEE. There is no notion that behavior (including the conversation of the original Turing Test) is a controlled consequence of simultaneous influences produced by the actor and the environment. Turing, like Harnad, was not a physicist; he had no idea how behavior worked. Like present day psychologists he assumed that the words he used to describe behavior were behavior; he perceived behavior at only one level; category.

I think the Turing Test is an excellent model of the basic misconception about behavior embraced by all the life sciences. According to the Turing Test (Total version), the only problem with clockwork simulations of behavior like those done in the 1700s was that they didn't mimic enough behavior and,

perhaps, didn't do it smoothly enough. The Turing Test shows very clearly that the goal of AI and like minded "sciences" is truly superficial -- they want to simulate how things "behave: not how they control.

Best regards Rick

Date: Thu Sep 10, 1992 10:34 am PST
Subject: Imagination objectives

[From Rick Marken (920910.1100)] Chris Love (920909.15:30)

> I read Bill's post on imagination
>[Bill Powers (920906)] and felt that something was either left un-
>explained or missing. This dealt specifically with the objective(s)
>of imagination.

This was an EXCELLENT post Chris. I think you're getting this modelling stuff down, particularly the idea that a good model does it all "on its own" with no help from the modeller. In the discussion of imagination in BCP Bill admitted that he had no idea what "threw the switch" that moved the system from imagination to regular control mode and back. Your post reminds us that whatever does throw the switch must have a reason for doing it -- and being a PCT type you probably (like me) imagine the reason is to control something. So the imagination switch is probably the output of a control system; I agree with your suggestion that this control system is probably involved in the control of some aspect of reorganization -- we imagine in order to learn. I also agree with your suggestion that it would be nice to have a system that alters the imagination connections through the perceptual and output systems so that new types of imaginings can occur. I agree that just playing back error only lets us imagine in the context of the existing structure of the control hierarchy. Your suggestion, I think, is designed to explain how I can put together perceptions I can already produce, but in novel and creative ways, to see if they might satisfy higher order goals. So I can produce in imagination perceptions that I have never controlled -- and probably would never be able to control, like a perception of myself jumping off the roof of my house and flying across the street to the neighbor's. I can imagine it; can't control it. But it was probably the ability to produce such imaginings that guided the reorganizations of the Wright Bros.

The problem now, of course, is to figure out how to do experiments that would guide our modelling of imagination. I can't imagine, right off hand, what those experiments would be, but I think they can be done; I think many have already been done -- though probably not exactly as we would have done them.

Best regards Rick

Date: Thu Sep 10, 1992 10:45 am PST
Subject: Missing mail From T. Baines?

[From: Chris Love (920910.1400)]
[To: Thomas Baines (920909)]

Hi Thomas,

Thanks for the response. However, for some reason or another I am missing your mail in response to Child care magic, which Bill posted a few days ago. Some other mail from my colleagues here, never appeared either, in particular my message yesterday. That is to say, it appears I may not be getting everything. If you could summarize what this topographic folding concerns I would be happy to comment on it if I can.

I did find two posting from you over the last few days; one concerns your child upbringing titled [920909.1020], Sciences and ideologies - I and another short one [920909.12:22] - Dealing with children, but neither talks about this folding topographies, etc. that you mention in your response posting to me [920909], Imagination objectives.

Hope to hear from you soon,
Best,
Chris Love. (DCIEM)

Date: Thu Sep 10, 1992 11:50 am PST
Subject: Imagination; Kinds of manipulation

[From Bill Powers (920910.1100)] Chris Love (920909.1530) --

>I read Bill's post on imagination [Bill Powers (920906)] and felt that >something was either left un- explained or missing. This dealt >specifically with the objective(s) of imagination. To me, an objective >is to assist ECS' that are having difficulty arranging their weights >to perceive the desired control variable (I think you guys call it CEV >now).

The imagination connection isn't in the theory as a way of helping the system do something. It's there because we imagine. The phenomenon to be accounted for is basically that of perceiving something that isn't coming in through the senses. This includes memory, hallucination, dreaming, and planning (by manipulating imaginary images and thoughts). Human beings do these things; the model has to provide a place for them. This is quite aside from asking why imagination exists -- what an organism with imagination can do that an organism without it can't do.

>Now I like Bill's proposal since it does away with having to *pair up* >each and every percept (input) with the corresponding *output* of an >ECS in the imagination loop....

Perhaps what you feel is missing is that pairing-up. In fact, it's still there, although without the complications that the previous version of this idea required.

Suppose a system of level n controls a perception that is a function of 3 lower-level perceptions. Normally, control is accomplished by adjusting the reference signals in all 3 lower-level systems. Those systems in turn work through still-lower systems, and eventually through effects on the environment that alter primary sensory information. Then the result is passed up the perceptual hierarchy until it gets back to the original system as three perceptual signals from the immediately-lower systems that received the three

reference signals. When all is well, those three perceptual signals match the three reference signals.

Let's say that some higher system decides that our level-n system is to operate with imagined perceptions rather than real-time ones. To do this, under the new arrangement, it does not throw the level-n system into the imagination mode, but the three level n-1 systems.

Now the error signal in each level n-1 system is short-circuited back to the comparator as a pseudo-perceptual signal. A copy of this pseudo-perceptual signal heads upward just as it does with real perceptual signals. This happens in each of the three level n-1 systems, so the level n system gets three perceptual input signals as usual. It combines them according to the form of its input function to generate the level n perceptual signal. So for the level n system (and all higher systems), control proceeds normally.

The great advantage of this method is that we don't have to have duplicate copies of all the weights that are applied in the level n perceptual function (or whatever form of computation exists), in order to get the right imagined signal. The level-n perceptual function applies its processing in the same way as usual. The level n-1 systems do no perceptual processing at all, so we don't worry about the weights there, either. So we get the best of all possible worlds: switching to imagination is an elementary process, while applying the correct perceptual function is done by the same means that is used in real-time perception.

I'll be interest in your implementation of an imagination loop that can adjust perceptual weights. This would definitely be part of planning or rehearsing a control process.

>So now the question arises as to how I will propose to test out these
>theories? I now want to make a pseudo control system, where most
>everything that can be preprocessed is, excluding a single ECS layer.
>Here I can isolate the effects of the various learning/reorg.
>algorithms that the group here at DCIEM (Martin Taylor's group) will
>develop. I hope this will provide a nice testbed for any proposals
>that come up on the net for percept learning and/or output learning.

Sounds wonderful. I will NOT ask stupid questions like "when?"

Greg Williams (920910)--

>Back to precision (Talmudicity?): Some kinds of manipulation, including
>examples of ALL FOUR types I enumerated in my last post, can SOMETIMES MAKE
>USE OF deception. They are NOT deception, just MAKE USE OF IT sometimes.

(from previous post)

>There are four kinds of manipulation (I claim): those which, if
>successful ... result in the manipulee's ... judging them ... harmful
>["exploitation"], neutral ["?"], good ["facilitation"], or mixed/who
>knows/can't say ["?"].

These aren't kinds of manipulation; they're judgements about the outcomes of manipulations. Any single manipulative act of any kind could be judged in all four of these ways by different people.

What I'm trying to sort out is the NATURE of manipulation, not its situation-dependent effects. Judging the effects gets us into the subject of "importance," which I still want to postpone for a while. I'm still concerned with establishing what kinds of manipulation are POSSIBLE.

We've established one kind of manipulation, which is controlling a person's action by applying a disturbing influence directly to the variable that the action is being used to control. For short, this is "rubber-banding."

Another kind involves the manipulator having control of part of the environment on which the manipulee's controlled perception depends, in a case where that part of the perception is NOT under direct control by the manipulee, but only contributes one variable to a multi-input- variable perception. For short, this is "manipulation of action through disturbance of uncontrolled perceptions" or just "manipulation through uncontrolled perceptions" (OK?).

In the final analysis, both of these forms of manipulation are manipulation through disturbance of controlled variables. In the "rubber-banding" case, the disturbance is applied directly to the controlled variable. In the "manipulation through uncontrolled perceptions" case, it is applied to a higher-level variable through disturbing an uncontrolled component of that variable. Both cases result in the manipulator being able to exert predictable and quantitative control over the actions of the manipulee, with the usual provisos about not violating other goals of the manipulee. I assume, of course, that the manipulator has correctly identified all pertinent controlled variables.

Can we agree, then, that we have established one method of manipulation, which we can call "manipulation of action through disturbance of controlled variables"? The nature of this method (although not its success) is independent of what the manipulee thinks of the results.

There is another type of manipulation that you've brought up previously. We could call it "manipulation by altering the properties of the environment." This is what Skinner referred to as altering a "schedule of reinforcement." To do this, the manipulator requires the ability to alter the part of the environment that lies between the manipulee's actions and the manipulee's perceived consequences of those actions. As the manipulee is trying to keep the perceptual consequences matching an inner reference state, the result will be a change in the actions that quantitatively compensates for the change of environmental properties. If the change in the environment is not too large, the manipulee can keep control without any change of internal organization. If the change is large, higher-order control may be needed, or reorganization may commence. In any case, successful manipulation requires that the manipulee keep or regain control by the time the process has ended.

We can call this "manipulation by altering properties of the environment."

Purposeful manipulation of actions is possible only when reorganization does not result from the manipulation. If there is reorganization, the outcome

can't be predicted. There is no possibility of purposeful manipulation unless there is a systematic relationship between the manipulative action and the behavioral outcome.

I can think of one exception: the case when the manipulator is using the E. coli method of acting. If the manipulator chooses a manipulative act randomly, but changes the action immediately if the observed behavior of the manipulee departs further from the desired action, eventually the manipulator will arrive at an effective manipulation (or die of old age). This, however, would not seem to be a markedly effective method, particularly if the manipulee is also reorganizing.

If we stick to cases in which systematic manipulative action has a systematic and reliable effect on the actions of the manipulee, we are left with

- (a) Manipulation by disturbance of controlled perceptions, and
- (b) Manipulation by altering properties of the environment.

Another type of manipulation can be called "giving new information." To give NEW information is not the same as disturbing an already-existing and already-controlled perception by either method (a) or method (b). New information is initially uncontrolled, and in fact may initially be unperceivable by the manipulee. Let's assume, however, that the manipulee can understand the new perception.

New information can alter the way behavior looks, in a predictable way. You see someone trying to pull a nail with a hammer, but holding the hammer-handle very close to the head. You say "The farther out on the handle you hold it, the more force you can apply to the nail." Assuming that the manipulee is not getting anywhere with pulling the nail out, and continues to want it out, you can predict that after a moment's thought, the manipulee will make the connection between the relationship you have just described and the task at hand. You can predict that the manipulee will change the grip toward the end of the handle. Of course you can't predict whether this will pull the nail, or whether the manipulee will make the connection and adopt your new information, but that's another matter.

This form of manipulation becomes far more reliable if the manipulee REQUESTS the new information. "Why can't I get this thing out?" Such a request for information practically guarantees that new information will at least be heard. Of course then giving new information is not an act of manipulation, but a response to a disturbance: it is the "manipulee" who is controlling for getting the information. This still leaves the manipulator, however, in the position of being able to choose the information in terms of the effect the manipulator wants. The manipulator could say "Because you're not holding the handle close enough to the head of the hammer," if the objective is to see the manipulee continue to fail, or even just to see if the manipulee will do any stupid thing you suggest.

Asking a question has already been covered under "manipulation by disturbance of controlled perceptions." Most people believe that hearing a question should be followed by hearing oneself or someone else emit an answer. It's like answering a telephone -- and the use of the term "answering" in that context is interesting. What do they say in other languages?

My last entry is "manipulation to see what the other person will do." This is not strictly purposeful manipulation, but only a statement of the purpose for which manipulations might be carried out. It can, however, lead to the other person doing things that satisfy the manipulator's purpose of understanding. Like the other methods, it depends for its success on doing nothing that prevents the manipulee from maintaining control of all variables for which the manipulee has reference levels.

Now we have three methods of purposeful manipulation, excluding the one above that really entails the first three:

- (a) Manipulation by disturbance of controlled perceptions.
- (b) Manipulation by altering properties of the environment.
- (c) Manipulation by giving new information.

Note that nothing has been said about whether the manipulee later judges the results to be good, bad, or neutral, or whether the application of a method of manipulation succeeds.

Can you think of any more?

David Goldstein (920909) --

Thanks for the post on Mahrer (and for showing me how to spell his name).

-
Best to all, Bill P.

Date: Thu Sep 10, 1992 2:04 pm PST
Subject: Barry Richard and STELLA

[from Gary Cziko 920910.1558]

I just had my first encounter with STELLA (version 2.1) and it sure looks like an easy way to do PCT-type modelling. Yes, Rick Marken would go bananas with this program (sorry, I should've said "MORE bananas")!

But I must say that I was distracted by the manual. I didn't expect to get a tutorial on control theory, but that's part of what's in the first chapter.

On page 18:

"This ongoing, circular control process in which "CONDITIONS LEAD TO ACTIONS WHICH RETURN TO INFLUENCE CONDITIONS" is known as a FEEDBACK LOOP. This closed-loop structure is well-known in engineering, the physical and life sciences, and in medicine. It is somewhat less-established in the social sciences, although it is every bit as applicable there. FEEDBACK CONTROL is an extremely important process; one which is absolutely essential to maintaining viability in a changing environment. We will revisit this concept several times throughout the remainder of the Guide."

On page page 37 and 38:

"The "internal feedback structure" viewpoint carries with it two additional perceptual reorientations which also are alien to many people.

"The first is a shift from a "discrete" to a "continuous view of underlying processes. Feedback loops, and goal-seeking activities, are CONTINUOUS in their operation. Thus, for example, you do not continue to ship product flat out until you have ZERO inventory--at which point you ship zero. Rather, as inventory falls below DESIRED levels, you would begin to restrict shipments (allowing your "lead time" to rise). Restrictions would become increasingly severe as inventory fell increasingly far below desired levels. A feedback viewpoint thus implies "continuous regulation", not flip-flop! Your STELLA models should reflect this viewpoint.

"The second reorientation in viewpoint, not unrelated to the first, is a shift from a static/equilibrium perspective to a dynamic/disequilibrium view. Much of our thinking is static. For example, it is not uncommon to hear people express sentiments like: "overwork causes stress". IN such expressions, causality runs only ONE-WAY: from overwork to stress. But in fact, it is just as legitimate to say that "stress causes overwork". If you're stressed, you're [sic] mind's not at ease. Maybe you're not sleeping too well either. Your diet (and exercise routines) may also be disrupted. All of these influences combine to result in a decrease in your work-related productivity. Lower productivity, in turn, means less work gets done, and more errors get created in the process. Slower work rates, and more rework, mean more work to do. More work means less sleep, more pressure, more fatigue, and still lower productivity--thereby reinitiating the cycle. We've all been there. And, the causality is NOT linear, it's CIRCULAR!"

The author of the manual and president of High Performance Systems (which publishes STELLA) is Barry Richard, Assistant Professor at Dartmouth (at least in 1988 when the manual was published). Does anyone know what his field is? He certainly would fit into the company on this network. Maybe we give him a free one-year membership in CSG in return for a steep discount on STELLA for CSG members!--Gary

P.S. Bill [Powers], I am dreaming of playing with the reorganization model you showed me in Durango on STELLA. But I still have a lot to learn. I tried to build a simple control system and got the message "no circular connections allowed" as I put in the last arrow! I think I forgot to add an "inflow" source to the model (STELLA makes a distinction between "input" and "inflow" which is not yet clear to me).

Gary A. Cziko

Date: Fri Sep 11, 1992 6:18 am PST
Subject: Kinds of Manipulation

From Greg Williams (920911) >Bill Powers (920910.1100)

GW>>There are four kinds of manipulation (I claim): those which, if
GW>>successful ... result in the manipulee's ... judging them ... harmful

GW>["exploitation"], neutral ["?"], good ["facilitation"], or mixed/who
GW>>knows/can't say ["?"].

>These aren't kinds of manipulation; they're judgements about the
>outcomes of manipulations.

Sounds fine to me. I was just trying to emphasize that my "manipulation"
doesn't ALWAYS have to have "bad" results (NOTE: FROM THE POINT OF VIEW OF THE
MANIPULEE!!!)

>Any single manipulative act of any kind
>could be judged in all four of these ways by different people.

Yes, but I think judgement by the MANIPULEE (sometime after a successful
manipulation, or after being apprised of all relevant facts about a failed
manipulation) is most telling. PCTers like to take the point of view of the
organism, don't they, for "true" explanation of behavior?

>What I'm trying to sort out is the NATURE of manipulation, not its
>situation-dependent effects.

Me too -- I want to sort out its "nature and limits" as revealed by PCT
notions. That's what I've been saying all along. Welcome aboard!

>Judging the effects gets us into the subject of "importance," which I still
>want to postpone for a while. I'm still concerned with establishing what
>kinds of manipulation are POSSIBLE.

Fine.

>We've established one kind of manipulation, which is controlling a
>person's action by applying a disturbing influence directly to the
>variable that the action is being used to control. For short, this is
>"rubber-banding."

Yes.

>Another kind involves the manipulator having control of part of the
>environment on which the manipulee's controlled perception depends, in
>a case where that part of the perception is NOT under direct control
>by the manipulee, but only contributes one variable to a multi-input-
>variable perception. For short, this is "manipulation of action
>through disturbance of uncontrolled perceptions" or just
>"manipulation through uncontrolled perceptions" (OK?).

I think so. You've spooked me about agreeing to anything with "control" in it
because I still don't know exactly how you mean it all the time. Manipulation
IS control in the sense that the manipulator is controlling some of his/her
perceptions, but I'm not so sure any more whether saying it is "control of
something else" makes anything less muddy.

>In the final analysis, both of these forms of manipulation are
>manipulation through disturbance of controlled variables. In the
>"rubber-banding" case, the disturbance is applied directly to the
>controlled variable. In the "manipulation through uncontrolled

>perceptions" case, it is applied to a higher-level variable through
>disturbing an uncontrolled component of that variable. Both cases
>result in the manipulator being able to exert predictable and
>quantitative control over the actions of the manipulee, with the usual
>provisos about not violating other goals of the manipulee. I assume,
>of course, that the manipulator has correctly identified all pertinent
>controlled variables.

Possibly. Just yesterday morning, Pat and I were discussing the possibility that "rubber-banding" is just a special case of the broader class of manipulation which I've been trying to convince you is ubiquitous (and not necessarily "statistical"/"unimportant"/always "bad"), lo, these many days. Perhaps the largest part of my problems were due to you not being able to generalize from the notion of "rubber-banding" to a broader class of manipulation? At any rate, I need to consider this more before coming to the conclusion that "rubber-banding" IS a special case of the broader class of manipulation I've been talking about.

>Can we agree, then, that we have established one method of
>manipulation, which we can call "manipulation of action through
>disturbance of controlled variables"? The nature of this method
>(although not its success) is independent of what the manipulee thinks
>of the results.

For now, I'm willing (provisionally) to accept your suggestion that applying a disturbing influence DIRECTLY to the variable that the action is being used to control is just a special case of manipulation as I see it. Ultimately, the issue of distinguishing the two might hang mainly on pragmatics (what does it get me to say that the one is a subclass of the other vs. that they are "different").

>There is another type of manipulation that you've brought up
>previously. We could call it "manipulation by altering the properties
>of the environment." This is what Skinner referred to as altering a
>"schedule of reinforcement."

Here you are being too specific. Skinner would have said that ANY (actual) manipulation of ANY kind would HAVE to be done by altering the properties of the other's environment. He, of course, did not believe in direct mental communication between individuals, and so (for him) it was obvious that the only way anybody can be influenced in ANY way by another is through changes in the influencee's environment. THEN he would have gone on to claim (rightly, wrongly, or (I think) confusedly) that altering a reinforcement schedule is ONE way an influencee's environment could be altered.

>We can call this "manipulation by altering properties of the environment."

I call it a broader view of how ANY manipulation could be accomplished.

>Purposeful manipulation of actions is possible only when reorganization does
>not result from the manipulation. If there is reorganization, the outcome
>can't be predicted.

I doubt this as a GENERAL statement. Of course, my first question would be: when does "correct" count when predicting? Ruling out "absolutely correct"

prediction as an absurd demand (even PCTers seem satisfied with 99.999999% accuracy at least some of the time!), we are in the realm of subjective "correctness." I suspect that some (good) con men wouldn't miss a beat in actually FOSTERING reorganization in a mark and still taking him/her. But we have yet to come to a serious consideration of what is and is not entailed by reorganization, right? So I must respectfully decline to answer at this time, because I can't do so in an intelligent way until I have more give and take from you on reorganization models. In particular, I am skeptical about whether the random part of reorganization is really very important (for manipulators) relative to the systematic part (when it stops): the solutions of some (many, I'll hazard!) problems require a very narrow range of actions.

>There is no possibility of purposeful manipulation unless there is a
>systematic relationship between the manipulative action and the behavioral
>outcome.

Yes, or a series of systematic relationships over time.

>I can think of one exception: the case when the manipulator is using
>the E. coli method of acting. If the manipulator chooses a
>manipulative act randomly, but changes the action immediately if the
>observed behavior of the manipulee departs further from the desired
>action, eventually the manipulator will arrive at an effective
>manipulation (or die of old age). This, however, would not seem to be
>a markedly effective method, particularly if the manipulee is also
>reorganizing.

I think you're on the trail of something big here. If you start endowing the manipulator with abilities going beyond the rather poverty-stricken ones of E. coli, then you might just find that allowing more human-like "second-guessing," "anticipation," "intuition" (no, I don't mean anything supernatural), etc., goes a long way toward reducing the "dying of old age" problem.

>If we stick to cases in which systematic manipulative action has a
>systematic and reliable effect on the actions of the manipulee, we are
>left with

(a) Manipulation by disturbance of controlled perceptions, and

(b) Manipulation by altering properties of the environment.

Why isn't (a) a kind of (b)? Be that as it could, I suspect, a third kind having to do with reorganization.

>Another type of manipulation can be called "giving new information."

Still a kind of (b), isn't it? In fact, still a kind of (a), isn't it. In fact, (hypothetically) still a type of my suspected third kind, too.

>To give NEW information is not the same as disturbing an already-existing and
>already-controlled perception by either method (a) or method (b).

To me, it sounds like a kind of (b) and also a kind of (a), but somewhat less direct than "rubber-banding."

>New information can alter the way behavior looks, in a predictable way.

Certainly.

>This form of manipulation becomes far more reliable if the manipulee
>REQUESTS the new information.

Or, more generally, if the manipulator has ANY kind of means of learning about what the manipulee wants to control (or will want to control, given altered context), i.e., he/she applies The Test.

>Asking a question has already been covered under "manipulation by
>disturbance of controlled perceptions."

So has "giving new information," I think!

>My last entry is "manipulation to see what the other person will do."

To my thinking, another kind of (b). But here I get a glimmer of (horrors!) recursion in manipulation: the possibility of manipulating in one way to better achieve manipulation in another way. Oh boy, what a (potential) can of worms. Just what I like to play with, intellectually. Well, nobody ever said that talking about human relationships isn't messy....

>Like the other methods, it depends for its success on doing nothing that
>prevents the manipulee from maintaining control of all variables for which
>the manipulee has reference levels.

Exactly my point over the past week about why (from the PCT standpoint) successful manipulation MUST not involve generating conflict. Nevertheless, I am still tantalized by the possibility that a manipulator might actually succeed by fostering conflict at some point(s) in an elaborate manipulative ruse -- even to the point of "inducing" reorganization.

>Note that nothing has been said about whether the manipulee later
>judges the results to be good, bad, or neutral, or whether the
>application of a method of manipulation succeeds.

Yes. Now, will you buy the notion that there can be SOMETHING "good" (for the manipulee!) in SOME successful manipulations? I am perfectly willing to change the name to "suggestion"; I'm just tired of saying over and over in my posts that MY definition of manipulation DOESN'T require "badness" for the manipulee.

>Can you think of any more?

My "long-term" manipulation idea, involving reorganization, as noted briefly above. But that's all besides your (a) I've been able to come up with so far. Still, there might be better ways of making subcategories.

Best wishes, Greg


Date: Fri Sep 11, 1992 7:50 am PST

Subject: Re: Barry Richard and STELLA

i believe it's actually barry richmond.

he's a system dynamics modeler/programmer. i'm not sure if he came out of the MIT group or not. as to particular specialty, i really couldn't say. he doesn't seem to be referenced in any of my system dynamics texts. what i can point out is that a lot of system dynamics folks (myself included), while they may have particular areas of concern and might therefore do a lot of publishing on a core topic (population studies, economics, urban studies, ecology, etc.), their paradigm causes them to see the world in terms of systems whose structures are essentially similar. the modeling method lends itself to viewing apparently different phenomena as fundamentally related. feedback is feedback, wherever it shows up.

as for the 'circular connections' error, it's a classic opening mistake. here's the canonical system dynamics structure (notice that it's essentially the same as what i showed before):



in which the only material flow is that into and out of the level variable. the rate of flow is controlled by equations which must, in some way, take into account information about that level. often, there is a pretty circuitous path, entailing many transformations and inclusion of extra variables. your 'circular' error arises from violating one of a very few basic rules of system building. i don't have forrester's "principles of systems" sitting at my desk right now, but here are three off the top that should help:

- 1) as just mentioned... any rate that controls flow for a level must in some way receive information about that level. this is the fundamental assumption of feedback.
- 2) no rate can control another rate except through its action on a level. this one can actually be physically violated both in dynamo and stella, and so requires a little care on the modeler's part. one of the skills of system dynamics modeling is determining what element of a system to actually quantify and call a level. if you can't build the system without violating this structural rule, then you've got a clue that perhaps you should rethink your approach toward its representation.
- 3) (here it is...) structural loops cannot be built that would result in simultaneously dependent equations. remember that the method involves integration in time by intervals. there *is* a calculation order. loops must involve something that will delay the information flow. instantaneous infinite feedback is disallowed. in dynamo, one can actually put a delay equation in the loop, and it will work. in stella, the compiler will not even allow you to do that, adhering to forrester's original dictum that any loop must contain a level.

Eric Harnden (Ronin)

Date: Fri Sep 11, 1992 8:49 am PST
Subject: VARIETIES OF MANIPULATION

CHUCK TUCKER 920911

RE: Some comments on "Varieties of manipulation"

I have enjoyed and learned from the conversation between Bill, Greg, Rick, Ed and several others on the topic of "control and manipulation"; it will be useful for my paper: "The Myth of Social Control" but I was somewhat startled by Greg's post 92090[8] in which he noted that he was concerned with several types of "manipulation". I can see what he is concerned with but I have seen these as types of "social transactions" and have tried to reserve "manipulation" to action which is designed "to manage or control (sic) artfully or by a shrewd use of influence especially in an unfair and fraudulent way" (my dictionary's definition). What Greg proposes is that what may be considered manipulation from A's perspective is not so considered from B's perspective. This would be possible if B's did not realize that A had done something "unfair or fraudulent" even though A intended to defraud B. This can happen if B is unaware of the rules or laws, the value of an object or service, believes that there is nothing amiss (e.g., the people in Homestead, Florida all were told that their houses were built to "code" to withstand a hurricane but the winds of Andrew revealed that they were told an untruth so until Andrew they believed the builder and inspector), B is ignorant (e.g., adults deceive children [incest!!] and the child doesn't become aware of it until later). I suppose there could be instances where A would transact with B and not realize that unfairness was involved but this may not fit my specification of "manipulation". But, whatever type of social transaction, I hope there is agreement that each person involved CONTROLS his/her own conduct and what transpires between A and B is NOT control.

I noticed today that Bill stated the issue of "control" and "influence" this way in 1974 ("The Illusion of Control" in LCS:II, pp. 55-66):

We must first distinguish between two sets of terms that are often used interchangeably when the speaker intends to say "control". Going with the word "control" are terms like "regulate, adjust, set, stabilize, direct, guide". All these terms apply directly to the behavior of a negative feedback control system. Another set of verbs is often used loosely when "control" is intended: "affect, influence, change, alter,

impinge upon, stimulate, initiate, drive, cause, force".

The first set necessarily implies the second: to "regulate," for example, one must necessarily "affect". BUT THE SECOND SET DOES NOT ENTAIL THE FIRST. We would agree we could INFLUENCE the behavior of a watch with a hammer, but we would not accept a hammer as a means of REGULATING the watch. . . . When one REGULATES the flow of water from a facet, he does indeed CHANGE the flow, but he does so relative to a PARTICULAR flow that he has in mind. Regulation implies a DIRECTED change, a change toward some predetermined state. (59) [some emphases supplied]

The above relates to what I noted last Spring (not found in the issue of CLOSED LOOP on "Social Control") that we ought to get our terms specified as best we can in the beginning of our conversations. PCT (or HPCT) has special (technical ?!?) specifications for the terms: Hierarchy, Perception, Control and Theory, i.e., for every major term of the formulation we have a specification that varies dramatically from the ordinary use of the term and markedly from every other formulation.

Much (but not all by any means) of the ambiguity in our conversations with others (and apparently ourselves) is that we falsely assume agreement on the specification of our terms. Yes, words are ambiguous in their "meanings" but with all of the other words that each of us has available to use it seems that we might be able, for at least a moment or two, to try to be more precise and specific in what we say to each other. Set those reference levels more precisely and narrow the gain thus increasing the error and reducing our range of allowable words; let's improve our control!

Best regards,

Chuck

Date: Fri Sep 11, 1992 9:19 am PST
Subject: Education as Manipulation

[Gary Cziko 920911.1000]

Greg Williams (920911) said in response to Bill Powers (920910.1100):

>I am still tantalized by the possibility that a manipulator might actually
>succeed by fostering conflict at some point(s) in an elaborate manipulative
>ruse -- even to the point of "inducing" reorganization.

Isn't this what education is? Isn't this what Socrates did with his students?

Here's a nice example. When time for my first sabbatical came up, I decided that I wanted my kids (then going into the 1st and 3rd grades) to learn another language. So I packed up our stuff and carted the family off to France. The kids wound up in a two-room country school in the Alps where nobody spoke English.

Talk about error and manipulation! They had no friends, couldn't understand a word of anybody, but had two very understanding and kind teachers (both males, by the way) and wanted very much to understand what was happening and be understood. I had predicted that they would reorganize and be speaking French well by the end of the school year. In fact, by Christmas my son (1st grader) was hardly distinguishable from native speakers of French and my daughter was not far behind.

This certainly required lots of error and reorganization, but I was nonetheless able to predict quite well the outcome. I suppose they could have both become autistic and catatonic, but they didn't.

So I think Greg is right when he considers education as (his nice) form of manipulation. The only way to push reorganization is to cause error in students. It doesn't always work, but in some situations it's pretty reliable.
--Gary

Date: Fri Sep 11, 1992 9:43 am PST
From: PAPANICOLAOU
EMS: INTERNET / MCI ID: 376-5414
MBX: PAPANICOLAOU@beach.utmb.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: coli

Dag,

The disk is on the way.
Let me know if it works.
Regards, Tom Bourbon

Date: Fri Sep 11, 1992 10:15 am PST
Subject: Re: MANIPULATION AND FORCE

[From Rick Marken (920911.0900)] Ray Allis (920910.0845)--

Thanks for the reinforcement posts. Obviously, people still assume that external events control behavior. This can become pretty silly at times, as in this comment:

>For negative reinforcement, probably the easiest example is that
>of putting up your umbrella in order to keep the rain off. The
>behavior is increased (which makes it reinforcement of some type),
>in order to keep away an aversive consequent (rain on your new suit).

A behaviorist would frown on the "in order to"'s in this description of negative reinforcement. But it shows how difficult it can be to ignore the actor's purposes when you have not been carefully trained to do so. Also, the

fact that reinforcement means "to strenthen" requires that the behaviorist view "putting up the umbrella" as a strenthened response. But what is stronger about it? The fact that it never happens except when it's raining? But this is only true when another condition is also met -- the actor wants to keep dry. And, besides, what would an even stronger version of "putting up the umbrella" be?

How much simpler it would be if behaviorists would see that reinforcements are just disturbances to controlled perceptual variables; positive reinforcements are disturbances that, because of the organism's setting of the reference for the perception, move the controlled perception toward the reference; negative reinforcements are disturbances that move the controlled perception away from the reference. This view explains why the same "reinforcement" -- like rain -- can be positive or negative at different times. When the reference for wetness is 0, rain is a negative reinforcement; when the reference for wetness is > 0 you get Gene Kelly "singin' in the rain". Same perception (wetness) -- different reference; all because of changes in the reference for that perception by that pesky, autonomous system deciding what it wants to perceive.

Simple but, somehow, alarming to establishment life scientists.

The war on autonomy

After making the mistake of watching the news last night I realized why PCT will always have a hard time; people just don't like to believe in autonomy for anyone but themselves. And they will apparently continue to wage war on autonomy even though the consequences of that war are precisely the opposite of what they hope to produce. I am speaking, of course, of the "war on drugs" -- the greatest and most sustained crime creation program in history. After watching the news it became obvious that there was little hope that this idiocy would end soon in this country. Here is a clear case of trying to do, at a societal level, what we have agreed is useless on an individual level -- society is trying to forcibly change the reference level of a controlled variable (drug usage) -- trying to force it to 0 for everyone. This is precisely like "rubber band" controlling -- society is pulling on one end, thinking that this SHOULD make the knot (drug usage) assume a new position. In fact, the knot changes very little but the means used by drug users to keep the knot where it was become as violent as the movements made by society on the other end of the rubber band in an effort to move the knot to where it "should be". I object to this idiocy, not because I want to take drugs (the usual assumption about those who want an end to this drug war idiocy) but because things I care about are endangered by the wild efforts of society and drug users as they tug more and more wildly on the ends of their rubber bands. The only solution is to go up a level ("what do you folks really care about?") or have a police state (a temporary "solution", at best). It looks to me like a solid majority would choose the police state in a second.

Ah well. Pessimistic regards Rick

Date: Fri Sep 11, 1992 10:36 am PST
Subject: Slightly altered copy of Varieties of Manipulation

CHUCK TUCKER 920911

RE: Some comments on "Varieties of manipulation"

I have enjoyed and learned from the conversation between Bill, Greg, Rick, Ed and several others on the topic of "control and manipulation"; it will be useful for my paper: "The Myth of Social Control" but I was somewhat startled by Greg's post 92090[8] in which he noted that he was concerned with several types of "manipulation". I can see what he is concerned with but I have seen these as types of "social transactions" and have tried to reserve "manipulation" to action which is designed "to manage or control (sic) artfully or by a shrewd use of influence especially in an unfair and fraudulent way" (my dictionary's definition). What Greg proposes is that what may be considered manipulation from A's perspective is not so considered from B's perspective. This would be possible if B's did not realize that A had done something "unfair or fraudulent" even though A intended to defraud B. This can happen if B is unaware of the rules or laws, the value of an object or service, believes that there is nothing amiss (e.g., the people in Homestead, Florida all were told that their houses were built to "code" to withstand a hurricane but the winds of Andrew revealed that they were told an untruth so until Andrew they believed the builder and inspector), B is ignorant (e.g., adults deceive children [incest!!] and the child doesn't become aware of it until later). I suppose there could be instances where A would transact with B and not realize that unfairness was involved but this may not fit my specification of "manipulation". But, whatever type of social transaction, I hope there is agreement that each person involved CONTROLS his/her own conduct and what transpires between A and B is NOT control.

I noticed today that Bill stated the issue of "control" and "influence" this way in 1974 ("The Illusion of Control" in LCS:II, pp. 55-66):

We must first distinguish between two sets of terms that are often used interchangeably when the speaker intends to say "control". Going with the word "control" are terms like "regulate, adjust, set, stabilize, direct, guide". All these terms apply directly to the behavior of a negative feedback control system. Another set of verbs is often used loosely when "control" is intended: "affect, influence, change, alter, impinge upon, stimulate, initiate, drive, cause, force".

The first set necessarily implies the second: to "regulate," for example, one must necessarily "affect". BUT THE SECOND SET DOES NOT ENTAIL THE FIRST. We would agree we could INFLUENCE the behavior of a watch with a hammer, but we would not accept a hammer as a means of REGULATING the watch. . . . When one REGULATES the flow of water from a facet, he does indeed CHANGE the the flow, but he does so relative to a PARTICULAR flow that he has in mind. Regulation implies a DIRECTED change, a change toward some predetermined state. (59) [some emphases supplied]

The above relates to what I noted last Spring (not found in the issue of CLOSED LOOP on "Social Control") that we ought to get our terms specified as best we can in the beginning of our conversations. PCT (or HPCT) has special (technical ?!?) specifications for the terms: Hierarchy, Perception, Control and Theory, i.e., for every major term of the formulation we have a specification that varies dramatically from the ordinary use of the term and markedly from every other formulation.

Much (but not all by any means) of the ambiguity in our conversations with others (and apparently ourselves) is that we falsely assume agreement on the specification of our terms. Yes, words are ambiguous in their "meanings" but with all of the other words that each of us has available to use it seems that we might be able, for at least a moment or two, to try to be more precise and specific in what we say to each other. Set those reference levels more precisely and narrow the gain thus increasing the error and reducing our range of allowable words; let's improve our control!

Best regards, Chuck

Date: Fri Sep 11, 1992 11:12 am PST
Subject: On using the same language

From Greg Williams (920911 - 2) >CHUCK TUCKER 920911

>... I was somewhat startled by Greg's post
>92090[8] in which he noted that he was concerned with several
>types of "manipulation". I can see what he is concerned with
>but I have seen these as types of "social transactions" and
>have tried to reserve "manipulation" to action which is
>designed "to manage or control (sic) artfully or by a shrewd
>use of influence especially in an unfair and fraudulent way"

I'd be happy to use "social transactions" (or "suggestion" or, my preference, "PURPOSIVE INFLUENCE") instead of "manipulation."

>But, whatever type of social transaction, I hope there is agreement that each
>person involved CONTROLS his/her own conduct and what transpires between A
and
>B is NOT control.

Yes! Back to square one (my "prolegomenon"), where I was trying to be careful to say that one CANNOT PCT-control another's perceptions. All one can PCT-control is one's own perceptions. How could I be more explicit about this?

>The above relates to what I noted last Spring (not found in the
>issue of CLOSED LOOP on "Social Control") that we ought to get
>our terms specified as best we can in the beginning of our
>conversations.

Amen. That's what I've been saying all along (and why I've needed to become "Talmudic" at times).

>... we [often] falsely assume agreement on the specification of our terms.

I agree completely.

>Yes, words are ambiguous in their "meanings" but with all of the other words
>that each of us has available to use it seems that we might be able, for at
>least a moment or two, to try to be more precise and specific in what we say
>to each other.

And to be able to try to do that without others complaining about the level of detail (sometimes) necessary! Do I wish!!!!

>Set those reference levels more precisely and narrow the gain thus increasing
>the error and reducing our range of allowable words; let's improve our
>control!

Right on, right on, right on. Best, Greg

Date: Fri Sep 11, 1992 11:20 am PST
Subject: Manipulation; Stella

[From Bill Powers (920911.0900)]

A goody to pass on from a murder mystery. A commercial announcer is warming up and doing a sound level check. Over the speaker system in the studio, his important voice says "Yes, friends, try Preparation H and kiss your hemorrhoids goodbye!"

Greg Williams (920911) --

>>Any single manipulative act of any kind
>>could be judged in all four of these ways by different people.

>Yes, but I think judgement by the MANIPULEE (sometime after a successful
>manipulation, or after being apprised of all relevant facts about a failed
>manipulation) is most telling. PCTers like to take the point of view of the
>organism, don't they, for "true" explanation of behavior?

That's what I meant -- "by different manipulees."

>>For short, this is "manipulation of action
>>through disturbance of uncontrolled perceptions" or just
>>"manipulation through uncontrolled perceptions" (OK?).

>I think so. You've spooked me about agreeing to anything with "control" in
>it because I still don't know exactly how you mean it all the time.
>Manipulation IS control in the sense that the manipulator is controlling
>some of his/her perceptions, but I'm not so sure any more whether saying
>it is "control of something else" makes anything less muddy.

I always mean control in only one way: the controller acts on the environment to make a perception in the controller match a reference level determined by the controller.

A manipulator works in two stages (maybe this will make it clearer). One stage is direct control of perception of an environmental variable. This is, shall we say, the _instrumental_ variable -- it will be instrumental in effecting manipulation, if manipulation succeeds. The manipulator does this in order to affect a second variable, which is the aspect of the action of the manipulee that the manipulator perceives, hopes to affect, and desires to control.

In the case of control by disturbance of perceptions, the manipulator controls (which is invariably short for "controls a perception of... ") the

instrumental variable as a means of controlling the action of the manipulee. As long as the manipulee maintains a constant reference level and succeeds in keeping the controlled variable at that level, the manipulator will have true control over the manipulee's action, and the manipulee will have NO control over that action (or the lower-level reference signal that specifies the action). The action is ENTIRELY under the control of the manipulator, not the manipulee -- as long as the manipulee is using this action to stabilize the manipulee's controlled variable. This is only to say that disturbances must be opposed by the action if control is to continue, which is true for all parties concerned.

Mostly for the sake of onlookers, I'll make this concrete by using the Rubber Band Game. The following analysis gets a little complex in spots, so it may be a good idea to demonstrate each point with actual rubber bands and players.

In the rubber band game, the instrumental variable is the position of the manipulator's end of a pair of linked rubber bands. The manipulee's controlled variable is the position of the knot. The manipulee's action that the manipulator perceives, affects, and hopes to control is the position of the manipulee's hand at the other end of the linked rubber bands. The manipulator controls the instrumental variable, the position of his/her own end of the rubber bands. Changing the reference position of the instrumental variable, and thus changing the actual position (because control is involved) tends to disturb the knot, so the manipulee moves the hand holding the other end of the rubber bands to counteract that effect and keep the knot over the selected spot on the table.

Thus the manipulator can affect the manipulee's hand by varying the intended (and actual) position of the instrumental variable. By watching the manipulee's hand, the manipulator can thus CONTROL the position of the manipulee's hand, placing it anywhere feasible that does not cause injury, cause the hand to dip into the soup, or create other problems for the manipulee.

If the manipulee changes the reference position of the knot and moves it by moving the hand position, the manipulator can compensate by altering the instrumental variable, and return the manipulee's hand to the original position. If a third party connects another rubber band to the knot position, and pulls on it, the manipulee will move his/her hand to counteract the disturbance of the knot. But the manipulator, seeing the movement of the hand, will change the instrumental variable so as to return the manipulee's hand to the original position intended by the manipulator.

From the manipulee's point of view, it seems that resisting a disturbance of the position of the knot now requires far less effort than it does when manipulator is not manipulating. Just beginning to move to counteract a disturbance seems to be enough to eliminate it -- in fact, the hand scarcely moves at all, and returns immediately to its original position. When the manipulee intends the knot to be in a new position, the mere thought, and just a suggestion of a movement of the hand, seems enough to get the knot to the new position. Again, the hand ends up where it began.

This uncanny ease of control is actually thanks to the manipulator, who is doing most of the work.

The manipulator, it should be pointed out, cannot freely choose the reference level for the instrumental variable. As long as the manipulator's goal is to keep the manipulee's hand in a specific position, the requirement for doing this is that the instrumental variable be in the position physically required to produce this end. Thus the manipulator's action (positioning the instrumental variable) must depend on the manipulee's choice of reference position for the knot, and on the amount and direction of any third-party disturbances of the knot. So in the final analysis, the manipulator's hand position (the instrumental variable) is not under the control of the manipulator, and the manipulee's hand position is not under the control of the manipulee. The manipulee can control the manipulator's hand position by varying the reference position of the knot; the manipulator can control the manipulee's hand position by varying the reference position of the hand. Each one controls the other's means of achieving the same other's goal.

 The King has arrived in the Queen's boudoir. The Captain of the Guards has hastily concealed himself behind the window drapes.

"My dear," says the King, "you seem a little flushed. Perhaps I should open the window..." The Queen leaps to her feet, rushes to the window, reaches between the drapes, and opens it wide. "You're right, dear Highness. I feel much better now that the breeze is coming in."

The King says "You seem troubled. I cannot see your face in this light. Perhaps if I just pull the drapes back..."

The Queen hastily turns up the dressing-table lamp. "No, Highness, I am quite untroubled; my brow is smooth as you can see."

There is a clink and a scuffling sound. The Queen relaxes. "Of course, Highness, if you will that the drapes be open, so they shall be." And she flings back the drapes, revealing to all but the King that the Captain has escaped through the window.

 >Just yesterday morning, Pat and I were discussing the possibility that
 >"rubber-banding" is just a special case of the broader class of manipulation
 >which I've been trying to convince you is ubiquitous (and not necessarily
 >"statistical"/"unimportant"/always "bad"), lo, these many days.

If we can enumerate all the methods of manipulation that exist, we will see what broader classes there are. I believe that the underlying mechanisms are far simpler than their numerous manifestations.

 -
 >There is another type of manipulation that you've brought up previously. We
 >could call it "manipulation by altering the properties of the environment."
 >This is what Skinner referred to as altering a "schedule of reinforcement."

> Here you are being too specific. Skinner would have said that ANY (actual)
 >manipulation of ANY kind would HAVE to be done by altering the properties
 >of the other's environment.

By altering "properties" of the environment I do not mean altering instrumental variables in the environment. I mean altering relationships among

the variables, and thus altering the effect that any action by the manipulee will have on the manipulee's controlled perception.

This kind of manipulation can be done without the act of manipulation itself disturbing the manipulee's controlled variable at all. Skinner could, by resetting the apparatus, change the number of bar-presses needed to get each reward. Doing so would not, in itself, immediately alter the rate at which rewards are obtained. But it would alter the effects of pressing the bar at any given rate. If the required number of presses per reward were decreased, the rat would find that the next reward appeared sooner. The hourly or daily rate of obtaining food would rise if the rat continued pressing at the old rate, but the rat would soon lower its pressing rate, bringing the obtained amount of food per unit time back toward the reference level. All this assumes that the schedule is not in the range of extreme deprivation, but in the range where normal control relationships hold: more food, less effort.

So by changing the schedule, it is possible to affect the rat's behavior in a known direction. The effect of the change in the properties of the apparatus is made visible only through the effect of the rat's actions on the outcome, the rate at which food is delivered. There is no direct disturbance of the reinforcement rate or the behavior rate. The only effect on reinforcement rate is that of the rat's behavior. It is still true that zero behavior produces zero reinforcement rate.

Skinner seemed to recognize ONLY this means of manipulating behavior. I don't think he ever used direct disturbances of the reinforcement rate.

A simpler example would be that of a person operating a lever to control the position of a spot of light in relation to a moving target. The experimenter can arrange for the ratio of spot movement to lever movement to change, effecting the change just as the lever crosses through the zero position. The change itself will cause no deviation of the spot, but as the subject's action continues it will be found that the spot now responds far more to the same amount of lever movement. As a result, the participant will decrease the lever movements, to keep the spot tracking the target.

This is clearly a different way of affecting behavior. It does not involve applying a disturbing force either directly or indirectly to a controlled variable. It entails altering THE FORM OF THE ENVIRONMENTAL FEEDBACK FUNCTION.

>We can call this "manipulation by altering properties of the environment."

I call it a broader view of how ANY manipulation could be accomplished.

The only thing this method shares with the method of disturbing controlled perceptions is that it allows the manipulator to control his/her own perception of the manipulee's actions. The actual method is entirely different.

In the rubber band game, an equivalent method would be to substitute two parallel rubber bands for the single one on the manipulee's end. Now the manipulee would respond half as much to disturbances at the manipulator's end. In operant conditioning, a second way of introducing this kind of manipulation would be to substitute a new lever that moves sideways rather than up and

down: by this means you cause the rat (eventually) to switch from a vertical to a horizontal action.

You can make an inexperienced public speaker speak less loudly into the public-address system by moving one of the loudspeakers closer to the orator, or obtain exactly the wanted loudness by giving the orator earphones with the volume-setting adjustable by you. You now control the loudness of the orator's speech as perceived by you. The orator continues to control the loudness of his/her own perception of the same speech at the same reference level, but the orator no longer controls the loudness as heard by others.

The environmental feedback function expresses the way the manipulee's perception depends on the manipulee's actions. Such a function can be expressed as a polynomial in ascending powers of the action variable, or as a nonlinear differential equation. If a constant term is included, then the most general means of purposefully manipulating the action of another control system could be described as "control by changing coefficients in the environmental feedback function." An externally-adjustable constant term is, of course, a direct disturbance of the controlled variable. So this last mode of manipulation covers all cases so far.

The method of "giving new information" still remains as an alternative.

OK?

Eric Harnden (920911) --

3) (here it is...) structural loops cannot be built that would result in simultaneously dependent equations.

I hope that you mean "linearly dependent." If Stella rules out cases that amount to analog solution of simultaneous equations, it's not going to be much good for HPCT. It could be used to simulate a single control system, but not two control systems controlling different aspects of the same environment.

Suppose you have two "levels", x and y. You want one control system that controls the perception $s = 2x + 3y$, and a second one that simultaneously controls the perception $t = 3x - 2y$. It is possible for reference signals s^* and t^* to be set and varied arbitrarily, so that, independently, $s = s^*$ and $t = t^*$. Rick Marken's spread-sheet program does a similar thing with SIX ways of perceiving six environmental variables, and at three hierarchical levels.

$$\begin{aligned} s &= 3x + 2y; & es &= s^* - s; & x &= x + gs*es; & y &= y + gs*es; \\ t &= 2x - 3y; & et &= t^* - t; & x &= x + gt*et; & y &= y - gt*et; \end{aligned}$$

If you set gs and gt to small values, this system will converge to the condition $s = s^*$ and $t = t^*$ for any settings of s^* and t^* . I think that x and y can be set up as "levels" with the "flows" being $gs*es$ and $gt*et$. Could you try this with Stella and see if it works? This is sort of crucial for any advanced uses of Stella.

-
Gary Cziko (920910.1558) --

Looks as if Forrester did understand a lot about control theory. George Richardson, author of Feedback thought in social science, is a system

dynamicist and a follower of Forrester's methods. George is friendly to the CSG (he got an article of mine published in the System Dynamics Review) but he has never joined. I'm pretty sure that Forrester never saw the whole picture; his concept of "levels" and "flows" is only a special case of possible relationships that will work in closed loops. But we, too, have found that integral output functions are pretty universally useful. He was mostly concerned with simulating social systems, I think.

Best to all, Bill P.

Date: Fri Sep 11, 1992 12:26 pm PST
Subject: Re: Barry Richard and STELLA

[from Gary Cziko 92091.1310]

Eric Harnden (920911)
i believe it's actually barry richmond.

You're right, it's Barry Richmond (and NOT Barry Richard) who is the man behind STELLA. It is also NOT George Richardson (whom Bill Powers mentioned in his last note.

Sorry.--Gary

Date: Fri Sep 11, 1992 12:35 pm PST
Subject: Language-learning manipulation

[From Bill Powers (920911.1200)] Gary Cziko (920911) --

> So I packed up our stuff and carted the family off to France.
>The kids wound up in a two-room country school in the Alps where nobody spoke
>English. Talk about error and manipulation!

It wasn't purposeful manipulation, but an open-loop attempt to have an effect. It also depended on your having greater physical resources than your children, so they couldn't object. You made an arbitrary change in their environments, but had no means of adjusting that change until the result you wanted -- learning French -- matched your reference level for it. If the children had been so miserable, or the teachers so unsympathetic and xenophobic, that your children just clammed up and gave up, you would have felt sorry that you used that method. Perhaps you had experience that told you your bet was a good one. But it was still an open-loop method, not control. Perhaps we could call this the "sink or swim" method of manipulation.

This method is necessarily statistical and imprecise. It can't be used reliably to produce a wanted effect, although sometimes it may work as you wish.

In case you think I'm criticizing, I'm not. I think you showed great respect for your children's autonomy, and considerable confidence in their ability to reorganize. You presented them with a situation and let them work out the solution. The only "control" involved was getting them to France and enrolling them in the school (with all the implied force in the background, such as

truant officers). If they accepted that as "just the way the world is," then you controlled nothing that violated any of their higher goals. All the rest was their doing; you only made it necessary for them to do it (swim) or fail (sink). A judgement call.

-

Ray Allis (920909) --

I echo Rick Marken's thanks for the posts on behaviorism.

This seems to be an excellent example of what I call "trying to learn the rules." This is what you have to do when you don't use generative models. You observe, classify, and try to memorize what class of event leads to what class of result. There is no explanation at all for why any of these pairings should occur; they just do, and you remember them. In most of the life sciences, this is called "research." When Skinnerians see lower life-forms learning about their environments in this way, they call it "superstitious behavior." The chicken asks not why it has to walk in figure-eights to make food appear. That's just nature's rule. Chickens don't make generative models either. Like Skinnerians, they deal strictly with the surface appearances of natural phenomena.

----- - Chuck
Tucker (920911) --

Your comments are on the money. Be patient. We are working our way toward conclusions you will find acceptable.

-

Best to all, Bill P.

Date: Fri Sep 11, 1992 1:01 pm PST
Subject: Re: Language-learning manipulation

Bill Powers (920911.1200) says:

> It wasn't purposeful manipulation, but an open-loop attempt to have an
>effect. It also depended on your having greater physical resources than
>your children, so they couldn't object. You made an arbitrary change in
>their environments, but had no means of adjusting that change until the
>result you wanted -- learning French -- matched your reference level for
>it. If the children had been so miserable, or the teachers so unsympathetic
>and xenophobic, that your children just clammed up and gave up, you would
>have felt sorry that you used that method.

If I knew nothing about what it takes to learn a language and just expected their presence in France to do the trick, then I guess you could call this open loop. But I was controlling for my perception that they were in fact doing the things that would lead to acquisition of the language. In fact, I spent part of the first couple of days of school with my daughter in her classroom, until I felt she had developed some basic understanding of what was going on. I would have hired a tutor if I had to, but didn't need to. We invited their schoolmates over for meals and sleepovers, etc. So it was anything but open loop. Very purposeful in fact. Very manipulative, in Greg's sense.

> Perhaps you had experience that told you your bet was a good one. But it
>was still an open-loop method, not control. Perhaps we could call this the
>"sink or swim" method of manipulation.

Not really. They could have sunk, I supposed, but not without first getting a
lot of lifesavers thrown at them by my wife and me. We were controlling for
their NOT sinking.

> This method is necessarily statistical and imprecise. It can't be used
>reliably to produce a wanted effect, although sometimes it may work as you
wish.

I think such manipulations can be quite reliable if one has an idea of what
the stages of reorganization look like and then control for those stages. A
math teacher controls for his students to be able to do addition before
requiring multiplication, etc. --Gary

Date: Fri Sep 11, 1992 5:09 pm PST
Subject: Phenomenon, Demo.

[From Dag Forssell (20911)]

Tom Baines, Penni Sibun, Ray Jackson and Uwe Schnepf:

Information packet never did get mailed last weekend. I dealt this past week
with some other (for me larger) error signals. Your packet IS being mailed
tomorrow, Saturday. Includes literature references. Hope you like it.

Dag

Date: Fri Sep 11, 1992 6:29 pm PST
Subject: Blathering

[From Bill Powers (920911.2230)] Gary Csiko (920911.1253) --

I knew while I was writing that thing about "sink-or-swim" language learning
that I was blathering. Sorry. No self-discipline.

Best, Bill P.

Date: Sat Sep 12, 1992 9:33 am PST
Subject: Dynamics of Hierarchical Systems: a reaction

[From Bill Powers (920912.0700)]

As I write this very sketchy and abortive review of Nicolis' Dynamics of
Hierarchical Systems (recommended by Martin Taylor), I feel much older than I
am, which is already old enough. I have this uncomfortable vision of a
crotchety wrinkled old person living in the past, dubious of everything he
doesn't understand, hanging onto simplicities of an older time and complaining
about new-fangled and incomprehensible developments beyond his grasp. This

does not generate much confidence. The fact that I've had much the same feelings about physics since about 1948 doesn't bolster my confidence a lot; perhaps I was as prejudiced then as I am now. Perhaps I have had one year's experience 40 times since my undergraduate days in physics.

On the other hand I've watched developments every year since that first course in "Atomic Physics" with the strong sense of *deja vu*. Nothing really important seems to have happened in physics since that Bomb went off. The main fallout of that explosion seems to have been a sense among physicists that they must now have the key to all the answers: after all, it worked, didn't it? In the light of that first nuclear dawn, physicists looked into each others' faces and saw a new race of superscientists. Mankind had progressed to a new level of understanding of the universe, or at least certain intellectually-favored members of mankind had. In physics lay all secrets of matter and energy, in whatever forms they might appear. The only voice of caution was that of Einstein and he was helpless before the wave of hubris that overtook his science. Suddenly abandoned, he fell silent.

In the preface to his book, Nicolis says this:

"My intention is to start from "scratch" if possible, keeping the reasoning heuristic and tied as closely as possible to physical intuition; I assume as prerequisites just basic knowledge of (classical) physics (at the level of the Berkeley series or the Feynman lectures), calculus, and some elements of probability theory."

This, to Nicolis, is starting from scratch. What he means, of course, is starting from the vast structure of suppositions and inventions that marks physics as different from other natural sciences. He is speaking to those who are already of the right faith. One knows immediately that he is NOT going to start from scratch in analyzing living systems, any more than the Pope is going to start from scratch in analyzing morals.

In Chapter 1 he says this (with allowance for typography):

"_Structural_ hierarchy theory ... aims at reducing the quadratically (N^2) rising number of switching elements as a function of the input- output number N needed to perform a given task to the theoretical minimum $N \log(N)$. This figure comes simply from the fact that $N!$ is the minimum number of states that a network should possess in order to handle N input-output connections. If this is implemented via S switches, the maximum number of states is 2^S ."

[** means exponentiation in Fortran]

So the mind-set behind quantum mechanics takes its place as part of "physical intuition." The cage begins to take shape; the bars are the 1s and the spaces between them are the 0s. The hardware of the organism is assumed immediately to consist of switches that are in states. Anyone familiar with the quantum mind-set knows that we are preparing a universe of categorical things and events in which the only choice is between "it occurred" and "it didn't occur" and in which the only concession to continuity arises from statistics. This, despite the fact that neurons are structures on such a gigantic scale that quantum effects are swamped by continuous relationships.

This book, however, is not about structural hierarchy theory; it simply assumes structural hierarchy theory.

"_Functional_ or dynamical theory, on the other hand, deals with evolution of large-scale systems from stable subassemblies, and studies the way these systems may _change_ through the exchange of energy and information with their environment."

These exchanges, you can be sure, will be quantized sequences, not continuous processes.

"In .. a communicative transaction, 'hierarchical' _coding_ and _decoding_ are again instrumental in reducing the exponentially (2^{*N}) rising number of decoding steps (as a function of the length N of the transmitted time series) down to an algebraic rate N."

"We will deal essentially with the problem of mutual _simulation_ between two dynamical hierarchical systems, and especially with the problem of _compressibility_ of the information which a given system receives from its partner. We proceed, however, in steps."

The systems of which he speaks are assumed, of course, also to "proceed in steps." Step by step, the intellectual cage is built up, the rules of conduct within it are laid out, and the nature of the possible conclusions is shaped to fit the confines of the argument.

If one happens to be looking for it, the organism itself will be seen to be strikingly missing. Starting from scratch clearly doesn't entail watching an organism doing something. Even more to the point, it doesn't involve an organism called Nicolis noticing what it, itself, is doing. We who do notice the organism can see that we are looking at the output of a brain. We are seeing a brain categorizing the continuum of experience into packages, putting the packages into sequence, and weaving the sequences into logical patterns that depend on imagined perceptions, otherwise known as premises. We can see that such patterns can indeed be built up to enormous complexity, and that if they are allowed to develop in isolation from the rest of experience, they can easily pass over into delusion. I think that physics has been transforming itself gradually, over the past 40 or 60 years, into a systematic delusion.

The main thing I was looking for in starting to read Nicolis' book was his concept of hierarchy. It's clear that he has one, but it's also clear that he doesn't know exactly what it is. If he did know, he would describe it unambiguously and plainly. Instead, he speaks from certain assumptions about hierarchy, but the assumptions are largely unconscious, so he speaks around and from these concepts without looking AT them and describing exactly what they are.

One hint comes from the concept of a system composed of switches, so that the whole system is in "a state" at a given instant. This implies that the collection of switches comprises the whole system at its most detailed level of physical organization. Whatever hierarchy exists, then, must exist within this collection as various separable aspects of its overall state. The whole system, apparently, participates in all hierarchical aspects of its organization.

Another hint comes from a distinction made in passing:

"In the first place, the system must be hierarchical, i.e., possess at least a hardware (H)(energetic-structural) level and a software (S)(cognitive) symbolic, functional level."

It does help to have a different place to stand. Nicolis' concept of hierarchy clearly has nothing to do with a hardware hierarchy. The hardware of the whole system is simply one level of the hierarchy. What the hardware DOES is another level. This is very different from the HPCT concept, in which the hardware itself is organized into levels, and what the system does is different at each level as dictated by the nature of the hardware at each level. This doesn't mean that the hardware determines content; only that it makes possible certain computations that could not be done at a different level of the hardware.

The above quote is immediately followed by this, where S stands for Software and H stands for Hardware:

"Information, in the form of discrete time-series emanating from the S-level of system II (S2) [systems I and II are interacting] impinges upon level H of system 1 (H1), where a set of convolutions (cross- correlations) takes place between the time series received and the intrinsic dynamics of system I at level H. The result of such convolutions is the emergence of a collective property possessing far fewer degrees of freedom than S2 or H1."

Hierarchical levels, therefore, seem to be "collective properties" of the physical system, separated not in space but in concept. And the commitment to a discrete universe is set in concrete.

By page 5 in Chapter 1, Nicolis is dwindling out of sight as he goes further and further into his imaginary universe, where I refuse to risk being trapped. By the bottom of page five he is talking about attractors and cascading bifurcations. At the top of page 6 he is telling us we will be considering how a signal is encoded into spherical waves and what happens when those waves, after passing through a stochastic medium, deliver their information to a finite aperture. He is pointing out that because perception is ambiguous, we must deal with the entropy of an information-carrying spherical wave, and that this of necessity takes us into statistical physics and formal information and coding theory. To Nicolis, this is evidently familiar and comfortable territory. He has no problem with these analogies, because he is not trying to explain any particular behavior of any real organism. He is exploring the properties of a world that exists in the imagination, that has only those connections to real systems that he assumes, through free subjective association, that they have. The real world, the world of direct experience, has become the dream, and this world of hypothetico-mathematical entities and relationships has become the reality. That way lies madness, and I will not follow.

I said that this review would be abortive. It aborted on page 6. In order to follow the argument further, I would have to brainwash myself, throw myself under the wheels of this blind machine, abandon all questions and adopt the faith. The thought of doing to myself what would be necessary in order to follow where Nicolis' brain wants to lead me fills me, in fact, with horror. There's no other word for it. I can see a slippery slope that descends at an

ever-steepening angle into an abyss. Venture too far down it, and there is no return. The mind can't free itself from such a trap. Or if one can, it would be a better mind than mine.

All this, of course, could be a feeble defense against admitting my own intellectual failings and the fact that the world is full of people smarter than I am -- considerations that are easy to document. What I feel, however, is not embarrassment or despair at not being able to follow further, but fear at knowing that I could. I've been on this slope before and scrambled back to level ground only by luck, not even knowing at the time that I had rescued myself (or that someone had rescued me).

Ideas are inside us, not outside us. What we think about the world isn't what makes it work, or even what makes it what it is. When we begin to let our ideas take precedence over direct experience, we start losing our grip on that tenuous connection with the real world, and take the first steps into a waking dream in which all is comfortingly consistent and solid -- but imaginary. Be Here Now, the man said. Whatever else he said or did, that one was a zinger.

Best, Bill P.

Date: Sat Sep 12, 1992 9:43 am PST
Subject: Autonomy--Referents

[FROM Dennis Delprato (920912)] Rick Marken (920911.0900)

>The war on autonomy

>After making the mistake of watching the news last night I realized why
>PCT will always have a hard time; people just don't like to believe
>in autonomy for anyone but themselves.

&

Bill Powers (920911.1200)

>[To] Gary Cziko (920911) --

>You made an arbitrary
>change in their environments, but had no means of adjusting that
>change until the result you wanted -- learning French -- matched your
>reference level for it. ...
>In case you think I'm criticizing, I'm not. I think you showed great
>respect for your children's autonomy

I recall Rick Marken, perhaps most vocally, taking a position that the "inner-workings" of the PCT class of models requires an individual who acts autonomously. From this, one postulate of PCT is individual autonomy.

What does this postulate imply? For starters, I see two basic states of affairs. One would be the traditional Western Judeo-Christian-inspired view first well-stated by Augustine whereby the individual is totally separate from the world. Augustine, the great egoistic apostle, interiorized and personalized the universe in the interests of personal salvation. Although

few histories reveal this, Augustine's internalism, with modifications, remains with us today. We continue to find one or another theorist suggesting that the answers as to why people do what they do, think what they think, experience what they experience are to be found solely under their skin (inside the Great boundary of the Skin). Great developments in scientific psychology represent (*partial*) movement away from extreme versions of internalism--dynamic psychiatry (Freud), environmentalism, behaviorism. Interactionism (of which there are unnumbered versions), a recent topic of CSG-L, also represents attempts to "get outside the skin." The least advanced interactionisms (e.g., Bandura's theory in psychology) bespeak their evolution when they offer a combination of inside (e.g., cognitive)-and-outside (e.g., behavioral) factors as always required for understanding psychological events.

I do not find that PCT calls for the sort of autonomy given us by cultural tradition. The equations are not restricted to internal variables in the sense of taking the skin as a Magical boundary. It is not that what in the vernacular we refer to as the "environment" or even "external environment" is always involved. It is that no specific, elementary control system is independent in that a hierarchy (>1) is always involved. No one control system (referring to a single collection of basic components-- reference level et al.) is self-contained. Furthermore, note that disturbances ("external" variables) must always be included (one of great contributions of PCT). Thus, I find PCT progressive on the venerable internal-external issue. It takes neither side *and* does not take a common interactionist position that there is a sequence of I-->E, THEN E-->I, THEN I-->E, THEN E-->I. Self-proclaimed progressive varieties of this add a mysterious feedback loop verbally or graphically (e.g., Kanfer).

I am questioning that the "inner workings" of PCT call for autonomy of the behavior--autonomy with referents going back to Augustine and still found in mainstream psychology.

Nor, do I suggest that Rick's and Bill's use of autonomy above calls for an autonomy postulate of the sort found in mainstream psychology. I find them saying (a) people have a right to be left alone unless they harm others and (b) a fundamental outcome of the way that people function psychologically is that given certain (unnecessary) disturbances they may adjust in ways that are unpredictable, not liked by others, disruptive to themselves.... Thus, the autonomy is more social than psychological. The organization and functioning of the psychological system that incorporates social variables does bear on what the outcome of politico- socio disturbances will be, but not because the individual operates autonomously as an internalist would mean.

So I suggest the second state of affairs (class of referents to autonomy) to be politico-socio. As a recommendation it might read something like: "We know enough about how people adjust to certain things that are done to them to suggest it is desirable to follow a policy of minimum intrusion. In fact, it might be useful to think of each individual as having a right to personal autonomy."

Instead of individual autonomy as a basic postulate of PCT, I find it useful to think of self-control--as opposed to external control. Thus, PCT's message here is that psychological events are under self-control with- out autonomy (in the traditional, internalistic sense).

Perhaps one test of whether or not autonomy is needed in PCT is to examine what PCT would be like without this assumption.

I was most immediately prompted to write this upon reading Rick's statement above and thinking that if Rick were autonomous in the Augustinian way (and how could Rick, above all, be an Augustinian), then he would blissfully go on his way not giving a hoot about what anyone else thinks. He would exercise his autonomy and thereby--to him--make everyone a 100% full-fledged perceptual control theorist. His Augustinian autonomy would free him from the pain and woe of having to be exposed to non-PCT advocates behind every nook and cranny. I am not being facetitious, for what I describe is just the sort of autistic thinking Augustinian autonomy justified and justifies--You don't like this world? Verbally create and maintain another having none of the spatiotemporal characteristic of the one in which we live.

I find that the recent intricate discussion on social control may relate to the autonomy as I have framed it. But I'd hate to have to put it all together.

A not unrelated addendum: For some time I have been touting "noninterventionism," which I believe to be a fundamental implication of naturalistic cybernetics. In contrast to all those who go around seeking to "fix" this and that involving humans, I suggest we LEAVE THEM ALONE. Take so-called psychotherapy. All 300 or so versions, more if one counts up "eclectic" attempts, are interventionist. Do any begin with the position that THE INDIVIDUAL, FAMILY, OR WHATEVER NEEDS NO "HELP"? I've suggested establishing a Center for Non-Therapy. Its function would be to show people they need no therapy. Trouble is I don't know under which DSM code this service could be billed. One thing, such a service would reduce iatrogenic effects.

So, clearly I am for politico-socio autonomy.

Dennis Delprato

Date: Sat Sep 12, 1992 6:53 pm PST
Subject: Morals and Malfunctions

[From Rick Marken (920912.1700)]

Well, thank goodness someone responded to my "autonomy" and "drug wars" post. Thanks Dennis.

Let me respond to one of your comments:

Dennis Delprato (920912) says:

> From this, one postulate of PCT is individual autonomy.

I think "autonomy" is just a word (with some existing connotations that seem appropriate, and, perhaps, some that are not) that is used to describe the operation of the HPCT model; at least, that's how I meant to use it. Instead of saying "autonomy" I could have loaded up my spreadsheet model, hit the F9 key to start it and stood pointing silently as higher level systems told lower

level systems what to perceive in order to satisfy the perceptual goals of the higher level systems.

>Nor, do I suggest that Rick's and Bill's use of autonomy
>above calls for an autonomy postulate of the sort found in
>mainstream psychology. I find them saying (a) people have
>a right to be left alone unless they harm others and (b)
>a fundamental outcome of the way that people function
>psychologically is that given certain (unnecessary)
>disturbances they may adjust in ways that are unpredictable,
>not liked by others, disruptive to themselves....

Speaking for myself, I meant to say something a lot more like like (b) than (a). I think it is good to point this out. My discussion of the drug war was not meant to make a moral point; I am not arguing (with Thomas Jefferson) that people have an inalienable right to be free (if they are not hurting anyone). People might perceive rights and they might try to control those perceptions. But that is not a moral point; it is a scientific one, which can be tested and rejected.

What I was discussing was malfunctions, not morals. If there is a fundamental postulate of PCT it is that organisms are control systems. A functioning control system is able to make its perceptual experience match it's references for that experience; I call this "autonomy" -- the normal operation of a control system. Anything that prevents normal operation is the cause of a malfunction. Conflict is an example of a control system malfunction; conflict prevents autonomy -- ie. the ability to control.

The drug war is an example of control systems IN CONFLICT. So the drug war is an example of control systems that are MAL- FUNCTIONING. There is no moral judgment here; that would imply that I LIKE the goals of one group (the drug warriors) better than I like those of another (the drug takers). In fact, I personally don't care for either of their goals -- but that is NOT why I don't like the drug war. I don't like it because there is CONFLICT between control systems; as I said, this conflict MIGHT have unpleasant side effects for me (I might get robbed by a druggie who has to pay ridiculously high prices for highly abundant substances or have my house broken into by an over zealous SWAT team that was off by a digit on the address of a crack house). But, in fact, the chances of those side effects are fairly low. I really object because conflict prevents the functional operation of the control systems involved; neither party (drug warrior and drug taker) is able to function as a full fledged hierarchical perceptual control system.

The conflict would be solved, of course, if the druggies decided to stop taking drugs OR the warriors decided to stop fighting drug takers (and suppliers). Either approach would end the conflict and people could start functioning again. I favor a solution to this conflict based on the drug warriors changing because they are the ones who created the conflict by trying to control other control systems. The other control systems (druggies) maintain the conflict by maintaining their references for the perception that the drug warriors want them to change. But, somehow (and it's hard for me to articulate it without becoming moralistic) it seems to me that its a lot easier for the drug warriors to stop controlling for what the druggies are controlling than it is for the druggies to change their own reference for what they are controlling. I don't know how to articulate this point in terms of

the model -- but it seems to me that expecting the druggies to solve the conflict by changing their reference for drugs is like expecting jews to solve their conflict with nazis by changing their reference for how alive they want to be (like, to zero).

Best regards Rick

Date: Sat Sep 12, 1992 6:55 pm PST
Subject: Morals and Malfunctions

[From Rick Marken (920912.1830)]

Well, of course. The reason the drug warriors are the problem is because they must push against another control system in order to control the variable they want to control. The victim (the druggie) could (and was until the drug warrior came along) control the variable s/he is controlling without creating conflict in another control system at all. So one set of control systems (the warriors) are CREATING conflict by trying to inhibit the autonomy (not consciously, but that is what they are doing) of others. Since the warriors don't understand PCT, they are creating this malfunction out of ignorance. So I still have no moral complaint here. As I said in my first post, the drug war is just malfunction-producing idiocy (stupidity) that results from a failure to understand the nature of autonomy. So problems like the drug war can be solved, not by trying to articulate better moral principles, but simply by understanding how control systems work. A person who understands control theory would simply shake their heads in dismay at the drug warrior -- just as a person who understands plumbing shakes their head in dismay at the person who pours grease down the drain; both are just watching people create malfunction.

Best regards Rick

Date: Sat Sep 12, 1992 6:58 pm PST
Subject: Autonomy

[From Bill Powers (920912.1400)]

Dennis Delprato (920912) --

RE: Autonomy.

You're coming very close to the position I think I have. As you say, the development is delicate. On the one side there is the pit of solipsism, with the world denied and made imaginary or subject to whim. Down the other side lurk the positivists, who must eventually advocate controlling people for their own good according to objective criteria. At the apex where these two slopes -- one slippery and the other studded with sharp rocks -- converge, there is, I think, a narrow ridge which, followed to its end, will show us a self-consistent view of the environment, of other people, and of ourselves.

Greg Williams is perfectly correct in wanting us to develop a theory of interaction. To cast the problem in terms of manipulation, however, is to beg a question, because we don't know, without analyzing the possible mechanisms, what can actually be manipulated by another and what can't.

The underlying question is what we mean by autonomy: in what respects are organisms autonomous and in what respects not? Organisms clearly depend on their environments (including other organisms) for sustenance of various kinds. Furthermore they have inherited needs about which they have no choice. Understanding autonomy is especially difficult in a hierarchical system, where combating a disturbance at one level entails altering goals at lower levels. Goal-directed behavior is not, per se, autonomous -- in a hierarchy.

There are certain processes in the human organism that are carried out simply because of the way the organism is put together, inside. Control itself is an example, as is reorganization. The environment contains no means of carrying out these processes for an organism, either to help it or hinder it. A human being must acquire perceptual functions that produce consistent perceptions, perceptions that vary in sensible ways and relate to each other without contradiction. The environment can't do that kind of making-sense; only a brain can. Once a perceptual signal exists, it is the brain that must carry out comparisons with reference signals, to generate an error signal. The environment does not inform the human organism of how it, the environment, should be, nor does it tell the organism what constitutes a discrepancy with the organism's goals. And it is the inner mechanisms of the organism that must find out how to convert error signals into those choices, amounts, and directions of lower-order actions that will make the errors smaller. The actual process of finding those means, as opposed to what means will be found, is strictly a product of internal mechanisms for change.

On the other hand, the organism can produce outputs, but it can't determine what the consequences of them will be. All the brain can do is drive the effectors. It is the environment that then serves up the consequences. The environment can generate consequences on its own; it can vary enough so that a given action seldom creates the same consequences twice, and it contains independent forces that affect perceptions. This is my best evidence for existence of a Boss Reality that exists and has properties independent of my experiences.

Because the environment has properties and contains independent sources of disturbance, the organism must learn by itself which output variations will control its perceptions and which will not. Once the organism has begun representing the external world in a particular perceptual way, and once it commits to a goal, to reproducing a particular state of that perception, it has no choice but to produce an output that is sufficient to create that perception. The forms of its perceptual functions determine how the environment will be apprehended, but given those forms, it is the environment, not the organism, that determines what actions will in fact be sufficient to control the perceptual result. The organism must contain means of discovering the effective actions, the means of control, but it has no say as to what those actions will turn out to be.

When an organic control system resists a disturbance, it does so relative to a goal-state of the affected perception. But that goal-state itself is adjusted

as it is because of other disturbances that have effects on higher-level perceptions, making them deviate from higher-level goal-states. There is no end until we reach the top level. Only there do we find the possibility of purely endogenous goal-generation. Below the top level, the way goals are set is determined by the perceptual structure that intervenes at lower levels, and by the properties of the environment, which dictate how resetting a lower-level goal will affect a higher-level perception.

Each level of control, therefore, comes into being through the action of internal mechanisms for change and development, but the final result, the control organizations that come into existence, must be designed to work through the properties of the world that actually exists, however we may perceive it.

Finally there are the intrinsic variables, their inherited reference levels, and the process of reorganization driven by intrinsic error. These are defined for the organism by its heritage. Thou shalt stay warm, but not too warm. Thou shalt breathe, and eat, and drink or find ways of acting that have equivalent effects. Thou shalt reproduce, and find the experience pleasant. Thou shalt avoid injury, and find injury painful. And so on through a list that is undoubtedly longer than any we can now write down.

The environment can affect intrinsic variables, but it can't say what will constitute an intrinsic variable, or a reference level for one. The process of reorganization, being capable of altering anything in the hierarchy of control, both perceptions and actions, overrides all other considerations. The environment may determine what must be done to affect a given perception in a given way, but it is the reorganizing system that decides what will constitute a perception, and whether any particular state of it is to be sought or avoided.

With respect, at least, to the current environment, these built-in systems define the true autonomy of an individual. The individual will learn to perceive the environment, and to control its perceptions relative to particular goals, in ways that use the properties of the environment and satisfy the requirements of the inherited control system that I call the reorganizing system. In an advanced organism, if the properties of the environment aren't sufficient to find a way of eliminating intrinsic error, the properties of the environment will be changed. The rain will not be allowed to fall, the wind will not be allowed to blow and freeze, the sun will not be allowed to beat down -- on this organism.

So autonomy, in any one lifetime, is awarded to the organism, in particular to its reorganizing system. But that is only in one lifetime.

One organism is a member of a species. There are processes of blind variation and selective retention that span the links between generations. The selective retention depends at least in part on criteria endogenous to the species; even the blind variation may be in part an act of the species rather than just a random effect of something else. The species contains mechanisms for change. And it must also contain reference signals that are passed down from generation to generation, which define the reason both for starting and for ending change.

I won't follow that trail any further; the point is only to show that the same principles of autonomy can be extended into the past, with the environment providing the stage and acting mindlessly to disturb, but being incapable of carrying out the processes of change, which continue to reside in the species. This trail leads all the way back to the first molecule that stabilized the conditions on which the accuracy of its replication depended, and thus gave birth to life.

Here and now, we live in a world whose properties are largely unknown to us and which determine for us the effects of our actions. We, however, choose our own goal structures, as our means of preserving ourselves in the state that our natures tell us is right. Between this ultimate personal autonomy and the impersonal events in the nonliving universe, there comes to be a hierarchy of control systems that reflects both our overriding inner needs and the conditions the environment places on meeting those needs.

This is the best I can do for now by way of laying out the meaning of autonomy and the relation of an organism to the world outside it. This still leaves open the question of interactions among organisms; organisms with similar organizations, and organisms asymmetrically related to each other. Whatever we decide to say about manipulation, I think it will turn out to describe the situation presented here: one organism learning how to affect its perceptions in the way it wants, by choosing actions from among those that the environment says are the only feasible ones.

Date: Sun Sep 13, 1992 4:09 am PST
Subject: Catching up on purposive influence

From Greg Williams (920913) >Gary Cziko 920911.1000

Nice example of how a purposive influencer can employ reorganization of the influencee!

>So I think Greg is right when he considers education as (his nice) form of
>manipulation. The only way to push reorganization is to cause error in
>students. It doesn't always work, but in some situations it's pretty
reliable.

It is nice to know that at least one netter understands how I am using
"manipulation."

>Rick Marken (920911.0900)

>After making the mistake of watching the news last night I realized why
>PCT will always have a hard time; people just don't like to believe
>in autonomy for anyone but themselves. And they will apparently
>continue to wage war on autonomy even though the consequences of that
>war are precisely the opposite of what they hope to produce. I am
>speaking, of course, of the "war on drugs" -- the greatest and most
>sustained crime creation program in history.

I think a much different hypothesis may be offered regarding the motivation for the "war on drug[ee]s": that, generally speaking, people are QUITE AWARE (in a non-technical way) that others are control systems like themselves, and

that extreme methods (threats and overwhelming physical force) are necessary to make others quit doing what they very much want to continue to do (i.e., not suffer from drug withdrawal pains), and so the "war" results from the influencers escalating their methods knowing full well that less ruthless techniques won't be effective against control systems.

The result, of course, is STILL not as the influencers desire, since the countermethods of the drugees tend to offset the (even nearly overwhelming) forces arrayed against them (grow dope indoors to avoid helicopter detection, start carrying guns to help prevent getting arrested, etc.). This is addiction on a grand scale, with "co-evolution" of a technology (including arms) race -- and what drives the escalation is both sides' knowing that the other will attempt to offset their disturbances.

I think this argues for attempting to figure out how purposive influence without direct coercion (that is, generation of conflicted control systems) might substitute for violence. I realize it's difficult for an influencer to make use of a potential influencee's controlling when the latter basically has one huge desire: keep blood concentration of such-and-such drug at so-and-so level, within a small tolerance. (When I was in the neurophysiology lab at MIT, one visiting professor who had used heroin one time advised me that "it feels so good you don't even want to try it.") I even suspect that effective influencers will NOT be case workers or other "outsiders," but neighborhood folks who have non-drug desires.

In sum, I argue that the PCT explanation for the continuing "war on drugs" is NOT that we tend to see others as NOT being control systems, but rather that we see ALL TOO WELL, and then want IDEOLOGICALLY ACCEPTABLE solutions. I also argue that PCT shows how the problem can be addressed without violence.

>Bill Powers (920911.0900)

>The manipulator, it should be pointed out, cannot freely choose the
>reference level for the instrumental variable.

Certainly. And, it should be pointed out, the manipulator's knowledge of the manipulee's (past) controlling -- in essence, by performing The Test -- can be decisive in making the appropriate choice of reference level.

>Thus the manipulator's action (positioning the instrumental variable)
>must depend on the manipulee's choice of reference position for the
>knot, and on the amount and direction of any third-party disturbances
>of the knot.

Exactly -- and this generalizes to more elaborate types of this kind of manipulation, as I've pointed out time and again.

>So in the final analysis, the manipulator's hand position
>(the instrumental variable) is not under the control of the
>manipulator, and the manipulee's hand position is not under the
>control of the manipulee.

In the final analysis, neither is controlling for his/her perception of his/her own hand position. Is that what YOU just said?

>The manipulee can control the manipulator's hand position by varying
>the reference position of the knot; the manipulator can control the
>manipulee's hand position by varying the reference position of the
>hand.

No, the manipulee doesn't give a whit where he/she sees the manipulator's hand, so he/she is NOT controlling it. If the manipulee started to control thusly, then the manipulation wouldn't work as the manipulator planned (on the basis of his/her previous knowledge of the manipulee's controlling); in this case, the manipulator would say he/she had a "bad" model of the manipulee.

>Each one controls the other's means of achieving the same other's goal.

This is an example of slippery use of the word control. Are you doing this on purpose, or by innocent error? Regardless, this is the sort of problem I've been up against repeatedly in this discussion. In this case, the manipulator is controlling his/her perception of the manipulee's hand position, but if the manipulator's model of the manipulee's control is "good," the manipulee is ONLY controlling his/her perception of the position of the knot. Sorry.

>The environmental feedback function expresses the way the manipulee's
>perception depends on the manipulee's actions. Such a function can be
>expressed as a polynomial in ascending powers of the action variable,
>or as a nonlinear differential equation. If a constant term is
>included, then the most general means of purposefully manipulating the
>action of another control system could be described as "control by
>changing coefficients in the environmental feedback function." An
>externally-adjustable constant term is, of course, a direct
>disturbance of the controlled variable.

I understand what you're talking about now...

>So this last mode of manipulation covers all cases so far.

... and that's what I had a more intuitive feeling about before.

>The method of "giving new information" still remains as an alternative.

I still think this might be subsumed under your second type.

>Bill Powers (920911.1200)

>Your [Chuck Tucker (920911)] comments are on the money. Be patient. We are
>working our way toward conclusions you will find acceptable.

Aren't we there already? I AGREED with every single one of Chuck's comments, too!

>Dennis Delprato (920912)

>We continue to find one or another theorist suggesting that the answers as to
>why people do what they do, think what they think, experience what they
>experience are to be found solely under their skin (inside the Great boundary
>of the Skin)....

>I do not find that PCT calls for the sort of autonomy given us by cultural

>tradition. The equations are not restricted to internal variables in the
>sense of taking the skin as a Magical boundary.

Exactly my sentiments!

>Thus, I find PCT progressive on the venerable internal-external issue. It
>takes neither side *and* does not take a common interactionist position that
>there is a sequence of I-->E, THEN E-->I, THEN I-->E, THEN E-->I.

Exactly my sentiments! May that progressiveness not be harmed by ideological
"autonomism"!!!

>I am questioning that the "inner workings" of PCT call for autonomy of the
>behave--autonomy with referents going back to Augustine and still found in
>mainstream psychology.

Welcome to the beef-of-the-month club.

>Nor, do I suggest that Rick's and Bill's use of autonomy above calls for an
>autonomy postulate of the sort found in mainstream psychology. I find them
>saying (a) people have a right to be left alone unless they harm others and
>(b) a fundamental outcome of the way that people function psychologically is
>that given certain (unnecessary) disturbances they may adjust in ways that
are
>unpredictable, not liked by others, disruptive to themselves.... Thus, the
>autonomy is more social than psychological.

I would say "more ideological than scientific."

>So I suggest the second state of affairs (class of referents to autonomy) to
>be politico-socio. As a recommendation it might read something like: "We
know
>enough about how people adjust to certain things that are done to them to
>suggest it is desirable to follow a policy of minimum intrusion. In fact, it
>might be useful to think of each individual as having a right to personal
>autonomy."

This can follow from PCT PLUS certain ethical postulates, but not from PCT
alone. PCT alone doesn't say that control-system conflict is "bad" -- not even
if it results in death of an organism. It doesn't say that one's death is
necessarily bad (or good, either)!

>Instead of individual autonomy as a basic postulate of PCT, I find it useful
>to think of self-control--as opposed to external control. Thus, PCT's
message
>here is that psychological events are under self-control with-out autonomy
(in
>the traditional, internalistic sense).

Now you're cookin' -- this is where I came in, with the Prolegomenon.

>Rick Marken (920912.1700)

>I think "autonomy" is just a word (with some existing connotations that seem
>appropriate, and, perhaps, some that are not) that is used to describe the

>operation of the HPCT model; at least, that's how I meant to use it. Instead
>of saying "autonomy" I could have loaded up my spreadsheet model, hit the F9
>key to start it and stood pointing silently as higher level systems told
lower
>level systems what to perceive in order to satisfy the perceptual goals of
the
>higher level systems.

I have no problem with this technical PCT-sense of "autonomy." Do you have a
problem with my saying that what I do now is NOT solely the result of what I
(and my ancestors) did in the past? That is, do you have a problem with my
current control being the result of BOTH my and my ancestors' pasts AND my and
my ancestors' environments?

>What I was discussing was malfunctions, not morals. If there is
>a fundamental postulate of PCT it is that organisms are control
>systems.

Here, too, I think we must begin to be more precise. Doesn't PCT say, more
precisely, that organisms constitute PARTS of ("natural") control systems?
Non-living parts of organisms environments can also constitute parts of these
control systems, which must ALWAYS have a loop through at least one organism.

>A functioning control system is able to make its perceptual experience match
>it's references for that experience; I call this "autonomy" -- the normal
>operation of a control system.

So far, so good.

>Anything that prevents normal operation is the cause of a malfunction.

But (maybe) here comes the moralizing. "Mal-""???

>Conflict is an example of a control system malfunction; conflict prevents
>autonomy -- ie. the ability to control.

But conflict allows "growth" via reorganization. (And "conflict" itself is
subject to the kinds of criticisms I've been subjected to from various netters
regarding my notion that "manipulation" needn't always be "bad" -- if
"manipulation" is supposed to be a loaded word, how much more so is
"conflict." The first definition of "manipulation" in my dictionary is value-
neutral; not so for "conflict.")

>The drug war is an example of control systems IN CONFLICT. So
>the drug war is an example of control systems that are MAL-
>FUNCTIONING. There is no moral judgment here; that would
>imply that I LIKE the goals of one group (the drug warriors)
>better than I like those of another (the drug takers).

There IS a moral judgement here: that control-system conflict is BAD and a
MALFUNCTION.

>Bill Powers (920912.1400)

>Greg Williams is perfectly correct in wanting us to develop a theory

>of interaction.

Thank you, thank you.

>To cast the problem in terms of manipulation, however, is to beg a question,
>because we don't know, without analyzing the possible mechanisms, what can
>actually be manipulated by another and what can't.

I'm not begging the question, I'm calling for answers to it. My "manipulation"
was defined, from the beginning, in as fundamental a way as I could see how.

>Understanding autonomy is especially difficult in a hierarchical system,
>where combating a disturbance at one level entails altering goals at lower
>levels. Goal-directed behavior is not, per se, autonomous -- in a hierarchy.

This can of worms is why I was ORIGINALLY trying to stick with implications of
PCT, not HPCT.

>There are certain processes in the human organism that are carried out
>simply because of the way the organism is put together, inside.
>Control itself is an example, as is reorganization. The environment
>contains no means of carrying out these processes for an organism,
>either to help it or hinder it.

The NON-ORGANISMIC environment, you mean, right?

>A human being must acquire perceptual functions that produce consistent
>perceptions, perceptions that vary in sensible ways and relate to each other
>without contradiction. The environment can't do that kind of making-sense;
>only a brain can.

The NON-ORGANISMIC environment, you mean, right?

>Once a perceptual signal exists, it is the brain that must carry out
>comparisons with reference signals, to generate an error signal. The
>environment does not inform the human organism of how it, the environment,
>should be, nor does it tell the organism what constitutes a discrepancy with
>the organism's goals.

The NON-ORGANISMIC environment, you mean, right?

>Each level of control, therefore, comes into being through the action
>of internal mechanisms for change and development, but the final
>result, the control organizations that come into existence, must be
>designed to work through the properties of the world that actually
>exists, however we may perceive it.

I suggest that the "coming into being" of levels of control depends on the
actions of BOTH internal AND external mechanisms for change and development.

>Finally there are the intrinsic variables, their inherited reference
>levels, and the process of reorganization driven by intrinsic error.
>These are defined for the organism by its heritage.

A heritage, I suggest, both hereditary AND environmental (think, for example, of physiological set-points changing when you move to a high altitude; and note that some of the "thou shalt's" can be overridden: one can choose self-respect which requires feeling much pain, for example, if one is being tortured to reveal secrets).

>The environment can affect intrinsic variables, but it can't say what
>will constitute an intrinsic variable, or a reference level for one.

I think it can't redefine types of i.v.'s, but it can affect their settings.

>The environment may determine what must be done
>to affect a given perception in a given way, but it is the
>reorganizing system that decides what will constitute a perception,
>and whether any particular state of it is to be sought or avoided.

I disagree.

>With respect, at least, to the current environment, these built-in
>systems define the true autonomy of an individual.

I have no problem with this definition of "autonomy." Just don't try to say that current controlling is solely the product of past controlling.

>So autonomy, in any one lifetime, is awarded to the organism, in
>particular to its reorganizing system.

No. I'll only award a moment-by-moment autonomy to the individual.

>I won't follow that trail any further; the point is only to show that
>the same principles of autonomy can be extended into the past, with
>the environment providing the stage and acting mindlessly to disturb,
>but being incapable of carrying out the processes of change, which
>continue to reside in the species.

The NON-ORGANISMIC environment, you mean, right?

>Here and now, we live in a world whose properties are largely unknown
>to us and which determine for us the effects of our actions.

Yes.

>We, however, choose our own goal structures, as our means of preserving
>ourselves in the state that our natures tell us is right.

What is this "choose"? PCT says that there is NO choice in the present -- one simply controls as one's control system allows. Are we getting back to the "free-will" argument? No thanks!

>Between this ultimate personal autonomy and the impersonal events in the
>nonliving universe, there comes to be a hierarchy of control systems that
>reflects both our overriding inner needs and the conditions the environment
>places on meeting those needs.

A nice way of putting the (PCT) science underlying the (partial) reconciliation of extreme environmentalism and extreme organismism, which I have been calling for. However, I'm not sure you really mean it as I can take it.

>This still leaves open the question of interactions among organisms;
>organisms with similar organizations, and organisms assymmetrically related to
>each other.

So you DID mean the non-organismic environment ONLY, all those places above.
OK.

Best wishes, Greg

Date: Sun Sep 13, 1992 4:11 am PST
Subject: one or two texts

I am a generally non-contributing member of this list. This is because my area of specialization is language (ESL) teaching and my academic background is in linguistics (though also psycholinguistics). I am writing this to request suggestions on one or two texts. That is, the old question of if you could only read one or two books about ..., what would you choose.

Thanks in advance.

Eileen Prince
Northeastern University

Date: Sun Sep 13, 1992 7:48 am PST
Subject: RE: one or two texts

[From: Dennis Delprato (920913)]

Eileen Prince asked for two recommended books. I'll recommend one, J. R. Kantor's *Psychological Linguistics*, 1977, Principia Press. You will not find PCT explicitly presented here. What you will find is a naturalistic re-examination of linguistics free from environmentalistic and mentalistic approaches that underlie all alternatives to how PCT handles linguistic acts. Unless one has rather thoroughly "digested" system theory and cybernetics (and bracketed the fluff from both), one will likely find Kantor's linguistics a valuable lead into detailed PCT handling of linguistic acts.

Date: Sun Sep 13, 1992 10:36 am PST
Subject: One or two texts, Progress report

As I post this, I see Dennis response. What subject did you mean to ask about Eileen? Here is my reply.

[From Dag Forssell (920913-1)]

>Eileen Prince (920913)

> I am writing this to request suggestions on one or two texts.
>That is, the old question of if you could only read one or two
>books about ..., what would you choose.

Without question, the first text I recommend is FREEDOM FROM STRESS by Ed Ford. Ask him for it! Or ask me for my intro/demo package with Ed's book included. Either way, your snail mail address is required.

Perhaps Gary Cziko could re-post the introductory message (usenet intro), which includes a brief PCT book list with commentary. It's been a long time, Gary.

Progress report:

You all know Christine and I got our first booking. When the Director of Personnel and Engineering Manager called with instructions from the Chairman to sign us up, we agreed on a pilot presentation to half the engineering department, about 15 people. They know they are our first commercial customer. I immediately sent FREEDOM FROM STRESS in addition to an extra set of all my intro information to the personnel manager, who was designated our contact person. That was two weeks ago. Now, the two are issuing the following memorandum to the WHOLE group of 30 engineers, including a few marketing people.

MEMORANDUM

DATE: September 10, 1992
TO: Distribution
FROM: (Personnel Mgr) / (Engr Mgr)
SUBJECT: Purposeful Leadership training

Every once in a while something new comes out and you feel it is worth the risk to give it a try. In our opinion, Purposeful Leadership is such an opportunity. We have read the background literature that Dag Forssell, the provider, has supplied and his line of thinking parallels many of the concepts we have been trying to engrain at [Company].

You have been selected to participate in a "beta test" of this new approach to management thinking that offers some excellent possibilities.

The course covers three full days (three consecutive Wednesdays, ...).
It is held...

If this course is as valuable as we feel it may be, we will be following up your training with training for managers and supervisors.

The attached pamphlet will give you an idea of the course content.

DISTRIBUTION:

THEY HAVE UNDERSTOOD!!!!!! (And I believe they will help me because they now understand what is in it for them).

Now, it is just the small matter of delivery on the promises.

As many of you know, I have been preparing for this as an enthusiast for several years and full time for a year and a half. We will deliver.

All the PCT literature has been very helpful to me. I thank all the members of CSG and CSGnet for support and encouragement along the way.

PCT *will* come of age and gain acceptance, and soon!

Dag

Date: Sun Sep 13, 1992 1:31 pm PST
Subject: Re: One or two texts, Progress report

I'd be most interested in the demo package. My snail mail address is:

Eileen Prince
English Language Center
Northeastern University
206BY
Boston, MA 02115 Fax: 617-437-8929 Plain old phone: 617-437-3925

Thank you for your prompt response to my query.

Best, Eileen

Date: Sun Sep 13, 1992 2:47 pm PST
Subject: Congrats, Dag; manipulation

[From Bill Powers (920913.1300)] Dag Forssell (920913) --

CONGRATULATIONS. Your management course is the second great step forward for PCT in bringing it to public attention (the first being Ed Ford's forays into education). Long may you wave.

-

Greg Williams (920913) --

>>Thus the manipulator's action (positioning the instrumental variable)
>>must depend on the manipulee's choice of reference position for the
>>knot, and on the amount and direction of any third-party disturbances
>>of the knot.

>Exactly -- and this generalizes to more elaborate types of this kind of
>manipulation, as I've pointed out time and again.

No. It's too soon to generalize. Your own ideology is showing through. I may end up agreeing with you, but at the moment it looks as though you have a

conclusion all prepared, and are interpreting what I say only in ways that support that conclusion.

Control that works strictly through disturbance of perceptions, with the manipulee's actions countering that disturbance and causing no errors in the manipulee by being performed, is by definition unimportant to the manipulee. So the manipulator is doing things that are satisfying to the manipulator, but irrelevant to the manipulee. In order to find some effect of manipulation that IS important to the manipulee, we will have to find some other form of manipulation than just disturbing perceptions.

When I say "important to the manipulee," I DO NOT mean important in some objective sense, like performing the act will get the manipulee killed or fired or hurt. Those outcomes are completely irrelevant to the control processes in the manipulee if the manipulee does not know about them at the time of the manipulation. Many of your arguments invoke what is OBJECTIVELY important about the consequences of manipulation. That is ideology, just as you say it is when interactions are rejected because they result in conflict. The only importance that counts in a functional analysis of interactions is importance to the control systems involved. Third-party judgements have nothing to do with the scientific analysis.

You may or may not have noticed that when I speak about conflict, I always speak about it conditionally. That is, control of another must not significantly alter any variable that the other is controlling, if conflict is to be avoided. Conflict must be avoided if skillful control is to be maintained. Skillful control must be maintained if the organism is to satisfy its intrinsic reference levels. Intrinsic reference levels must be maintained if the organism is to survive. There is not a single moral or ethical judgment in any of those statements. When I speak of a malfunction, it is always in relation to the idea that proper functioning results in survival. What other definition of "proper" functioning could there be?

I admit to a value judgment when I say that I prefer, and work toward, survival of most human beings. I do, however, eat hamburgers.

>>So in the final analysis, the manipulator's hand position
>>(the instrumental variable) is not under the control of the
>>manipulator, and the manipulee's hand position is not under the
>>control of the manipulee.

>In the final analysis, neither is controlling for his/her perception of
>his/her own hand position. Is that what YOU just said?

Yes, but more prolixically.

>>The manipulee can control the manipulator's hand position by varying
>>the reference position of the knot; the manipulator can control the
>>manipulee's hand position by varying the reference position of the
>>hand.

>No, the manipulee doesn't give a whit where he/she sees the
>manipulator's hand, so he/she is NOT controlling it. If the manipulee
>started to control thusly, then the manipulation wouldn't work as the
>manipulator planned (on the basis of his/her previous knowledge of

the >manipulee's controlling); in this case, the manipulator would say >he/she had a "bad" model of the manipulee.

I said CAN. If the manipulee decides it is more important to control the manipulator's hand position than to maintain a fixed reference position for the knot, the manipulee can vary the reference position for the knot while watching the manipulator's hand, and put that hand anywhere the manipulee pleases. At the same time, the manipulator can keep the manipulee's hand where the manipulator wants it. Try it.

An example is the police riots at the Democratic convention in Chicago in 1968. The police were trying to manipulate the behavior of the protestors by a show of force. They disturbed, and the demonstrators pushed back and went right on protesting. But the demonstrators didn't want the police just to THREATEN force; they wanted them to USE force. So the demonstrators began running around creating incidents in different places which the police tried to suppress, until finally the police lost it and degenerated into brutality. Nominally, the protestors had the goal of defending their right to protest. But they manipulated HOW they were protecting that right, so that the police were maneuvered into controlling them by the means the protestors wanted the police to use right on camera.

>>Each one controls the other's means of achieving the same other's goal.

>>This is an example of slippery use of the word control. Are you doing >>this on purpose, or by innocent error?

On the contrary, this is a precise use of the term control. Each one can sense the other's action, and alter his/her own action in such a way as to make the other's action match the controller's reference level for it. This is accomplished without conflict because neither cares where his or her own hand is, so control can be as good as the parameters of the systems allow. Each one, by varying a reference signal, can make the other's action vary exactly as the reference signal varies. Extraneous disturbances of the other's actions will be counteracted. Why is this a slippery use of the word control? Did you actually test my claims with real rubber bands and players as I suggested?

>>Understanding autonomy is especially difficult in a hierarchical >>system, where combating a disturbance at one level entails altering >>goals at lower levels. Goal-directed behavior is not, per se, >>autonomous -- in a hierarchy.

>This can of worms is why I was ORIGINALLY trying to stick with >implications of PCT, not HPCT.

I now think this can't be done in talking about manipulation.

>>There are certain processes in the human organism that are carried out >>simply because of the way the organism is put together, inside. >>Control itself is an example, as is reorganization. The environment >>contains no means of carrying out these processes for an organism, >>either to help it or hinder it.

>

>The NON-ORGANISMIC environment, you mean, right?

I see I didn't manage to explain clearly what I meant. Organisms perceive. There is nothing the environment (animate or inanimate) can do to help or hinder the processes by which neurons respond to energy inputs by creating output signals; there is nothing any external agency or thing can do to aid or oppose the combining of such signals through neural computations to produce new perceptual signals. Not, that is, without surgical or chemical intervention. There is nothing the environment, live or dead, can do to help carry out the processes of perception, comparison, and acting.

Perceptual functions are created and changed by mechanisms that alter neural connections in perceptual functions. Nothing external to the organism has an influence on these mechanisms -- that is, on how they manage to alter connections, or even on exactly what connections they will alter, and how. All the external world can do is present sensory stimuli to the receptors. All else that takes place is done by processes working entirely inside the behaving system.

A manipulator can't alter the perceptions of the manipulee; there is no pathway to mediate such an alteration. All the manipulator can change is part of the external physical world in the vicinity of the manipulee, as the manipulator perceives it. Whatever happens after that is done by processes inside the manipulee.

>>A human being must acquire perceptual functions that produce
>>consistent perceptions, perceptions that vary in sensible ways and
>>relate to each other without contradiction. The environment can't do
>>that kind of making-sense; only a brain can.

The NON-ORGANISMIC environment, you mean, right?

No, the environment, organismic or not. Only the organism can find ways of making sense of its raw intensity signals. That process is an endogenous function, part of the physical and functional makeup of the organism. There is no way anything in the environment can suggest weightings, forms of functions, or anything of that sort that will be of the slightest help in creating a new perception or modifying an old one. The mechanisms that do these things operate independently of the outside world.

>Once a perceptual signal exists, it is the brain that must carry out
>comparisons with reference signals, to generate an error signal. The
>environment does not inform the human organism of how it, the
>environment, should be, nor does it tell the organism what
constitutes >a discrepancy with the organism's goals.

The NON-ORGANISMIC environment, you mean, right?

No, nothing in the environment can be of any help or hindrance in the process of comparing two neural signals, nor can it reach inside an intact organism and tell it what state of the environment should be indicated by the reference signal. Each part of the process of control arises from the functioning of the organism's nervous system and its mechanisms for internal change. The environment can alter input stimuli, but that is as far as its direct influence reaches. Nothing at all would happen beyond that point but for the internal processes of the organism.

>>Each level of control, therefore, comes into being through the action
>>of internal mechanisms for change and development, but the final
>>result, the control organizations that come into existence, must be
>>designed to work through the properties of the world that actually
>>exists, however we may perceive it.

>I suggest that the "coming into being" of levels of control depends on the
>actions of BOTH internal AND external mechanisms for change and development.

Levels of control come into being only when neurons become physically organized into new layers of control. The processes that do that can't be carried out by any mechanism outside the organism. You're talking about making it NECESSARY for such processes to occur. Making it necessary doesn't make it possible, not does it select what actual changes inside the organism will cope with the necessities imposed by the external world. All the mechanisms that change an organism are inside it.

>>Finally there are the intrinsic variables, their inherited reference
>>levels, and the process of reorganization driven by intrinsic error.
>>These are defined for the organism by its heritage.

>A heritage, I suggest, both hereditary AND environmental (think, for
>example, of physiological set-points changing when you move to a high
>altitude; and note that some of the "thou shalt's" can be overridden:
>one can choose self-respect which requires feeling much pain, for
>example, if one is being tortured to reveal secrets).

Once it is decided by DNA that blood volume will be a controlled variable, nothing in the lifetime of the organism can dictate that it NOT be a controlled variable. A true intrinsic variable is defined strictly by heredity -- defined, I said, not controlled. There is a built-in sensor and a built-in reference signal for oxygen tension in the blood. You can move to any altitude you like, and oxygen tension will still be an intrinsic variable controlled relative to a genetically-fixed reference level. Clearly, neither breathing rate/volume nor number of red blood cells is an intrinsic variable, because both change when you move to a new altitude, as a way of maintaining the required oxygen tension. They change BECAUSE there is an error in oxygen tension, and nothing outside the organism and nothing the organism does can make that error into anything other than an error.

When an organism feels pain, it begins reorganizing. It has no choice in the matter, even if it has itself created the circumstances that led to the pain. It will keep reorganizing until the pain disappears, unless some other intrinsic error shows up. That is why torture sometimes works, and dieting doesn't. Torture doesn't always work because its results are random, and there may be no way to reorganize that will stop the pain (if you're being asked for information that you don't know and can't fake). I'll bet that even J. Gordon Liddy doesn't find many occasions, any more, to show how brave he is by holding his hand over a candle flame. Reorganization will conquer even right-wing fanaticism, given time. Although it won't necessarily eliminate the most obvious intrinsic error. Look how JGL turned out after all that self-administered intrinsic error.

>>The environment can affect intrinsic variables, but it can't say what
>>will constitute an intrinsic variable, or a reference level for one.

>I think it can't redefine types of i.v.'s, but it can affect their settings.

If it can alter the reference settings, then we aren't talking about intrinsic reference levels. See control of oxygen tension above. This is a matter of definition. An intrinsic variable is by definition one that is specified as important by DNA. Its reference level is by definition fixed by DNA. There are lots of biochemical variables that aren't in this category. Perhaps you would like to argue that there are no true intrinsic variable and reference levels.

>>The environment may determine what must be done
>>to affect a given perception in a given way, but it is the
>>reorganizing system that decides what will constitute a perception,
>>and whether any particular state of it is to be sought or avoided.

>I disagree.

Do you claim that it is the environment that specifies what function of a set of input signals will be computed? Are you saying that the organism does not have to discover FOR ITSELF what settings of reference signals will result in consequences that it prefers?

>>With respect, at least, to the current environment, these built-in
>>systems define the true autonomy of an individual.

>I have no problem with this definition of "autonomy." Just don't try to
>say that current controlling is solely the product of past controlling.

I don't even know what that means. I've tried to explain how a reorganizing system has the ultimate say over what inner states will and will not be preferred by the organism, and therefore over what will and will not become controlled by the organism. It is not the environment that says pain hurts and pleasure feels good. The environment knows nothing about what is good for the organism, at the level of reorganization.

>>So autonomy, in any one lifetime, is awarded to the organism, in
>>particular to its reorganizing system.

>No. I'll only award a moment-by-moment autonomy to the individual.

But that's precisely what I have been giving up. If disturbances control your actions, and if even reference levels can be controlled from outside by disturbing higher-level perceptions, where is there any moment-by-moment autonomy of action OR of goal-selection? The only place left is at the highest level, or beyond that in the intrinsic reference levels which BY DEFINITION are genetically set.

>>I won't follow that trail any further; the point is only to show that
>>the same principles of autonomy can be extended into the past, with
>>the environment providing the stage and acting mindlessly to disturb,
>>but being incapable of carrying out the processes of change, which
>>continue to reside in the species.

>The NON-ORGANISMIC environment, you mean, right?

Wrong. Even the organismic part can influence the organism only through altering the non-organismic environment in the vicinity of an organism. And the signals coming in through the sensors do NOT provide the machinery that carries out behavior or changes its organization. All changes inside the organism require a judgment by the organism, and an internal action by it. Nothing outside the organism is in a position to do those things.

>>We, however, choose our own goal structures, as our means of
>>preserving ourselves in the state that our natures tell us is right.

>What is this "choose"? PCT says that there is NO choice in the present
>-- one simply controls as one's control system allows. Are we getting
>back to the "free-will" argument? No thanks!

I mean that it is a mechanism inside us that varies our goal structures and stops varying them when the result is pleasing to that little man who lives in our pineal glands.

There's no argument about free will. It exists, as anyone can verify from direct experience. You can voluntarily and consciously step into the operation of many of your systems at any level and cause arbitrary changes in their operation (the result is usually not great). This phenomenon is, I agree, difficult to account for within any current theoretical framework. But it occurs. I, too, say "no thanks" to any explanation of it that anyone has offered so far. Including theories that say it doesn't exist.

>>This still leaves open the question of interactions among organisms;
>>organisms with similar organizations, and organisms assymmetrically
>>related to each other.

>So you DID mean the non-organismic environment ONLY, all those places
>above.
>OK.

No, I didn't. I meant that we have yet to analyze interactions among persons, beyond the simple interaction of controlling actions by disturbance of perceptions.

RE: an earlier comment:

>Exactly my sentiments! May that progressiveness not be harmed by
>ideological "autonomism"!!!

I see that ideological autonomism is a disturbance to some sentiment, and contrary to progress in some direction. Would you care to elaborate on what that sentiment is, other than by giving everything an interpretation that supports it? I see lots of arrows pointing in some direction, but I don't yet see what's there at the convergence.

What if ideological autonomism turns out to be the right approach? If so, so what?

Best, Bill P.

Date: Sun Sep 13, 1992 4:53 pm PST
From: Ken Mann / MCI ID: 507-3685

TO: * Dag Forssell / MCI ID: 474-2580
Subject: MCI Mail News 9/92

Dear MCI Mail Subscriber:

If you'd like more information on any of the below listed topics, just send a message to my Agency's toll-free mailbox 501-6408 and I'll have it sent to you.

If you'd like to be taken off this list, just let me know.

- NEEDANOTHER MAILBOX?
AUTOFORWARD
CREATE
CANCEL
DELETE
AUTOFORWARD FOR REMS ACCOUNTS
TRY OUT DOW JONES FOR FREE!
FREQUENTLY ASKED QUESTIONS
Faxing
Lotus Express
X.400/Interconnections
Communications
LANS/Gateways
Creating Messages
Security

Regards: Ken Mann CAI Agency

Date: Sun Sep 13, 1992 7:58 pm PST
Subject: Re: Turing

(ps 920913.2000)

[From Rick Marken (920910.0900)]

There was also an article in comp.robotics that dealt with the Turing Test. There is apparently an article on the subject by S. Harnad

I don't know if I'll waste my time reading it (NB-- Penni Sibun; I judge that I will learn nothing from it because I have read Harnad before. Might

that's ok; i don't usually read you since i know what you're going to say. however, you did pique my curiosity by mentioning the turing test--even though i knew what you were going to say ;-).

i think the fact that it still more or less bears in mind the turing test is the only thing that really keeps ai coherent and intellectually respectable.

be a mistake but I'll take the risk). But I do think that the Turing Test is an EXCELLENT example of the behavioristic basis of AI (and cognitive science) etc. Harnad's contribution makes it even clearer -- he suggests a Total Turning Test meaning its not enough to get a simulation to answer questions like a real person -- you must also get it to behave in all ways like a real person -- ie -- brush teeth, play soccer, build model airplanes, etc.

have you read turing's article (recently)? (i mean ``computing machinery and intelligence,' ' _mind_ LXI:236, Oct. 1950, pp433-460.) the question turing considers is ``can machines think?' ' the method he proposes to answer it is to play a game in which an interrogator communicates via teletype w/ both a human and a machine and has to decide which is which. if the interrogator can't decide, or decides wrong, this suggests that the machine thinks (or the human doesn't). who better to judge a thinker than a thinker? what else to compare a potential thinker to than a thinker?

turing based his test on what was apparently a sort of parlor game in which the interrogator has to decide which is male and which is female. maybe the analogy will be helpful. how do we know if someone is male or female? that's easy, right? we're scientists! we just check the primary sexual characteristic: a male mammal has an x chromosome and a y chromosome; a female has two xs. (of course, not always: some people are xyy, some are xxy, and maybe other interesting combinations.) what? you've never checked anyone's chromosomes? then how do you *know*?? well, there are the secondary sexual characteristics, having to do w/ shape of genitalia, distribution of body hair, subcutaneous fat, etc.. but what if you can't see enough of these to make a judgement or it would be rude to ask? what if in this particular person they're ambiguous, or the person is actually a hermaphrodite? what if the secondary sexual characteristics are convincing, but the person is a transsexual, and thus has 2ary characteristics that don't match the primary ones? gosh, it's a toughie, isn't it? WE DON'T KNOW WHAT ANYBODY'S SEX IS!!

this hasn't been keeping you up at nights? why not? you say you know the sex of everyone you deal with? even though you don't know their chromosomes? howzat?? you mean you base it on their behavior--how they dress and walk and talk and name themselves and interact with you? that's pretty dicey don't you think? here everyone is walking around and applying the sex test--based on ``superficial' ' characteristics, w/o looking inside for the true facts--and sometimes they are WRONG--if you go by the facts of the chromosomes.

as it happens, we don't go around worrying about our ability to tell what sex others are. occasionally we make mistakes and are embarassed. but we usually decide w/ complete confidence that someone is female or male. it is this *judgement* of someone's sex, based on our perceptions of them, that constitutes what is female or male. our judgements are surely highly correlated w/ what a check of the chromosomes would tell us, but the

judgements aren't determined by the chromosomes, and they don't always match. but i think most people would trust their judgements, w/o giving it a thought.

if *sex* is decided in such a ``superficial'' way, then why not the ability to think?

cheers. --penni

Date: Mon Sep 14, 1992 4:04 am PST
Subject: More science, more ideology

From Greg Williams (920914)

>Bill Powers (920913.1300)

>Control that works strictly through disturbance of perceptions, with
>the manipulee's actions countering that disturbance and causing no
>errors in the manipulee by being performed, is by definition
>unimportant to the manipulee. So the manipulator is doing things that
>are satisfying to the manipulator, but irrelevant to the manipulee.

"Irrelevant to my controlling (now)" might or might NOT be "important" to me
LATER.

>In order to find some effect of manipulation that IS important to the
>manipulee, we will have to find some other form of manipulation than just
>disturbing perceptions.

Why should we be needing to do this? I can grant that, in your sense, ALL workable purposeful influence MUST be irrelevant to the influencee's controlling -- except for the possibility of setting the stage for reorganization -- and still preserve MY point that successful purposeful influence is often considered important AFTERWARDS by the influencee. You don't think this is an important point? However, that said, I still think that one could claim "relevance" (differently defined than yours) to the influencee because the DETAILS of what the influencee is controlling (and therefore is "relevant" to the influencee) depend in part on the influencer. If that mark had never met that con man, the mark wouldn't have (with high loop gain!) wagered on that "sure-thing" horse.

>When I say "important to the manipulee," I DO NOT mean important in
>some objective sense, like performing the act will get the manipulee
>killed or fired or hurt.

Neither do I -- I mean how the influencee judges the outcome for him/herself. You can't get much more subjective than that! As I said in my post, some folks have preferred death to dishonor.

>The only importance that counts in a functional analysis of interactions is
>importance to the control systems involved.

With all due respect, this is beginning to sound positively Skinnerian. Nevertheless, I am willing to proceed along the "functional analysis" route if

you think it best. I just think that sometimes it makes sense (contra Skinner, and maybe contra you) to ask the parties what THEY think.

>Third-party judgements have nothing to do with the scientific analysis.

I'll let it go, but I still think that if the "third party" is the influencee, it makes sense to consider his/her notions when conducting an analysis. Again, your tack has, to my thinking, a Skinnerian restrictiveness which I think might compromise the fruits of a more open scientific analysis.

>When I speak of a malfunction, it is always in relation to
>the idea that proper functioning results in survival.

When I speak of a "manipulation," I am talking about purposive influence. Neither "malfunction," as you use it, nor "manipulation," as I use it, exactly fit all (or major?) dictionary definitions of the words. Tell you what... I'll not complain about your use of "malfunction" if you'll accept my use of "manipulation."

>What other definition of "proper" functioning could there be?

"Proper," of course, also has lay connotations of "good."

BP>>>So in the final analysis, the manipulator's hand position
BP>>>(the instrumental variable) is not under the control of the
BP>>>manipulator, and the manipulee's hand position is not under the
BP>>>control of the manipulee.

GW>>In the final analysis, neither is controlling for his/her perception of
GW>>his/her own hand position. Is that what YOU just said?

BP>Yes, but more prolixically.

Good. I agree. The manipulator is controlling his/her perception of the manipulee's hand position, and the manipulee is controlling his/her perception of the position of the knot.

BP>>>The manipulee can control the manipulator's hand position by varying
BP>>>the reference position of the knot; the manipulator can control the
BP>>>manipulee's hand position by varying the reference position of the
BP>>>hand.

GW>>No, the manipulee doesn't give a whit where he/she sees the
GW>>manipulator's hand, so he/she is NOT controlling it. If the manipulee
GW>>started to control thusly, then the manipulation wouldn't work as the
GW>>manipulator planned (on the basis of his/her previous knowledge of the
GW>>manipulee's controlling); in this case, the manipulator would say
GW>>he/she had a "bad" model of the manipulee.

>I said CAN. If the manipulee decides it is more important to control
>the manipulator's hand position than to maintain a fixed reference
>position for the knot, the manipulee can vary the reference position
>for the knot while watching the manipulator's hand, and put that hand
>anywhere the manipulee pleases. At the same time, the manipulator can
>keep the manipulee's hand where the manipulator wants it. Try it.

OK. Sorry. Now I understand that you are noting the possibility of a SPECIAL CASE (involving mutual influence) of rubber-banding. I had thought you were trying to say that this is what can happen in "standard" rubber-banding.

Here we have a simple example of the (sometimes complex) "dance" of (some) manipulations which I spoke of in earlier posts. As I said then, it is perfectly possible for the mark to be "in charge" of the con at times -- but the con attempts to stay "in charge overall." If it suited the con's purposes, he/she would ENCOURAGE the mark to (in analog) "control the manipulator's hand position" for a while -- one example is "letting" the mark win the first few games of cards to "build the mark's confidence," before cleaning out the mark on a high-stakes hand.

>Why is this a slippery use of the word control?

It isn't. My mistake. I apologize.

>Organisms perceive. There is nothing the environment (animate or inanimate)
>can do to help or hinder the processes by which neurons respond to energy
>inputs by creating output signals; there is nothing any external agency or
>thing can do to aid or oppose the combining of such signals through neural
>computations to produce new perceptual signals. Not, that is, without
surgical
>or chemical intervention.

Right, but beside the point.

>There is nothing the environmentalive or dead, can do to help carry out the
>processes of perception, comparison, and acting.

Here's why it is beside the point. I am trying to argue that one's current controlling is the result of one's history. That history includes BOTH internal and external influences. While it is true that the external environment (living or dead) does NOT "aid or oppose the combining of... signals through neural computations to produce perceptual signals," I claim that the one's external environment significantly influences how one will control in the future. An example is Gary's: his children wouldn't be speaking in French today, but for him (as necessary but not sufficient condition).

>No, the environment, organismic or not. Only the organism can find
>ways of making sense of its raw intensity signals. That process is an
>endogenous function, part of the physical and functional makeup of the
>organism. There is no way anything in the environment can suggest
>weightings, forms of functions, or anything of that sort that will be
>of the slightest help in creating a new perception or modifying an old
>one.

Again, I think this overstates the case. But as long as you aren't saying that one's environment now CANNOT affect the "trajectory" of one's controlling in ANY way, I won't argue.

>The mechanisms that do these things operate independently of the outside
>world.

So, are you saying that "raw intensity signals" can NEVER affect the trajectory of control? If so, I think you need to go back to diff. eqs. 101 (well, 201 -- the equations themselves are time-varying) -- all we need is ANY sort of connection between the equations' (of control) parameters and inputs (intensity inputs will due nicely) to the equations to have a causal relationship between the equations' structure and the environment.

>No, nothing in the environment can be of any help or hindrance in the
>process of comparing two neural signals, nor can it reach inside an
>intact organism and tell it what state of the environment should be
>indicated by the reference signal.

If we're NOT talking about (presumed purposeful) "help or hindrance," but CAUSAL relationships, I can say explicitly that I claim causal efficacy between inputs to an organism and the trajectory of control of that organism.

>The environment can alter input stimuli, but that is as far as its direct
>influence reaches. Nothing at all would happen beyond that point but for the
>internal processes of the organism.

OK, we agree on that. The important point that you seem to be obscuring is that those "internal processes" (now) are the products of BOTH previous internal processes AND inputs from the environment. Do you disagree with this?

>Levels of control come into being only when neurons become physically
>organized into new layers of control. The processes that do that can't
>be carried out by any mechanism outside the organism. You're talking
>about making it NECESSARY for such processes to occur. Making it
>necessary doesn't make it possible, not does it select what actual
>changes inside the organism will cope with the necessities imposed by
>the external world. All the mechanisms that change an organism are
>inside it.

All the mechanisms that PURPOSIVELY change an organism are either inside the organism OR in other organisms. They are NOT in the non-living environment -- I agree with that much. The environment (non-living OR living) does NOT determine EXACTLY what the trajectory of control will be. But it influences the trajectory.

>That is why torture sometimes works, and dieting doesn't.

Dieting worked for me. (:-->)

>Do you claim that it is the environment that specifies what function
>of a set of input signals will be computed? Are you saying that the
>organism does not have to discover FOR ITSELF what settings of
>reference signals will result in consequences that it prefers?

No.

>I've tried to explain how a reorganizing system has the ultimate say over
>what inner states will and will not be preferred by the organism, and
>therefore over what will and will not become controlled by the organism.

And I've tried to explain how environmental inputs have SOME say in reorganization.

>But that's precisely what I have been giving up. If disturbances
>control your actions, and if even reference levels can be controlled
>from outside by disturbing higher-level perceptions, where is there
>any moment-by-moment autonomy of action OR of goal-selection?

There is none, according to PCT. But if you take "autonomy" to mean unconflicted current operation of one's control system, then we ARE autonomous (unless conflicted) in our moment-to-moment behavior. If you take "autonomy" to mean that our current controlling has no relationship to any of our past environmental inputs, then I claim that we are NOT autonomous (in this "historical" sense). So that is what I mean by reconciling Skinner's ideology of "historical" non-autonomy with an ideology of "current" autonomy. PCT says that we are autonomous in the current moment (if not conflicted) AND, I think, NOT autonomous vis-a-vis our histories. PCT shows that Skinner's extreme environmental determinism is wrong on two counts: historically, the environment is not the SOLE influence on the trajectory of controlling; currently, environmental inputs are highly circumscribed and subject to current controlling.

>There's no argument about free will. It exists, as anyone can verify
>from direct experience. You can voluntarily and consciously step into
>the operation of many of your systems at any level and cause arbitrary
>changes in their operation (the result is usually not great). This
>phenomenon is, I agree, difficult to account for within any current
>theoretical framework. But it occurs. I, too, say "no thanks" to any
>explanation of it that anyone has offered so far. Including theories
>that say it doesn't exist.

(Against my better judgement -- whatever that means!?!?) Unfortunately, there ARE (raging!) arguments about "free will." Especially the "free" part. Will certainly occurs, I agree.

>I see that ideological autonomism is a disturbance to some sentiment,
>and contrary to progress in some direction. Would you care to
>elaborate on what that sentiment is, other than by giving everything
>an interpretation that supports it? I see lots of arrows pointing in
>some direction, but I don't yet see what's there at the convergence.

See my previous post where I attempted to spell out my ideology with regard to promotion of PCT, and the above comment marked with ***.

>What if ideological autonomism turns out to be the right approach? If so,
>so what?

I think it conflicts with PCT science. But nobody ever said that ideologies had to jibe with science. There is something to be said for giving the masses what they want, even if it is at odds with what you think scientifically. Even the Buddha opined that individuals should be taught in ways they can appreciate, given where they are on the karmic road. In the case of PCT, vague claims that it is the basis for anti-environmentalism might attract quite a few folks. Eric Hoffer called them "True Believers."

Best, Greg

Date: Mon Sep 14, 1992 5:42 am PST
Subject: stella for pct, cont'd

from: eric harnden (920914.0900)

i plugged the two level system suggested by mr powers into stella over the weekend, and can report that it does just fine. s and t converge to s* and t*. i need to bring binhex home so i can send a copy of it to mr cziko, as a modeling example. on the subject of simultaneous equations...

there is a slight terminology difference here (as usual, it seems). stella is perfectly capable of being set up to model solutions to systems of equations such as those provided by mr powers. the compiler effectively treats the existence of a level variable as a delay of one dt interval. variables dependent on level state are calculated first, and then the level (which is dependent on those values) is calculated for the end of the timestep. what is disallowed (and i believe is the source of mr cziko's error) is the following:

```

source/sink<-----valve----->level-----
                /\
                                     \/
                                ----->v1
                                     \/
                                rate<-----v3<-----v2

```

where v1,v2, and v3 are mutually dependent variables which must be calculated at the same time. in that loop there must be a delay of some kind (preferably a level).

one thing i can suggest to anyone trying to model with stella is to at least obtain "introduction to system dynamics modeling with dynamo", by george richardson and alexander pugh. the stella docs do not go into the principles very deeply, and getting a look at the underlying assumptions of what constitutes a 'good' model is very worthwhile. even though the book is written as a dynamo text, it is immediately applicable to stella (since stella basically just compresses two steps of dynamo programming into one, at the cost of some flexibility).

Eric Harnden (Ronin)

Date: Mon Sep 14, 1992 5:49 am PST
From: Uwe.Schnepf MBX: Uwe.Schnepf@gmd.de

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Phenomenon, Demo.

Many many thanks. Uwe

Date: Mon Sep 14, 1992 7:33 am PST
Subject: Re: Turing

From Tom Bourbon [920914]

Penni Subin [920913] gave what I think is a very nice reply to Rick Marken's [920910] post concerning the Turing Test. The example of judging another person's sex (gender) is elegant.

In nearly all of my recent presentations and demonstrations of PCT tasks in which people, or hands, or models, or combinations of those interact, I have opted to show the model interacting with a person, to make the point that you can scarcely tell the difference between that condition and two people interacting. That seems to me to be precisely the sort of test-demonstration Turing had in mind in his 1950 article. I have come to think of such demonstrations as The Second Test (along with the test for the controlled variable) that PCT can off as an alternative to "business as usual" in the behavioral, life and cognitive sciences. AI people (at least some of them) do still accept a Turing Test as evidence concerning a model or theory. That gives us a great opening to show the wide range of phenomena in which the PCT model passes the test.

Date: Mon Sep 14, 1992 10:33 am PST
Subject: Imagination construction!

[From: Chris Love (920914.1300)]
[To: Bill Powers (920910) and Rick Marken (920910)]

Thanks for the encouraging posts!

Bill, I have already started working on this *new* model. And I don't know either when it will be complete, but I can't see it taking too long (< 1 month) since all the *control objects* have already been created, such as linking objects, connection routines, the basic ECS object, etc.. But in development work, as I'm sure your aware this is never a sure thing!

>Perhaps what you feel is missing is the pairing up. In fact, it's still there...

I see Bill. This really does appear that, for the purpose of that n-level ECS, it is imagining without even knowing it. Kind of like drifting in the middle of a lecture - it's not until you bring yourself back to attention (reality) that you realize you ever left! I say this because that n-level ECS is being provided all information in a regular manner as if nothing unusual was happening to it, although the information (percepts) are really imagined ones.

The problem: By allowing those n-1 level ECS' to provide the imagined percepts to your (focus) n-level ECS you are getting (undesired?) influence from neighboring n-level ECS'?

How? This is occurring because those n-1 level ECS' are also receiving outputs from the neighboring n-level ECS' in the process of generating the net input reference sum.

Do you want this? If you wish to evaluate the effectiveness of that n-level ECS, I think not since you now have external uncontrollable influence from neighboring n-level ECS' through an indirect path. If this was your intention then I do not understand the reason, or possibly overlooking something. Please explain.

In my proposal, you short circuit the ECS that you want in imagination mode. No influence from neighboring n-level ECS' occurs. I know it does in the *global* sense through the environment **eventually** as received references, but this is far enough removed, I feel, to be of little effect. When that n-level ECS imagines you want it to try to perceive doing well, that's the objective, right? So, in achieving, or attempting to achieve this objective we must vary it's direct weights in relation to it's imagined percepts such that it reduces its error. This results, of course, in a better controller. This is my hypothesis.

Bill, if you removed the influence of those other n-level ECS' contributions to those other n-1 level ECS reference sums then I think we have agreement on the imagination objectives.

>I'll be interested in your implementation of an imagination loop that can adjust perceptual weights...

Me too! I'm open to any suggestions. My baseline proposal is to use some form of Hebbian style learning (also known as the Widrow-hoff rule). As applied to PCT it goes something as follows:

Delta weight = learning rate * connection wt. * error signal
New wt. = Old wt. + Delta wt.

Comments:

The learning rate is typically set to a small value between 0.1 and 0.3.
The connection weight is the existing weight (wt.) of the connection of interest.

The error signal is that signal emanating from the comparator.

What does this do?

Well, as I see it it correlates particular references with particular percept signals. In wording the Hebbian rule says something like, "... if two connecting units are correlated then strengthen that connection." This lowers the threshold necessary to *fire* that receiving unit in the future under similar circumstances.

Is this the right thing to do? That depends on whether there is correlation between the references and the percepts? And whether this correlation is desirable to learn, i.e., will it make controlling better?

If anyone has any comments, improvements, either to prove or disprove the above proposal or anything else - great! Let's have at 'em!

Rick Marken-

>Whatever throws the switch must have reason for doing it...
>the reason is to control something.

>... it would be nice to have a system that alters the imagination
connections through the perceptual and output system...

>the imagination switch is probably the output of a control system..

Sounds good to me too Rick.

I also thought that another possibility would be to allow the ECS to determine that it needs to imagine. This would happen, like me, when I work too hard and get (temporarily) *brain dead* and my mind automatically switches to imagination mode to try to solve *the problem*, I guess subconsciously.

When I say "allow the ECS to determine", I should specify that the *switch* is internal rather than coming externally from another ECS. I would like to avoid this path of exploration first if I can since it requires, I feel, a dedication to function, i.e., that ECS' objective is to *only* know when to send a *switch to imagination* signal to other ECS'. I, as a designer, now have to worry about how to make such an ECS do just that. I would prefer to develop, if I can, a *almighty* ECS that is VERY GENERIC in the sense that everything that must happen in an ECS or to an ECS can be handled *EXCLUSIVELY* by that ECS. It does not require outside influence to make decisions on what its actions should be.

As an engineer, if it turns out that your proposal works ok, great I'll implement it tomorrow. I'm in no way biased to either approach - just one that works and is reliable, reproduceable, and explainable (is that a word?)!!! I think at this point, however, that the almighty ECS may be the way to go. What do you think about this?

>figure out how to do experiments that would guide the modelling of
imagination...

I can't imagine these expts. now either but you can't always see the end of the tunnel; you just have to believe that it will appear before you drop!

Thanks, Chris. (DCIEM)

Date: Mon Sep 14, 1992 10:34 am PST
Subject: turing test; autonomy & reorganization

[From Bill Powers (920914.0900)]

Penni Sibun (920913.2000) --

That was a delightful essay on the Turing test for sex. I think you have a great sense of humor to go with a sharp mind.

I would add another requirement to the Turing test, although it doesn't alter the basic criterion of telling the machine from the real person. The Turing test is basically a test for humanity at the level of symbolic interaction. One of the kinds of interesting tests that an AIer is not likely to think of applying is a test for control. A human being, in a conversation, has (may have) a point he or she is trying to make. In that case, when one person fails to understand, fails to agree, shows that the point has been missed, and so on, the other person will try different methods of correcting the perceived error -- use alternative arguments, paraphrases, illustrations, and so on. So I would look for the thing at the other teletype to do this sort of error-correcting communication. If it says "I am a mechanical engineer," and I say "who made you?" I would look for it to correct my mistaken understanding of "mechanical" instead of telling me who its parents are.

With communication limited to a teletype, we can only test the other entity at the higher levels of organization. At the strictly rule-driven level, I doubt whether we could distinguish a program from a person if the program were complex enough. But what would happen when we try to see the principles or system concepts that the other "person" is working from? Is it possible, for example, to write a program that follows principles of humility, arrogance, caution, and so on? I don't doubt that a program could make use of vocabularies constructed from the speech of people judged to exhibit these principles, but could a program in itself generate any new principle? And if it could, would its principles be indistinguishable from the ones human beings adopt and control for? Could a computer program think up something like the Turing test? Could a computer program be devised to carry out the Turing test with another entity?

Would you be surprised to learn that I am really an 80486/33 with 4 Megs of memory?

-

Greg Williams (920914) --

It may not seem so, but I think we're approaching some sort of understanding. You're forcing me, at least, to clarify my thinking.

I have still not managed to get across the point I'm making about abilities that lie (autonomously) inside an organism. Actually I've had the same failure in getting this point across to S-R types. The question I'm trying to answer isn't "what effects on the behavior of another person can you (empirically) demonstrate?" but "on what internal properties and processes in the other person are you depending in order to have those effects?"

This question arises when a radical behaviorist talks about the effects of reinforcement. It can be shown experimentally that making reinforcement contingent on behavior according to a schedule of reinforcement results (after a little trial and error and rejecting some potential test-organisms) in predictable patterns of behavior in almost all test organisms. So the behaviorist says, "See? I have complete control of this animal's behavior. It is obviously the contingency of reinforcement that is causing it to behave as it does, and that contingency is under my control."

The question I would want to ask in return is "But doesn't this apparent control depend on properties and processes inside the organism over which you

have no control?" I would really want to ask, "What if the organism doesn't want the reinforcer you offer?" But that wouldn't go across big with a behaviorist.

Even under the S-R postulates, the experimenter is depending on many autonomous functions inside the organism, starting with sensing and recognizing the reinforcer. The internal connections from the stimulus to the muscles then have to be established. Then the connections have to be adjusted so that when the muscles act, they have an effect on the environment (forget about disturbances) that produces a pattern that either an apparatus or the experimenter recognizes as the one that merits a reinforcement. If these things inside the organism weren't present, and didn't somehow become organized in just the right way, the contingency would have no effect at all, or at least it would have an effect quite different from the expected one. And there is NOTHING the experimenter can do to provide the organism with any of these internal elements, should any of them prove to be missing. How do you control an animal that has no sensory organs? The only way is to pick it up and run it across the floor like a toy car. Of course such an organism couldn't live, but this is a thought experiment.

In a hierarchy of control that accretes through reorganization, this situation is compounded in complexity. The external agent can produce expected behaviors from this hierarchy, but only if they don't violate the conditions that the reorganizing system must, by its nature, maintain.

On the surface, the external agent is dealing only with a single control system, the one whose actions the agent is trying to control. But as soon as this control has any effect that significantly disturbs the control hierarchy at the same or a higher level, either direct resistance will develop or the reference signals in the original system will begin to shift. So in order to retain control in general, the external agent has to give up control of that action and pick a higher-level aspect of behavior to control.

It's inevitable that the general desire to control another must lead to applying disturbances at higher and higher levels, which of course becomes more and more difficult. And in the background, at any level of interaction, the reorganizing system sits watching for any consequences of these interactions, known to the hierarchy or not, that have adverse effects on the fundamental variables that signify the viability of the organism. As soon as any such violation takes place, the reorganizing system will come into action. Ideally it will act to prevent exactly the manipulation that is causing the problem, but more probably it will begin trying new organizations at random. Either way, the external agent must start over because now its map is obsolete.

One way to think of the hierarchy of control is as a series of layers of defense erected by the reorganizing system. When a systematic control system is acquired that has, as a side-effect, the ability to prevent changes in the environment that cause intrinsic error, the organism has created a defense, over a limited range, against external events that previously caused intrinsic error. Unless some marked change in the properties of the environment occurs, those events are now controlled by the behavioral hierarchy and prevented from causing intrinsic error again. As the hierarchy builds up, the number of events of a kind still capable of creating intrinsic error must decrease; the effects of the events must become more subtle and indirect. Control behavior

becomes more and more sophisticated as the levels are added, to the point where in an adult there are few things that can upset the milieu interieur. Either the adult organism directly cancels such effects or it knows how to avoid exposing itself to them.

I define the reorganizing system in a particular way. Its input function is genetically designed to monitor some set of critical variables (whose nature is still undetermined). For each such variable there is a genetically-given reference level. The reorganizing system contains a comparator that detects the sum of the absolute values of all these errors, resulting in an intrinsic error signal. The intrinsic error signal drives the processes of change that unsystematically alter the organization of the hierarchy. Obviously this is just the outer envelope of a system that may have much more detail in it that I can provide. The effects of reorganization may m^uend to biochemical systems, too, such as the immune system. I don't know.

The main point is that if there is a reorganizing system of the kind I propose, then it is autonomous with respect to the events of a single lifetime. Its criteria for acting are given in its set of intrinsic reference levels, which are fixed for life and independent of particular experiences and particular organizations of the behavioral hierarchy.

It is this reorganizing system that sets the limits on what an external agent can get the organism to do or not do. The entire hierarchy is constructed around the unchangeable requirement that intrinsic error shall be held near zero. What constitutes intrinsic error is defined by DNA, so it represents the judgment of the entire species, acquired over geological time, as to what is good and bad for any individual of this species. This judgment is unaffected by the events of a single lifetime.

I have postulated, as you know, that awareness is somehow connected with the reorganizing system. Apparently, awareness can adopt the point of view of any level of perception in the hierarchy, so that the conscious world consists of the perceptual signals in that level and is shaped by the perceptual interpretations that give rise to those signals. This is the basis of the Method of Levels, the HPCT design for psychotherapy.

Apparently, also, awareness can direct the locus of reorganization -- that is, reorganization is concentrated on the systems with which awareness is most closely identified at a given time, at least in the neural-behavioral hierarchy.

This is my theory of "free will," which defines both "free" and (conscious) "will." I do not know the nature of awareness, nor does anyone else. But anyone who or anything that lacks it clearly gives the impression of "nobody home." This is a working hypothesis, of course, nothing more. But it does seem to have noticeable explanatory power.

The point in bringing this up is that nobody can have any freer a will than this -- the ability to change point of view, and to direct reorganization to a place in one's own system, with the same apparent randomness that we attribute to reorganization itself.

If this is true, then what we call manipulation of one person by another is simply the interaction of two systems, each attempting to reorganize itself to suit its own internal definition of what is good or bad for it. We do NOT have the situation where one all-knowing omnipotent system of unlimited intellect is manipulating a dimly conscious and defenseless system of some inferior kind -- that is the situation only in the relation between a human being and an emphatically lower order of life. It is not the situation that holds between normal adult human beings. The manipulator's "will to manipulate" is not different in kind or effect from the manipulee's "will to defend."

The term "manipulation" is unfortunate, because it implies power and will on the side of the manipulator, and weakness and automaticity on the part of the manipulated: exploiter and exploited, controller and controlled, agent and effect. This implies that the manipulator has some abilities or powers that are lacking in the manipulee, an awareness that is denied to the manipulee, a natural advantage over the manipulee. This is simply not the case in the HPCT concept of behavior. All human beings have the same basic abilities, save for those that are crippled in some striking way. Whatever conceptions of social life we may have developed, in truth we all have the same equipment, which is designed to allow us to control what happens to ourselves -- and especially what happens to our ability to continue living.

Manipulation is just a one-sided way of looking at interaction.

Eric Harnden (920914.0900) --

Thanks much for trying out that two-system control thing. I am relieved that it turned out as I predicted. I think this settles the question as to whether Stella can be used for significant PCT modeling.

Can you tell me a bit about Stella's ability to create time plots of the variables in a simulation? Is there any way for Stella to accept real-time input from, say, a mouse or a joystick? Would be possible to provide Stella with results from an experimental tracking run if the data were presented as a list of successive values? Does Stella have any provision for user-defined functions (which could be used to do input-output even if Stella doesn't have that ability)?

Finally, how Big a Mac is required to run Stella? If one of the Little Macs could run it, I would consider getting a used Mac just to be able to communicate through Stella with Mac users.

I hope you're willing to be the Stella guru for a lot of PCTers who have Macs. You would be helping a whole lot of people beside yourself to get a real grip on PCT.

I should point out that the Forrester/System-dynamics concepts of "levels" and "flows" are somewhat narrow concepts. Some rather unnatural translations would have to be made, for instance, to describe how a position-control system works. But that's just a matter of words. I think Stella should be able to do the job.

--

Best to all, Bill P.

Date: Mon Sep 14, 1992 11:31 am PST
Subject: Re: turing test; autonomy & reorganization

Penni Sibun (920914.1200) --

[From Bill Powers (920914.0900)]

With communication limited to a teletype, we can only test the other entity at the higher levels of organization. At the strictly rule-driven level, I doubt whether we could distinguish a program from a person if the program were complex enough. But what would happen when we try to see the principles or system concepts that the other "person" is working from? Is it possible, for example, to write a program that follows principles of humility, arrogance, caution, and so on? I don't doubt that a program could make use of vocabularies constructed from the speech of people judged to exhibit these principles, but could a program in itself generate any new principle? And if it could, would its principles be indistinguishable from the ones human beings adopt and control for? Could a computer program think up something like the Turing test? Could a computer program be devised to carry out the Turing test with another entity?

right. without having read it, i think this is the kind of thing the article that prompted rick's post was saying. a program that can actually pass the original turing test would be impressive, but we probably still wouldn't want to ascribe intelligence to it when push came to shove. we'd want to see more, we'd want to see versatility.

the first ``official'' (it's got a lot of publicity and money attached) turing test was held last year. i've seen some transcripts and the system that won seemed pretty credible (it ws judged human about half the time). however, the task was limited not only to teletype conversation, but to a particular topic. all very interesting, though. (see ``AI magazine'' 13(2), summer 92, pp80-95 for a report.)

Would you be surprised to learn that I am really an 80486/33 with 4 Megs of memory?

very. ;-}

cheers. --penni

Date: Mon Sep 14, 1992 12:22 pm PST
Subject: Re: Turing

[From Rick Marken (920914.1200)]

penni sibun (920913.2000) says (to me):

>that's ok; i don't usually read you since i know what you're going to say.

ouch!

> however, you did pique my curiosity by mentioning the turing
>test--even though i knew what you were going to say ;-}).

>turing based his test on what was apparently a sort of parlor game in
>which the interrogator has to decide which is male and which is
>female. maybe the analogy will be helpful. how do we know if someone
>is male or female?

>if *sex* is decided in such a ``superficial'' way, then why not the
>ability to think?

This analogy rests on the assumption that behavior is to thinking as gender is to chromosomes. Apparently, you assume that the "superficial" signs of gender (maleness vs femaleness) are an output of chromosome structure:

```
xy --> maleness
xx --> femaleness
```

So, the "sex test" looks at observed maleness or femaleness and guesses at the "real" sex of the person (their chromosome structure). The above assumptions make it reasonable to think that you can "work backwards", inferring chromosome structure from appearance.

I agree that this is just how Turing looked at his test for mind. He assumed that "mental behavior" was an output of mental structure, and "non-mental behavior" was what resulted when there was no mind:

```
mind --> mental behavior
no-mind --> non-mental behavior.
```

So the Turing test is based on the same idea as the sex test; look at observed behavior and work backwards to infer whether or not the behavior was generated by a system with a mind. If you see "mental behavior" then guess mind; if you see non-mental behavior then guess "no mind".

The point of my post (if you are still reading) was precisely this: the Turing test is Behavioristic in the sense that it assumes that mental behavior is the OUTPUT of a mental system. This assumption is also the basis of cognitive psychology. The idea is that you can infer something about the mind by looking at behavior because behavior is an OUTPUT of mind.

I brought this up because, according to PCT, the Turing test (as Turing described it) cannot work; behavior (results that are intentionally produced by the organism) is NOT OUTPUT -- it is controlled INPUT (ie. perceptions. The PCT model of mind produces a 3rd alternative to the two above, namely:

```
mind -->input< -- actions
      | "behavior"      |
```

The mind (reference signals) specifies intended levels of input. The system acts to keep these inputs at their intended levels (but it does not cause these actions; these actions are caused by the disturbance - resistant characteristics of the system).

An observer might call either the actions or the observed correlate of the intended input the "behavior" of the system. In PCT, we reserve the word "behavior" for the intended inputs; actions are not really the behavior of the system since the system doesn't really determine what those actions will be. But actions (or their side effects as seen by the observer) are what are most likely to be seen as behavior. These actions, however, don't really reveal what Turing is trying to get at -- whether or not he is dealing with a system with a mind.

What Turing really needs to do (if the PCT model of the system is correct) is figure out whether or not he is dealing with a system whose actions are aimed at protecting the state of an input variable (a variable that is determined by a mind). To do this, Turing would have to know that there might be a controlled variable involved (the "input" in the diagram) and test to see if that variable is being protected from disturbance. This process is called "the test for the controlled variable". It differs from the Turing test only inasmuch as it clearly specifies QUANTITATIVELY how a system with a mind differs from a system without one; the system with a mind is busy trying to control input variables (at least if it is the kind of mental system assumed by PCT -- a purposeful mental system). The no-mind system is not controlling an input.

You cannot tell whether or not a system is controlling an input by simply watching to see if its behavior "looks" intelligent. But this is what Turing suggested WAS possible. You "poke" at the system (with questions) and watch to see how it responds. If the response is judged "intelligent" then you guess that it was made by a mental system. This approach is not only unnecessarily subjective (as you hint at in your "sex" example) -- it is bound to give the wrong results when you are dealing with a mental system that is organized as a closed loop control system. The failure of the Turing test is demonstrated by my "mind read" and "find mind" programs (since you don't read my stuff I guess I can count on your not reading about the mind read program in my Mind Readings book). The "find mind" program, which you will be happy to learn is NOT written up anywhere, is a good example of the Turing Test. One of five numbers is actually controlling it's position but all five are "doing" pretty complex movements. I could have each one trace out the form of a character punched into the keyboard -- their movement in response to the input is like their response to a Turing question. I could even have only one consistently answer correctly. But, still, all but one are generating behavioral output. Only one is actually controlling an input -- and this is the only one with a "mind" -- a reference for it's position. This one is easily detected by disturbing it's position (with the mouse).

I doubt that you are still there, penni, but for the benefit of anyone who still is, the point is the the Turing Test (as it is described in Turing's article -- yes, I read it) is based on a behaviorist conception of behavior -- ie. that behavior is an output generated by mind events. PCT says that this concept of behavior is demonstrably false. So there is no way to distinguish mind generated from non-mind generated behavior by just LOOKING at characteristics of the behavior itself. You CAN, however, distinguish mind generated from non-mind generated behavior using the Powers Test (the test for the controlled variable) which involves:

1. identifying a possible controlled variable

2. produce disturbances to the variable (events that should change the hypothetical controlled variable if it is NOT controlled)
3. monitor state of hypothetical controlled variable under continuous disturbance
4. look for lack of effect of disturbance
5. if there is no effect, make sure system being tested can perceive the variable (control stops when variable is obscured) and can effect it

Turing was a smart, swell guy and all but, unfortunately, he didn't understand control (who did at the time?) so his test (like everything else in the life sciences) was based on the idea that behavior is OUTPUT. An understandable (but in 1992 a quite unnecessary) misconception.

Best regards Rick

Date: Mon Sep 14, 1992 12:45 pm PST
Subject: Re: Turing

penni sibun (920914.1300):

[From Rick Marken (920914.1200)]

>female. maybe the analogy will be helpful. how do we know if someone
>is male or female?

>if *sex* is decided in such a ``superficial'' way, then why not the
>ability to think?

This analogy rests on the assumption that behavior is to thinking as
gender is to chromosomes.

no, it rests on the assumption that, like intelligence, sex is something we
make judgements about when we consider people or other organisms.

Apparently, you assume that the "superficial"
signs of gender (maleness vs femaleness) are an output of chromosome
structure:

xy --> maleness
xx --> femaleness

i said no such thing. how could you possibly have got that from my post? i
listed several counterexamples to your ``outputs.'' i think you come to
things w/ such preconceptions that i barely understand what you're talking
about when you summarize what you think i said. some of this is undoubtedly my
fault; i'm hardly perfect at presenting my ideas. but i find you especially
difficult to communicate w/.

So, the "sex test" looks at observed maleness or femaleness and guesses
at the "real" sex of the person (their chromosome structure).

no. the sex test looks at a person and guesses whether they are female or male. period. the sex test doesn't care about what's inside. it only has access to what you consider ``superficial'' things.

The above assumptions make it reasonable to think that you can "work backwards", inferring chromosome structure from appearance.

no. the test doesn't work backwards; it doesn't have anything to work backwards to.

I agree that this is just how Turing looked at his test for mind.

well, we're obviously not agreeing at all.

He assumed that "mental behavior" was an output of mental structure, and "non-mental behavior" was what resulted when there was no mind:

mind --> mental behavior
no-mind --> non-mental behavior.

no. when you go on like this, i see so clearly preston's point that some people are so locked into the behaviorist/mentalist mindframe that they can't talk about anything else. (did you ever get her paper?)

i think you're a smart and articulate person, btw. i just think you're terrifically wedged in this area. it's not my job to unwedge you; but maybe one day one of my stories will make you look at things just a little differently.

cheers. --penni

Date: Mon Sep 14, 1992 2:24 pm PST
From: Control Systems Group Network EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.BITNET@pucc.princeton.edu TO: *
Dag Forssell / MCI ID: 474-2580 TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767 TO: Henry James
Bicycles Inc / MCI ID: 509-6370 TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414 MBX:
CSG-L%UIUCVMD.BITNET@pucc.princeton.edu Subject: Imagination; turing test
[From Bill Powers (920914.1530)]

Chris Love (920914.1300) --

> The problem: By allowing those n-1 level ECS' to provide the >imagined
percepts to your (focus) n-level ECS you are getting >(undesired?) influence
from neighboring n-level ECS'?

Let's see, when could this cause a problem? I guess it would be when there are environmental disturbances that actually cause trouble at the level n-1 systems in the imagination mode, which of course those systems don't report to the higher level as a change in the perceptions. So the systems at level n that think they are dealing with real inputs are getting partially-made-up

inputs, and therefore fail to take corrective actions they should be taking. Of course that's true in spades for the level-n system that wants to be imagining, but it's a mistake to cause other systems at that level to deal with false information.

Perhaps there has to be some "context" such that a whole group of level-n systems has to use imagined information at the same time. Or perhaps -- and this is just as realistic -- real systems DO get into this kind of trouble! This thing didn't come with an instruction manual.

```
> Delta weight = learning rate* connection wt. * error signal > New wt. =
Old wt. + Delta wt. > > Comments: > The learning rate is typically set to a
small value between >0.1 and 0.3. > The connection weight is the existing
weight (wt.) of the >connection of interest. > The error signal is that
signal emanating from the comparator.
```

I think there's something to the GENERAL Hebbian idea of learning, but the concept that more strength in a connection is better just doesn't hold up. When you get around to designing control systems that have to behave in an environment with somewhat physical properties -- in which the effects of actions don't happen instantly -- you'll find that "more" output is a good thing only up to a certain point, after which more still will lead to instability. Also, trying to use Hebbian concepts in adjusting weights will work only as long as the optimal condition is all weights at maximum or zero. For computing many perceptions, some weights have to be adjusted upward and some downward to get the optimal result. I think in general that some particular weight is the correct one for the situation, not maximum weight.

I've seen some work (Harry Klopf, on the "selfish neuron") in which all the working models that used this Hebbian sort of learning failed, because their only stable state occurred when all weights went either to maximum or zero. To get them to settle down anywhere else, the modelers had to introduce all sorts of ad-hoc limits and nonlinearities. Just the right ones, of course, to produce the results they wanted. When you have to fiddle with a model at that level of detail to make it work, something is probably basically wrong with the model. You end up with a "right-result-producer" that only works when it's tweaked just so.

```
> What does this do? > Well, as I see it it correlates particular
references with >particular percept signals.
```

These words don't make any sense to me. Why would we want to correlate a reference [signal?] with a percept signal? Do you mean that the resulting error signal will cause lower level system to alter the reference signal toward the value of the reference signal? I just want to make sure we're talking about the same organization.

As to your Almighty ECS that you mentioned to Rick, you're going to have to be almighty careful about what kind of knowledge you give it. After all, it doesn't know whether it's controlling categories or intensities. Whatever you build into it has to work exactly the same way no matter what it's controlling. Also, if you try to build in all levels of functioning of all kinds into every ECS, you're going to have a tremendous duplication of function, sort of like giving every cell that uses oxygen a pair of lungs.

----- Penni
Sibun (920914.1200) --

>the first ``official'' (it's got a lot of publicity and money >attached) turing test was held last year. i've seen some transcripts >and the system that won seemed pretty credible (it ws judged human >about half the time). however, the task was limited not only to >teletype conversation, but to a particular topic.

Even with just a particular topic, it shouldn't be too hard to tell whether the other guy was a machine if you were free to ask questions. For example, you could take any two statements by the other entity and ask "Can you show me that these statements really relate to this topic?" or "Are these statements consistent with each other?" Actually, it would be interesting just to respond to any particular statement by saying "No, you're wrong about that." Anything to elicit some signs of higher-level perceptions or an attempt to control for something.

----- In your
reply to Rick you say

>no. the sex test looks at a person and guesses whether they are >female or male. period. the sex test doesn't care about what's >inside. it only has access to what you consider ``superficial'' >things.

I guess that as usual big disagreements rest on little words, in this case "are." When you say you're trying to guess whether they "are" male or female, the normal interpretation would be that you're trying to guess what they REALLY are, not just how they look. In other words, you can only tell whether your guess is "right" by (blush) peeking.

If you mean ONLY judging by the readily-visible characteristics, there isn't much guessing about it, is there? You just classify your perceptions the way you're used to classifying them.

----- -- Best
to all,

Bill P.

Date: Mon Sep 14, 1992 8:00 pm PST From: sibun EMS:
INTERNET / MCI ID: 376-5414 MBX: sibun@parc.xerox.com TO: *
Dag Forssell / MCI ID: 474-2580 TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767 TO: Henry James
Bicycles Inc / MCI ID: 509-6370 TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414 MBX:
CSG-L%UIUCVMD.BITNET@pucc.princeton.edu Subject: Re: Imagination; turing test
(ps 920914.2100)

[From Bill Powers (920914.1530)]

In your reply to Rick you say

>no. the sex test looks at a person and guesses whether they are
>female or male. period. the sex test doesn't care about what's >inside.
it only has access to what you consider ``superficial'' >things.

I guess that as usual big disagreements rest on little words, in this case "are." When you say you're trying to guess whether they "are" male or female, the normal interpretation would be that you're trying to guess what they REALLY are, not just how they look. In other words, you can only tell whether your guess is "right" by (blush) peeking.

i've been trying to establish that ``what they REALLY are'' has two answers, and that these answers are distinct (sex-by-chromosomes doesn't always match sex-by-judgement). i've further insinuated that the answer that is used 99.99% of the time is the one by judgement. and i've left the conclusion to be drawn that the judgement answer is in fact the really real one (unless you're a doctor or scientist and have some reason to muck about in somebody's chromosomes).

so the ambiguity you perceive is deliberate. i'm hoping that rather than thinking i've been sloppy, y'all will think i'm making a point.

cheers. --penni

Date: Tue Sep 15, 1992 3:52 am PST
Subject: More on my ideology; models -> asymmetries of interaction

From Greg Williams (920915)

>Bill Powers (920914.0900)

First, let's "go up a level." In our current debate, and in all previous (sustained) ones between us on the net, I have been motivated by two fundamental problems.

One: I have observed that various PCTers put forth highly technical usages of terms such as "control" and "autonomy" in contexts where the fact that they ARE being used technically, in "PCT-senses," is not made clear explicitly. By doing this, PCT-novices and laypersons are misled into thinking that PCT supports notions which it doesn't support. For example, our debate about what PCT says about "free will" was specifically motivated by casual conversations with PCTers who appeared to (in my view, mistakenly) presume that PCT provides a basis for something closely akin to the traditional Western religious notion of "free will." Only by making explicit the specific implications of PCT can such misunderstandings be resolved/prevented.

Two: I have observed a tendency among PCTers to treat PCT/HPCT as a monolith, as if all its aspects were equally well-supported. It seems undeniable to me that the basic tenets of PCT are virtually unassailable (there is no other even plausible candidate "overall" mechanism for accomplishing ends consistently in a disturbance-filled world). But I find no such rock-solidity in HPCT (particularly at higher levels) and reorganization theory. These are plausible, but so, I think, are other theories (WHICH REMAIN CONSISTENT WITH THE BASICS OF PCT; examples: changing gains rather than reference levels and "directed" (non-random) reorganization attempts). Only by understanding that

some of Bill's ideas are more impregnable than others can PCTers assess the degree to which various notions "which follow from" particular aspects of PCT/HPCT are worthy of being espoused as "gospel."

I believe that these two problems stand in the way of more widespread adoption of PCT ideas (which I take to be desirable). That's why I've been trying to bring these problems to the fore.

Back to the debate at hand.

>Even under the S-R postulates, the experimenter is depending on many
>autonomous functions inside the organism, starting with sensing and
>recognizing the reinforcer.

Certainly. Skinner's problem is not having a model of organismic mechanics. PCT doesn't have that problem. Nevertheless, the experimenter DOES HAVE AN INFLUENCE ON WHAT THE "CONTROLLED" ORGANISM DOES (just as the organism has an influence). The organism doesn't "work" (thanks, Dennis) autistically!

>In a hierarchy of control that accretes through reorganization, this
>situation is compounded in complexity. The external agent can produce
>expected behaviors from this hierarchy, but only if they don't violate
>the conditions that the reorganizing system must, by its nature,
>maintain.

Given a PARTICULAR (current) hierarchy, yes. No question about it. That's the constraint on successful purposeful influence: no conflict generation (with the possible exception of using the influencee's reorganization to get what the influencer wants).

>On the surface, the external agent is dealing only with a single
>control system, the one whose actions the agent is trying to control.
>But as soon as this control has any effect that significantly disturbs
>the control hierarchy at the same or a higher level, either direct
>resistance will develop or the reference signals in the original
>system will begin to shift.

Yes -- that's why the external agent (has to be living, and rather skilled; in fact, needs to have a reasonably good model of the control system!) tries to NOT significantly "disturb" the hierarchy. I claim (empirically) that all of this is indeed possible, and that successful purposeful influence without appreciable conflict generation is ubiquitous.

>So in order to retain control in general, the external agent has to give up
>control of that action and pick a higher-level aspect of behavior to control.

No, the external agent just has to be skilled enough to not NEED to pick a higher-level aspect.

>It's inevitable that the general desire to control another must lead
>to applying disturbances at higher and higher levels, which of course
>becomes more and more difficult. And in the background, at any level
>of interaction, the reorganizing system sits watching for any
>consequences of these interactions, known to the hierarchy or not,
>that have adverse effects on the fundamental variables that signify

>the viability of the organism.

No, not inevitable. If the external agent's model is good enough so that negligible conflict is generated, there is no going to higher and higher levels. Anyway, the purposeful influencer might be "disturbing" at a low level in such a way that the intrinsic reference signals are being BETTER satisfied than otherwise (i.e., the example of feeding tasty AND healthy food to kids).

>The main point is that if there is a reorganizing system of the kind I
>propose, then it is autonomous with respect to the events of a single
>lifetime.

Well, maybe that is just a tautological statement. BTW, I see empirical evidence against it. An individual can ACQUIRE in a single lifetime a "reference signal for honor among peers" which can OVERRIDE the intrinsic reference signals related to survival. Even granting you that the KINDS of intrinsic reference signals an organism has don't change within a lifetime (although their levels might), the functional result is that ACQUIRED (within a lifetime) reference signals can, in effect, alter their functioning.

>It is this reorganizing system that sets the limits on what an
>external agent can get the organism to do or not do.

It does set some limits. It does not set everything about what an organism does. Ditto for environmental influences.

>The entire hierarchy is constructed around the unchangeable requirement that
>intrinsic error shall be held near zero.

No, I think intrinsic error can be OVERRIDEN in some cases by error associated with reference signals "acquired" in a single lifetime.

>Apparently, also, awareness can direct the locus of reorganization --
>that is, reorganization is concentrated on the systems with which
>awareness is most closely identified at a given time, at least in the
>neural-behavioral hierarchy.

This begs the questions of whether "directing the locus" is "free" of environmental influences and whether it is "random."

>This is my theory of "free will," which defines both "free" and
>(conscious) "will."

>...

>The point in bringing this up is that nobody can have any freer a will
>than this -- the ability to change point of view, and to direct
>reorganization to a place in one's own system, with the same apparent
>randomness that we attribute to reorganization itself.

Did you really just answer the begged questions? Why should it be random? What evidence do you have that it is? My awareness seems to go to "problem" areas, which are (I claim) anything BUT randomly distributed. It also seems that I can "direct" my attention to various "places," but I never seem to except for HISTORICAL REASONS (which seem environmentally related).

>If this is true, then what we call manipulation of one person by

>another is simply the interaction of two systems, each attempting to
>reorganize itself to suit its own internal definition of what is good
>or bad for it.

REGARDLESS of the potential problems I see with your argument above, I see this statement as being correct.

>We do NOT have the situation where one all-knowing omnipotent system of
>unlimited intellect is manipulating a dimly conscious and defenseless system
>of some inferior kind -- that is the situation only in the relation between a
>human being and an emphatically lower order of life. It is not the situation
>that holds between normal adult human beings.

I see this statement as correct, also.

>The manipulator's "will to manipulate" is not different in kind or effect
>from the manipulee's "will to defend."

Here, we begin to differ. There MIGHT not be a difference -- in the case of symmetric purposive influence. But purposeful influence CAN be asymmetric. The asymmetry lies in the use (or more appropriateness) of a model of the other's controlling by one of the parties but not the other. The successful influencer employs a sufficiently sophisticated model of the influencee's controlling. The influencee might not employ a model of the influencer's controlling AT ALL (he/she might not care) or might not be using a sophisticated model of the influencer's controlling -- in this case, the interaction would be asymmetric. So use of models of others' controlling is the important (PCT-derived!) underlying mechanism which explains asymmetry of interactions. Pretty heady stuff! (Note that in everyday interactions, the asymmetries of interactions generally go BOTH ways -- simultaneously and consecutively -- between dyads, not exclusively one-way (which we might (dare to?) call "domination").)

>The term "manipulation" is unfortunate, because it implies power and
>will on the side of the manipulator, and weakness and automaticity on
>the part of the manipulated: exploiter and exploited, controller and
>controlled, agent and effect.

If you continue to use "manipulation," you'll be on your own. I'm using "purposive influence." Consider me converted... now, if we can just convert you!

>This implies that the manipulator has some abilities or powers that are
>lacking in the manipulee, an awareness that is denied to the manipulee, a
>natural advantage over the manipulee. This is simply not the case in the HPCT
>concept of behavior.

It is simply not the case in the PCT concept of behavior. (Occam's razor!)

>All human beings have the same basic abilities, save for those that are
>crippled in some striking way. Whatever conceptions of social life we may
>have
>developed, in truth we all have the same equipment, which is designed to
>allow
>us to control what happens to ourselves -- and especially what happens to our
>ability to continue living.

Some are more skillful at developing and using models of others' controlling, and that makes a great deal of difference in the symmetry (lack thereof) of social interactions.

>Manipulation is just a one-sided way of looking at interaction.

Purposive influence is the PCT way of explaining asymmetric interaction.

Best, Greg

Date: Tue Sep 15, 1992 7:58 am PST
From: g cziko
MBX: g-cziko@uiuc.edu
TO: * Dag Forssell / MCI ID: 474-2580
Subject: CSGnet Intro Document

Dag (direct): You said:

>Perhaps Gary Cziko could re-post the introductory message (usenet intro),
>which includes a brief PCT book list with commentary. It's been a long
>time, Gary.

I don't think for a moment I don't know it! I've been swamped with the start of the academic year and related duties in addition to trying to finish my book. I'm starting to think that I am not the right person to be in charge of the documentation to be made available via CSGnet and Bill Silvert's fileserver.

Is this something for which you would like to take primary responsibility? Your talent in interest lie in making PCT attractive to newcomers, so this duty would fit in. I realize that you are at somewhat of a disadvantage since you have to pay for your access to CSGnet, but I would certainly be willing to help out here where I could.

What I think we should have is a not-too-long posting sent to new subscribers and monthly to the net which serves as an introduction and which refers to other documents and programs on the fileserver. You already started to create such a document last summer (spring?) which I have on file and I could send back to you. I could then fill in the info on the documents and files on the fileserver.

Should I start the ball rolling again by sending back what I have from you with inserted commets?--Gary

P.S. Woops, I just checked and I don't seem to have your document any longer on my current CSGnet floppy. Do you still have a copy? If not, I should be able to pull it out of the archives if you can give me an idea of the of the date it was sent.

Gary A. Cziko Telephone: (217) 333-8527
Educational Psychology FAX: (217) 333-5847
University of Illinois E-mail: g-cziko@uiuc.edu
1310 S. Sixth Street Radio: N9MJZ

210 Education Building
Champaign, Illinois 61820-6990
USA

Date: Tue Sep 15, 1992 8:02 am PST
Subject: Re: turing test

The fascination with the Turing test has always intruiged me. To me, the following are obvious. Perhaps I need to get educated.

1) At first consideration, X passing some test must only be a NECESSARY, and not a SUFFICIENT condition for X to have the property tested for. Thus X passing the TT is not determining that X is intelligent: surely it could be the case that something that is not "really" intelligent can be made to look like it to any finite approximation; indeed, that seems to be very common. Thus passing the TT is at best non-determining EVIDENCE for the INDUCTIVE INFERENCE that X is intelligent.

2) Point (1) depends on a PRIOR DEFINITION of the property being tested for, i.e. intelligence, so that the TT is performed relative to it. Therefore it would be possible in principle to distinguish between the "really" intelligent and the simulations. But that just begs the question: this "new" test just replaces the TT. In the end, we must KNOW, or at least SAY, what intelligence is, and then construct the TT to measure X's satisfaction of the definition. Of course, in the case of intelligence, this is exactly NOT the case, whereas in control, the Powers school (as I'll call it) DOES have a prior definition of control.

3) Lacking a prior definition of intelligence, the group of conditions necessary to pass the TT simply BECOMES the ad hoc definition. This house of cards has been the methodological basis of AI for fifty years! Turing has not really advanced the argument any: How do we know something is intelligent? because it acts like us. What is intelligence? Well, just whatever it is that we do. Of course, we hasten to observe that much of what we do and say is exactly NOT intelligent, but lets not let that clutter the issue (a Turing Test for life itself, anyone?).

I'm glad that Control Theory seems to be founded on less shaky methodological ground than traditional AI.

O-----
>
| Cliff Joslyn, Cybernetician at Large, 327 Spring St #2 Portland ME 04102 USA
| Systems Science, SUNY Binghamton NASA Goddard Space Flight Center
| cjoslyn@bingsuns.cc.binghamton.edu joslyn@kong.gsfc.nasa.gov
V All the world is biscuit shaped. . .

Date: Tue Sep 15, 1992 9:54 am PST
Subject: Re: Turing

(eric harnden: 920915.1230)

ok, i think i'm coming a little closer, here. mr marken's post on the turing test contained a thumbnail of the pct perspective which makes it a little easier for me to understand why his insistence that ai should not focus on modeling behavioral outputs has been so strongly stated, and am finding myself more in agreement. i think that part of my resistance to this point has been a degree of confusion over the importance of the specific terminology. frankly, i don't think it matters where one puts the 'behavior' label. my own approach has always been a cybernetic one, and the shifting of a label doesn't change any of my (mental or experimental) model structures. i also think that such a (perhaps misplaced?) emphasis on nomenclature may be partly responsible for an earlier interaction between mr marken and ms sibun. given the interaction between entity and environment as the basis for discussion, the question of the locus of mind is largely philosophical. as i told ms sibun in a private memo, i found some of the material that she presented interesting in as far it related to ideas about mind that i have picked up from bateson. but so what? or so what about objections to those views? i model the mechanism, not the philosophy.

one way in which i am now more confused, however... mr marken said something about the organism not being the cause of behavior. that behavior is 'caused' by environmental disturbance of the organism's perceptions. on the strictly structural level i'm willing to accept this, given the local definition of behavior. but, since the purpose of behavior is adjustment of perceptions toward reference values, then that which is caused by environmental disturbance must be the organism's actions... which is precisely the 'normal' definition of behavior. is there a contradiction here? sorry if the answer to this question is made glaringly clear by closer reading of some part of BCP which i have skimmed/forgotten. also, btw, note that this question falls in the category just discussed. i don't really see how the distinction affects the model.

Eric Harnden (Ronin)

Date: Tue Sep 15, 1992 9:55 am PST
Subject: Really real sex; HPCT; purposeful influence

[From Bill Powers (920915.1000)] Penni Sibun (920914.2100) --

>i've been trying to establish that ``what they [sexes] REALLY are'' has
>two answers, and that these answers are distinct (sex-by-chromosomes
>doesn't always match sex-by-judgement). i've further insinuated that
>the answer that is used 99.99% of the time is the one by judgement.
>and i've left the conclusion to be drawn that the judgement answer is
>in fact the really real one (unless you're a doctor or scientist and
>have some reason to muck about in somebody's chromosomes).

A theorizer about human nature lives a strange life. If not 99.99% of the time, then 90% of the time it lives life exactly as you describe. What appears to be the case is taken without question as if it is the case. Judgment, in other words, is based on unanalyzed perceptions and is trusted. So 90% of the time, the theorizer is just as sane as anyone else.

But the other 10% of the time the theorizer is mad. It's only when you begin viewing your own experiences in a way that departs significantly from the norm

that you begin to see beneath the surface, to question what normal people take for granted or don't even notice.

Take the category "sex." If you base your notions of sex strictly on external appearances -- habits of speech or dress, mannerisms, length of hair, use of makeup, pitch of voice, attitudes toward children, roles in relationship to people in a disjoint category, and so on -- you will base your expectations and wants on those appearances. If, for scientific or other reasons, you decide to base your concepts of sex on chromosomes, skeletal conformity, organ configuration, and the like, you will base your expectations and wants on these appearances instead.

Which one of these ways of categorization is "right?" To me it's obvious that this is a non-question. "Sex" is just a word. We can attach it to any perceptual category we like. We could even use it to mean a number, 6. It's the category, the way of lumping different perceptions together and treating them as "one thing," that makes the difference, not the word. When you strip away the word, it becomes obvious that what the scientist is looking at is one collection of perceptions, while what the normal person is looking at is a different collection.

What the world does is different depending on how you categorize its components. Even the relationships among categories will be different if you draw the boundaries differently. And what you DO in the world -- what you control for, what you try to make happen and try to prevent from happening, is entirely shaped by how you perceive the world. You can interact intentionally with the world only AS YOU PERCEIVE IT.

So NONE of the ways of perceiving sex is the "really real" way. There's no arbiter sitting up on a high stool beside the net interpreting the rules and informing the players whether the ball was REALLY "in" or "out." The players have to work out a common set of rules between themselves for the purposes of the game. Of course all the players are subject to the laws of physical dynamics, gravity, elasticity, aerodynamics, and other things -- or however you want to express the nature of That Itself, the Immanent Order. All the rules that anyone can possibly invent must be played out on that stage. The only "really real" aspects of the game are the consequences of our actions that occur whether we think about them or not, whether we want them to occur or not. The mad theorist is trying to figure out mostly the nature of the stage, and only incidentally the particular rules that two or three generations of players have brought into the game with them. In a century they'll all be dead and all those rules will be different anyway.

In the meantime, or 90% of it, the theorist is just like anyone else.

--

Greg Williams (920915) --

>One: I have observed that various PCTers put forth highly technical
>usages of terms such as "control" and "autonomy" in contexts where the
>fact that they ARE being used technically, in "PCT-senses," is not made
>clear explicitly. By doing this, PCT-novices and laypersons are misled
>into thinking that PCT supports notions which it doesn't support.

I agree that this is a problem, particularly in a public discussion where we may understand each other but others will misinterpret. Not that we necessarily agree on everything, which is another complication.

>Two: I have observed a tendency among PCTers to treat PCT/HPCT as a >monolith, as if all its aspects were equally well-supported. It seems >undeniable to me that the basic tenets of PCT are virtually >unassailable (there is no other even plausible candidate "overall" >mechanism for accomplishing ends consistently in a disturbance-filled >world). But I find no such rock-solidity in HPCT (particularly at >higher levels) and reorganization theory.

Again, agreed. When we talk about higher levels, however, there are certain threads that follow through from what we know more solidly about the lowest levels. One is the concept of viewing all experience AS PERCEPTIONS rather than in the naive realist fashion. The HPCT model doesn't totally deny the underlying concepts of realism, because it uses models of the external world just as if they were a true representation of that world. So an independent external world is assumed. The main difference is that the HPCT model shows perception of that world as an activity in a brain, with (as far as we know so far) completely arbitrary perceptual transforms creating each new level of representation. So the brain literally "makes sense" of the world -- creates a self-consistent network of representations that depends on the world but has no one necessary relationship to that world.

We can see this relationship of perception to our physics-models of the world very clearly at the lower levels of the nervous system. This provides a plausible justification for the general concept of experience as experience of perceptual signals, not of the world directly. By applying that principle to higher levels of perception (without, of course, the relatively detailed circuit-tracing that we can do in the spinal-cord systems) we can come to understand the experienced world at those levels in a new way. We can see that the principle of controlling perception rather than reality applies at the higher levels as well, and indeed makes more sense of those kinds of experiences than any other point of view can provide. I think this counts as reasonably "solid" knowledge, on any practical scale of solidness.

There are other isolated islands of solidness in the HPCT concept. We can pick any particular behavior that seems to be of a kind that belongs in the higher levels -- behaviors relating to morals, for example -- and set up experimental conditions under the assumptions of PCT (not HPCT). And generally we find exactly what we expect: variations in behavior that oppose disturbances and maintain constant some aspect of a variable like a moral aspect of behavior. So we can demonstrate that the PCT model works with a great variety of behaviors, varieties that range over what we vaguely see as "levels." Control of morals is clearly of a different level than control of muscle tension, which is clearly of a different level from control of posture, and so on.

The "H" part of HPCT then shows up when we examine the MEANS by which "higher" control processes are carried out. Invariably, they are cast in terms of variables which themselves can be tested as controlled variables in the usual PCT way. And they, too, prove to be resistant to disturbance through actions of a still more detailed sort. Although investigations like these have never been carried out in a systematic and exhaustive way (that will take organization and money), we can see that when such research is finally

started, there will be organization there to find, and that it will almost undoubtedly resemble the hierarchy in HPCT.

There's one final aspect of HPCT that seems reasonably solid: private experience. When you look at your own life in terms of a hierarchy of controlled perceptions, you can make sense of it in a way that's intuitively satisfying. The HPCT meaning of a goal structure shows up clearly in at least some segments of private experience: I want this in order to get that, and when I get that I will satisfy something even more important, and so on. With the HPCT model in mind, you can put these goals into some kind of order that seems real; they really do seem to be related as HPCT claims, in general if not in detail. This isn't "objective" verification, but if you think about it, when many different people agree that they are observing the same phenomenon, does it matter where they are observing it? We can't really PROVE that two people are seeing the same meter-reading; we take the agreement itself as indication of objectivity if there aren't too many dissenters.

Back to the argument.

I don't want to stretch this post out much more so I'll be brief

Whenever I hear you saying something like "the experimenter DOES HAVE AN INFLUENCE ON WHAT THE 'CONTROLLED' ORGANISM DOES (just as the organism has an influence).. " I become uncomfortable. It appears to me that you're classing two things together that are not alike. One of the things is simply "having an influence," which just means bumping into someone or getting in the way or otherwise disturbing the other's perceptions. That naturally leads to actions that keep the effect of the disturbance small, in those respects that matter to the influencee. The other thing is "VARYING an influence so as to produce a CONSISTENT EFFECT on the other's behavior." That's a completely different matter. That's control, and only control can be called purposeful. Of all the influences that affect what a person does, only a very tiny fraction of them is used intentionally by someone else to produce a preselected effect.

In its first meaning, "having an influence" is a sure bet. Practically anything that is done by a person or a natural event influences perceptions, and results in changes of action that counteract incipient changes in controlled variables. Just being in the world means being constantly influenced by external processes. That's the whole reason that control exists.

The second meaning, however, brings in purposeful influence, which means influence that is organized not around the physical actions taken by the influencer, but around the behavior of the influencee. Purposeful control means having a goal for the behavior of another person, and taking whatever action is required to make that desired behavior appear. No particular action can be influential in the sense of control; control produces consistent ends by VARIABLE means. There is no way to tell whether a given disturbance is part of someone else's control process, not just by looking at the disturbance. There is no difference between an intentional disturbance and an unintentional one. It's just a disturbance.

It often seems to me that you are lumping together "having an effect" with "having a particular intended effect." But simply showing that one person's action results in a change of action by another is not sufficient to show purposeful control, even if the person doing the disturbing wanted to see the

action that the other took, even if that action satisfied a goal in the disturber. That wanting, that goal, is completely irrelevant -- it's just a lucky guess -- unless the controller can VARY the disturbance and KEEP the other's action in the form desired, either for a long time or on repeated occasions.

You say " The organism doesn't 'work' (thanks, Dennis) autistically!" But this means nothing with respect to PURPOSEFUL influence. If an organism is in an environment, it must control in order to keep its perceptions as it wishes them to be, whether the disturbances that interfere were generated purposefully or naturally (that is, not by control systems). Autism, behavior without regard to environmental happenings or properties (as I understand your meaning), has no bearing on the subject of purposeful influence. It isn't true that in the absence of purposeful influences there are no influences, nor that behavior will be autistic unless the influences are purposeful. There are far more influences unrelated to anyone's goals than influences aimed at producing a particular intended change in behavior.

The question before us is not whether disturbances of perceptions can alter actions, but the extent to which an external agent can use disturbances in a systematic way to make other people behave to suit the disturber's goals. You claim that this extent is far greater and far more significant than I am yet willing to admit (absent coercion, which I do see as a serious problem). That's what we're arguing about, in my opinion.

I'll comment on only one more thing from your post -- this is getting too long, and anything I miss will come up again.

>>So in order to retain control in general, the external agent has to
>>give up control of that action and pick a higher-level aspect of
>>behavior to control.

>No, the external agent just has to be skilled enough to not NEED to
>pick a higher-level aspect.

That's impossible. If you're controlling at one level, there is no way you can keep a higher system from changing the reference signal in the system you're taking advantage of. It's impossible from another standpoint, too. To be "skillful enough" to avoid disturbing any variable of the same or higher level that matters to the controllee, the controller would have to understand the controllee in complete detail, and also predict every independent disturbance that is going to happen. I don't think this will ever happen, but more to the point, I don't think it is likely to HAVE HAPPENED. Neither can you get out of this by saying that a person can have "sufficient" knowledge to do this -- that's merely asserting that such control has happened, so it must be the case that the knowledge was sufficient, then using the supposed sufficiency of knowledge to "prove" that such control happens. And lastly, I think that anyone with enough understanding of human nature to know another person's hierarchy in the necessary detail (not to mention predicting all events in the natural world that could disturb the process) would be wise enough to have given up the goal of controlling other people.

One last general observation. I've seen the term "controlling for" being used inappropriately on the net. If your enemy trips and falls down, you oughtn't say you're happy because you're controlling for misery for your enemy. Wishing

and controlling aren't the same thing. If you're controlling for A, then you're acting in a way designed to bring A to some reference level A*. You're not just hoping A will turn into A* all by itself. If you mean hoping for or wishing for, say hoping or wishing, not controlling. Let's not dilute our technical language.

Best to all, Bill P.

Date: Tue Sep 15, 1992 10:09 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: Gary (Ems)
MBX: G-CZIKO@UIUC.EDU
Subject: Starter Document

Gary (direct)

You are certainly doing plenty already. I am not clear what exactly you have in mind with "primary responsibility".

Let us start the ball rolling again by picking up where it was left off.

My original post to you was April 28 9:32 am.

Your edited post was June 11, 11:12 am.

I posted a few nitpick comments the same day @ 3:10pm.

I would think this is a very good document as it stands. I do not think it should be shortened further. (Suspect few netters actually download from Bill Silvert's server. I should try it myself)!

Why don't you go ahead and send this (with a few corrections as you see fit) monthly to the net and to new subscribers. I do not think a renewed creative effort is warranted at this point.

Netters may provide feedback after a few postings. None answered your call in June.

Ask if anyone wants to suggest improvements or try their hand at writing an introductory essay. The effort will clarify the writer's own thoughts if nothing else.

I will always be there. My interest is well established and I will serve as consultant.

Christine and I are now starting to get customers. A second one will solidify in a week. 99% probability. Before long, we will be on the road for two weeks at a time. I will download 500 kb at a time, and will fall way behind the net.

If that scenario is compatible with primary responsibility, fine. But we should still issue the open invitation. Anyhow, primary responsibility seems alien to the kind of distributed leadership we practice on this net. Anyone will speak up when an error signal arises, as I did to create the starter

document when the usenet discussion got underway. You may be the only one with a clear primary responsibility, by virtue of list ownership.

For now, I hope the present document can serve well, with minimum effort for either of us (I have to prepare my seminar and will not volunteer to re-consider anything for at least a month).

Control well, Dag

Date: Tue Sep 15, 1992 10:29 am PST
Subject: turing; perception of model structures

[From Bill Powers (920915.1200)]

Cliff Joslyn (920915) --

Right on the button, as usual. Your second point, about defining intelligence, is the most fundamental in my opinion. If I were constructing an intelligence test, I would make sure that everything I know know to do well would score high, and the things I've never mastered would contribute nothing. Intelligence, of course, means understanding control theory, so Dizzy Gillespie would score low, allowing me to get even for being a lousy jazz musician.

You'll notice that when I discussed the Turing test, I expressed it as trying to find out whether the entity at the other end was a human being or a program.

-

Eric Harnden (920915.1230) --

A lot of the things I've written about the slipperiness of words, about trying to see beneath the words to the perceptions that they mean, were attempts to express what you say so well:

>frankly, i don't think it matters where one puts the 'behavior' label.
>my own approach has always been a cybernetic one, and the shifting of a
>label doesn't change any of my (mental or experimental) model structures.

The model itself is the perception that matters, isn't it? I've always claimed that the way to prove you understand control theory is to explain it to three different people without using any of the same words for what the system does. When you have a picture of the organization in your mind, it really doesn't matter how you describe it.

For convenience, I think it's worthwhile differentiating actions from consequences, because there are some phenomena relating them that are interesting. When the controlled consequence is disturbed, assuming the reference level remains the same, the action HAS to change in just the way required to counteract the disturbance. It's in that sense that the disturbance determines the action. But this is very much an AS IF situation; you can't explain this relationship without grasping the way the whole loop behaves. The appearance is that the disturbance causes the action; that's what led to S-R theory.

Best, Bill P.

Date: Tue Sep 15, 1992 1:00 pm PST
Subject: Influence/purposive influence/controlling others

From Greg Williams (920915 - 2) >Bill Powers (920915.1000)

>It often seems to me that you are lumping together "having an effect"
>with "having a particular intended effect."

I think the confusion arises largely because TWO arguments are going on in parallel here: one about whether purposive influence is possible and, if so, common and significant; and one about whether organisms in some important sense are completely self-determining (not counting heredity). There are common threads in both arguments, but they ARE different. In the former argument, I am mainly concerned with purposive influence; in the latter, with more general sorts of environmental influence (including influence due to non-living parts of the environment).

>But simply showing that one person's action results in a change of action by
>another is not sufficient to show purposeful control, even if the person doing
>the disturbing wanted to see the action that the other took, even if that
>action satisfied a goal in the disturber. That wanting, that goal, is
>completely irrelevant -- it's just a lucky guess -- unless the controller can
>VARY the disturbance and KEEP the other's action in the form desired, either
>for a long time or on repeated occasions.

I agree. I am claiming that purposive influence as you have just defined it is common and results in important effects as judged by both influencer and influencee. I do NOT claim that it ALWAYS works. That depends not on luck, but on the influencer's having an adequate model of the influencee's controlling, so that the influencer is able to keep perceiving what he/she wants to perceive, as the influencee's actions change.

>You say " The organism doesn't 'work' (thanks, Dennis) autistically!"
>But this means nothing with respect to PURPOSEFUL influence. If an
>organism is in an environment, it must control in order to keep its
>perceptions as it wishes them to be, whether the disturbances that
>interfere were generated purposefully or naturally (that is, not by
>control systems).

Right. Here I was (or thought I was) considering the second argument.

>There are far more influences unrelated to anyone's goals than influences
>aimed at producing a particular intended change in behavior.

Sure.

>The question before us is not whether disturbances of perceptions can
>alter actions, but the extent to which an external agent can use
>disturbances in a systematic way to make other people behave to suit
>the disturber's goals. You claim that this extent is far greater and

>far more significant than I am yet willing to admit (absent coercion,
>which I do see as a serious problem). That's what we're arguing about,
>in my opinion.

That's argument 1. Argument 2 is about the significance of environmental influences in general, and here, I also claim that it is "far greater" than you -- I think.

BP>>>So in order to retain control in general, the external agent has to
BP>>>give up control of that action and pick a higher-level aspect of
BP>>>behavior to control.

GW>>No, the external agent just has to be skilled enough to not NEED to
GW>>pick a higher-level aspect.

>That's impossible. If you're controlling at one level, there is no way
>you can keep a higher system from changing the reference signal in the
>system you're taking advantage of.

You don't HAVE to, if you have a good model of the influencee. The kids WANT good-tasting food. If they don't WANT to change that reference signal, Pat will never need to worry about the possibility. She doesn't MAKE them not change the signal, she just guesses that they won't want to. She could be wrong, but hasn't been to date.

Pat implores me at this point to make the point that most day-to-day interactions involve purposive influence which works without the influencer having a very sophisticated, detailed model of the influencee. Simple instances like "pass the salt, please." Give-and-take. The success of the influence is presumed -- such behavior is "civilized." Won't work if the influencee doesn't know English, true. But it works almost always, otherwise.

>It's impossible from another standpoint, too. To be "skillful enough" to
>avoid disturbing any variable of the same or higher level that matters to the
>controllee, the controller would have to understand the controllee in complete
>detail, and also predict every independent disturbance that is going to
>happen. I don't think this will ever happen, but more to the point, I don't
>think it is likely to HAVE HAPPENED.

Perhaps we have reached an impasse. This is an empirical issue: how detailed a model of someone's controlling is necessary for, and how important are independent disturbances in interfering with, purposive influence. I look at the evidence and say, "in many cases, not much detail, not very important." You -- looking at the evidence? sometimes I wonder -- say, "always, too much detail, overwhelmingly important." I can cite cases from family life, education, advertising, politics, crime, etc., etc., all day and I doubt you'll change your opinion. (In fact, I have cited several examples to no avail.) From the other direction, your theoretical arguments strike me as unpersuasive. So, where to from here? It certainly makes no sense to go on as we have, although I suspect that some in our (rather deafeningly silent) audience have been swayed to a degree one way or another by what we've said.

>And lastly, I think that anyone with enough understanding of human nature to
>know another person's hierarchy in the necessary detail (not to mention

>predicting all events in the natural world that could disturb the process)
>would be wise enough to have given up the goal of controlling other people.

Ah. Now I see that you are STILL laboring under the misconception that purposive influence has "the goal of controlling other people." That would be impossible, Bill. One can control another person only by employing physical force (or threat of same), and we've been disallowing that in our discussion of the last several posts. A purposive influencer INFLUENCES an influencee so that the influencer can perceive what he wants to perceive. PCT says that the best way to try to do this is to do The Test and thereby figure out what the influencee wants, which, with the appropriate context provided by the influencer, will result in the influencer perceiving what he/she wants to perceive. (Sometimes you'll figure out that nothing the influencee wants has much chance of giving the result you want, and you give up on purposive influence -- and THEN maybe pick up a gun, send the kid to "experts," or just walk away and give up your goal). The purposive influencer who is aware of the PCT basis of what he/she is attempting HAS INDEED GIVEN UP THE GOAL OF CONTROLLING OTHER PEOPLE. In fact, it could be argued (on the basis of pervasive ideologies, including the Judeo-Christian heritage) that purposive influence is an ethically superior alternative to control of others.

>One last general observation. I've seen the term "controlling for"
>being used inappropriately on the net. If your enemy trips and falls
>down, you oughtn't say you're happy because you're controlling for
>misery for your enemy. Wishing and controlling aren't the same thing.
>If you're controlling for A, then you're acting in a way designed to
>bring A to some reference level A*. You're not just hoping A will turn
>into A* all by itself. If you mean hoping for or wishing for, say
>hoping or wishing, not controlling. Let's not dilute our technical
>language.

I agree.

Best, Greg

Date: Tue Sep 15, 1992 4:08 pm PST
Subject: Interactionism

[From Wayne Hershberger 920915]

Apparently, I failed to get this on the net when I posted it 920912 so I am trying again. I appologize for being 2 or 3 threads late.

Tom Bourbon:

I am sorry to hear your funding has been cut. What about your contacts in France this summer? Any prospects there?

RE: INTERACTIONISM

Penni Sibun, Tom Bourbon, Martin Taylor, Avery Andrews, Rick Marken, Bill Powers: (Please read this at least twice; it took me so long to write.)

I am inclined to agree with Martin Taylor that Agre's idea of "leaning on the environment" is not necessarily inconsistent with Bill Powers' hierarchical control model, because many of the control loops comprising Bill's model have an environmental part. However, this environmental part of Bill's model is a conceptual realization expressed in a language we call physics. It is not the environment that we experience perceptually (e.g., perceptually, the sun rises; conceptually, the earth rotates).

The perceptual environment is a realization occasioned by a perceptual process that we conceive or conceptually realize (e.g., Bill's model) as involving physiology and physics, but this is not to say that the perceptual environment is the offspring of that conceptualization--for instance, consider this nonsense: when a conceptualization is revised (say, by Copernicus), which conceptualization gave birth to those perceptualizations (i.e., sunrises) that preceded the revision, the original or the revised version? Then, what about the perceptualizations that came later?

We may say that a perceiver realizes a perceptual environment, but not that s/he perceives a realized environment. That is, s/he does not represent (re-present) to her/himself a Real environment that has been realized a priori in the form of Platonic ideals, or things in themselves, or even scientific conceptualizations--and certainly not the naive conceptualizations (i.e., perceptualizations) of naive realism where the perceived environment is considered to be a re-presentation of itself (cf., Agre's "the world is its own best representation"). This means that a perceiver can not be said to lean on a realized perceptual environment.

What our perceptual and conceptual realizations both seem to lean on (albeit, like a house of cards) is the immanent order that those constructs themselves realize or model. That is, although our perceptual and conceptual realizations appear to be inventions, the immanent order which they realize is not; rather, it appears to be "given."

Fundamentally, perception is a process involving realization not representation. Therefore, perceptual theorists who concern themselves with the representation of reality seem to me (and to Agre?) to be wasting their time, because they are ignoring the fundamental realization process and addressing the superfluous representation process. As for Bill Powers' HCT, it seems clear to me that the input functions in his theory comprise the core of a perceptual- realization process, not a representational process (e.g., Bill has observed that the sourness of lemonade is a perceptual realization that is not to be found in the conceptual world of physics).

Although I sympathize fully with Agre's scoffing about the need to re-present realizations of the world, he does not seem to be addressing the realization process itself, and, hence, I can understand Tom's, Rick's and Bill's reluctance to tender a blanket endorsement of Agre's notion of "leaning on the environment."

Martin Taylor: (re: Martin Taylor 920831 1330)

I believe that the question you are raising about whether the controlled variable is an environmental variable or a perception of an environmental variable is ambiguous for precisely the same reason that Agre's "leaning on the environment" is problematic. That is, what I said above applies here as

well. The answer to your question depends upon the realization you are talking about. When you conceive of my driving (steering an auto) as the control of one of your conceptual realizations (i.e., distinguishing between physiology and physics), then the controlled variable is physiological rather than physical (i.e., it is not an environmental variable). However, since my driving necessarily involves the control of one of my perceptual realizations, and since that controlled variable is in the environmental part of our common perceptual realization (I trust that you do not perceive the car as being in my body any more than I do) the answer to your question is: in the environment. (This ambiguity is what accounts for "the 3 stages of Satori.") If two individuals share a common perceptual realization of the world, it should be possible for them to control the same environmental thing--at least in the perceptual sense of the term. Is that part of what you are driving at?

Martin, could you repeat the reference to Georgopoulos?

Rick Marken:

I think Mary Henle may have published the sort of criticism of Behaviorism you are looking for--it would have been many years (decades?) ago. Also, I seem to recall Skinner himself addressing the matter and explaining his own attempts to control others by controlling their environments; of course, he said that he finds it reinforcing to do so. Daniel Dennett lampooned this use of the term, reinforce, in a critique titled "Skinner Skinned," reprinted in Dennett's book, *Brainstorms*?

I second Gary Cziko's great idea that you put your computer demos in a menu driven package.

Bill Powers: A belated happy 66th birthday (920829) to you!

Warm regards, Wayne

Date: Tue Sep 15, 1992 9:41 pm PST
Subject: Purposive/nonpurposive influence

[From Bill Powers (920915.2000)]

Greg Williams (920915-2) --

>There are common threads in both arguments, but they ARE different. In
>the former argument, I am mainly concerned with purposive influence; in
>the latter, with more general sorts of environmental influence
>(including influence due to non-living parts of the environment).

Well, good. Yes, these are different arguments. Let's dispose of one of them so we can get to the other.

With respect to "more general sorts" of environment influence, I think we need to draw a sharp line separating non-purposive from purposive influence. Purposive influence is not just a subset of non-purposive influence: it is a totally different phenomenon. Let's put it aside for a while.

Non-purposive influence results, in the final analysis, from the non-sentient environment applying undirected and goalless disturbances to variables that are critical to a living system's survival. These disturbances are not aimed specifically at an organism; they fall equally on any object or process, living or nonliving, that gets in their way. They are simply the normal unfolding of physical processes as entropy winds up toward its maximum. As the environment does not care whether such influences have any effect or not, the organism can win; but it can win only if it is predisposed to acquiring control systems, or comes into individual existence already controlling (as most of the organisms on this Earth do).

Organisms must learn to control the variables important to them by learning how to act in opposition to these undirected disturbances. They must learn what actions are needed to prevent disturbances from affecting the variables critical to their existence, which means simply that they must learn to affect those variables themselves. That's the genesis of control: learning what actions have an effect on critical variables, then building organized systems that can act on those variables to maintain them in specific states. The organism does not need to know the source of disturbances; it acts directly on the consequence that matters to it. Learning that connection closes the loop and creates a control system.

Clearly, it's the environment that determines what the organism must learn, because the effects of the organism's actions on its own critical variables is mediated by the environment. But the mere fact that the environment makes learning necessary and possible does not mean that the environment guides learning. The environment doesn't care; it makes no difference to the environment whether the organism learns or not. Whether the organism learns or not depends on the existence inside it of a mechanism for learning, which has to predate the first occasion on which it's needed.

This means that this mechanism for learning can't take advantage of any of the organized methods of control that eventually will result from its actions. Given an error signal, it has no a priori way of determining what action will have an effect on the error. Neither does it have any sophisticated hierarchy of perceptions and computations to aid it; that ultimate product of its action lies far in the future. As far as I can tell, the MINIMUM machinery needed to acquire control from this primitive start is an error-driven output that affects action in a random way.

With respect to food-seeking, that is as far as the evolution of *E. coli* got. Every time it seeks food (or avoids noxious substances) it has to reorganize. It starts from scratch every time it tumbles. This, apparently, is all it needs to find enough food and to avoid poisons.

What lies behind *E. coli*'s behavior? It is an inherited ability to sense substances that are good for it and substances that are bad for it. Each *E. coli* individual comes into existence knowing these things. It knows what to sense (has the equipment for sensing) and it knows whether the time rate of change of each input signal should be positive or negative. The current environment did not teach it that: evolution did. In fact, the current environment never does teach *E. coli* anything about seeking food and avoiding poisons. *E. coli* never learns to control these variables any better during its whole life as an individual.

Yet we can certainly say that the environment influences E. coli. What it influences is the direction in which E. coli, on the average, progresses through 3-space in relation to sources of diffusing molecules. From an outside point of view, we can use substances of various kinds to push E. coli around in any way we please. We have, it would seem, total control over E. coli. Even the nonpurposive environment has total determinism of its actions. E. coli must move according to where the molecules are coming from, every random shift in the locations and concentrations of various substances rotting away causing E. coli's path to change.

But have any of these influences actually changed E. coli in any way? No. Through all of these effects on its direction of locomotion, E. coli continues to maintain a positive time rate of change of beneficial substances at its sensors, and a negative time rate of change of harmful ones. From inside E. coli, these external effects, its own movements in space, don't even exist.

And any external influence that DID materially affect, for any length of time, the variables E. coli senses and controls -- would kill it.

E. coli can't use its reorganizing capacity to build control systems. But higher organisms can. The particular systems they build depend on the environment that happens to exist outside. But those nonpurposive environments still don't care what organized systems result; whether the result is to seek or to avoid creating any particular state of the environment. The criterion for what is learned is still the inherited one: whatever is learned, it must preserve the critical variables in their inherited reference states. To an external observer, it may seem that the environment has enormous effects on the organism, because the organism behaves differently with every change in the environment. But the net result of all these apparent effects is NO effect -- on what matters to the organism.

Of course the organism may fail to deal with the challenges of natural disturbances. If so, it will die. Then the environment can't even influence its behavior. Its behavior has stopped. No external event can get it going again. Only autonomous systems can irrevocably die.

----- For
well over 300 years, science has looked at all the things organisms do and all the environmental events that seem to cause these doings. Common sense has gone along with science, or perhaps it was the other way around. Our social interactions are perceived and understood from the same point of view. Both science and common sense have become obsessed with what the world can do to organisms, from evolution to everyday stimuli. Most of our concerns revolve around this concept of behavior-as-effect. When we want to change our behavior or someone else's, we look for something in the world around us to change. It's impossible to overestimate the pervasive effects of this mind-set on the way each individual sees and understands all living systems, including the observer.

It is very difficult to find that point of view from which it is obvious that the environment's influences on an organism are side-effects of what is actually going on. Even the behavior of organisms is a side-effect, misleading in the extreme. What we see organisms doing is of only passing interest; it will change with every change in the environment. All that matters is what the organism is accomplishing, and to know what that is, we will have to become

very much better acquainted with organisms than anyone is today -- and from a radically different point of view.

I don't know if that "disposes" of non-purposive influences, but it's a start.

Purposive influence is entirely different from non-purposive influence, because the causal influence is no longer of interest. The causal influence is just like any other; its effect on the organism is the same as if non-purposive sources were generating it. What the influencer is concerned with is not influence-as-cause, but influence- as-accomplished-result, as effect. Not "applying an influence" but "having an influence." These two meanings of influence, which our language mashes together into a single fuzzy concept, are different phenomena.

When we are talking about influence-as-effect, only the reasons for which influence-as-cause is applied matter. An influencer who wants to have a particular effect on another person must attend to the effect, not the cause. Because the environment, including the total system of the person to be affected, is always in flux, producing a given disturbing action will never have the same effect twice in a row. So there can be no set formula of action for having a desired effect. All that can be done is to monitor the aspect of the other's behavior that you want to affect and vary your actions until you see what you want or fail.

Having an effect is most easily done when the effect you want to see is of no consequence to the influenced person. You don't have to understand control theory to find such effects. All you have to do is experiment and remember. Some kinds of actions will elicit particular actions from the other person. You may have to vary your action under different circumstances to continue having the same effect, but you can learn that, too, by experimenting. Applying the Test is a perfectly natural procedure that people use all the time. They just don't put their conclusions into a theoretical framework, so they usually miss seeing what the other person is controlling for. What they see is action and reaction. That is quite valid as an empirical observation, even if it does give a false picture of cause and effect.

This is done all the time. Almost invariably, it leads to conflict. It is rare, I think, that any protracted use of this method on an individual can avoid creating situations in which the controllee is caused an inconvenience or worse by these attempts at control. The seething violence that we see all around the world is the direct result of failures of this method of controlling people. Most of these attempts at control are done because of what the controller sees as high-minded reasons, for the good of the other or of the society. Of course there are other motives too, but I would guess that most of the violence-producing control is done in service of an ideal. But what causes trouble is that the controller's high-minded goals are not those of the controllee.

What makes this method of control fail is that it is invariably done for the good of the controller, not for the good of the controlled. It is done to satisfy the controller's concept of the right way for other people to behave, the best foods for them to eat, the substances they should avoid, the way they should interact with others, what they should do for a living, how they should dress and speak and appear, what they should know. The controllers are trying

to arrange the world and the people in it to satisfy their picture of how things should be.

Things can be rearranged without any problem. But not people.

Theoretically, it's possible for one person with unlimited resources and time to learn enough about another person's control structure to elicit any desired behavior from that person without inducing errors that would create resistance. The controller, in effect, would become the means by which all the controllee's wants and needs are satisfied. So in theory, half of the human race could devote its time to controlling the other half's actions. The controlled half would find that all of its goals were met while each person performed an unending series of inexplicable acts that had no bearing on their own goals. Their actions would satisfy the whims, prejudices, beliefs, and convictions of their controllers, while the controllers made sure that the controllees' goals were not violated, whatever the effort or expense.

This is a totally unrealistic scenario. It has never been even remotely approached. What has happened is that every attempt at controlling another eventually arrives at the point where avoiding inconvenience to the other person is just too inconvenient. The controller, after all, has a life to live, too. Unending solicitude for the welfare of others is impractical, not to mention that it demands more knowledge than anyone has.

>Pat implores me at this point to make the point that most day-to-day
>interactions involve purposive influence which works without the influencer
>having a very sophisticated, detailed model of the influencee. Simple
>instances like "pass the salt, please." Give-and- take. The success of the
>influence is presumed -- such behavior is "civilized." Won't work if the
>influencee doesn't know English, true. But it works almost always, otherwise.

Yes. But no model at all is needed for this. All you need to know is that if you ask someone for something trivial, you will probably get it. It isn't necessary to realize that the person is controlling for something, and that NOT complying with the request would cause that person an error. You may think that your request was what caused the other to pass the salt. You may come to think that the request should be SUFFICIENT to get you the salt. That's the pervasive mind-set I'm talking about. All that the request did was set up a situation. The salt would not have been passed if the other didn't have a sense of error until it was passed. The result you see, which you interpret as an influence and its effect, was really generated by the other in order to prevent an error, not in order that you get the salt.

This will not work for requests that demand anything important of the influencee. Please change your job. Please hand me all your money. Please cut your hair off. Please speak respectfully to me. Please let me ahead of you in line. Please don't tell my insurance company that I got drunk and ran into your car. Civilized give-and-take works only for trivial things, things that are really optional, that matter only in a sort of hypothetical way, like manners.

So how do we take care of our young, get along with our friends, teach people things that will be useful to them, help people in trouble, persuade others to do things our way, all without starting down that road that ends in violent antagonism and chaos? I think there are ways. I think that PCT can help us

find them. But to find those ways we have to break with the old mind-set, and we have to stop looking for justifications of the way we're already doing things. The way we're already doing things, no matter how nice and logical and right it may seem, has given us the world we live in, which is approaching a state of total failure. This is why I have so little motive to establish links between PCT and more traditional approaches. My attitude toward the traditional approaches is that the smartest people in the world have tried them, and look what it got us.

Best, Bill P.

Date: Tue Sep 15, 1992 11:42 pm PST
From: Jackson
EMS: ATTMail / MCI ID: 414-0940
MBX: GI=Ray
MBX: IN=RL
MBX: DDA=ID=rljackson

TO: * Dag Forssell / MCI ID: 474-2580

Application message id: MAC-1.4-306690-rljackson-98
Grade of Delivery: Normal

Subject: Received Your Packet 9/15

Dag,

Many thanks for sending the Purposeful Leadership packet; although I've only had a little time, I listened to the tape in my truck and read the literature while hiding out in an elevator. Your interpretation of PCT in the workplace is exceptional, and it's sorely needed in America today. This is exactly what I've been striving to do in the factory I work at (I'm kind of a one-man-show: in-house organizational development, training and education).

For the longest time, I felt like I was the only person in the world who viewed manufacturing environments as collections of control systems to be integrated; I really thought that I was the only one doing this kind of work. At that point, I fellowshiped with Tom Hancock and he showed me the reason to my madness (PCT). Tom then introduced me to Ed Ford, who further enlightened my budding understanding of PCT. Both Tom and Ed encouraged me to join the CSG and sign on to the net, and now (although I'm still a PCT intellectual midget) I've found someone (you) who personifies the combination of what I do and what I believe; but I still lack the knowledge to be truly effective.

You've done an EXCELLENT job with "Purposeful Leadership". I want to give you some specific feedback in the future, but your approach is understandable and gripping. It's so hard to get people to want to break their paradigm in the first place without giving false expectations, but you certainly come off as promising, yet realistic.

Do you find it difficult to cover your one-day modules in one day? One of my specialties is in course construction and delivery, so that was one of my questions.

I've worked so very hard in the area of workplace satisfaction; I'm an outsider (background in education) in the manufacturing environment, but I knew it COULD be a better place...fortunately, my current paradigm is closely related to the PCT paradigm: in other words, it's easy for me to 'see' things, because I already believe them. I'm turning to PCT (and you) to help me define, explain, and teach it to others.

I can't overstate how impressed I was with your material...you probably know how many corporate ed packages I have to review and assess; and generally, I end up constructing my own. I hope this is a beginning of a long and fruitful relationship for both of us, as I hope to do whatever I can to support the dynamic and critical work you are doing.

Please stay in touch and let me know if there's anything I can do for you. By the way, I'm thinking about doing some research this spring on PCT in the workplace...any suggestions?

Thanks, and Best Regards,

Ray

Ray L. Jackson 602-963-6474
3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson

Date: Tue Sep 15, 1992 11:46 pm PST
Subject: Dag and The Phenomenon of Control

Hello All!

Rec'd Dag's packet today; he's done some tremendous work here. In my world -- and my world is organizational education and renewal -- there is absolutely nothing like it. It crystallizes all that is positive in current corporate training, and presents it (through PCT) in a single, understandable vehicle.

Please continue to support Dag and his work -- despite the sense of the ideas or the critical degree that they're necessary, it's incredible how many people turn a deaf ear (or a blind eye?) to a new paradigm. Just ask Ed Ford or Jim Soldani.

Corporations across the country are looking for the 'magic' of empowerment; they say, how do we "empower" people? They don't realize that the magic lies in seeing each person as an individual control system...that understanding empowerment means understanding PCT. Real leaders (the minority) can sense that...but, it's unfortunate that learning a new mindset is too much work for the others (the majority) who are looking for the 'quick fix'.

Best to All,

Ray

Ray L. Jackson 602-963-6474

3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson

Date: Wed Sep 16, 1992 3:55 am PST
Subject: Influence (not cause) for optimism

From Greg Williams (920916) >Bill Powers (920915.2000)

[on influence in general]

>Clearly, it's the environment that determines what the organism must
>learn, because the effects of the organism's actions on its own
>critical variables is mediated by the environment.

So the environment "determines" what the organism must learn? I thought you earlier said that the organism (ALONE) determined it. If you really mean what you say here (and I would be happy with your broadening it to "the environment and the organism co-determine" what the organism must learn), I'm happy. This is truly a Big Step away from organismism.

>Yet we can certainly say that the environment influences E. coli.

That's what I've been saying PCT says! And that provides a path to bridge-building to the ideology of environmentalism, eventually (I propose) resulting in a new PCT-based ideology of NEITHER radical environmentalism NOR radical autonomism, but a middle road which is supported by common sense notions as well as PCT science.

>But have any of these influences actually changed E. coli in any way? No.

Oh, come on. They haven't influenced SOME of properties of E. coli, but they HAVE influenced others. That's OBVIOUS. You can argue, of course, about which properties are more "important" than others from various standpoints, but your statement as it stands is simply incorrect.

>Of course the organism may fail to deal with the challenges of natural
>disturbances. If so, it will die. Then the environment can't even influence
its
>behavior. Its behavior has stopped. No external event can get it going again.

That, of course, is irrelevant to the idea that even an organism's non-living environment has significant influences on the LIVING organism -- significant enough to sometimes (temporarily or even permanently) prevent control by the organism.

>Both science and common sense have become obsessed with what the world can do
>to organisms, from evolution to everyday stimuli.

I agree. But I think the solution to this problem isn't swinging to the opposite ideology of radical individualism.

[on purposive influence]

>Because the environment, including the total system of the person to be
>affected, is always in flux, producing a given disturbing action will never
>have the same effect twice in a row. So there can be no set formula of action
>for having a desired effect.

Often, the "flux" isn't very important, and so actions needn't be altered very much to achieve the same purposive influence repeatedly.

>All that can be done is to monitor the aspect of the other's behavior that you
want to affect and vary your actions until you see what you want or fail.

Before that, you can try to understand how the other controls; with a good -- not necessarily detailed -- model of the others' controlling, it won't be such a difficult, trial-and-error process as you make it out to be.

>Having an effect is most easily done when the effect you want to see
>is of no consequence to the influenced person.

Actually, many times it would be easier if the influencee WANTS (a lot) what you offer!

>You don't have to understand control theory to find such effects. All you
>have to do is experiment and remember.

PCT provides a theoretical underpinning and, I suspect, as we get into greater details (of HPCT), might provide practical tips to enhance the success of purposive influence...

>Applying the Test is a perfectly natural procedure that people use all the
>time. They just don't put their conclusions into a theoretical framework, so
>they usually miss seeing what the other person is controlling for. What they
>see is action and reaction. That is quite valid as an empirical observation,
>even if it does give a false picture of cause and effect.

... and PCT could help to clear up such misunderstandings...

>This is done all the time. Almost invariably, it leads to conflict.

... and PCT could help prevent conflict by showing the way toward improved success in human interactions WITHOUT conflict.

>It is rare, I think, that any protracted use of this method on an individual
>can avoid creating situations in which the controllee is caused an
>inconvenience or worse by these attempts at control. The seething violence
>that we see all around the world is the direct result of failures of this
>method of controlling people.

To the contrary, I think no-conflict, successful instances of purposive influence are the rule, not the exception. Perhaps violence is most often the result of would-be purposive influencers being rather poor at influencing, and then giving up and pulling out the bazooka. An understanding of PCT principles underlying the requirements for SUCCESSFUL purposive influence (first and

foremost, give the influencee what he/she wants) could help to remedy this unfortunate problem of inept purposive influence.

>What makes this method of control fail is that it is invariably done
>for the good of the controller, not for the good of the controlled.

Exactly -- to be successful, it must be done for the good (as perceived by each, respectively) of both. That's what PCT says. When will would-be purposive influencers come to realize this fundamental fact? It will help if they understand PCT basics. (An aside: for POSITIVE purposive influence -- "facilitation" -- to work, the influencee must want what is "offered" by the influencer WITHOUT "catches"; for NEGATIVE purposive influence -- "exploitation" -- to work, the influencee must want what is PUTATIVELY offered, not what is offered IN TOTO. I don't think knowing PCT basics will stop some purposive influencers from trying to be exploiters, but it might help to warn marks about what to watch out for.)

>Theoretically, it's possible for one person with unlimited resources
>and time to learn enough about another person's control structure to
>elicit any desired behavior from that person without inducing errors
>that would create resistance. The controller, in effect, would become
>the means by which all the controllee's wants and needs are satisfied.

This, of course, is NOT what I've been talking about, or advocating.

>This is a totally unrealistic scenario. It has never been even
>remotely approached.

I agree. I'm not talking about putting you in a Skinner box.

>What has happened is that every attempt at controlling another eventually
>arrives at the point where avoiding inconvenience to the other person is just
>too inconvenient.

Impasse time, again. Your "every"s and "all"s are just too much for me to buy, because I don't see them being supported either empirically or by PCT science.

>Civilized give-and-take works only for trivial things, things that are really
>optional, that matter only in a sort of hypothetical way, like manners.

Well, I suppose non-violent family interactions, education, politics, advertising, interactions in the workplace, etc., etc., ARE in the last analysis "optional." What isn't optional is dying if your head is blown off by some guy who escalates from purposive influence to violence because he never heard of PCT.

>So how do we take care of our young, get along with our friends, teach people
>things that will be useful to them, help people in trouble, persuade others
to
>do things our way, all without starting down that road that ends in violent
>antagonism and chaos? I think there are ways. I think that PCT can help us
>find them. But to find those ways we have to break with the old mind-set, and
>we have to stop looking for justifications of the way we're already doing
>things. The way we're already doing things, no matter how nice and logical
and

>right it may seem, has given us the world we live in, which is approaching a
>state of total failure.

I agree completely.

>This is why I have so little motive to establish links between PCT and more
>traditional approaches.

This is why I have so great motive to establish what PCT science says about purposive influence -- a pervasive, significant phenomenon in human interaction which all too often is unsuccessful, setting the stage for violence. This also leads to one final thought: successful purposive influence among humans might be NECESSARY (if violence is to be avoided) simply because it is virtually impossible that everyone who interacts will naturally have the same reference levels -- thus, we each must try to come to an (always partial, but good enough to avoid violence) understanding of what others want, and "give" them some of what they want in order to get some of what we want. All along, I've not been trying to justify purposive influence, but to explain it. Now it looks like our radical individualities might justify promoting positive purposive influence as a way to reduce violent attempts at controlling people. An alternative to the traditional approaches of thinking about purposive influence -- namely, PCT -- could be the key to showing people how to be more successful purposive influencers, so they wouldn't be compelled to use violence.

Optimistically, Greg

Date: Wed Sep 16, 1992 5:09 am PST
Subject: DEMO; DRUGS; BECOMING SOCIAL


CHUCK TUCKER 920916

TOPICS: Anti S-R Demo; "War on Drugs"; Becoming Social

ANTI S-R DEMO

Without modesty I would refer each of you to the paper that appears in CONTINUING THE CONVERSATION (Greg - could you put a copy of that on our list of available documents, please??) that I wrote using statements from Bill demos which appear in BCP (241-244). I merely transformed his demo into several sets of instructions (as I am wont to do on most occasions) so they could be used with a group (e.g., a class or conference). I have found that instead of using a chalkboard with the rubber bands (sometimes I use a single rubber band knotted in the middle) I use 3' X 5' drawing paper (which I attach to the board with masking tape) to make a record of the trace of pencils or markers (each person has a different color tracing instrument) then I take down each paper after the demonstration and after all of the demonstrations I can compare the tracings (you could also reduce the tracings to distribute to the group and also use the tracings in papers you may write about the demonstrations). One way to make a record of the movement of the knot (as well as the tracings) is to video tape the entire demonstration making certain that the center of the frame is the

knot. I would also suggest making the knot a different color from the background and drawing a grid around the target. The target (+ below) could look like this:



then with "slow mo" you could see (and even crudely measure) the movement of the knot. You could also use the video to show to other groups as well improve upon your procedures. (NB: make certain the the lens of the video camera in on the same plane as the knot so as not to introduce errors of "perspective" [like the person in the passenger side of the car always notes that the driver is going slower than the driver notes]). I look forward (of course) to seeing the videos at the next CSG meeting !?!

"WAR ON DRUGS" [Re: Marken 920911.0900]

This so-called "war" is, in my view (and some analysis), designed to show the voters that the members of the Republican party are not as "soft" on crime as are the Democrats (especially the Liberals) who coddle, protect and make their hearts bleed for the criminal. It has very little to do with actually trying to reduce the usage of those substances which most laws make illegal. If there was a serious interest in decreasing the use of these special illegal substances (I keep writing this to note that caffiene, nicotene and most certainly, alcohol [the favorite drug in the USA] are all drugs as are all of the billions of drugs prescribed by MD's and dispensed by pharmacists) one of the obvious ways to do it is to make them legal and dispense them through official means (some of these illegal drugs were once dispensed in this manner). But this would take away the political and moral issue that can be used in this crusade.

Another way is to not try to decrease the use of these drugs at all but suspend all laws (and law enforcement) with regard to them this is not making them legal but rather alegal) and watch the price drop to all time lows. These is the position of Thomas Szasz (remember THE MYTH OF MENTAL ILLNESS) who says we (by which he means our government officials and our hired guns) should not be in the business (literally and figuratively) of trying to alter some behaviors, e.g., killing oneself, using drugs, prostitution, and other classified as "victimless" (which means to me that only the person him/her self, being autonomous, is harmed by such acts). The main problem I see with this position (outside of the obvious one that it would not be accepted by those who want to force others to comply with their moral rules) is that some of the acts may be harmful to others, e.g., prostitution should be arranged so that others aren't robbed, physically harmed or acquire deadly diseases

[by the way, contrary to popular knowledge, only about 8% of the cases of STD can be traced to paid prostitutes; STD is bad for business and professionals know it]. If these sorts of problems could be regulated without force (in most European countries prostitution is regulated by the state as a service provided to the citizens, by comparison in Columbia yesterday the County Council was arguing over a law about whether "nude" dancers should wear "pasties" and "G-strings"; next there will be a law for beachwear!) then this might be the best solution from a PCT perspective.

Unfortunately, most of the disturbances that become transformed into political issues, are very unlikely to be solved to my (and I think, our) satisfaction because political solutions focus on attempting to force people to alter their behavior rather than assisting them to find a way to get them to reorganize their own conduct. So, in an important way you are correct Rick (920911.0990), political "solutions" do not respect the self-regulating and autonomous "nature" of the human being.

BECOMING SOCIAL

[Re: previous posts on manipulation and force; Dennis (920912); WTP 920912.1400; 13.1300; GW 920913 & 14 (last post)]

I think that these recent posts on "autonomy" (which are masterpieces and definite candidates for CLOSED LOOP) have brought the central issue (or problem) of Perceptual Control Theory into sharp and definite focus. The problem was stated by Comte (as paraphrased by McPhail): How can the Individual be both the cause and consequence of Society? By Hobbs, Locke, Jefferson, Simmel, Mead, Cooley, Blumer (I am being expansive here!): How is society possible? A paraphrase of Tucker and Stewart: Given that human beings are self-regulating, self-controlling and autonomous organism, how are societies constructed and maintained?

This problem seems to be beyond solution since we claim that there is no way that A can get "inside" of B or B can get "inside" of A to have any direct effect on each other's operation. There are no chains that link A and B to make it possible for each to effect the other's physical movement. There are no "hardwired intrinsic variables" that can even come close to accounting for the behavior of the organism. There just seems to be no way to solve the problem.

But we know (Bill claims it is obvious) that people do come to regulate and guide themselves in ways such that they produce children, construct and maintain relationships, build shelters, grow and process food, create and maintain networks of friends and acquaintances, learn how to accomplish a vast variety of tasks and all of the activities that each of us observes that we and others are doing everyday. So how is this done by self-regulating, self-controlling and autonomous human beings? How do we become social and remain an individual simultaneously?

The answer to these questions is, of course, a model of social conduct which is based (at least, I believe) on the PCT model of individual conduct. The central idea of the social model was stated by Clark and I in our piece on "purposive collective action" which is that individuals generate similar perceptual signals independent of each other by interpreting their experiences, generate similar perceptual signals interdependently or do so by both taking the actions of another and using it to generate similar perceptual signals. These "ways" are fairly obvious (even somewhat routine and ordinary) and what is not fully modeled is how this "generation process" works. The set of ideas we have found useful we find in the work of Dewey, Mead, James and Blumer and stated most simply in that human being "take" the actions of others and make them their own and use them to control their own conduct. To "take" them and make them their own they must transform them but in most instances their transformation is such that another can "take" the action and make it their own and still find it useful enough to successfully interact with the other. We can see this "taking" most clearly (but perhaps too clearly) in posts where Bill will actually reproduce a statement from Greg and then Greg will reproduce his previous statement, Bill's reply and then make another comment thereto. Mead used the phrases "taking the attitude of the other" and "taking the role of the other" to mean that I control my own conduct by "taking and using" those transformations I make of your conduct. Blumer spoke of "taking others into account" in much the same way.

I can not detail it for you at this point but we do have some very convincing evidence from experiments we have done which shows quite clearly that this "taking the attitude(role) of another" does happen in a systemic manner for a variety of actions. We, of course, don't have all of the evidence that we would like to have but I doubt that we even will so we just keep working on it trying to find out the flaws in the model. Someday I will try to lay this out for all of you but in the meantime we can keep up the conversation on these questions.

CODA

I found the quotes provided by Ray Allis (920910.0845) quite scary and realized for the first time that when behaviorism refers to the inside as a "black box" they mean that there is nothing under the skin of the human organism that is of interest to them if they are to be true to their "model"; there are no internal organs, no brain, no central nervous system - nothing since the stimulus just passes thru without any problem. Am I correct about this? Please point out where there is a serious discussion by a behaviorist about the "intrinsic variables" of the human organism.

Regards, Chuck

Date: Wed Sep 16, 1992 6:11 am PST
Subject: Imagination comments

[From: C. Love (920916.9:30)] To: B.Powers (920914.15:30)

I have been thinking about your imagination model more Bill and now feel that I better understand what your after. This came about after a brainstorming session with a fellow contractor here at DCIEM, with whom is also interested in both my work and PCT, etc. Your imagination model involves three levels. The top level (n+1) is the trigger level. It tells it lower level (n) neighbors(s) to begin imagining because it hasn't been *happy* for sometime! (This could mean persitently bad error, who knows?) Anyways, in turn, these level (n) ECS' (how many - one, two, all of them???) recieve a signal (how is it received - details???) saying that it should begin imagining. Ok, so now these level (n) ECS' tell all the level (n-1) ECS' to short-circuit their error signal (output) into their percept output, i.e., the signals that are treated as percepts by the level (n) ECS'.

This is clear to me now, and I hope to everyone who reads this.

Now what happens. We get this closed loop involving three levels. THE objective: AS I see it, it is to satisfy the ECS's that requested imagination to occur - thhe level n+1 ECS's. Ok, what I want to happen here is for something better to occur as a result of this process. How can that happen unless those percept weights are changed, right?

Which percept weights you ask? The percept weights connected to the ECS' in the (n+1) level. Why these? Because it is these weights that are causing the problems since we must assume that the reference signals are correct, therefore only our perceptions are wrong, and since we haven't been able to *learn* the right perception to see, it must appear that something like conflict exists. This means that the only OTHER way of resolving this is to bias the percept weights. The solution is to change the weights, either more negative or more positive. By the way when I referred above to the reference signals I mean the ones entering into the (n+1) level ECS(s). The refernce signals must be relied on as being correct or this entire argument falls apart.

We can't say that we must change one vble (percepts) according to the references if the references are just as unreliable - we'll end up going in circles and never get anywheres. At the highest level, at least, the references are those conceptual objectives, "match tarhet to finger", which we KNOW to be correct. As such we therefore must assume that all lower level ECS' will also be receiving like objectives in their own way. It is the percepts that must be changed to perceive the right signals to match the references.

Bill, about his almighty ECS. I think you misunderstood me. I don't wish to place lungs into my ECS model. THE ECS model may have all these capacities but the model is set up in a manner that the ECS will only use those methods that it should know about. Remember that I'm doing object orientated programming and all these *methods* are held within my version of the *ECS* object. Many different methods may be there, but may not always be used. I think of it as a person having the ability to do anything but that person only does a subset of those *possible* tasks due to their environmental restrictions, whereas if another person (object instance) is created and placed in a different environment, this person will be able to do a different set of tasks. Do you follow me?

About the Hebbian stuff? This is more just shots in the dark right now. I'm not going to defend any of this just yet. I'm floundering around as much as anyone else in search of good methods to achieve control in such a model.

Suggestions/criticism is always welcome because if you learn from it, it is worth listening to.

Best, Chris. (DCIEM)

Date: Wed Sep 16, 1992 9:51 am PST
From: Dag Forssell / MCI ID: 474-2580
TO: owen (Ems)
MBX: RSCH287E@cl.uh.edu
Subject: Reply

Got your cryptic message of 9/8.

Please introduce yourself. How did you get my name? What do you already know? Have you read introductory messages from CSGnet? What is your profession? What is your organizational affiliation?

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

Date: Wed Sep 16, 1992 10:33 am PST
Subject: A shot at imagination =>

[From: C. Love (920916.13:30)]
[To: Whomever is following these proceedings]

I have come up with a simplistic approach to imagination based primarily on the way Bill proposes imagination to occur, i.e., the three levels of interaction of ECS'. Now this says nothing of how or why imagination is triggered by that level (n+1) ECS, but it does say what to do when it happens.

Let's set up a scenario...
OK, the level (n+1) ECS, we'll call ECS top.
Next there are three lower level (n) ECS', which we'll call ECS M1, ECS M2, and ECS M3, (M=middle).
Lastly there are also three even lower level ECS', which we'll call ECS B1, ECS B2, and ECS B3, (B=bottom level).

So the ECS (B1 to B3) are the ones that are short circuited according to Bill's model, and the ECS' (M1 to M3) are the ones imagining. (For the moment let's assume that those other level (n+1) ECS', which did not request imagination to occur are simply shut off during this process for simplicity's sake.

That's the scenario. Now what happens...

One last point, all weights are continuous value and fully interconnected between layers and upon initialization have small random weighted values between $\{-1,+1\}$.

Let's start by assuming the the error of ECS top is >0 and error is calculated by the rule $E=R-P$, where

R =net weighted sum of references

P =net weighted sum of incoming percepts.

$ECS\ top > 0 \implies R - P > 0 \implies R > P$ (\implies = implies)
 the necessary action here is therefore to
 increase P and decrease R .

Which is better to do? What are the implications? I'm not sure at this point. We must remember WHY we are doing this - it is because we haven't been able to control our percepts entering into ECS top.

The course of action I'm proposing to alleviate this mismatch is the following:

To increase Net Percept of ECS top do the following -

1) Consider an arbitrary percept weight, W_n , entering into ECS top along with the percept signal, P_n .

2) If $W_n * P_n > 0$
 If $W_n > 0$ and $P_n > 0$
 $W_n = W_n + \Delta$
 Else
 $W_n = W_n - \Delta$
 Else
 If $W_n > 0$ and $P_n < 0$
 $W_n = W_n - \Delta$
 Else
 $W_n = W_n + \Delta$

3) What's Delta?

Delta is, I feel, something related to a portion of the error.

I propose the following:

$\Delta = L * ABS(E)$
 L = proportion factor, say 0.2
 E = error of ECS top.

This same sort of process would be applied to the outputs of ECS top also.

I think we would have stability in finding the equilibrium due to the proportion of the error, L , being less than one.

Any comments?

(Keep in mind this is more or less a little idea that popped into my head, but I thought it was worth sharing anyways.)

Chris.

Date: Wed Sep 16, 1992 10:40 am PST
Subject: Realizations; purposive influence=control

[From Bill Powers (920916.0900)]

Last night I tried to be brief. So maybe today I should try to be lengthy. Something is working backward.

Wayne Hershberger (920915) --

Perception as realizations of the immanent order? I think I like that.

>We may say that a perceiver realizes a perceptual environment, but not
>that s/he perceives a realized environment.

In other words, a perception IS a realization of the immanent order? Yes:

>What our perceptual and conceptual realizations both seem to lean on
>(albeit, like a house of cards) is the immanent order that those
>constructs themselves realize or model. That is, although our
>perceptual and conceptual realizations appear to be inventions, the
>immanent order which they realize is not; rather, it appears to be
>"given."

This is very much like my concept, perhaps identical to it. What you call a "realization" of the immanent order I have referred to as a "perceptual function" of the immanent order. There is only one immanent order, but there are many potential functions of it that amount to perceptual realizations. Perceptual realizations may differ from each other in different people, but once they are defined or acquired, the resulting experience of the world depends on the "given" processes and entities of the immanent order in a regular way. Whatever orderliness exists in the immanent order is reflected as a transformed orderliness in the perceptual realization. Does that sound consistent with your idea?

I think that when Agre said that the environment is its own best model, he may have meant that it's easier to use actual outputs and sense their actual effects, letting the immanent order determine how the outputs alter the effect, than it is to try to capture all the immense detail and exactness of the immanent order in a symbolic model and compute the effects. If that's what he meant, I emphatically agree. If I could BUILD a realistic arm, I certainly wouldn't bother modeling it as equations of motion.

Realization, in the sense of "making real," is a better term than representation. Representation always implies that corresponding to the inner representation there is something of like form or nature outside, to be represented. Realization captures the idea that while the result is a perceived real experience, the immanent order is not necessarily organized in just the same way, so the experience doesn't correspond to a mirror image of itself in the immanent order. The concept of representation makes it difficult to understand why people supposedly in the same environment don't come up with identical representations of it.

Thanks for the hipy papi bththdy (if the reference isn't too obscure).

-

Ray Allis (920915) --

I got vicarious joy out of your post to Dag Forssell.

Greg Williams (920916) --

>Clearly, it's the environment that determines what the organism must
>learn, because the effect of the organism's actions on its own
>critical variables is mediated by the environment.

>So the environment "determines" what the organism must learn?

It determines how the organism must learn TO ACT, given what the
organism wants TO ACCOMPLISH.

>I thought you earlier said that the organism (ALONE) determined it.

The organism alone determines what it wants to accomplish (i.e., a higher
level sets the reference signal). If I want to open a door, that's up to me.
It's up to the door whether I have to push it or pull it to get it open. What
I have to learn is which way to exert my effort, not what I want to happen as
a result. I know what I want. I just don't have any a priori information about
the act that will bring it about.

>If you really mean what you say here (and I would be happy with your
>broadening it to "the environment and the organism co-determine" what
>the organism must learn), I'm happy. This is truly a Big Step away
from >organismism.

Avoiding an "ism" is just as limiting as adopting one. All you get from that
is negativism. The implication is that somewhere in the background you have a
better ism.

>>But have any of these influences actually changed E. coli in any way?
>>No.

>Oh, come on. They haven't influenced SOME of properties of E. coli, but
>they HAVE influenced others. That's OBVIOUS.

I don't think they have influenced ANY properties of E. coli. They've only
influenced where it goes in space and the values of some of its variables,
which are not "properties." They haven't altered E. coli's mode of operation
or organization a bit.

Perhaps we're using "property" in different ways. Behavior is not a property
of an organism, as I use the word. The relationship of behavior to other
events may provide evidence of properties, but a property is not a specific
behavior. It's a relationship between variables, not a variable. If a
behavioral response is expressed as an equation, the properties of the
behaving system are described by the coefficients and exponents, not the
variables. The coefficients and exponents remain the same no matter what the

variables are doing. To change the properties of a system you have to change the coefficients and exponents.

>>Both science and common sense have become obsessed with what the world
>>can do to organisms, from evolution to everyday stimuli.

>I agree. But I think the solution to this problem isn't swinging to the
>opposite ideology of radical individualism.

What's "radical" about individualism other than that it denies radical environmentalism? You seem to be looking for a theoretical stance that can encompass enough environmentalism to leave loopholes in control theory, thus keeping it from being "radical" control theory.

>>All that can be done is to monitor the aspect of the other's behavior
>>that you want to affect and vary your actions until you see what you
>>want or fail.

>Before that, you can try to understand how the other controls; with a
>good -- not necessarily detailed -- model of the others' controlling,
>it won't be such a difficult, trial-and-error process as you make it
>out to be.

You don't need ANY model to exert control of another person's behavior. All you need is to find an action that has a reliable effect on the other's behavior. You don't need to know WHY it has that effect. You must monitor the other's behavior, and if it deviates from the state you want to see, alter your actions in ways you can learn will work. Knowing control theory won't make you any better at that.

The test for the controlled variable begins with finding actions that reliably affect the actions of another organism. That's what we do anyway.

If you happen to pick a method of control that avoids conflict in the other person, you simply control that person's actions, no problem. It's true control, with no violence. Of course if you do happen to disturb something else, then you run into resistance. What happens next depends on how badly you need to control what the other is doing.

>>Having an effect is most easily done when the effect you want to see
>>is of no consequence to the influenced person.

>Actually, many times it would be easier if the influencee WANTS (a lot)
>what you offer!

If the influencee already wants what you have to offer, then you're simply disturbing the influencee's control loop, or becoming a means of completing it. If the influencee is already trying to get something, your help will result in a relaxation of the influencee's own efforts as the error gets smaller. That is opposition to your disturbance. You end up doing all the work.

A little old lady is trying to pick up a 50-pound suitcase, but she can lift only 40 pounds. She's 10 pounds short. You help her lift it, and end up

supporting 50 pounds, not 10 pounds. She applies a negative 40 pounds of effort against your positive effort.

Try helping somebody who wants to thread a needle, without just doing it yourself.

>>You don't have to understand control theory to find such effects. All >>you have to do is experiment and remember.

>PCT provides a theoretical underpinning and, I suspect, as we get into >greater details (of HPCT), might provide practical tips to enhance the >success of purposive influence...

Only in the sense of learning how to avoid disturbing other variables. If you pick a method of controlling behavior that doesn't disturb other variables, PCT won't make your control any better. You can already exert perfect control of the other's behavior (the particular action in question) if you like.

>>This is done all the time. Almost invariably, it leads to conflict.

>... and PCT could help prevent conflict by showing the way toward >improved success in human interactions WITHOUT conflict.

The way to improved success is through understanding that purposive influence (control) of other people is not a practical mode of interaction. The basic problem is that people want some sort of lever that removes other people's ability to choose their behavior freely. They can call it anything they like: help, influence, advice, guidance, correction, or facilitation. It's still something THEY can do that makes OTHERS behave more to their liking. What they want is control of others, by definition of control. Calling it by a different word doesn't change what is going on.

>To the contrary, I think no-conflict, successful instances of purposive >influence are the rule, not the exception.

I'm not denying that people interact, nor that most of the time the disturbances that arise are taken care of automatically, with no difficulty. But that is accidental influence, not purposive influence. When influence is purposive, its effects on the other's actions are specifically intended by the influencer. That is, you ask someone to pass the salt not just to get the salt within your reach, but specifically for the satisfaction of seeing the person pick up the salt and hold it toward you. That is the person's action. If all you want is to get the salt, you don't care how that person gets it to you, or whether someone else closer to the salt hands it to you. But when control is what you're after, you specifically want to have an effect on the other's actions, not just on the consequences of the actions.

It's this specific desire to perceive having an effect on another's actions that leads to conflict. It doesn't necessarily lead to conflict, because it's possible to pick an action of someone else that the other person really doesn't care about (made-ya-look). But when that's the case, the only one getting off on the control is the controller; to the controllee it's a matter of indifference. The "influencing" action is just another disturbance, being taken care of automatically.

But controllers aren't satisfied with having effects that are of no importance to anyone else. They aren't satisfied with trivial interactions. They want the rest of the world off their backs. They want others to stop being stupid or acting helpless or getting in the way. As soon as they try for control of something important, they meet resistance. And if they don't understand about HPCT, they'll just increase their efforts and overcome the resistance, and so on to the inevitable end.

>>What makes this method of control fail is that it is invariably done
>>for the good of the controller, not for the good of the controlled.

>Exactly -- to be successful, it must be done for the good (as perceived
>by each, respectively) of both. That's what PCT says.

No, PCT says that the controller can control only for what the controller considers good. The controllee can control only for what the controllee considers good. The trouble starts when the controller claims to know what is good for the controllee. Then the action is based on a reference signal in the wrong head.

>When will would-be purposive influencers come to realize this
>fundamental fact?

It won't work any better when the purposive influencer tries to make the influencee behave in a way that is good for the influencee. The conflict will still come up if any of the influencee's goals are disturbed. A helping disturbance is opposed just as energetically as an opposing one. The little old lady thankfully lets go of the handle as soon as you start helping. She doesn't go on exerting her maximum 40-pound lift. "Good" control of other people's behavior doesn't work any better than "bad" control. Look at the welfare mess.

>What isn't optional is dying if your head is blown off by some guy who
>escalates from purposive influence to violence because he never heard of PCT.

You're making a distinction between purposive influence and control by violence. These are both true control: that is, only the controller gets to determine the state of the controlled variable. Purposive influence, when it works, is simply control. When it doesn't work it's meaningless; it's just one disturbance among many. When it disturbs something important to the controllee, and the controller insists on continuing, it leads to violence. We're just talking about control over a continuum from control of trivial things to control of important things.

I think you're looking for something that doesn't exist: controlling without controlling. The word "influence" has that sound. It doesn't seem to mean grabbing someone and forcing the person to behave. It connotes sort of nudging, hinting, suggesting, pointing, calling attention, gently urging. It implies that the controller uses only the mildest of means and leaves lots of room for variations. And of course it's assumed that the controller is just trying to help, to do what's good for the controllee as well as the controller.

But if influence really worked that way, it wouldn't work. Little nudges will produce little effects which the controllee will shrug off. The controller

will keep the influences small and gentle only until it's seen that the result isn't what the controller wants. Then the efforts will get larger, and larger, and larger, until the effect IS seen. That, after all, is what PURPOSEFUL influence -- control -- is about: to produce the effect you want to see. If a little effort doesn't do the job, and you still want to control, you'll use more effort. In the end, if you think that purposeful influence is for a good end and very important to achieve, you'll end up doing whatever it takes.

From all the things you've written, I get the impression that you think PCT might provide a "nice" way of controlling people's actions, so that important influences can be achieved without the nastiness of direct force, deceit, and so on. I'm afraid that there isn't any such answer in PCT or HPCT. Control is control. If you use gentle means and leave room for variation and choice, all you get is poor control. The issue is not whether to do a poor or a good job of controlling. It's whether controlling anything about other people is the way to build a world we would want to live in.

 I think it might be profitable to turn to how we CAN have influences on the social world that are useful and that could cure some social problems. I'm convinced that NO kind of control of another's behavior, nice or nasty, will work. But we're far from exhausting the possible ways in which people can interact, and in which they can get what they want from or through other people. The real challenge from PCT is to see what is left by way of social interactions if we simply remove the desire to make people behave in any particular way (unless they're directly interfering with us, of course). Would this mean just lying back and becoming the world's doormat? Since that's not what anybody wants, obviously not. Does it mean something like unilateral disarmament? Again, obviously not: the goal is not suicide.

There are actually lots of things left, once we get over the idea that controlling others is going to help. Once we stop devoting so much attention to finding clever nice ways of controlling others, we can put all those loose brain cells to work on finding different approaches to the real problems.

-

6 pages. Well, it's still working backward.

Best to all, Bill P.

Date: Wed Sep 16, 1992 3:02 pm PST
 Subject: Purposive influence... uh... control

From Greg Williams (920916 - 2) >Bill Powers (920916.0900)

>The organism alone determines what it wants to accomplish (i.e., a
 >higher level sets the reference signal).

The current controlling of the organism alone determines what the organism wants currently, yes. That is the PCT-demolition of traditional "free-will." But I claim that the current controlling depends significantly on the history of environmental (both living and non-living) influences on the organism.

>Avoiding an "ism" is just as limiting as adopting one. All you get
 >from that is negativism. The implication is that somewhere in the

>background you have a better ism.

That's correct. I think I do have a better "ism," and have discussed it more than once in previous posts (also what "better" means to me, with regard to spreading interest in PCT science): a middle ground between Skinner's extreme environmentalism and your extreme autonomism. Perhaps I could call it co-determinism, if you need a name.

>I don't think they have influenced ANY properties of E. coli. They've
>only influenced where it goes in space and the values of some of its
>variables, which are not "properties." They haven't altered E. coli's
>mode of operation or organization a bit.

Oh. I see. But higher organisms can reorganize, and even the non-living environment can ("blindly") set them problems they need to solve; if and when they solve them (via reorganization), the nature of their new (post-reorganization) control systems results in part from the environment doing that, since it was the nature of the problem that "selected" from among the reorganization "trials" to give the new control systems. The non-living environment can't INTENTIONALLY select in that way, but (some of) the living environment can.

>What's "radical" about individualism other than that it denies radical
>environmentalism?

It is "radically" extreme, in the opposite direction from radical environmentalism. I think neither ideology is supported by PCT science. Skinner didn't care about what PCT science supports, but I do. Do you? Do you think clinging to an ideology which contradicts theories you believe in is acceptable? If you simply deny radical environmentalism, I think that is perfectly consistent with PCT science. To say that an organism's current controlling is independent of that organism's past environmental influences is not consistent with PCT science.

>You seem to be looking for a theoretical stance that can encompass enough
>environmentalism to leave loopholes in control theory, thus keeping it from
>being "radical" control theory.

No loopholes. Radical PCT science is just fine with me; it's the radical ideologies I have trouble with -- BECAUSE they conflict with PCT science (and empirical observations).

>You don't need ANY model to exert control of another person's
>behavior. All you need is to find an action that has a reliable effect
>on the other's behavior. You don't need to know WHY it has that
>effect.

True, you can get lucky. But, usually, it helps to have a model for predicting how to influence others to achieve what you want.

>Knowing control theory won't make you any better at that.

To the contrary, The Test can be quite useful, I think. Furthermore, I think lots of people have already figured that out, in a folk sort of way. So they

didn't need to know PCT science to get that far; I suspect some would be even more successful if they did know PCT.

>The test for the controlled variable begins with finding actions that
>reliably affect the actions of another organism. That's what we do
>anyway.

You needn't actually present "concrete" disturbances. Asking about what another likes is often sufficient, if done in fair detail. Just observing passively can be sufficient in some cases. The Test is more general in scope than actual trial-and-error.

>If you happen to pick a method of control that avoids conflict in the
>other person, you simply control that person's actions, no problem.
>It's true control, with no violence. Of course if you do happen to
>disturb something else, then you run into resistance.

Exactly what I've been saying all along. But I don't think that successful purposive influence is usually the result of "luck." Rather, the large number of instances of successful purposive influence which I can see happening incessantly implies to me that some folks have pretty good models of other folks' wants. I DON'T see a lot of attempts at trial-and-error purposive influencing.

>What happens next depends on how badly you need to control what the other is
>doing.

No, what happens next depends on the accuracy of the influencer's model of the influencee's controlling. If the influencer makes the mistake of trying to get the influencee to do something he/she doesn't want to do (given his/her "take" of the situation), look out!

BP>>>Having an effect is most easily done when the effect you want to see
BP>>>is of no consequence to the influenced person.

GW>>Actually, many times it would be easier if the influencee WANTS (a lot)
GW>>what you offer!

>If the influencee already wants what you have to offer, then you're
>simply disturbing the influencee's control loop, or becoming a means
>of completing it.

Which includes "shifting" control (at the lower levels) to wants which the influencee would never have had the opportunity to satisfy if the influencer hadn't provided the opportunity.

>If the influencee is already trying to get something, your help will result
in
>a relaxation of the influencee's own efforts as the error gets smaller. That
>is opposition to your disturbance. You end up doing all the work.

Not necessarily. I think you are confusing wants at a high level with wants at lower levels. The influencer can aid (in fact set up the possibility of novel) lower-level wants which are means to an UNCHANGING higher-level want of the influencee. Besides, you can "offer" what someone wants without doing it for

the person. I can't read for my son Cam, but I can provide the opportunity for his reading (he likes car magazines a lot, and I supply them to him).

>The way to improved success is through understanding that purposive
>influence (control) of other people is not a practical mode of
>interaction.

There are many things it might not be, but I can't see that "practical" is one of them. I look around and see it working all the time, repeatedly over long time periods, in everyday life. Now, THAT'S practical!!

>The basic problem is that people want some sort of lever that removes other
>people's ability to choose their behavior freely.

Purposive influence WON'T WORK if it restricts other people's ability to choose (in the PCT sense, NOT the traditional sense, of "choose"), so that their control systems are conflicted. (With the possible exception of the influencer fostering conflict, and hence reorganization, in the influencee to achieve the influencer's aims.)

>I'm not denying that people interact, nor that most of the time the
>disturbances that arise are taken care of automatically, with no
>difficulty. But that is accidental influence, not purposive influence.
>When influence is purposive, its effects on the other's actions are
>specifically intended by the influencer. That is, you ask someone to
>pass the salt not just to get the salt within your reach, but
>specifically for the satisfaction of seeing the person pick up the
>salt and hold it toward you. That is the person's action. If all you
>want is to get the salt, you don't care how that person gets it to
>you, or whether someone else closer to the salt hands it to you. But
>when control is what you're after, you specifically want to have an
>effect on the other's actions, not just on the consequences of the actions.

Not necessarily (in reply to the last sentence above). By and large, the SPECIFIC actions of the influencee don't matter to the influencer, as long as they are resulting in what the influencer wants. Really, Bill, you can answer my posts standing on your head if you want to; I just care about the answers as they appear on my CRT in front of me. And I really don't mind if my son sometimes throws the salt shaker at me -- just as long as it gets to me. (It's plastic, nearly unbreakable.)

>It's this specific desire to perceive having an effect on another's
>actions that leads to conflict.

Yes, I DO see how wanting a person to act "exactly so" could easily lead to conflict. This is NOT the only kind of purposive influence, and it is probably not the most pervasive kind BECAUSE it often fails.

>But controllers aren't satisfied with having effects that are of no
>importance to anyone else. They aren't satisfied with trivial
>interactions. They want the rest of the world off their backs. They
>want others to stop being stupid or acting helpless or getting in the
>way. As soon as they try for control of something important, they meet
>resistance. And if they don't understand about HPCT, they'll just increase
>their efforts and overcome the resistance, and so on to the inevitable end.

This goes WAY too far, in several directions at once. Once again, I see successful purposive influence all around. The successful influencers are those who are satisfied with whatever they can achieve by (most of the time) NOT causing conflicts in the influencees' control systems. Empirically, it looks like that amounts to quite a lot. And it includes (but isn't limited to) giving influencees opportunities to control what is important to the influencees. Yes (in reply to the last sentence above), knowing about PCT indeed might make for more successful purposive influencers, who are concerned about NOT generating conflict in the influencees' control systems. Generating conflict is a time-waster, even if nothing more sinister.

>No, PCT says that the controller can control only for what the
>controller considers good. The controllee can control only for what
>the controllee considers good. The trouble starts when the controller
>claims to know what is good for the controllee. Then the action is
>based on a reference signal in the wrong head.

The purposive influencer makes a model of what the influencee wants. The influencer (via the Test in various forms) essentially ASKS the influencee, "What do you want?" The resulting model certainly could be wrong. But (and this I take to be foundational for PCT science) having even a crude model is better than having no model at all (if only because the crude model can be tested and improved where found wanting).

>It won't work any better when the purposive influencer tries to make the
>influencee behave in a way that is good for the influencee. The conflict will
>still come up if any of the influencee's goals are disturbed.

Having a model of the influencee's control is for the explicit purpose of trying to avoid conflict. Empirically, it appears that "good-enough" models are very common.

>I think you're looking for something that doesn't exist: controlling
>without controlling.

No. I'm looking for what PCT has to say about CONTROLLING WITHOUT CAUSING CONFLICT. You USED to say that "control is impossible without overwhelming physical force (or threat of it)." NOW you say that "control is possible without overwhelming physical force (or threat of it), but not likely or able to persist for long or be important (in SOME sense) or not result in eventual violence." I say that "purposive influence (control, if you like) is possible without overwhelming physical force (or threat of it), and often persists indefinitely, has important results as judged by both influencers and influencees, and need not result in eventual violence." Further, I say that the success of purposive influence (or control) can often be enhanced by applying some form of The Test. Finally, I note (for the nth time) that there is a lot of reciprocal purposive influencing going on in everyday life, as well as some quite asymmetric purposeful influencing, and that purposive influence can be judged (afterwards) by the influencee as good, bad, neutral, or I-don't-know.

>But if influence really worked that way, it wouldn't work. Little nudges will
>produce little effects which the controllee will shrug off. The controller will

>keep the influences small and gentle only until it's seen that the result isn't
>what the controller wants. Then the efforts will get larger, and larger, and
>larger, until the effect IS seen. That, after all, is what PURPOSEFUL influence
>-- control -- is about: to produce the effect you want to see. If a little
>effort doesn't do the job, and you still want to control, you'll use more
>effort.

Not if you're wise (should I say "PCT-wise?"), you won't. You'll try a DIFFERENT approach, NOT more of the same. (You might decide to apply The Test some more, first.) You want what you want to see, and you (being PCT-wise) don't want to give rise to conflict in the influencee's control system, since that would surely NOT produce what you want to see. So you alter your influencing in a sophisticated way -- no simple "pedal-to-the-metal" PCT-ignorant approach. That said, there is still the possibility of FOSTERING conflict to achieve purposeful influence -- but I suspect that is much rarer (and harder to accomplish) than garden-variety purposive influence, which requires NOT generating conflict in the influencee's control system to achieve success.

>In the end, if you think that purposeful influence is for a good end and very >important to achieve, you'll end up doing whatever it takes.

"What it takes," says PCT, is not generating conflict in the influencee's control system. If one ends up trying to control via overwhelming physical force because The Test suggests that there is no non-violent way to get what one wants without generating conflict in the other's control system, that fact still does not call into moral question the facts of the matter, as seen from a PCT viewpoint. PCT doesn't say that violent control is "bad"; you and I do.

>I think it might be profitable to turn to how we CAN have influences on the >social world that are useful and that could cure some social problems.

Go for it. I'm all ears.

>There are actually lots of things left, once we get over the idea that >controlling others is going to help.

So let's hear some. If the only good control of others is no control (with or without generation of conflict in the controlled), I can suggest a logistically important place to start. Please tell me how parents of young children can keep their kids from starving or meeting an untimely end due to a household "accident" without purposeful influence. Or does your "no control" philosophy extend only to your adult peers? (Note that I've seen the cat door at your house -- consider the insidious attempt at influencing that poor cat to use the ramp YOU designed to meet HIS desires. :->)

Patiently, Greg

Date: Thu Sep 17, 1992 5:37 am PST
Subject: stella stuff

(eric harnden - 920917.0930)

answers to mr powers' stella questions:

1) variables can be plotted with respect to time, plotted with respect to each other (comparitive mode, they call this), and tabulated. the newest version allows time plots to be constrained as to their time window. the version i have plots the entire time sequence, always. the new version also places time in a column of the table, which my current one does not. tables can be saved and loaded into spreadsheets with a minimum of fuss. plots are automatically scaled on the vertical axis, although this can be overridden.

2) realtime input does not appear to be possible. the newest version does allow one to halt the simulation, change values, and resume, for gaming uses. but no joysticks or mouse tracking kind of stuff.

3) data for table functions must be entered by hand, as far as i can tell. this annoys me, since a large national model might require a fairly large database. this data can either be typed in, or drawn in on a graph. one of the tenets of system dynamics is that really precise input data doesn't affect the dynamics of the system much. certain cusps may take longer to evolve, threshold may rise or fall a little, but the theory is that it's the structure that defines the behavior, much more than the exact values of the parameters. hence, graphic input is deemed appropriate, since getting the shape of the function is considered the critical part. i will withold comment.

4) no user functions. there are alot of really useful math and other functions built in, and again, output can be saved and analyzed with other tools. but you can't link to exotic homebrew tools. that's one of the reasons i like dynamo. it links quite easily to C routines. (one can also call on external data files).

5) it'll run on anything se30 or above. although it is a bit of a pig. smaller machines should not be running much of anyhting in the way of inits or such, or stella will run out of memory quickly. but that's really the only problem. it's fast enough, even on the small machines, and sims are cheap. you can put in your 4meg and not worry about it.

Eric Harnden (Ronin)

Date: Thu Sep 17, 1992 1:52 pm PST
Subject: giving up; Stella

[From Bill Powers (920917.0900)] Greg Williams (920917) --

If anyone is trying to make an ideological point here, it is you, not me. For the past several posts, I've been answering your counterproposals by reasoning as strictly as I could from the model. I've been pointing out that you've made mistakes in applying the model, that you're shifting the definitions of your terms to suit your immediate objectives, and that control theory itself, regardless of ideology, doesn't agree with your claims. None of this has succeeded in shaking your position that the world is full of people purposively and successfully controlling the behavior of other people. You seem to want me to change the model so its conclusions will agree with your interpretations.

I've repeatedly pointed out things you say about what a control system would do are not what it would do, according to the theory. You never come back and say that my analysis was wrong, and that the model would behave as you said, not as I said. Instead you just look for a different way of making your point, as though the strict rightness or wrongness of a particular argument involving the model is of no account. This gives the impression that protecting the conclusion from disturbance is more important than any particular argument by which this is done.

Taking a tip from my own model, I must admit defeat. I can see what you're controlling for, and I can see that arguments based on the model are not going to affect the outcome. Your world is full of people successfully controlling the behavior of other people, and will remain that way. I can't convince you that this is a misinterpretation.

I don't know where to go from here. Answers to all your points are boiling around in my head, but I guess I've stopped believing that they make any sense to you.

eric harnden (920917.0930) --

RE: Stella.

Thanks for all the answers. It does seem that using Stella either for interactive experimenting or for matching models to real behavior would be pretty difficult, perhaps impossible. About the best that could be done would be to read out tables of results from Stella, and write another program to compare them with experimental results. How do the System Dynamics people handle the comparison of model behavior obtained with Stella with real-system behavior?

Would you mind posting once more the name of the guy who's in charge of Stella, and if you have it, an address? I think I'll write to see if he has any interest in PCT. If he could be persuaded to include just user-written functions, a lot of needs would be met.

Best, Bill P.

Date: Thu Sep 17, 1992 2:38 pm PST
Subject: STELLA & Barry Richmond

[from Gary Cziko 920917.1600]

Bill Powers (920917.0900) asks of Eric Harnden (920917.0930):

>Would you mind posting once more the name of the guy who's in charge
>of Stella, and if you have it, an address? I think I'll write to see
>if he has any interest in PCT. If he could be persuaded to include
>just user-written functions, a lot of needs would be met.

Bill, Hugh Petrie recently wrote me about Barry Richmond:

> I THINK I sat next to Barry Richmond

>(according to his card) on a plane trip about a year ago. He showed me STELLA

>in a demo form. He did indeed, seem very PCTish as I recall. Maybe you >could remind him of his trip in first class a year ago with a dean of education

>and see if he wants to join CSG-L. I have the following information.

> Barry Richmond
> High Performance Systems
> 45 Lyme Road
> Hanover, NH 03755
> Fax: 603-643-9502
> phone: 603-643-9636
> AppleLink: X0858

That AppleLink address translated to Internet should be X0858@applelink.apple.com. I sent a test message to this address yesterday but haven't gotten a reply yet. I will let you know if I do.

I was going to try to get discounts on STELLA for CSG members. Bruce Hannon here on my campus gets it for \$65 for his students. I'll gladly let you write to him about this instead of me (of course, if he's good he already knows about you and won't know about me).--Gary

Date: Fri Sep 18, 1992 4:06 am PST

Subject: Giving up: a smokescreen?

From Greg Williams (920918) >Bill Powers (920917.0900)

Two short points:

1. I fail to see how ANY of your "reasoning... from the [PCT] model" contradicts my main point that, while in present time one CANNOT alter what another WANTS, one CAN (sometimes) influence HOW the other acts to satisfy (some of) the other's wants so as to satisfy (some of) one's wants. I believe that strictly follows from the model. The PCT model also says that the chance of success of this kind of influence is enhanced if the would-be influencer uses The Test to make a model of whether the potential influencee would want to act in ways such that the influencer would get what he/she wants IF the influencer provided the opportunity (not otherwise possible) for such actions. I believe that that, too, strictly follows from the model. I have attempted to meet your objections without appealing to ideology, and I deny that there is any need to change the PCT model.

So we agree to disagree, but I don't agree regarding the reasons for the disagreement. I claim that you try to distort and read into the PCT model to suit your own ideology. Ideologically, I believe that such distortion is unfortunately compromising the acceptance of PCT ideas by others.

2. Again I invite you to present alternative means (grounded in PCT science) for individuals to achieve their respective goals via interactions which don't involve purposive influence. Your doing this would show that you aren't giving up simply because you have nothing constructive to offer, regardless of your posturing about my "ideology." I am certainly willing to work with you in exploring ANY implications for social interactions of the PCT model. But I am unwilling to allow ANYONE to run a bluff which makes false accusations simply

because it is a convenient way for them to get out of having to produce what they have promised but can't provide.

If that sounds harsh, ditto your post to my ears. And I'm still all ears.

Soberly, Greg

Regarding differential and difference equation simulators, my life might have just been made easier. The TutSIM folks have announced that they are working on a (nonclunky looking) interactive graphical interface for release this fall (in beta testing now). TutSIM can run fast enough (though no speed demon) on a bare-bones IBM-PC compatible with no math coprocessor from a floppy. With the graphical interface, user-defined blocks, real-time input-output, and ASCII system-definition files (can be passed around via e-mail), I estimate the price will be around \$200 (perhaps less with educational discounts). IBM-compatibles capable of running the program go for ca. \$500 (even less if used) and speed demons go for ca. \$1000.

Nevertheless, perhaps we could do better by rolling our own program. I suspect that STELLA is limiting in numerous pain-in-the-buttocks ways, simply because it came from the rather limited approach of Industrial Dynamics (apparently, Forrester decided early on that management types had to have everything handed to them on a rates-and-accumulators platter -- modeling physiology can get VERY complicated when you have to build up everything from capacitors).

Might somebody be able to photocopy the basic specs for STELLA from ad flyers or the manual and USPS it to G. Williams, 460 Black Lick Rd., Gravel Switch, KY 40328? Thanks!

Greg

Date: Fri Sep 18, 1992 8:57 am PST
Subject: Influence is not control

[From Bill Powers (920918.0800)] Greg Williams (920918) --

I'm sorry about the harsh tone of my last post. I was frustrated. I really like you and Pat a lot, and I always will.

"Purposive influence" is an oxymoron if it doesn't mean simply "control." If you have only an influence on what someone else does, and not control, then you can't predict what the other will do, and you can't vary what you do in order to make the other do any particular thing. In that case it can't be a purposive influence. An "influence" is simply one among many independent variables. It does not have a determining effect on the result. Over the long run the effect of an influence can be detected only statistically, by averaging out the effect of all the other independent variables that act at the same time, also influencing the same outcome.

To convert a mere influence into control, you have to monitor the influenced effect, compare it with a reference level, and adjust your influences whenever

there is a deviation from the effect you desire. Now you compensate for variations in all the other independent variables, not by knowing what they are but simply by monitoring the final effect you want. If you can't monitor the effect you want to influence, and vary your actions to oppose changes in it, then you have only influence, which is necessarily statistical, and not control.

In an example such as feeding children healthful food, you may be influencing their health, but you're not controlling it. You also may not be influencing their health. In principle, according to statistical theories of nutrition, "people" are better off eating some foods than others. If your children are healthy, however, there is no way to tell whether that is the result of what they eat, or the effect of powerful biochemical control systems inside them that can convert a wide variety of inputs into sufficiently useful compounds and reject a wide variety of useless or deleterious inputs without harm, or simply a result of the fact that the food you feed them is as good as any other kind. At best you can say you're not making them unhealthy, for the simple reason that they're healthy and therefore NOTHING is making them unhealthy. Unless you're continuously monitoring their biochemistry and correlating it against variations in their diet, you have no knowledge of the actual effects of their diet. If you feed them only one kind of diet, you have satisfied one of your own reference levels, but have no information on the effects of this diet as opposed to others on your children. You are certainly not controlling their state of health. You won't know whether you're even controlling what they eat until they demand to eat something else. Then you'll find out.

You've cited your nutritional customs several times as an illustration of purposive influence of another. The only purposive influence in this case is on your own perceptions of the effects you hope to have on your children, in relation to your own beliefs about nutrition. Those perceptions you can control, because they're your own. Those effects you can't actually control, you can imagine.

There is no problem with making sure your children eat the things you want to see them eating, because they are too young to find their own food and they can't leave to find a wider variety from which to choose for themselves. You have physical control over them, as your parents had over you.

But you have another advantage, which is that you perceive in ways and at levels where your children are not yet competent and don't feel competent. As long as your children don't have their own goals for what they eat, they will simply accept what they are given and believe what they are told about it. They won't ask to go to MacDonald's, so you won't have to prevent them from wanting and eating Big Macs, or prevent them from having the money and freedom to buy Big Macs if they want to. They may adopt the food they are given as a reference level (they have few other experiences from which to choose, I presume). After a time, when they do begin to state their preferences for food and want them taken seriously, you can shift to a higher level and explain to them your reasons for choosing a particular diet. If they've reached the point where they can understand and appreciate logical reasoning, they will accept your reasoning because it's all they know (I presume you don't call in a representative of a fast food manufacturer to present the other side). And when they begin to try out other modes of reasoning, you can fall back to

principles, and explain how they make sense of some kinds of reasoning but not others.

When you're supplying information, activities, attitudes, and so on at the level where it has an effect, you are the only control system involved at that level. You present a world exactly in the form you want, which the child is only dimly beginning to perceive. You can't control what the child is going to acquire from this new world; it won't be exactly what you think, and it will include many things you were not aware of presenting. The child will often take a disliking to something you're trying to get across, and there's nothing you can do about it, short of force.

When children reach the point where they have adopted their own reasoning, their own principles, and their own system concepts, you have finished raising them and can have no further effects on them that they do not ask for and adopt. They are now your equals and will demand being treated as such -- unless they've developed their own attitudes toward parents that make them go on treating you as they did when they were little.

I think it's clear that the goal structures children actually end up with bear only a general resemblance to those of their parents, despite the parents' best efforts to educate them "properly." Real learning requires reorganization, and the outcome of reorganization is not predictable by an outsider. The criteria for starting and stopping reorganization are internal, unreachable from outside. That's why you aren't much like your parents.

Parents can purposively influence their children as long as they don't interfere with levels of control that have already developed or appeal to levels that are still inconceivable to the child. There's a moving window, a leading edge of development, where children soak up everything they're exposed to (including everything you hope they will have overlooked). What control you can exert is within this moving window. It isn't nearly as much control as you think or wish and it's gone before you feel that you've finished. A long long time later you will look back and see that this is a good thing: that you were not actually as wise as you felt then -- or, on occasions, as stupid as you felt then. You will realize that your children learned a lot more from how you WERE than what you DID, and that the impetus for that learning came from them.

My cat does something that looks like controlling my behavior. Through long interactions, it has come about that the cat jumps up onto the bathroom counter, puts his forefeet down in the sink, and waits. I then come over and turn on the water. The cat has plenty of water in his own dish; he also has a cup with water in it right on the counter. So the cat doesn't need me to give him water. He just likes to drink out of different things in different ways for the fun of it. This amuses me, and I like the sense of communication between us when I turn on the water and he looks up at me and then turns to drink. We obviously get very different things out of my act of turning on the water. But there's also something in common. Two entities, not controlling each other but acting to hold up two ends of something that can't be carried by one end, and each of them knowing it and counting on the other. Doing something together that neither one could do alone.

Let's by all means explore the new TUTsim. But let's also lay a snare for Barry Richmond and Stella, because that brings in the Mac world. Of course you're perfectly free to say you'd rather develop your own system, which would

undoubtedly surpass either of the others -- some day, on a PC. I agree that Stella is designed around Forrester's "stocks and flows" concept, which is limited. But there are ways to get around that bias, as we've found out already. At least Stella seems a promising tool for teaching basic principles (we can write our own user's guide). And I repeat, it's accessible right now to non-programmer Mac users. Best, Bill P.

Date: Fri Sep 18, 1992 10:10 am PST
Subject: Re Chuck Tucker's last post

From Greg Williams (920918 - 2) >CHUCK TUCKER 920916

>Without modesty I would refer each of you to the paper that appears
>in CONTINUING THE CONVERSATION (Greg - could you put a copy of that
>on our list of available documents, please??) that I wrote using
>statements from Bill demos which appear in BCP (241-244).

Chuck's article, titled "Demonstrating Control Theory," appeared in CONTINUING THE CONVERSATION (12), Spring 1988, pp. 16-19. For a copy of CC (12), send \$2 for delivery within the U.S. or \$3 for delivery elsewhere (via air mail) to: HortIdeas, 460 Black Lick Rd., Gravel Switch, KY 40328. U.S. funds only.

>I found the quotes provided by Ray Allis (920910.0845) quite scary
>and realized for the first time that when behaviorism refers to the
>inside as a "black box" they mean that there is nothing under the
>skin of the human organism that is of interest to them if they are
>to be true to their "model"; there are no internal organs, no
>brain, no central nervous system - nothing since the stimulus just
>passes thru without any problem. Am I correct about this?

Partly. At least some (I wouldn't presume to speak for ALL) behaviorists believe that they can accomplish what they want to achieve with their endeavors (sometimes sloganed as "predict and control [others'] behavior") WITHOUT NEEDING TO MODEL ORGANISMS' INNARDS. They DON'T think that there AREN'T any innards, and some admit that sometime in the (far distant, they opine) future, there will be sufficient information about the innards (discovered by physiologists) to PERHAPS aid "prediction and control of behavior." But the behaviorists say trying to model the innards now or in the foreseeable future is mostly just a distraction from the (they claim) highly efficacious methodology which they use, disregarding the innards. I think they are fooling themselves -- are HANDICAPPING themselves -- by refusing to model the innards, but that's another story.

>Please point out where there is a serious discussion by a behaviorist
>about the "intrinsic variables" of the human organism.

In ABOUT BEHAVIORISM (1974), B.F. Skinner says the following:

pp. 51-52 - "Mentalistic terms associated with reinforcers and with the states in which reinforcers are effective make it difficult to spot functional relations. For example, the statement 'The term "aggression" should be restricted to behavior motivated by the wish to injure' is intended to make a useful distinction between behavior which is merely aggressive in form and any part of such behavior which is emitted because it injures another person, but nothing is gained by speaking of the wish to injure or, in particular, of

being motivated by a wish. When the Utilitarians held that pleasure and pain were 'the motives influencing human behavior' they were referring to feelings associated with consequences rather than motives. The experimental analysis of contingencies of reinforcement puts these matters in better order."

pp. 161-162 - "Turning from observed behavior to a fanciful inner world [by the ignorant, untutored masses, I think he wanted to say!] continues unabated. Sometimes it is little more than a linguistic practice. We tend to make nouns of adjectives and verbs and must then find a place for the things the nouns are said to represent. We say that a rope is strong, and before long we are speaking of its strength.... The extraordinary appeal of inner causes and the accompanying neglect of environmental histories and current setting must be due to more than a linguistic practice. I suggest that it has the appeal of the arcane, the occult, the hermetic, the magical..."

pp. 164-165 - "There are, of course, reasons why a fluid flows slowly, and a molecular explanation of viscosity is a step forward. There are physiological reasons why a person behaves in a manner we call cautious; and the physiologist will, we assume, eventually tell us what they are.... The exploration of the emotional and motivational life of the mind has been described as one of the great achievements in the history of human thought, but it is possible that it has been one of the great disasters. In the search for internal explanation, supported by the false sense of cause associated with feelings and introspective observations, mentalism has obscured the environmental antecedents which would have led to a much more effective analysis." [Didn't he EVER talk to a physicist about the utility of modeling underlying mechanisms, as opposed to just describing the observed phenomena and then hoping that the future will be "like" the past?]

pp. 219-220 - "What [behaviorism] says about consciousness is this: (a) Stimulation arising inside the body plays an important part in behavior. (b) The nervous systems through which it is effective evolved because of their role in the internal and external economy of the organism. (c) In the sense in which we say that a person is conscious of his surroundings, he is conscious of states or events in his body; he is under their control as stimuli.... (g) No special kind of mind stuff is assumed...."

Best, Greg

Date: Fri Sep 18, 1992 2:06 pm PST
Subject: I give up, too, Turing, Stella

[From Rick Marken (920918.1330)]

Well, I give up. Psychological Science is not going to publish the "Blind men" paper as a commentary and I'm not going to try to get it (or any other paper) published in any scholarly journal any more. I will only publish stuff in the Journal of Living Systems if it ever exists. The brief comments of the person who reviewed the "Blind men" paper made it clear to me that attempts to get psychologists interested in PCT are just futile -- they just don't see anything wrong with the current state of psychological "science" (as Dag is always pointing out to me). The fact that their data is junk, that they have no working models that match the behavior of real organisms, that their models cannot behave in real environments, that their concepts of the nature of

behavior are based on skewed views of a pervasive phenomenon that they don't notice or understand (control); all this is of no interest. They are happy with their statistical tests; I'm happy with my models. Apparently, there is going to be no dialogue between conventional psychology and PCT -- other than the one going on right here on CSGNet. Uncle. Uncle.

penni sibun (920914.1300) says:

>no. the test doesn't work backwards; it doesn't have anything to work
>backwards to.

If this is true, then the Turing test doesn't seem to "test" much. As you describe it, the test involves judging whether the behavior I see seems intelligent or mindlike or whatever. How can that judgement be wrong? What's the test for? If this is really the Turing test then it would probably be better described as the Turing experience. I just recently had this experience in an airport where there were some Coke cans dancing to the music coming out of a radio. Those cans seemed like they were really bright; they passed the Turing test. Do you have any reason to suspect that they didn't?

>no. when you go on like this, i see so clearly preston's point that
>some people are so locked into the behaviorist/mentalist mindframe
>that they can't talk about anything else. (did you ever get her paper?)

I received and read Preston's paper. It was VERY long. I don't know if I am locked in the behaviorist/mentalist mindframe. I am locked in the mindframe of thinking I want to understand human nature. If the Turing test is actually what you describe then I have no argument with it; I just don't see what the big deal is. The Turing "test" is apparently nothing more than the claim that people can have the experience of perceiving something called "intelligent behavior". Well, the Marken test goes way beyond that -- it says that people can experience "moral behavior", "spiritual behavior", and much much more. I thought Turing (and ai types) were interested in understanding the mechanisms that produce what is called "intelligent behavior" -- something which, most folks agree, is produced by humans. Thus, a measure of the success of the mechanism you build (or postulate) for producing intelligent behavior is whether it produces intelligent behavior. And, since the only standard for intelligent behavior is that which is produced by humans (and described as such), then the way to test the mechanism is to compare it's behavior (blindly) to that of a human. If the behavior of the mechanism is indistinguishable from that of the human, then the mechanism is "artificially intelligent". If this is not what the Turing Test is about then what is it for? How is it relevant to ai?

The Turing test (as I understand it) would be just fine if behavior were what Turing (and everyone else) imagined it to be -- actions and consequences of those actions. But there is reason to believe that behavior is controlled perceptual experience. So looking at actions and their consequences gives a misleading picture of the behavior of the organism. In order to test whether a hypothesized mechanism behaves like the "real" thing you must test to see whether it controls perceptions like the real thing. Thus, the only problem with the Turing test (from a PCT perspective) is that it uses the wrong criterion to evaluate the success of the mechanism proposed as a model of behavior. The Turing test focused on actions and their consequences -- PCT

says look at controlled perceptions. And it is only possible to look at controlled perceptions using the test for the controlled variable.

If there is another way to look at this, that's fine with me. I don't think I want to stay locked in a behaviorist/mentalist mindframe. I would like to try to see things your way. But, quite frankly, I want to do science and to me that means observing phenomena and testing models of those phenomena -- to see if the model acts like the phenomenon. I admit that I am locked in this mindframe and it seems to me that this is the mindframe that you believe I should break. It's hard for me to imagine how I can break it -- even with your help -- because in the small domains in which I have used this mindframe, it has been most successful. But please keep trying to convince me. Maybe we could get more concrete and discuss your perspective in terms of experiments and models.

Chuck Tucker says:

> "WAR ON DRUGS" [Re: Marken 920911.0900]

> If there was a
> serious interest in decreasing the use of these special illegal
> substances ... one of the obvious ways to do it is to
> make them legal and dispense them through official means

It sounds like you're saying that there IS a way to control drug use but the drug war is the wrong way. Maybe I am misreading you but I want to clarify my post about the drug war anyway. My point was not that the drug war is a "poor" way to control drug use. My point was simply that it IS an approach to control drug use (a behavior) -- with the predictable consequences -- ie. conflict. If you tried to control drug use by legalizing drugs you would run into the same problem; conflict. What is important about control of behavior is not the actions that are taken in an effort to produce the intended result (in this case, lower drug use) it is simply having the intention that another person will behave in a particular way -- ie. it is having a reference for a particular perception of the behavior of another person. If a person is trying to control this perception (of another's behavior) then he or she will act in whatever way results in movement of the controlled perception toward the intended level. It is the intention to control behavior that causes the problem -- not the actions that are used in an effort to carry out this control. An observer might judge some of the actions that are used to control a behavior (like drug use) to be very nice -- like giving the user alternative drugs or putting the user in a real nice "rancho relaxo" environment to kick the habit. But as long as it is the reference for drug use of the controller and not that of the controllee that is being forced as the one to be adopted by the controllee, then there is conflict and, possibly, violence.

Re: Stella: I have a copy at work. It looks fine to me. But what I do can be done just fine in pascal or basic or lotus. But I'll fiddle with it.

Best regards Rick

Date: Fri Sep 18, 1992 3:09 pm PST
Subject: not much argument

From Greg Williams (920918 - 3)

>Bill Powers (920918.0800)

>I'm sorry about the harsh tone of my last post. I was frustrated. I
>really like you and Pat a lot, and I always will.

I understand. I was frustrated, too. Perhaps we would be less so in this "lovely argument" -- as you termed it in a private post just a couple of days back -- if we each don't frame it as trying to convince the other of anything, but as trying to see how far we can go regarding control-theory implications for social interactions. So, back to the lovely argument....

>"Purposive influence" is an oxymoron if it doesn't mean simply
>"control." If you have only an influence on what someone else does,
>and not control, then you can't predict what the other will do, and
>you can't vary what you do in order to make the other do any
>particular thing. In that case it can't be a purposive influence.

I agree. Explicitly, influence, if "purposeful," is control.

>An "influence" is simply one among many independent variables. It does not
>have a determining effect on the result. Over the long run the effect of an
>influence can be detected only statistically, by averaging out the effect of
>all the other independent variables that act at the same time, also
>influencing the same outcome.

OK, a non-purposeful influence (i.e., not control) is as you say. And I'll add that the non-living portion of an organism's environment (but, now that I think of it, I don't know whether that includes at least some artifactual non-living feedback control systems -- but that's a detail to be pondered later) can ONLY influence organisms and CANNOT control ANYTHING, simply because that portion of the environment has no reference signals.

>To convert a mere influence into control, you have to monitor the
>influenced effect, compare it with a reference level, and adjust your
>influences whenever there is a deviation from the effect you desire.

Yes.

>Now you compensate for variations in all the other independent
>variables, not by knowing what they are but simply by monitoring the
>final effect you want. If you can't monitor the effect you want to
>influence, and vary your actions to oppose changes in it, then you
>have only influence, which is necessarily statistical, and not
>control.

I agree.

>In an example such as feeding children healthful food, you may be
>influencing their health, but you're not controlling it.

That is correct. What Pat is doing is controlling her perception that the kids eat food which she THINKS is "healthful," according to her definition of the term.

>You also may not be influencing their health.

Right. In controlling her perception that the kids eat food which she considers "healthy," the food she provides might not alter their actual health relative to their health if Pat didn't control that perception.

>In principle, according to statistical theories of nutrition, "people" are
>better off eating some foods than others. If your children are healthy,
>however, there is no way to tell whether that is the result of what they eat,
>or the effect of powerful biochemical control systems inside them that can
>convert a wide variety of inputs into sufficiently useful compounds and reject
>a wide variety of useless or deleterious inputs without harm, or simply a
>result of the fact that the food you feed them is as good as any other kind.

You said it far better than I did above.

>At best you can say you're not making them unhealthy, for the simple reason
>that they're healthy and therefore NOTHING is making them unhealthy.

In fact, I think, Pat actually COULD be making them unhealthy, but "luckily," something else is compensating for that. She really has no way to monitor all of the sources of the kids' health unless -- MAYBE -- she could keep them penned up some way.

>Unless you're continuously monitoring their biochemistry and correlating it
>against variations in their diet, you have no knowledge of the actual effects
>of their diet. If you feed them only one kind of diet, you have satisfied one
>of your own reference levels, but have no information on the effects of this
>diet as opposed to others on your children. You are certainly not controlling
>their state of health. You won't know whether you're even controlling what
>they eat until they demand to eat something else. Then you'll find out.

Right. Pat will know BY DEMONSTRATION rather than just BY GUESS whether she is controlling for her perception that the kids are eating what she considers "healthy" only if there is an error signal in her perception, which she then able to correct.

>You've cited your nutritional customs several times as an illustration
>of purposive influence of another. The only purposive influence in
>this case is on your own perceptions of the effects you hope to have
>on your children, in relation to your own beliefs about nutrition.

Being careful here, I would want to rephrase this and say that Pat GUESSES that she is controlling her own perception of the kids' eating food she considers healthy if she HAS NOT SEEN an error signal in that perception and managed to correct it, but that she has WITNESSED successful control of her own perception of the kids' eating food she considers healthy only when she HAS SEEN an error signal in that perception and managed to correct it. However, assuming that Pat IS successfully controlling her perception of the kids' eating food she considers healthy, she is accomplishing that control by INFLUENCING the kids. Possibly you don't buy this last claim. If so, please explain your problem with it.

>Those perceptions you can control, because they're your own.

Yes.

>Those effects you can't actually control, you can imagine.

At least SOME effects one can't control can be imagined -- there might be some such effects one can't control because one is ignorant or stupid in certain ways.

>There is no problem with making sure your children eat the things you
>want to see them eating, because they are too young to find their own
>food and they can't leave to find a wider variety from which to choose
>for themselves. You have physical control over them, as your parents
>had over you.

There WOULD be a problem if Pat gave our kids food which she considers healthy AND that tastes bad to the kids. Only because they WANT to eat good-tasting (to them) food AND because Pat gives them good-tasting (to them) food is there no problem. Pat has made (occasional) mistakes about whether particular kinds of food are tasty to one or the other kid, and, believe me, then there are PROBLEMS, regardless of our having (some) physical control over them and it being inconvenient for them to get food independently of Pat. If the food tastes VERY bad to them, they'll just not eat, period.

So the crux of what Pat must do to control her perception of the kids' eating food she considers healthy is to give them food which is BOTH "healthy" (to her) AND "tasty" (to the kids). It is her purposive COMBINING of the two food-properties which allows her control of her perception to be successful.

>But you have another advantage, which is that you perceive in ways and
>at levels where your children are not yet competent and don't feel
>competent. As long as your children don't have their own goals for
>what they eat, they will simply accept what they are given and believe
>what they are told about it.

That is correct. If the kids adopt the goal of, say, eating not-very-tasty "cool" food (raw fish soaked in horseradish? who knows??) to impress their peers in preference to eating tasty food, Pat would be well advised to start dishing up "cool" AND "healthy" (to her) food -- if such food exists. And if the kids decided to go on a fast, there is simply NO food she could serve to control for her perception of them eating "healthy" (to her) food.

>When you're supplying information, activities, attitudes, and so on at
>the level where it has an effect, you are the only control system
>involved at that level. You present a world exactly in the form you
>want, which the child is only dimly beginning to perceive.

I don't follow you here. Can you expand on this? It seems to me that Pat is NOT presenting the world "in the form she wants" to the kids, but tacitly acknowledging their wants in order to successfully control her perceptions which she wants to control. No way can she disregard that they want a certain kind of food (tasty, to date). The "trick" of her achieving control of her desired perception is in her pairing WHAT SHE WANTS with WHAT THE KIDS WANT.

>You can't control what the child is going to acquire from this new world; it

>won't be exactly what you think, and it will include many things you were not
>aware of presenting. The child will often take a disliking to something
you're
>trying to get across, and there's nothing you can do about it, short of
>force.

Sure. Through time, the necessary "pairing" as per above will alter in ways Pat cannot predict. That's where using The Test comes in. When her attempts at control of her perceptions (which depends for success on having a good model for what the kids want at any time) fail, she will, for example, ask the kids what kind of food they like NOW. And then try giving them food which is both "healthy" (to her) and is now liked by them. Of course, her model could be wrong (perhaps the kids lied about what they like) -- and she tries again to use The Test, perhaps this time by not asking the kids, but simply observing them in various situations.

>Parents can purposively influence their children as long as they don't
>interfere with levels of control that have already developed or appeal
>to levels that are still inconceivable to the child.

Yes. They also must take into account (as per the "pairing" discussed above) what the kids want. Actually, that goes for purposeful influence in general.

>You will realize that your children learned a lot more from how you WERE than
>what you DID, and that the impetus for that learning came from them.

And how -- Bateson's "deuterolearning" is a prime source of embarrassment for parents, years later!!!!

Not much argument in this post, I think... but I hope you still find it lovely.

>Two entities [Bill and his cat], not controlling each other but acting to
>hold up two ends of something that can't be carried by one end, and each of
>them knowing it and counting on the other. Doing something together that
>neither one could do alone.

Symmetrical (more-or-less) purposive influence?

>Let's by all means explore the new TUTsim. But let's also lay a snare for
>Barry Richmond and Stella, because that brings in the Mac world. Of course
you're perfectly free to say you'd rather develop your own system...

All three possibilities sound interesting to Pat and me at this stage. If any IBMers want more details on the current version of TUTsim, which starts at \$130 -- much less in student versions which are limited in various ways -- I'd be happy to provide ordering details.

I suspect that it wouldn't take much to convince the STELLA folks to add some bells and whistles specifically desired by PCTers.

Optimistic again, Greg

Date: Fri Sep 18, 1992 3:15 pm PST
Subject: Turing test: can it identify a person?

[From Bill Powers (920918.1600)] Rick Marken (920918) --

I think there's one aspect of the Turing test you may be overlooking. It's not just an examination of output spontaneously created by a hole in the wall behind which either a machine or a person might be lurking. It's an interactive test; that is, you send messages through the hole and get messages back. Your messages can be anything you like. So if you like you can send messages that amount to disturbances, under the hypothesis that the other entity is controlling for something by using its verbal outputs. You can try to discover controlled variables by clever pushing and pulling on ideas, using words. Like we do on the net all the time.

Of course this isn't what the AI types would do. I don't think they really know what they would be looking for. Didn't Penni say that the best program fooled about half the people, at the last big contest? That's about chance, isn't it? In other words, the best the judges could do with a machine that didn't give obvious clues about its machine origins was to guess whether it was "intelligent," yes or no. And they guessed at random, which means they thought a human being was a machine half the time, too. They couldn't even identify intelligence for sure when they were interacting with a human being. How did they expect to tell when it was a machine behind door number 2? Maybe that wasn't how it came out, though. Penni, how many people were misjudged as machines?

Intelligence is a SILLY idea. Suppose you read a paragraph from a book to a machine, and asked it to repeat it back. A person with a good memory might do it, and a person with a bad memory wouldn't be able to. The machine could easily repeat it back perfectly, so to fool you it would have to deliberately forget some lines some of the time, or pretend to. Or suppose you asked it to prove a theorem. If it started transmitting the proof back the instant your transmission ended, you would have to conclude that it was a machine on the other end, because it was too good at theorem-proving. Unless the proof were full of mistakes.

And so on. No matter how you define intelligence, it comes down to doing something well or fast or both. The particular abilities on which you judge intelligence will naturally be those that you admire. But being able to do anything too well or too fast will reveal that HUMAN intelligence isn't responsible -- or rather, that a human being has had too much time to prepare a machine to handle all possible contingencies in a way impossible for a human being, thus at best reflecting a non-human intelligence.

But the Turing test, just being a free interaction between two entities, can be used for any purpose you like. Under PCT we'd look for evidence of something other than "intelligence." We'd look for control. And if we found it, we might not be able to tell whether the other entity was really human just because it's able to control, but we could be sure that if it's a machine, it wasn't programmed by an AI aficionado.

--

As to Stella, you might be able to accomplish the same thing with Pascal, BASIC, or a spreadsheet, and so might about six others in the CSG, but what about the rest? Maybe Stella isn't the answer, but I think a block-diagram based system would be a lot easier for non-programmers to get the hang of than the other ways.

Best, Bill P.

Date: Fri Sep 18, 1992 6:54 pm PST
Subject: Blind Men Woes

From Gary Cziko (920918.1925)

Rick Marken (920918.1330) attempting to cop out says:

>Well, I give up. Psychological Science is not going to publish
>the "Blind men" paper as a commentary and I'm not going to try to
>get it (or any other paper) published in any scholarly journal any
>more. I will only publish stuff in the Journal of Living Systems if
>it ever exists. The brief comments of the person who reviewed the
>"Blind men" paper made it clear to me that attempts to get psychologists
>interested in PCT are just futile -- they just don't see anything wrong
>with the current state of psychological "science" (as Dag is always
>pointing out to me). . . .

I can understand your frustration, particularly since you think this is such a good paper (as I do). But I can also understand why reviewers won't like it.

Have you stopped to think that any reviewer who is not familiar with PCT is being called (at least partly) "blind" in your paper? Do you expect him or her to respond positively to that?

You may well have to move outside the mainstream psychological world to get this published, but I would like to encourage you not to give up. Not yet anyway.

I know that there are journals out there that specialize in non-mainstream and more radical approaches to psychology. Isn't there something called New Ideas in Psychology or something like that. And didn't someone mention on CSGnet not too long ago that there was a new journal that might be more friendly to PCT ideas? Let's get some suggestions from CSGnetters where Rick's article could go.

If this were my article, I wouldn't give up yet. If anything, I would make the article longer and talk about how progress in science is marked by new theories explaining what the old theories could not and showing how what are initially considered to be separate phenomena to be special cases of a more general phenomena. You could then make the case that this has NOT yet happened in psychology, where essentially new fads are invented which may become popular for a while and then disappear (or perhaps just fade a bit into the background; behaviorism is certainly still alive with its own journal). You could explain why this is so, how cognitive psychology and behaviorism share the same basic architecture of input causing output. Maybe you should

even scrap the blind men metaphor to make the article appear more "serious." You could even get into a little bit of philosophy of science and stuff like that.

If you don't feel up to this yourself, I propose that you draw on the expertise represented on CSGnet. I'd certainly be willing to help where I could. I think that this article could become a classic if you get it published, although perhaps not in your lifetime (think how proud your grandchildren will be!).

What do other CSGnetters think? Should Rick give up? (Boo!) How about some ideas on parallels in other fields of science. The only one that I can think off the top of my head is Galileo's terrestrial mechanics and Kepler's planetary physics being combined into Newton's more general theory (or at least I've been told--I'm no physicist). Could this be said for Einstein as well? What about from the fields of biology, geology, etc.?

I know that Lakatos talks about the progress of scientific research programs in terms that could be incorporated in the argument, and there is no shortage of people writing about how psychology is not shown itself to a "real" cumulative science like the physical and biological sciences. If you could get input on this kind of stuff from people like Hershberger, Taylor, Powers, Delprato, Williams, Tucker (just to name a few on CSGnet) it could be quite put into quite a powerful package with you original insight as its core.

Reactions?--Gary

P.S. Rick, could you share the review(s) with us? Remember you once said that you submit your articles to mainstream journals so that you can chuckle about how such smart people can be so benighted. I like to chuckle, too.

Gary A. Cziko

Date: Fri Sep 18, 1992 8:42 pm PST
Subject: Turing, Giving Up (sort of)

[From Rick Marken (920918.2100)]

Well, this is my last post from home. I'm giving my kid my Mac; he's going off to college tomorrow. I was going to get a replacement before he left -- but the anticipated "trickle down" failed to materialize. But this should help me buckle down to some real work on the weekends.

But here are some quick replies to swat around until monday.

Bill Powers (920918.1600) says, to me:

>I think there's one aspect of the Turing test you may be overlooking.
>It's not just an examination of output spontaneously created by a hole
>in the wall behind which either a machine or a person might be
>lurking. It's an interactive test; that is, you send messages through
>the hole and get messages back.

I know that. I should have been clearer, I guess. My point was that the Turing test is NOT aimed at the discovery of controlled variables; but that's all that's wrong with it. And that is really all that is wrong with conventional psychological research also. But that is a BIG "all". Skinner tested the effect of "schedules" on behavior until his ears turned blue; he just failed to consider the possibility that the organism was controlling some aspect of the reinforcement; so he forgot a few, little aspects of "the test" -- mainly, he forgot to hypothesize what variable might be controlled. He did introduce disturbances ("schedule changes") but he didn't monitor the hypothetical controlled variable under disturbance and change hypotheses if it turns out that the variable is not protected from the effects of the disturbance. The differences between the "behavioral" approach and the PCT approach to behavioral research are fairly small -- but crucial. I think the Turing test provides a good framework for pointing out these crucial differences (as you did in your post).

I also believe that idea of "intelligent behavior" is silly -- but I think that's the way many AI types talk about it.

Gary Cziko (920918.1925) says:

>Rick Marken (920918.1330) attempting to cop out says:

Not just attempting -- succeeding.

>Have you stopped to think that any reviewer who is not familiar with PCT is
>being called (at least partly) "blind" in your paper? Do you expect him or
>her to respond positively to that?

Good point. But the one very brief review was actually fairly positive; it said some nice things. There was no evidence that the reviewer thought that he/she was being referred to as "blind" -- actually, I think that the paper might have had a better chance if this were noticed. The reviewer basically said "so what" -- not "what do you mean 'I can't see the true nature of behavior'". The reviewer was clearly blissfully blind to my imputation that he/she was blind.

>I know that there are journals out there that specialize in non-mainstream
>and more radical approaches to psychology.

Why try? Hardly anyone pays attention to my published stuff anyway. The CSGNet forum is plenty for me. See my reply to Estes below.

>If this were my article, I wouldn't give up yet.

You haven't done this as much as I have. I've published enough articles for the time being. If I do another unbelievable experiment maybe I'll try again. Right now, it just doesn't seem worth it. What's the point of publishing anyway? For me, the goal is to get some colleagues to help with the research and modelling; to show what's been done and hope some people will help carry the ball forward. Extra heads really help. But I don't get extra heads from my published stuff;; the conventional psych audience just ignores it because it is simply orthogonal to their goals. So, forget em.

>If you don't feel up to this yourself, I propose that you draw on the

>expertise represented on CSGnet. I'd certainly be willing to help where I
>could. I think that this article could become a classic if you get it
>published, although perhaps not in your lifetime (think how proud your
>grandchildren will be!).

If you want to work together to try to make it publishable I'd be honored to have you as a joint author. I think all of your ideas about how to express the ideas in the paper are great. I'm just pretty much through with trying to impress people who don't even know what the show is about.

>P.S. Rick, could you share the review(s) with us?

The review was VERY short so I think you can glean its contents from this letter I am sending to Estes (the editor). I wrote it as my swan song to publishing in conventional venues -- but if you want to enter the fray as Marken and Cziko or vice versa, that would be great.

18 September 1992

William K. Estes
Department of Psychology
Harvard University
33 Kirkland Street
Cambridge, MA 02138

Dear Dr. Estes:

I give up. Apparently conventional psychology does not want a dialogue with control theory psychology so I'll retreat back into my lunatic fringe. It is clear from the reviewer's comments on my paper 'The blind men and the elephant: Three perspectives on the phenomenon of control' (MS 92-61) that psychologists have no fundamental problem with the basic assumptions of their discipline. I can see that there is no use pointing out that 'the emperor has no clothes' when everyone knows this and seems to accept it as evidence of imperial splendor. I will just have to continue this conversation with the few behavioral scientists who see that there is a conversation to have.

I will note that the reviewer made a few factual mistakes about the paper. First, Powers' model does not say that all behavior is best viewed as an attempt to "maintain sensory stimulation constant". The model says that it is intentional behavior that is best viewed as an attempt to maintain sensory input at constant or varying reference levels that are specified by the organism itself; other (non-intentional) behavior is best viewed as the result of ordinary physical laws. The reviewer suggests that the control theory view of behavior is 'obvious'. But if the input control view is obvious then why is virtually all psychological research based on the assumption that this obvious view is false. If organisms are input control systems then the independent-dependent variable approach to research reveals almost nothing about the organism but plenty about the environment in which it lives (as shown by equation 5 in the paper). This is the 'behavioral illusion'. If the input control view of behavior were obvious this illusion would be well understood and experimental psychology would already be based on the methodology implied by control theory -- the test for controlled input

variables. This is the only methodology that can lead to a correct understanding of the nature of the behavior of closed loop control systems.

The reviewer asks how control theory can make a real difference. I explained how, in the very last paragraph of the paper, viz. by changing the goal and methodology of psychological research. Psychologists should look for controlled inputs, not caused outputs. I forgot to mention that by doing this they can start shooting for relationships between variables on the order of $r = .99$ instead of $r = .56$ or so. The reviewer asks if control theory can do better at explaining puzzling phenomena. From my perspective, nearly all of the purported phenomena of psychology are puzzling because they cannot be accounted for precisely (less than 2% error, say) and reliably by the behavior of working models. Control theory can account for several phenomena with this level of precision; But, with only a small number of people doing this work, it's tough to make progress quickly. Nevertheless, I think it's better to have a high quality account of a few phenomena than a low quality account of many phenomena.

Trying to publish articles on control theory is like pulling teeth -- only worse; in the case of control theory the teeth are actively pulling back (because, of course, they are control systems). And when these articles are finally published nobody pays attention to them anyway. Worst of all, nobody seems to try to publish experimental results showing that we (control theorists) are wrong. Instead, we get approving nods (the reviewer noted that "it is undoubtedly helpful to remind oneself from time to time that the organism is usually part of a feedback loop"), misinformed proclamations ("in sensitization, for example, the feedback can be positive") and, finally, dismissal ("it would all be much better if the author could show how his ideas make a real difference"). If demonstrating the fallacy of a basic assumption of behavioral research (the assumption that behavior is a dependent variable) does not make a "real difference" then I surely don't know what does. Clearly, we are presenting solutions to problems that do not exist for conventional psychologists. C'est la vie.

Regards Richard Marken

Date: Fri Sep 18, 1992 9:49 pm PST
Subject: Marken: "I give up"

From Ray Jackson 921809.2230

>>[from Rick Marken (921809.1330)] >>I give up...Uncle, Uncle.

>[from Gary Cziko (921809.1925)]

>What do other CSGnetters think? Should Rick give up?

Rick,

You know I'm new to the discipline and the net... I can't console you with stories of similar incidents, nor genuinely empathize with your feelings due to my inexperience. I can understand your frustration because you take pride in what you and your peers do -- you're dealing with an understanding of behavior which can significantly impact the way we consider human nature, and

you obviously work tirelessly to theorize, define, test, confirm, then share your remarkable conclusions. Turn-downs must be terribly disheartening.

I can tell you that despite my lack of knowledge, I'm working hard to thoroughly understand PCT to enable me to assist adults to change their work environment to a fulfilling and successful culture. Eventually, I hope to follow in the (rather large) footsteps of Ed Ford and Dag Forssell. And, in the future, I also plan to contribute to the understanding of PCT as they, and you, have.

I'm sure you were probably being facetious about "giving up," but I thought a note of support would be appropriate. I find your work, and everyone else's on the net, inspiring -- even though I often find it hard to follow. I'm actively looking for the "thread" Bill told me I'd find, and in the meantime, I'm following the CSGers as you fight to develop, define, and express this theory despite the politics in the academic and scientific communities. Remember, that we in the field support and need you, feeling comfortable and confident that you are doing the homework while we slug it out in the trenches. But, also keep in mind that it's imperative that intellects such as yours fight it out in those academic arenas to legitimize what we do.

If other people can't see the value in PCT(for now), hell with them. This concept is too vital; but, then again, everyone will see that in 2, 3, or maybe 10 years, but it will only be due to the work you are doing now.

Keep at it, Rick, there are too many of us now (and in the future) relying on you.

Respectfully Yours,
Ray

Ray L. Jackson 602-963-6474
3613 W. Saragosa St.
Chandler, Az 85226
attmail.com!rljackson

Date: Sat Sep 19, 1992 4:46 am PST
Subject: Next CLOSED LOOP

From Greg Williams (920919)

Time for me to start working on the next CLOSED LOOP, due out in mid-October. Any suggestions for threads to be included? "Standards" is going in; what else do YOU think should go in?

I still don't have permission to reprint edited versions in CLOSED LOOP of the posts of some active netters (such as P. Sibun -- yes, I'm pointing at you). You folks are missing a glorious opportunity to have your thoughts immortalized in print! CLOSED LOOP is distributed to all members of the Control Systems Group; back issues are available from Mary Powers. Please fill out and SIGN a hard copy of the permission slip below and USPS it at your earliest convenience to Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328. I use posts in CLOSED LOOP only from netters who have explicitly given

me permission to do so. I edit the posts minimally (I hope!) for continuity and grammatical (NOT political!) correctness. Feel free to write in additional requirements on the permission slip if you wish. THANKS for helping to make this experiment in publishing a success!

TO GREG WILLIAMS:

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

SIGNED _____
 DATE _____
 PLEASE PRINT:
 NAME _____
 ADDRESS _____

Best wishes, Greg

Date: Sat Sep 19, 1992 7:17 am PST
 Subject: Influence; misc

[From Bill Powers (920919.0900)]

Greg Williams (920919.0800)--

>Being careful here, I would want to rephrase this and say that Pat
 >GUESSES that she is controlling her own perception of the kids' eating
 >food she considers healthy if she HAS NOT SEEN an error signal in that
 >perception and managed to correct it, but that she has WITNESSED
 >successful control of her own perception of the kids' eating food she
 >considers healthy only when she HAS SEEN an error signal in that
 >perception and managed to correct it. However, assuming that Pat IS
 >successfully controlling her perception of the kids' eating food she
 >considers healthy, she is accomplishing that control by INFLUENCING the
 >kids. Possibly you don't buy this last claim. If so, please explain
 >your problem with it.

Also being careful, I want to reserve control to refer to situations in which the loop is actually closed. The reason is that we often do things because we hope that they are going to have some effect, a situation that easily transforms into one where we imagine that the effect is occurring and forget that we have no actual information about the results. Until Pat finds a way of perceiving the actual effects of her nutritional scheme, she is operating open loop.

I don't mean to suggest that she should be operating closed-loop. That's impossible. The technology and the understanding of the biochemical systems

just isn't there. In areas like this, and there are many, the best we can do is use the old statistical approach. We can look at the data about lots of people, assume that the present situation is typical, and act accordingly, with no proof that our actions are in fact having the "typical" effect. We will, of course, have the usual success rate that comes out of this approach. But what else can we do?

In this case, then, Pat is applying an influence to the kids, but she does not know what influence she is having. If we put this as you suggest, saying "Pat is INFLUENCING the kids," then it isn't clear that she is only manipulating a cause, while the effect is the resultant of many independent variables which change without her permission or knowledge. All she can control is what she can perceive: namely, what the kids are eating. That's true control; what they eat, or at least have available to eat, is adjustable by her. We don't need to call that an influence, and indeed shouldn't, because it's a controlled outcome of the actions she takes.

>There WOULD be a problem if Pat gave our kids food which she considers
>healthy AND that tastes bad to the kids.

Right. The problem would not be with the giving, but the eating. The kids will eat if it's tasty to them, no problem. So you can control what they're eating and they won't push back. If they don't like the taste, they'll push back against the taste, not the eating.

>>You present a world exactly in the form you
>>want, which the child is only dimly beginning to perceive.

>I don't follow you here. Can you expand on this? It seems to me that
>Pat is NOT presenting the world "in the form she wants" to the kids,
>but tacitly acknowledging their wants in order to successfully control
>her perceptions which she wants to control.

I don't know the developmental stage of your kids right now, but there must have been a time when they simply ate what they were given because they were hungry and it tasted fine to them. From their point of view that's all that was going on. But to Pat, something else was going on. She was planning, choosing, and preparing meals controlled along dimensions that the kids didn't perceive. Tastiness was only one dimension, picked for practical reasons having to do with the kids' goals. The other dimensions mattered much more: it had to be food that was free of poisons, that was nutritionally balanced, and that fit an aesthetic view of how we should live in relation to nature. None of these considerations meant anything to the kids, yet by eating the food Pat gave them, they were fitting into her picture of how to raise children and how to be in the world. As they get older, they may gradually come to perceive what's going on in similar terms, and then what she's doing (and what they themselves are doing at the table) will take on added meaning in their perceptions. They may, of course, decide on different reference levels for those new perceptions.

>The "trick" of her achieving control of her desired perception is in
>her pairing WHAT SHE WANTS with WHAT THE KIDS WANT.

In the context, however, the important thing is that she was controlling perceptual variables at a higher level than they were. That means that she can

work around their wants as a means of controlling for what she thinks is important, because their wants are at a lower level, and many different wants would be consistent with the higher goal. This isn't just a "pairing," because Pat has much more flexibility than the kids do. As you pointed out -- if they want "cool" food, Pat can probably think up something acceptably cool that still exemplifies her own higher-level perceptions and goals.

>>Doing something together that neither one could do alone.

>Symmetrical (more-or-less) purposive influence?

Not of each others' behavior. It's control of a variable of a special kind: one that needs two independent control systems to be controlled. Carrying a bed upstairs is an example. Both want the bed upstairs. But neither can do it alone. So each one picks one end of the bed to control, and with just a little mutual interference the bed gets where they both want it.

When four people try to carry the bed upstairs it gets physically easier -- less weight for each to carry -- but organizationally harder. The potential for conflict goes up steeply with the number of people trying to "help." If each person tries to control for position, then only three of the people can control at the same time without conflict; three points determine the location and orientation of a plane. At least one person has to control for lifting force instead of position, contributing only a force and not trying to make the bed be in a given location.

If you want to help the little old lady who can lift only 40 pounds pick up her 50-pound suitcase, you will lift with a force of 35 pounds, and let her supply the remaining 15 as well as the reference- level for position. Then only one of you is controlling for where the suitcase is. You keep your upward force constant; she varies her upward force and creates the desired position. No conflict.

This isn't a case of symmetrical purposeful influence of the other's behavior. Whatever mutual influences remain, they are only disturbances which the other's control system can handle without difficulty.

-

Glad you're open to all possibilities for our CSG simulation language.

>If any IBMers want more details on the current version of TUTsim, which
>starts at \$130 --

It may start there, but doesn't it go up pretty high when you get the options that make it useful? And does it have a Mac version?

Ray Jackson (920918) --

First, my apologies for attributing your nice comments about Dag to Ray Allis.

And my thanks for your encouraging words about keeping up the struggle, which have meaning to many of us, not just Rick Marken. To paraphrase an older work on human behavior:

Thunder over the mountain. Perseverance furthers. The superior man carries an umbrella.

Best to all, Bill P.

Date: Sat Sep 19, 1992 1:44 pm PST
Subject: Re: I give up, too, Turing, Stella

> Well, I give up. Psychological Science is not going to publish
> the "Blind men" paper as a commentary and I'm not going to try to
> get it (or any other paper) published in any scholarly journal any
> more. I will only publish stuff in the Journal of Living Systems if
> it ever exists.

I presume that BBS is out of the question? Otherwise, please do consider the systems literature. Perhaps /Systems Research/?

O-----
>
| Cliff Joslyn, Cybernetician at Large, 327 Spring St #2 Portland ME 04102 USA
| Systems Science, SUNY Binghamton NASA Goddard Space Flight Center
| cjoslyn@bingsuns.cc.binghamton.edu joslyn@kong.gsfc.nasa.gov
V All the world is biscuit shaped. . .

Date: Sat Sep 19, 1992 4:37 pm PST
Subject: Why not "influence"?

From Greg Williams (920919 - 2) >Bill Powers (920919.0900)

>Also being careful, I want to reserve control to refer to situations
>in which the loop is actually closed.

Agreed.

>Until Pat finds a way of perceiving the actual effects of her nutritional
>scheme, she is operating open loop.

No. From the first time I gave this example, I have said that Pat is trying to control for her perception that the kids are eating healthy food. The "healthy" is her definition. I have made no claim that she is CONTROLLING for the kids' being healthy because they eat the "healthy" (her definition) food. Pat is operating CLOSED-LOOP: she can see whether or not the kids are eating the food she offers. If they don't, she'll alter her actions in an attempt to regain control of her perception that she wants to control.

Lest you think this is a cop-out, consider the case where Pat is attempting to control for her perception that the kids are healthy. Again, the definition of healthy MUST be hers -- as the slogan goes, it's all perception! How she came up with her definition of health is not an issue here. In this case, if she sees the kids NOT being healthy (her definition), she performs actions which result in that error signal being corrected. If she performs an action and sees the kids becoming LESS healthy (her definition), she performs different actions.

I see no fundamental difference in control which involves other persons (of whatever developmental stages) and control which involves non-living entities, EXCEPT that in the former kind of control, (some of) the other persons' wants must be taken into account for the control to succeed without generating conflict in the others' control systems. Both kinds of control might be unsuccessful because of ignorance or ineptness on the part of the would-be controller. Both kinds of control might require very complicated monitoring processes. Both kinds of control might (practically) need to involve statistics (if the controlled perception is a function of the perceived activities of a POPULATION of entities, living or not). And both kinds of control might have continuous and/or discontinuous monitoring of the entities' activities.

The essential distinguishing feature between the two types of control, as far as I can see, is that if the entities are alive and the controller doesn't want to give rise to conflict in their control systems, then he/she must take (some of) the entities' wants into account, as I've discussed previously, using The Test (repeatedly, if necessary) so as to determine how to "link" something the entities want with what the controller wants (i.e., "tasty" (to the kids) linked with "healthy" (to Pat) food).

For Pat to successfully control for her perception that the kids are "healthy" (her definition), she would need to figure out how to "link" whatever promotes their health with something(s) the kids want. She could be wrong about what is health-promoting; then, when she saw less healthy kids, she would need to do something else. It is possible that NOTHING she thinks is health-promoting will actually result in the kids becoming more healthy, in which case her control would be impossible. But this is no different from wanting to pick up a rock, deciding to lift it by hand, and finding that the rock is too heavy to lift -- and then finding that it is too heavy (in turn) for a shovel, for a winch, and even for a bulldozer, which exhausts your abilities to attempt to control -- and so control fails.

However, if Pat is wrong about WHAT THE KIDS WANT, then she must revise her model of the kids' controlling and try again. If rocks, rather than kids, were important in her controlling certain of her perceptions, the problem would never arise. Rocks don't have wants.

And if Pat can't figure out a way to "link" what she wants to perceive to something the kids want (i.e., the kids' fasting example in the last post), she simply CANNOT control for her perception without resulting in conflict in the kids' control systems. Such a problem does not arise when dealing with rocks, rather than kids.

>All she can control is what she can perceive:
>namely, what the kids are eating. That's true control; what they eat,
>or at least have available to eat, is adjustable by her. We don't need
>to call that an influence, and indeed shouldn't, because it's a
>controlled outcome of the actions she takes.

But it is successful (without causing conflict in the kids' control systems) ONLY if Pat succeeds in "linking" what she wants to perceive with something the kids want. Not just Pat, but the KIDS, too, must act in a PARTICULAR way (eat healthy (Pat's definition) food) rather than in other ways (eat any old

tasty food) for Pat to control her desired perception. Pat MUST "influence" the kids to achieve their wants in a manner which ALSO results in her achieving control. Pat must purposively ... hmmm... not "control," so it must be "influence," right?... the kids. She controls her perception by (in part) influencing the kids. Do you have another candidate term for providing an opportunity for another to control for what he/she wants in a way which also results in your being able to get what you want, too? Pat is certainly AFFECTING the kids -- if Pat didn't do what she does, they'd eat tasty (to them) but not healthy (to Pat) food. Pat is PURPOSEFULLY AFFECTING the kids. Is she not "purposively influencing" the kids? If there is a better terminology, I'm open to it. Maybe we should say that Pat is "channeling" the kids' controlling along the path of interest to her. Similarly, teachers, therapists, advertisers, politicians, con men, etc., etc. "channel." A lot of the time, I claim. With much success, I further claim. Judging the MORALITY of those successes is not something which PCT science can do, but PCT science explains what is NECESSARY for such successes: "linking," as described above. And The Test is the means to successful "linking."

>The kids will eat if it's tasty to them, no problem. So you can control what
>they're eating and they won't push back. If they don't like the taste,
>they'll push back against the taste, not the eating.

Right. Pat must "link" her want to THEIR "tasty." Similarly, I contend, teachers try to "link," advertisers try to "link," politicians try to "link," therapists try to "link," etc., etc. Maybe even cats try to "link." And it looks like the "linking" is often successful. (An aside: One reason for the "drug wars" is that no LEGAL "linking" is possible. The addicts want drugs, period, and the cops aren't allowed (methadone excepted) to "link" ANYTHING to that want. So the cops must resort to the only kind of control possible: the kind that produces conflict in the addicts. Of course, some cops go overboard, controlling for their own perception that they MAXIMIZE the conflict by beating the crap out of the "offenders.")

>I don't know the developmental stage of your kids right now, but there
>must have been a time when they simply ate what they were given
>because they were hungry and it tasted fine to them. From their point
>of view that's all that was going on. But to Pat, something else was
>going on. She was planning, choosing, and preparing meals controlled
>along dimensions that the kids didn't perceive. Tastiness was only one
>dimension, picked for practical reasons having to do with the kids'
>goals. The other dimensions mattered much more: it had to be food that
>was free of poisons, that was nutritionally balanced, and that fit an
>aesthetic view of how we should live in relation to nature. None of
>these considerations meant anything to the kids, yet by eating the
>food Pat gave them, they were fitting into her picture of how to raise
>children and how to be in the world. As they get older, they may
>gradually come to perceive what's going on in similar terms, and then
>what she's doing (and what they themselves are doing at the table)
>will take on added meaning in their perceptions. They may, of course,
>decide on different reference levels for those new perceptions.

I agree with all you say here, except that you miss the point that Pat's model of the kids' goals (here, with regard to what kind of food they want) is CRUCIALLY IMPORTANT for Pat's successfully controlling for her perceptions.

GW>>The "trick" of her achieving control of her desired perception is in
GW>>her pairing WHAT SHE WANTS with WHAT THE KIDS WANT.

BP>In the context, however, the important thing is that she was
BP>controlling perceptual variables at a higher level than they were.
BP>That means that she can work around their wants as a means of
BP>controlling for what she thinks is important, because their wants are
BP>at a lower level, and many different wants would be consistent with
BP>the higher goal. This isn't just a "pairing," because Pat has much
BP>more flexibility than the kids do. As you pointed out -- if they want
BP>"cool" food, Pat can probably think up something acceptably cool that
BP>still exemplifies her own higher-level perceptions and goals.

I don't follow this reasoning. I think it's the SAME story if Pat wants to see
ME eating healthy (her definition) food. If she wants to succeed, she'll try
to figure out something I want which she can "link" to healthy (her
definition) food. My "linkable" want could be high- or low-level. The
possibility of linkage has nothing to do with levels of her or my controlling
relating to each other as higher/lower, so far as I can see. The question is
whether she can find ANY way to "link." I might have very-high-level
perceptions which I am trying to control, which could be as readily "linked"
by Pat to healthy (her definition) food as could the kid's "tasty" perception.
Advertising writers appeal to consumers' patriotism to sell them dinnerware at
gas stations; who is controlling the higher-level perception? A particular
consumer might be controlling for self-as-patriot while the writer is
controlling for maintaining a single number (namely that on his/her monthly
paycheck) constant.

BP>>[Bill and his cat] Doing something together that neither one could do
alone.

GW>>>Symmetrical (more-or-less) purposive influence?

BP>Not of each others' behavior.

Why not? Why are you not controlling for being amused by seeing your cat drink
from the faucet, and in order to successfully control for that, PURPOSIVELY
providing the opportunity for the cat to get what you think (model) him to
want, namely drinking from the faucet. And why isn't the cat controlling for
drinking from the faucet, by PURPOSIVELY providing an opportunity for you to
provide him with the opportunity to get what he wants?

BP>It's control of a variable of a special kind: one that needs two
BP>independent control systems to be controlled.

It's that, also. IN ADDITION to being (I think) an example of simultaneous
purposive influence by two interacting organisms.

I need to think some more about your suitcase-lifting example. Can you expand
a bit on your final remarks?

No Mac version of TUTsim, but I just had a thought: IBMers could probably
build systems with identical blocks (amps, integrators, delays, etc.) as those

used in STELLA, so what could be passed from IBMers to MACers and back could be lists of the blocks and their connections and parameters. The basic TUTsim "personal" version allows up to 1000 blocks and has a vast variety of block types, so it probably can do everything STELLA can do. It also comes with ONE real-time input (joystick)/output (kludged) block; unlimited i/o (including full support for A-to-D boards) costs extra, as does user-defined blocks (which I think wouldn't be needed to mesh with STELLA). What's really needed is an improved STELLA, with real-time i/o and user-defined blocks; then it would be a match for the deluxe "personal" TUTsim, which probably would run less than \$300 even with the new graphical interface. And I think the TUTsim folks would be generous with discounts for PCTers.

Best wishes, Greg

Date: Sun Sep 20, 1992 11:30 am PST
Subject: Re: Why not influence?

[From Bill Powers (920920.0900)] Greg Williams (920919-2) --

>>Until Pat finds a way of perceiving the actual effects of her
>>nutritional scheme, she is operating open loop.

>No. From the first time I gave this example, I have said that Pat is
>trying to control for her perception that the kids are eating healthy
>food.

We're talking about several levels of control here. One level is the control of what is put on the table before the kids. This level is completely under Pat's control. She's in control of whatever she wants to put there, whatever she calls it.

While still satisfying the serving-food goal, she can vary the kind of food from healthful to unhealthful; that's an independent dimension of perception. She wishes to serve food she perceives as healthful, and she can also control for that because she can see the kind of food she's serving. But unless she can monitor the actual effects of the food, she can't control for its effects on health. Without a means of perceiving the way the state of health varies as the type of food varies, there's no way she can control health itself. All she can perceive is that she's doing something that theoretically should be good for the kids. To say that she serves "healthy food" is to assert that the effect of the food is known. What's really going on would be much plainer if you'd say that she serves food that she believes ought to be good for the kids.

This is a language problem, not a theoretical problem. Language lets us sneak in assertions without coming right out and stating them. Before you can apply PCT to any situation, you have to take the description apart to make all the hidden assertions explicit.

>Lest you think this is a cop-out, consider the case where Pat is
>attempting to control for her perception that the kids are healthy.
>Again, the definition of healthy MUST be hers -- as the slogan goes,
>it's all perception! How she came up with her definition of health is
>not an issue here. In this case, if she sees the kids NOT being

>healthy >(her definition), she performs actions which result in that
>error signal being corrected. If she performs an action and sees the
>kids >becoming LESS healthy (her definition), she performs different
>actions.

Now you're talking about actual control. If Pat has ways of perceiving the state of health of the children (her way, but what else have we?), and knows how to vary the offered diet to counteract changes in the indicator variables, then she isn't "trying" to control her perception of their health; she's controlling it. The implication is that she could make their health match any reference level for it by varying their diet. If she can actually see how varying the kind of food changes the kids' state of health, then she can control the state of health. Unless, of course, they're already controlling it themselves.

>I see no fundamental difference in control which involves other
>persons (of whatever developmental stages) and control which involves
>non->living entities, EXCEPT that in the former kind of control,
>(some of) the other persons' wants must be taken into account for the
>control to succeed without generating conflict in the others' control
>systems.

I agree that control is control, whatever is being controlled. But that little difference, I claim, makes all the difference.

We're pulling in different directions here. You're looking for generalizations in which differences don't make a difference. From my point of view, all this does is obscure the differences that DO make a difference. I'm interested in how control of people differs from control of inanimate things. To me, that is more "fundamental" than ways in which such relationships can, by ignoring details, be viewed as the same. When you say you see no fundamental difference here, to me this looks like going UP the ladder of abstraction, which leads AWAY from the "fundamentals" and in the direction of less and less useful generalizations.

Anyway, it isn't sufficient to know SOME of the variables under active control by the other person. You must know ALL of them. You can control a person's reaching out by holding up something you have reason to believe that the person wants. But you have to understand that in reaching out, that person is shifting the center of gravity of the body, so you must be sure that there is room for the person to lean backward or shift a foot forward, and you must take care that in leaning backward the person won't touch some wet paint, or that in shifting a foot forward, the person won't pitch headlong down a flight of stairs. You have to know how the person feels about having to reach out to get what you're offering; the King, for example, might take it amiss if you make him reach too far for the wine glass he wants.

Of course you may be unaware of the extent to which you're disturbing another's hierarchy by trying to control some innocent-seeming aspect of that person's behavior. To appreciate the full extent of the effects, you'd have to use the Test systematically, with a far more organized and extensive effort than anyone knows how to mount now. There are more effects than people realize, in no small part because in our culture, it's considered demeaning or weak to allow others to see just how much their efforts at control disturb you. A person may be perfectly aware of having his or her actions manipulated,

but will go along anyway simply to keep the manipulator from knowing the murderous impulses that the controllee is holding in check. Robbers and con-men often seem puzzled at what seems to them an overreaction to what they did, when they're caught. They know a little, but not enough, about what the victim wants. It's their ignorance of human nature that gets them into such professions in the first place.

>I think it's the SAME story if Pat wants to see ME eating healthy
>(her definition) food. If she wants to succeed, she'll try to figure
>out something I want which she can "link" to healthy (her definition)
>food.

Again we have that language problem: "healthy food." Calling food "healthy" is classifying it by its hoped-for effect. There's no question that you can control for giving another person a particular kind of food, while avoiding kinds the other person rejects. But if you can't monitor its effects, there's no way you can say you're controlling for the effect you hope it will have. As I commented a few posts back, not entirely at random, hoping for something or wishing for something is not the same as controlling for it.

I have a problem with terms like "linkable wants." What good are such terms? They just make us try to interpret them in PCT terms. Why not simply express what you mean in PCT terms in the first place? Is "linking" a new phenomenon that requires a new explanation? To me, it seems a step backward into vagueness and ambiguity. The general term "linking" can apply in a lot of specific situations. You could "link" my eating what you put in front of me to treating me nicely; i.e., make your nice treatment of me contingent on my wanting what you want. If you didn't mean THAT kind of linking, then you shouldn't use "linking."

>Advertising writers appeal to consumers' patriotism to sell them
>dinnerware at gas stations; who is controlling the higher-level
>perception?

The consumers.

>Why are you not controlling for being amused by seeing your cat drink
>from the faucet, and in order to successfully control for that,
>PURPOSIVELY providing the opportunity for the cat to get what you
>think >(model) him to want, namely drinking from the faucet.

Providing an opportunity is like serving healthy food. What I can provide on purpose is a physical situation under my direct control that I think of as an opportunity to drink. But it doesn't matter to the cat how I think of it. What matters is whether the cat uses this situation as a means of drinking. I have no control over whether my physical manipulation amounts to an opportunity or not. Half of the time the cat switches its tail at me, ignores the sink, and drinks out of his cup. It's not as if I can control what the cat's going to do. That's not the relationship. I can't control for my amusement, because I can't make the cat "take advantage of the opportunity." There's a critical part of the process that's under the cat's control, not mine. The relationship between me and the cat is not one of control.

BP>It's control of a variable of a special kind: one that needs two
BP>independent control systems to be controlled.

>It's that, also. IN ADDITION to being (I think) an example of
>simultaneous purposive influence by two interacting organisms.

Show me, in PCT terms. How are the two people carrying the bed upstairs making each others' actions be in a desired state? What are the controlled variables? The reference levels? The actions taken to affect the controlled variables?

(from my 920219)

>>This isn't a case of symmetrical purposeful influence _of the
>>other's behavior_. Whatever mutual influences remain, they are only
>>disturbances which the other's control system can handle without
>>difficulty.

>I need to think some more about your suitcase-lifting example. Can
>you expand a bit on your final remarks?

I'm controlling an upward force at a reference level of 35 pounds, independently of position. The LOL is controlling for position by _varying_ an upward force, which need be only 15 pounds to move the 50-pound suitcase, as I am now cancelling 35 pounds of its weight. As she moves the suitcase she will be disturbing my control system somewhat, and my system will respond by bringing my upward force back to the required level. The LOL will bring the suitcase to whatever height above the ground she wishes; I'm not controlling for that.

At a higher level, I'm helping her with the suitcase because she let me know she wanted help, and I decided, for numerous reasons, to comply. I've got my mitzvah; she has the suitcase where she wanted it. I lent her some lower-level systems to use in controlling the variable she wanted to control -- not me, but the suitcase's position. It was I who rearranged my momentary wants to allow her to use my muscles in her control action. I could have said no.

>IBMers could probably build systems with identical blocks (amps,
>integrators, delays, etc.) as those used in STELLA, so what could be
>passed from IBMers to MACers and back could be lists of the blocks
>and their connections and parameters.

That's a definite possibility. I haven't heard anything from Hammond yet. Stella would have to be modified to accept a setup language. And I think the language should be both generated by and interpreted by the programs, to avoid errors and for the sake of people just starting to learn.

Best, Bill P.

Date: Sun Sep 20, 1992 3:59 pm PST
Subject: Re: Next CLOSED LOOP

Greg, I vote for a thread on purposive influence, some carefully edited version of the recent exchange between you and Bill.

Hugh G. Petrie, Dean

Phone: 716-636-2491

Graduate School of Education
367 Baldy Hall
University at Buffalo
Buffalo, NY 14260
USA

FAX: 716-636-2479
BITNET: PROHUGH@UBVMS

Date: Sun Sep 20, 1992 4:22 pm PST
Subject: More purposive influencing; STELLA

From Greg Williams (920920) >Bill Powers (920920.0900)

>What's really going on would be much plainer if you'd say that she serves
>food that she believes ought to be good for the kids.

What I said is that she controls for her perception of the kids eating "healthy" (her definition) food. She isn't just controlling for seeing the "healthy" (her definition) food SERVED, which depends ONLY on HER actions, but also for seeing the kids EAT the "healthy" (her definition) food, which depends BOTH on (some of) her actions AND (some) actions of the kids. In order to control her seeing the kids eating the "healthy" (her definition) food, she must serve food which is both "healthy" (her definition) and wanted (on SOME basis) by the kids. So she makes a model of what the kids want with respect to food, and serves them food which is "healthy" (her definition) and which she believes also has the property of being wanted (currently) by the kids. The belief is developed via her using the Test. She might be wrong in her belief, and will then see an error; to correct and regain good control of her desired perception, she will try the Test again. It is possible that she will never be able to come up with a property of food which the kids want and which can be combined (in the same food) with "healthiness" (her definition), in which case she is not able to control her desired perception. So far (in real life) Pat has been able to repeatedly see both the kids and me eat the "healthy" (to her) and "tasty" (to me and the kids) food she serves us, with a few exceptions where her belief about the "tastiness" of the food she served was (I can attest!) quite incorrect. She never served that stuff again, and regained good control of her perception that we are eating "healthy" (her definition) food.

I call this sort of control, where one's desired outcome depends on outputs of others, "purposive influence," because it is being done purposefully (it is control of one's perception) and because, to succeed, it requires influencing the way others control. It does not involve overwhelming physical force or threat thereof. Generation of conflict in the others' control systems is minimized if the influencer can sufficiently ascertain (via The Test) what the prospective influencee wants in order to provide the opportunity for those wants to be satisfied in a manner which allows the influencee to see what he/she wants to see. Empirically, I find that purposive influence is extremely common, and often successful in the sense that the influencer retains good control of his/her desired perceptions, which depend on the influencees' controlling, over long periods of time. Rubber-banding is a special case of purposive influence, which is a special case of controlling (namely controlling where the desired outcome depends on other organisms' controlling).

>If she [Pat] can actually see how varying the kind of food changes the kids'
>state of health, then she can control the state of health. Unless, of course,
>they're already controlling it themselves.

To avoid generating conflict in an influencee's control system, the influencer must not attempt to control a perception (of the influencer, of course) which is incompatible with possible actions which could be taken by the influencee to enable control by the influencee of the influencee's desired perceptions. Nevertheless, the empirical evidence suggests that (1) often there are SEVERAL possible actions which can enable control by the influencee of the influencee's desired perceptions, (2) often there is at least ONE possible action which ALSO is compatible with the influencer's desired perception, and often the "compatible" action/s is/are "don't-care" for or actually embraced (because of perceived "convenience," lack of perceived "drawbacks," or other reasons) by the influencee, and so there is no conflict generated in the influencee's control system. The Test (in some form, and applied repeatedly if necessary, perhaps in different forms) is useful to predict whether conflict generation is likely to be a problem in any given case.

GW>>I see no fundamental difference in control which involves other
GW>>persons (of whatever developmental stages) and control which involves
GW>>non->living entities, EXCEPT that in the former kind of control,
GW>>(some of) the other persons' wants must be taken into account for the
GW>>control to succeed without generating conflict in the others' control
GW>>systems.

>I agree that control is control, whatever is being controlled. But that
>little difference, I claim, makes all the difference.

No argument here. It makes all the difference. I was trying to make the point that needing to take (some of) the other persons' wants into account IS INDEED the crux of the difference.

>I'm interested in how control of people differs from control of inanimate
>things.

I'm interested in how (successful) control of one's perceptions must differ (according to PCT science) when (a) the perceptions depend on other peoples' actions and (b) when the perceptions do not depend on other peoples' actions. Is that the same interest as yours?

>Anyway, it isn't sufficient to know SOME of the variables under active
>control by the other person. You must know ALL of them.

You must know ALL of them to do what? I claim that, for successful purposive influence, SOME -- but by no means ALL -- variables under active control by the influencee must be correctly determined by the influencer, so that the influencer can provide the opportunity for the influencee to control in such a way that the influencer will see what he/she wants to see. Pat need only correctly figure out that the kids want "tasty" (to them) food, serve them the food, and then watch them gobble it up. If they don't gobble it up, she has incorrectly figured out ONE of their reference signals (or maybe the laws of physics changed, or the dining table blew up -- but even control SOLELY involving inanimate entities relies on a faith in certain constancies, like

the continued existence of the universe). Sure, sometimes there are details unforeseen by the influencer (suddenly one of the kid's friends bursts in with a box of popsicles -- cherry, even!!), but that's possible (Murphy's law says likely!) with inanimate entities, too. Pat isn't trying to see the kids eat the food JUST SO; JUST EATING it will suffice. And if they fall out of their chairs and can't eat the food for some strange reason, Pat will work to correct that problem.

>Of course you may be unaware of the extent to which you're disturbing
>another's hierarchy by trying to control some innocent-seeming aspect
>of that person's behavior.

Yes, you might. But I look around and see that this doesn't appear to be a big problem in the real world, where successful purposive influence is common.

>To appreciate the full extent of the effects, you'd have to use the Test
>systematically, with a far more organized and extensive effort than anyone
>knows how to mount now.

I claim there's no need to appreciate it.

>Robbers and con-men often seem puzzled at what seems to them an overreaction
>to what they did, when they're caught. They know a little, but not enough,
>about what the victim wants. It's their ignorance of human nature that gets
>them into such professions in the first place.

Is this just opining, or do you have some evidence to back up what you say here?

>I have a problem with terms like "linkable wants." What good are such
>terms? They just make us try to interpret them in PCT terms. Why not
>simply express what you mean in PCT terms in the first place? Is
>"linking" a new phenomenon that requires a new explanation?

I've tried at the start of this post to eliminate all but PCT jargon from my explanation of what I'm talking about.

GW>>Advertising writers appeal to consumers' patriotism to sell them
GW>>dinnerware at gas stations; who is controlling the higher-level
GW>>perception?

>The consumers.

In your last post, you seemed to be claiming that a requirement for successful purposive influencing is that the influencer control for (his/her) perceptions at a HIGHER level than the perceptions being controlled for by the influencee. Have you dropped that claim, or perhaps never actually made it?

>Providing an opportunity is like serving healthy food.

Yes.

>What I can provide on purpose is a physical situation under my direct control
>that I think of as an opportunity to drink. But it doesn't matter to the cat
>how I think of it. What matters is whether the cat uses this situation as a

>means of drinking. I have no control over whether my physical manipulation
>amounts to an opportunity or not.

No control if you have a poor model of your cats' controlling.

>Half of the time the cat switches its tail at me, ignores the sink, and
>drinks out of his cup.

So you DO have a poor model. Models of non-human animals' controlling is made
difficult by the fact that they don't speak, so don't feel too bad.

>It's not as if I can control what the cat's going to do.

That's right. (You and I mean without overwhelming physical force. Maybe even
WITH overwhelming physical force, for a cat!) But I claim you MIGHT be able to
apply The Test and figure out what your cat wants sufficiently to improve the
repeatability of your amusement.

>I can't control for my amusement, because I can't make the cat "take
>advantage of the opportunity."

Providing the RIGHT SORT of opportunity would require a better model of your
cat's controlling than you now have. But I can't speak about the likelihood of
being able to find the right sort of opportunity, because I don't have the
opportunity to observe a lot of successful purposive influencing with cats.
With people, however, I observe that "good-enough" models don't seem very
difficult to construct -- even by influencers who know nothing about PCT and
apply The Test intuitively.

>There's a critical part of the process that's under the cat's control, not
>mine. The relationship between me and the cat is not one of control.

Yes, yes, yes! You can only control YOUR perceptions; your cat can only
control HIS perceptions. But BOTH of you can (not necessarily WILL) control
(some of) your OWN perceptions by purposive influence.

>Show me, in PCT terms. How are the two people carrying the bed
>upstairs making each others' actions be in a desired state? What are
>the controlled variables? The reference levels? The actions taken to
>affect the controlled variables?

I want to see this bed upstairs. (reference level: bed upstairs; controlled
variable: bed position)

I ask you, "Would you like to help me move this bed upstairs?" (action:
applying The Test to attempt to determine whether I can purposively influence
you so I can control for the bed being upstairs (it's a big, heavy bed, and I
can't move it alone; you just happen to be here)

You say, "Sure!" (PUTATIVE reference level EXPRESSED by you: bed upstairs
(note shared reference level); controlled variable: bed position)

I think you get the gist. Now I'd tell you how I want the movement of the bed
to go in some -- BUT NOT MICROSCOPIC -- detail: "I'll go first; you keep it
level; etc., etc." I'd want to hear either "OK" from you ("asymmetric

successful purposive influence") or better ideas about how it should be done (to which I would no doubt agree -- you're on the bottom -- "symmetric successful purposive influence"). Our purposive influencing would continue throughout the moving, as you and I each deem necessary.

>At a higher level, I'm helping her with the suitcase because she let
>me know she wanted help, and I decided, for numerous reasons, to
>comply. I've got my mitzvah; she has the suitcase where she wanted it.
>I lent her some lower-level systems to use in controlling the variable
>she wanted to control -- not me, but the suitcase's position.

Now you've got it! You've been successfully purposively influenced. Didn't really hurt, did it?

>It was I who rearranged my momentary wants to allow her to use my muscles in
>her control action. I could have said no.

You didn't. That LOL sure psyched out some (enough!) of your wants, didn't she?

GW>>IBMers could probably build systems with identical blocks (amps,
GW>>integrators, delays, etc.) as those used in STELLA, so what could be
GW>>passed from IBMers to MACers and back could be lists of the blocks
GW>>and their connections and parameters.

>Stella would have to be modified to accept a setup language. And
>I think the language should be both generated by and interpreted by
>the programs, to avoid errors and for the sake of people just starting
>to learn.

All we'd need is the format for setup files for STELLA; then we could write a pre-processor (translator) to convert from TUTsim's ASCII setup files to STELLA setup files and vice-versa. The program could be entirely separate from STELLA and TUTsim, and wouldn't be very complex or hard to write -- unless the STELLA folks are proprietary about their file formatting, and we have to figure it out for ourselves! A MACer could send a STELLA setup file to an IBMer who would run it through the translator and end up with a regular TUTsim setup file, and vice versa.

ADVERTORIAL: Maybe if PCTers are serious about having a common programming language (i.e., block diagrams), they are willing to pay a little something for the translator program, as well as for TUTsim or STELLA?

Best wishes, Greg

Date: Mon Sep 21, 1992 8:00 am PST
Subject: Giving Up

[From: Delprato (920921)]

Given Rick's "giving up" stance, I find it interesting that in response to my communication with a former student who is currently in a doctoral

program--where it seems his arm is being twisted to keep him in line with venerable tradition--I received the following:

"It was then, in fact, high time to bid farewell, and I soon received proof. Richard Wagner, who seemed all-conquering, but was in reality only a decayed and despairing romantic, suddenly collapsed, helpless and broken, before the Christian Cross. Was there not a single German with eyes in his head and sympathy in his heart for this appalling spectacle? Was I the only one he caused suffering? In any case, the unexpected event illuminated for me in one lightning flash the place that I had abandoned, and also the horror that is felt by every one who is unconscious of a great danger until he has passed through it. As I went forward alone, I shuddered, and not long afterwards I was ill, or rather more than ill -- weary: weary from my ceaseless disappointment about all that remained to make us modern men enthusiastic, at the thought of the power, work, hope, youth, love, flung to all the winds...."

Friedrich Nietzsche
September, 1886

By asking the hard questions we gain the opportunity to further scrutinize ourselves. The alternative lies in the safety of a group, whose sole purpose is to suppress the asking.

Mark A. Swain
SWAIN@UNSSUN.SCS.UNR.EDU

Date: Mon Sep 21, 1992 12:12 pm PST
Subject: giving up

[From Francisco Arocha, (920921 13:07) Rick Marken (920918.2100):

>What's the point of publishing anyway? For me, the goal is to get
>some colleagues to help with the research and modelling; to show
>what's been done and hope some people will help carry the ball
>forward. Extra heads really help. But I don't get extra heads from
>my published stuff;; the conventional psych audience just ignores it
>because it is simply orthogonal to their goals. So, forget em.

Sorry for the rejection of your paper. I thought it was quite a good exposition of scientific PCT in relation to pre-scientific approaches. But do not give up, please. There are some other journals that can publish your work. As Gary mentioned, New Ideas in Psychology is one of them. It may not be read by mainstream psychologists, but at least it is read by people who are dissatisfied with mainstream psych. (it has even published 3 papers by Mario Bunge!). It is dissatisfied researchers the ones you should reach, not people who are committed already to standard psychology. Remember, Skinner and the skinnerians were not convinced a bit by information processing psychology (not that they should). So, please, give it a try. Again.

Targeting mainstream psychology is very difficult, and as you said, even if the paper gets published, no one will pay attention, anyway. Has not Paul Meehl been publishing papers on hypothesis testing in psychology for almost 30 years, demonstrating that it is a useless exercise? That it works in exactly

the opposite way hypothesis testing in works physics? That in physics, the more a theory is tested the harder the tests, and therefore, the better a successful theory is? That in psychology is EXACTLY the reverse? And, mind you, Meehl is not a revolutionary. Are psychologists worried about poverty of their testing procedures? Of course not. They say this is good enough. They believe this is the way science works. Where the only thing they have to do to test their hypothesis is to show a statistically significant result How can you expect them to abandon such a piece of cake?

What can one expect of a discipline for which "laws" (if there is such a thing in psychology) are just statistical generalizations and not theoretical statements, as in the physical science? Of a discipline for which a theory is nothing more than a bunch of vague verbal statements and not a mathematically formulated body of propositions, as it is in physics? Are information processing psychologists worried about the ambiguity of their basic terms, such as information? They are not and they do not care. They are content with saying that people (and machines) process information, whatever that means. They like to talk about the great advances of the last 30 years in understanding how the mind works.

What PCT needs is more experimental and theoretical work, maybe reaching other audiences. I have always thought that educational researchers are more open to new ideas than experimental psychologists. Maybe an education-oriented paper would succeed.

Also, think of the other people in this net (like myself) who eventually would like to see our papers published. We are counting on you to clear the path.

Cheers, Francisco

Date: Mon Sep 21, 1992 12:16 pm PST
Subject: Giving up -- reconsidered

[From Rick Marken (920921.1030)]

Thanks to everyone who posted encouraging words in response to my "giving up" post.

Let me just say that I am not giving up the idea of publishing PCT ideas. What I am giving up (maybe) is trying to publish these ideas in the standards psychological journals. After thinking about this a bit over the weekend, I realize that my main reason for "giving up" is not the difficulty of getting PCT papers published in the standard journal. Heck, it's hard (or, at least, should be hard) to publish ANY paper in a journal. Journal space is precious and the job of the editor is to make sure that only the most valuable work gets in. I sure don't mind working my butt off trying to get a paper to meet the demands of editors and reviewers (even when I don't agree with their critiques) -- in fact, I think it's kind of fun. And I have successfully published a number of papers -- some in reasonably prestigious journals. And it almost always took one to three years of rewrites, sometimes more studies, etc. I don't get frustrated because of the difficulty of getting PCT studies published.

The reason I am considering giving up on the publishing biz is because (as I said in my letter to Estes) nobody pays attention to these articles anyway, once they are published. And when the articles DO get some attention, it is not very useful attention. For example, the person who reviewed my "Blind men" paper was familiar with Powers' work --s/he must have read something that Bill had published. But s/he had only a vague understanding of what Bill was talking about. The reviewer had read Bill's article, nodded "uh huh" and figured that s/he already knew what was going on.

There is a reason why PCT stuff is not going to get the attention or consideration of readers of psych journals -- it's because PCT is not about what the people who read these articles care about. PCT isn't even a disturbance to these readers. These readers are not looking for evidence that the foundations of their discipline are wrong; that they are asking the wrong questions using the wrong methods. No indeed. These readers want to know the state of the art in testing cognitive abilities, how people determine the optimal interval between medical check-ups, the effect of "meaning frequency" on processing of lexically ambiguous words, predicting health from stress, etc (these topics are culled from the latest issue of Psychological Science -- the journal that rejected the "Blind men" paper). The readers of these journals are not interested in hearing that these are misguided research questions based on the wrong assumptions. So, even if I manage to get stuff published in the standard psych lit it is only good for archival purposes anyway.

I do want to see PCT ideas promoted and disseminated; PCT work should be made known to those who might be interested and shared with those who ARE interested. In the old days the only way to do that was to publish in the standard literature and hope that someone would pay attention -- sometimes it works (apparently it did with Einstein) and sometimes it doesn't (Gregor Mendal published his genetic work, I believe, and received about the same reception as PCT does now). Nowadays we have computer networks (like this one) that not only make the information available but allow QUICK replies from those who oppose or question the ideas -- that's what makes this medium really useful. The reader of a journal who disagrees with something in a paper on PCT is most likely (I suspect) to just ignore it, But if s/he is in contact with the author, and in a public arena, the differences can be immediately clarified, refined and, hopefully, translated into testable models.

So, what I am giving up is trying to publish in the standard journals -- not because it's hard, but because it's useless. The network, I think, is the way to go from now on -- not just CSGNet but other related networks as well. Now that I think of it, I think I'll submit a revised version of the "Blind men" paper to the "Psychology" net. What the hey?

Hasta luego Rick

Date: Mon Sep 21, 1992 12:16 pm PST
Subject: More Purposive Influencing; misc

[From Bill Powers (920921.0800)] Greg Williams (920920) --

We agree up to a point, perhaps farther along than before.

While NOT altering your control for taste, and while NOT altering your control for eating, Pat can vary the food you eat in the remaining dimensions that aren't under control by you: that is, her definition of "healthiness" of the food. Assuming that these uncontrolled dimensions are really don't-care dimensions for you, Pat can control what goes into you as long as it doesn't disturb the perceptual variables you are controlling for.

OK so far?

>I call this sort of control, where one's desired outcome depends on
>outputs of others, "purposive influence," because it is being done
>purposefully (it is control of one's perception) and because, to
>succeed, it requires influencing the way others control.

I can see why you call it purposeful, and why it is truly control, but why do you say it requires influencing "the way others control?" It can't be done if it DOES influence the variables others are controlling for (eating and taste). How does supplying a particular kind of food influence how you eat it (other than the mechanics)? It seems to me you would eat food that Pat considers unhealthy in exactly the same way you'd eat food she considers healthy. What has she influenced other than the composition of what you put in your mouth (without altering its tastiness)? What would you do differently if she substituted squash for potatoes?

You say

>the empirical evidence suggests that (1) often there are SEVERAL
>possible actions which can enable control by the influencee of the
>influencee's desired perceptions,

But when Pat chooses food to feed you that does not disturb your perception of eating or your perception of tastiness, how can you say that this "enables" control by the influencee of the influencee's desired perceptions? Do you mean that if Pat weren't putting SOME kind of food down, you'd be unable to eat tasty food? What's to prevent you from making your own tasty food if she doesn't make it? Why isn't her presenting tasty food just an aiding disturbance that allows you to relax your own efforts to get tasty food? Your relaxation is, of course, an opposing action. We always resist disturbances of things we're controlling for, even if they help our efforts. We can't help it; that's how control works.

> (2) often there is at least ONE possible action which ALSO is
>compatible with the influencer's desired perception, and often the
>"compatible" action/s is/are "don't-care" for or actually embraced
>(because of perceived "convenience," lack of perceived "drawbacks,"
>or other reasons) by the influencee, and so there is no conflict
>generated in the influencee's control system.

This begins, albeit in a most convoluted way, to get into more of the influencee's hierarchy. It's convenient to have Pat prepare the food, so you agree to eat what she provides, if it tastes good to you. That agreement is under your control as well as hers. An agreement, like a conversation, is a member of a class of social variables that can't be controlled by one person. You can freely make any agreement with another that you like, as long as the other is willing to make the complementary part of the agreement. In BCP I

referred to a similar relationship, "bargaining," a trade in which each party offers to provide something the other wants in return for getting what he/she wants. Bargaining can involve use of the other's physical actions: I'll help you carry the bed upstairs if you'll lend me your truck tomorrow. Often this kind of bargaining is so deeply a part of a culture that all parties concerned understand that a bargain has been made simply by the act of helping another. If Mrs. Pacourek comes visiting Mrs. Hynek and brings a couple of jars of preserved mushrooms as a gift, Mrs. Hynek will not be able to rest until she can take something to Mrs. Pacourek and erase the debt. In other circles, dinner parties fulfil the same function.

 The distinction between environmentalism and organismism has to take into account the fact that for a given organism it makes a lot of difference whether the environment in question is living or nonliving, and if it is living, whether its organization extends to the same levels of control as in the given organism.

It's clear that the environment that is fully described by the laws of physics and chemistry has no natural control systems in it, and thus ranks zero on the scale of control capabilities. This environment poses no particular problems for any organism; organisms have never exhausted the possibilities of the physico-chemical environment. So if we divide the world into organisms and non-organisms, there is no contest. The world of organisms, or life in the aggregate, determines what will happen on the local scale, for only living systems have purposes and the capacity for realizing them. The non-organismic environment is passive; it doesn't aid or resist being pushed into various forms. It can affect organisms, even kill them, but when organisms find ways of controlling the effects, the environment does not retire to think up a new strategy. It is, in principle, utterly predictable and therefore completely controllable. It is the exact opposite of autonomous.

The question of autonomy is, I think, best understood in terms of organisms versus nonorganisms. The purposes carried within organisms and the machinery for carrying them out render organisms autonomous in all the practical physical regards that matter to their survival.

When organisms that are autonomous in that sense encounter each other, however, autonomy leads to conflict, and becomes a much more relative concept. To retain autonomy of some sort, organisms interacting with each other must give up major areas of autonomy that would not be any problem if the opponent were not alive. Organisms that are equal in control capabilities have to develop social systems to avoid destroying each other. I claim that these social systems are developed by striking bargains: I won't murder you if you won't murder me. One-sided bargains -- you can't murder me, but I can murder anyone I please -- can't work, because bargains require agreement, and in making agreements each person is trying to maintain control of what matters to that person. The underlying physical autonomy remains. Only force, which strikes through all the organizational defenses against disturbance and reaches to the core life-support system on which all the others depend, can maintain an unfair agreement.

----- I think that coequal organisms constantly use the Test for the Controlled Variable on each other. But they don't do this in order to discover ways of controlling the other organism. They do it to avoid trying to control something the other organism is already controlling. The Test is easier to apply as a way of

finding out what the other ISN'T controlling than as a way of finding out what it IS controlling. All that most organisms want or need to know is whether controlling some aspect of the world that matters to them is going to meet with inconveniently energetic resistance that will make control difficult. If not, they'll go ahead and control it. They don't have to stop and find out what IS being controlled first.

Trying to discover what another person IS controlling, using the Test, is a much more elaborate and difficult process. Your first attempt at applying a systematic disturbance is much more likely to reveal a variable that isn't controlled than one that is. If you're concerned only with controlling for your own perceptions, that's all you need to know. Searching for the other's controlled variables is important only when you have some positive reason to want to manipulate the other's actions, the way the other moves, to suit yourself. Or when you simply want to understand how the other system works.

Pat can feed you any food she likes as long as you don't resist it. If you don't resist it, she doesn't need to search for food that you would resist eating. If there is such a food, it will turn up in the natural course of things anyway. The Test will naturally occur; why spend a lot of time finding out all the things you won't eat, when she would probably never think of feeding most of them to you in the first place -- such as stewed puppy or chocolate-covered ants? What revolts her would probably revolt you, too.

Just ask yourself, what if Pat did an extensive survey of foods and discovered all the things you resist eating. What would she then know about the controlled variable that's being disturbed when you resist all these different foods? She wouldn't know much; you probably don't know what it is, either. Certainly, SOMETHING about these various foods is disturbing one or more important controlled variables. But it's possible that each one disturbs a different controlled variable. You won't eat a cup of stewed puppy for one reason; you won't eat a cup of pure salt for a different reason. Trying to generalize from all the foods you won't eat to the "common factor" that reveals the controlled variable is probably futile. You control for dozens of variables having to do with food, some of which may be at so low a level that you're unaware of them. It took my son 15 years to realize that he became a vegetarian not for any of the high-flown reasons that others give or because he didn't like eating meat, but because eating meat always upset his stomach.

The Test for the Controlled Variable is useful in a scientific setting, when you have reason to want to know what is in fact being controlled, at least at the time you use the test. But even then, all it does is identify the variables that an external agency can't simultaneously control without creating conflict. In most human interactions, the Test doesn't need to be carried much past the first step, finding out whether there's resistance. If I want you to pass me the salt because I want the salt, I'll just ask for it. If I get it, that's the end of the matter -- you weren't, at that moment, controlling for the position of the salt-shaker for your own purposes. If you were, you'll just say "Half a sec," shake it on your own plate, and then pass it to me. You cope with the disturbance and so do I. You get what you want; I get what I want. It's just one of those little understood bargains.

The only time I'd be concerned about controlling your behavior would be if I were trying to teach you a lesson in manners. Then I wouldn't want you to use the salt yourself before passing it, because that's supposed to be selfish and

impolite. HOW you passed the salt to me would then be the main point, and my getting the salt secondary. And I certainly wouldn't have to find out all the things you would resist doing with the salt shaker in order to see if you would pass it to me before using it yourself. I'll bet you wouldn't shake the salt in your eye, but so what?

>ADVERTORIAL: Maybe if PCTers are serious about having a common programming
>language (i.e., block diagrams), they are willing to pay a little something
>for the translator program, as well as for TUTsim or STELLA?

How about it, folks? Sounds good to me.

Gary Cziko (920921)--

I haven't heard back from Richmond, whom I keep wanting to call Hammond. Next week I suppose he'll be Wurlitzer.

Maybe Wurlitzer doesn't read his email very often.

Best, Bill P.

Date: Mon Sep 21, 1992 12:38 pm PST
Subject: influence, again

CHUCK TUCKER 920921 - THIS IS THE THIRD ANNIVERSARY OF HUGO WHICH OCCURRED
AT OUR LAST MEETING HELD IN WISCONSIN - MIXED EMOTIONS!

I have not read the recent posts on influence so an answers to my queries below might be found in them but I was struck by several statements made by Bill (WTP) and would appreciate some assistance in understanding them.

In your WTP 920916.0900 post you state:

"There are actually lots of things left, once we get over
the idea that controlling others is going to help."

I have some ideas about what those "things" might be but I would like to have
yo Thanks

Then in your WTP 920918.0800 post you state:

"Two entities, not controlling each other but acting to hold up
two ends of something that can't be carried by one end, and each
of them knowing it and counting on the other. Doing something
together that neither one could do alone."

You may notice that this statement is about you and your cat but could it not
be applied to actions between humans? Is this the type of relationship that a
recognition and application of PCT should bring about?

Just refer me to other posts if these "issues" have been covered.

Regards, Chuck

Date: Mon Sep 21, 1992 1:20 pm PST
Subject: Re: Next CLOSED LOOP

(ps 920921.1200)

I still don't have permission to reprint edited versions in CLOSED LOOP of the

posts of some active netters (such as P. Sibun -- yes, I'm pointing at you).

You folks are missing a glorious opportunity to have your thoughts immortalized in print! CLOSED LOOP is distributed to all members of the Control Systems Group; back issues are available from Mary Powers. Please

humph. i bet i wasn't listening when i was told about this. my inclination is to say no, but maybe if you told me more about it, like who actually reads it and how much editing you do, i'd change my mind.

--penni

Date: Mon Sep 21, 1992 6:42 pm PST
Subject: Summaries, anyone?

Something just occurred to me.

Of course, we all know that the output of this list is prodigious. On the other hand, it is also crucial. What's a poor boy to do? I log on every day, and wade through the first screenful of every message. Should I read that one? What's he saying there? Should I bother paging down for 6 screenfuls to get to the third item mentioned in the subject line? What if he says something wonderful that makes all cybernetics obsolete? I might miss it! But I've got to get some work done today!

Other times, the subject is interesting (I WANT to know people's views on influence vs. control), but there's no way in hell I'm going to go through the whole thing with you people.

Further, most of these threads are actually excellent ARGUMENTS. Now you know the structure of an argument: it usually ENDS by either deciding a point, or at the very least clarifying a point of disagreement. This is easily LOST in all the verbiage.

I KNOW that there are other subscribers who share my feelings.

So, a request: as a thread is coming to a close, SUMMARIZE! Just something simple: Cziko said X, Bill said no, X', Marken retorted Y, and they all came around to see that Bill was right (or whatever).

Indeed, /Closed Loop/ would actually be USEFUL to me if it contained such summaries. I can always download the logs.

Think "highlights"! Just a suggestion.

Cliff Joslyn

Date: Tue Sep 22, 1992 4:51 am PST
Subject: Next CLOSED LOOP; "Purposive disturbance"???

From Greg Williams (920922)

(Re: Next CLOSED LOOP)

>P. Sibun (920921.1200)

>humph. i bet i wasn't listening when i was told about this.

Maybe you weren't told previously. (OOOPS!) I'm not sure whether there is anything in the CSGnet introductory materials about CLOSED LOOP. We should add a brief explanation if it isn't there already. Who do I talk with about this? Gary?

>my inclination is to say no, but maybe if you told me more about it, like who >actually reads it and how much editing you do, i'd change my mind.

First, I want to make it clear that I twist no arms to get permission to reprint. To date, I've been able to work around posters who haven't given permission (i.e., "In reply to a post from such-and-such [a non-permission-giver], Bill said: ... "), although it gets a little awkward at times.

CLOSED LOOP is "free" with membership in The Control Systems Group. It is published quarterly. Its main purposes are to give folks who are interested in PCT ideas but aren't on the net some insight into what is going on here, and to provide an accessible record of some significant threads on the net. It might also provide a springboard to a refereed JOURNAL OF LIVING CONTROL SYSTEMS. A few (?) nonmembers have ordered back issues (from Mary Powers, \$6 each, I think). Each issue is 72 pages long, laser-typeset. I try to edit for continuity and grammar, period. So far, I've made one fee pee which was called to my attention (by Dag F., who didn't request an correction in a subsequent issue). If I screw up, I'd be happy to print a retraction/apology/correction/rewording in the next issue, and the person offended may cancel permission for future us of his/her posts.

Ed, why don't you send a sample copy of CLOSED LOOP to Penni. She can look it over and consider her decision at leisure. I am far enough along on the next issue to know that I don't want to use any of her posts in it. I think the whole issue will be devoted to "Standards" and closely related topics.

>Bill Powers (920921.0800)

>I can see why you call it purposeful, and why it is truly control, but
>why do you say it requires influencing "the way others control?" It
>can't be done if it DOES influence the variables others are
>controlling for (eating and taste). How does supplying a particular
>kind of food influence how you eat it (other than the mechanics)? It
>seems to me you would eat food that Pat considers unhealthy in exactly
>the same way you'd eat food she considers healthy. What has she
>influenced other than the composition of what you put in your mouth

>(without altering its tastiness)?

All an organism can affect DIRECTLY (both you and I assume) is that organism's physical environment. So Pat is DIRECTLY "influencing" (my dictionary says: "the power of persons or things to affect others, seen only in its effects") her physical environment, which is also the kids' physical environment. The changes in the kids' physical environment which are (at least partly) due to Pat then influence (or "disturb," the technical term in PCT, which has the unfortunate lay connotation of "harm," whereas "influence" is more neutral) the operation of the kids' control systems such that, given their unchanging reference signals at high levels (here, perceiving their eating "tasty" food), some of their lower-level reference signals will be altered so as to enable actions which satisfy the high-level reference signals. So the way the kids' control systems are working is different than it would be in the absence of Pat's purposive influencing. (Would you prefer "purposive disturbing"?) Many of the kids' lower-level reference signals are different WITH Pat than they would have been WITHOUT Pat. That is what I meant by Pat influencing the way the kids are controlling. Of course, as I said in my last post, Pat's wants (to be met with generating conflict in the kids' control systems) must NOT be contradictory to the kids' high-level reference signals for eating "tasty" food.

Perhaps a clearer example is convincing a person, who wants to be correct (in general -- a high-level controlled perception) about matters of "fact," that in a particular case, he/she is wrong. If the person is convinced that he/she is indeed wrong, then his/her control processes could be vastly different from what otherwise would be. In fact, the person might even "override" intrinsic reference signals and choose "death before dishonor." Or do you claim that the perception of "dishonor" is intrinsic?

Getting into even more contentious territory, if the purposive influencer fosters reorganization in the influencee and then influences when the reorganization stops (i.e., by posing a problem having a determinate solution), the influence can result in PROFOUND differences in the operation of the influencee's control system, versus the operation in the absence of influence.

>But when Pat chooses food to feed you that does not disturb your
>perception of eating or your perception of tastiness, how can you say
>that this "enables" control by the influencee of the influencee's
>desired perceptions?

I mean that it enables control in the manner which satisfies the influencer's desired perceptions, AS WELL AS the influencee's desired perceptions.

> Do you mean that if Pat weren't putting SOME kind of food down, you'd be
>unable to eat tasty food?

I mean that if Pat weren't serving "healthy" (to her) food, the kids wouldn't eat (except occasionally, by accident) "healthy" (to Pat) food.

>Why isn't her presenting tasty food just an aiding disturbance that allows
>you to relax your own efforts to get tasty food?

It is. But I personally prefer the neutrality of "purposive influencing" to the negative-sounding "purposive disturbing." Nevertheless, I can bow to precedent. Whatever we call it, we can relate it back to folk/sociological categories of manipulation, dominance, bargaining, facilitation, exploitation, etc. The important thing is to show what those categories amount to as explained by PCT. Why exploiters have to do certain things, but not others, and all that.

>The distinction between environmentalism and organismism has to take
>into account the fact that for a given organism it makes a lot of
>difference whether the environment in question is living or nonliving,
>and if it is living, whether its organization extends to the same
>levels of control as in the given organism.

Yes.

The non-organismic environment is passive; it doesn't aid or resist being pushed into various forms.

Certainly.

>It can affect organisms, even kill them, but when organisms find ways of
>controlling the effects, the environment does not retire to think up a new
>strategy.

Certainly.

>The question of autonomy is, I think, best understood in terms of
>organisms versus nonorganisms. The purposes carried within organisms
>and the machinery for carrying them out render organisms autonomous in
>all the practical physical regards that matter to their survival.

If "autonomous" simply means "successfully PCT-controlling," then I have no problems with this. But in questioning the ideology of extreme autonomism, I am calling into question the claim that the "trajectory" through time of a organism's controlling is not influenced by the organism's environment. I claim that REGARDLESS of whether the environmental influences are purposive or not, controlling IS influenced by the environment. In the "short-term," PCT says that actions (and hence, in general, low-level reference signals) are influenced by environmental disturbances AS WELL AS current high-level reference signals. In the "long-term," PCT says that where reorganization stops is influenced by the nature of the problem set by the environment (which triggered reorganization in the first place) AS WELL AS by the pre-reorganization control-system organization and (totally or partly) random or pseudo-random generation new ("trial") control-system organizations.

>The underlying physical autonomy remains.

The control systems remain. If they are not conflicted, the autonomy (successful PCT-controlling) remains. Is that all you are saying? Sometimes in the past it has sounded like so much more -- as if control systems (or even Selves therein) are masters of their destiny in a manner opposing Skinner's extreme environmentalism (the environment is master of the organism's destiny,

and there is no efficacious Self/Will"). PCT says that neither extreme is correct, but that each has an element of truth in it.

>Only force, which strikes through all the organizational defenses against
>disturbance and reaches to the core life-support system on which all the
>others depend, can maintain an unfair agreement.

That is, can maintain an agreement PERCEIVED AS UNFAIR by a party to the agreement. If there is deception, a party might NEVER decide that an agreement is "unfair" (to him/her), even though he/she would claim it to be "unfair" if the deception were revealed. So this is another purely verbal problem: we must be clearer about whether who is making the "unfair" judgement, and when. The possibility of deception obviates simplistic arguments that an individual should NEVER be taken seriously when he/she complains that "somebody else caused my problems." He/she might be justifiably complaining that "somebody else deceived me."

>I think that coequal organisms constantly use the Test for the Controlled
>Variable on each other. But they don't do this in order to discover ways of
>controlling the other organism. They do it to avoid trying to control
>something the other organism is already controlling.

I think they do both. For example, the con man must try to find out whether the prospective mark is controlling for wanting a lot of money AND not controlling for avoiding illegal activities. Pat must try to find out what the kids want about food as well as that they don't mind if it is "healthy."

>The Test is easier to apply as a way of finding out what the other ISN'T
>controlling than as a way of finding out what it IS controlling.

Why? In either case, with people, often you can simply ASK, either way.

>Searching for the other's controlled variables is important only
>when you have some positive reason to want to manipulate the other's
>actions, the way the other moves, to suit yourself.

Which happens ubiquitously! Hey, YOU called it manipulation, I didn't!

>It's [getting salt passed] just one of those little understood bargains.

Yes, one of the multitude which deserve PCT explanations! We're well on the way to that goal, I think, even though it took a while to stop spinning our wheels.

Best,

Greg

Date: Tue Sep 22, 1992 6:39 am PST
Subject: *really* naive questions

(from: eric harnden 920922.0940)

i told you these would be coming. mostly, they stem from my effectively zero knowledge of psychology.

1) mr marken, in his letter to estes, reasserts the essential argument that study should be directed not toward behavioral outputs but toward controlled inputs. further, he indicates that the reaction to this and other assertions is 'yeah, so what... that's obvious', and goes on to say that it doesn't seem like its obvious, as demonstrated by the experimental modes of mainstream behavioral researchers. now, i was under the impression that behavioral studies often involved the presentation of stimulus to an organism, and an observation of response. i was further under the impression that these stimuli were presented under a not entirely unfounded set of assumptions of what, in the local terminology, the organism is controlling for. please explain to me why it is invalid for a mainstream behaviorist to say 'yes, well of course the organism is going to avoid pain, and prefer food. in this supposedly new jargon, i am being told what i already know: that it is controlling for lack of pain and hunger. what interests me is the characterization of its outputs, since it is the nature of these that ultimately affects itself, its environment, and other organisms.' now, i understand that this demonstrates, on my part, a basic ignorance of both classical psychology, and (so far at least) pct. but i am hoping that, by framing the question so baldly, i can elicit a response strong enough to clear a few of my cobwebs.

2) mr powers said something about 'reserving the term "control" for the closed loop situation.' why is the experimenter not seen as part of the loop? not to imply that this then gives the experimenter the 'control' against which mr powers is arguing. but it seems pushing it to then imply that the experimenter is then 'controlled'. but s/he definitely must fit in somewhere... i think.

Eric Harnden (Ronin)

Date: Tue Sep 22, 1992 11:49 am PST
Subject: *really* good questions

[From Rick Marken (920922.1000)]

eric harnden (920922.0940) asks some VERY GOOD questions:

>1) mr marken, in his letter to estes, reasserts the essential argument
>that study should be directed not toward behavioral outputs but toward
>controlled inputs.

We're all friends here on the net -- you can call me Dr. Marken (just kidding -- please call me Rick).

>now, i was
>under the impression that behavioral studies often involved the presentation
>of stimulus to an organism, and an observation of response.

yes.

> i was further
>under the impression that these stimuli were presented under a not entirely
>unfounded set of assumptions of what, in the local terminology, the organism

>is controlling for.

Not formally. This is certainly true in operant conditioning studies where the researchers "know" that the organism is controlling for some aspect of reinforcement. Of course they wouldn't say it that way. But they sometimes break down and say things that imply that they know that their experimental manipulations will work only if the animal WANTS the food, water, or whatever is the "reinforcer". The problem is that they don't really know what "wanting" is (it is a reference for a particular level of a perceptual variable) so they don't know how to study it. Instead, they are content to study the noticeable side effects of achieving wanted perceptual states -- the disturbance resistance process. They never think of trying to really nail down WHAT perceptual variable(s) is (are) being controlled. For a good example of the process of testing for controlled perceptual variables in an operant conditioning situation, see Powers article in Behavioral Science (I think in 1971). You will see in that article that it is possible to be far more precise about a controlled variable than, say, "wants to avoid pain". Bill shows quantitatively how it is possible to tell that the rat in a "shock avoidance" experiment is controlling (as best it can) the interval between shocks (as I recall). He also showed that this variable (interval between shocks) is a better representation of the controlled variable than at least one reasonable alternative -- the probability of a shock. So "the test" is a process of honing in on the best representation of a controlled variable -- where "best" means a representation that allows you to predict EXACTLY how the organism will react to disturbances of that variable.

>please explain to me why it is invalid for a mainstream
>behaviorist to say 'yes, well of course the organism is going to avoid pain,
>and prefer food. in this supposedly new jargon, i am being told what i
>already know: that it is controlling for lack of pain and hunger.

Because the behavioral psychologist does NOT already know what is being controlled without doing the test. PCT is not just a change of jargon; it is a working model of behavior. It says that behavior is organized around the control of PERCEPTUAL VARIABLES. An observer cannot see what perceptual variables are being controlled just by looking. Well, perhaps the observer can get some COARSE idea -- like "avoid pain". But that's too vague for PCT modelling. As Bill's Behavioral Science paper (also reported in the last chapter of his book, BCP) shows, it is possible to get quite precise about exactly what perceptual variables are being controlled. The "coin game" described in BCP is also a good illustration of the how to of get a good idea of what perceptual variable(s) another organism (in this case, a person) is controlling -- and how difficult this can be. I also report a study (in my Mind Reading collection) where I show that it is possible to tell that a person is controlling the area rather than the perimeter of a quadrangular shape. The determination of controlled perceptual variables is not necessarily easy to do -- there is a lot of skill and creativity involved. And it has not been done AT ALL in conventional psychological research (actually, I think Piaget got close to doing a version of the test).

Once you know what is being controlled, you can understand how the organism will act in MANY different situation -- because it MUST act certain ways if it is to keep the perceptual variable under control. But the fundamental thing to understand about organisms, according to PCT, is what they are controlling -- what kinds of perceptual variables they control. Such a catalogue of

controlled variables is not yet available in the literature of conventional psychology; this is what we need in PCT. Developing such a catalogue is one piece of work that we PCT researchers can "get on with" -- once we stop kvetching about conventional psychology.

> what

>interests me is the characterization of its outputs, since it is the nature
>of these that ultimately affects itself, its environment, and other organisms.

Yes, this is the interest of conventional psychologists. But such studies reveal nothing much about the nature of the organism that produces such outputs. The relationship between inputs (usually disturbances) and outputs will always be puzzling, different and usually statistical because the outputs that are observed depend on 1) other influences that are also acting on the controlled variable 2) secular variations within and between organisms in the level at which they intend to keep a controlled perception and 3) differences across organisms in the perceptual variable that is actually controlled. These three factors make it CERTAIN that the results of the conventional methodology (which looks at arbitrary measures of action -- arbitrary with respect to the perceptual variables that are actually controlled) will reveal nothing about what the organism is actually doing (controlling certain perceptual variables). The outputs may be fun to watch but they are truly superficial -- and will always be puzzlingly different across organisms, situations and times.

Best regards Rick

Date: Tue Sep 22, 1992 1:22 pm PST
Subject: Piaget and THE TEST

[from Gary Cziko 920922.1835 GMT]

Note. I am going to TRY to remember to stamp my posts with Greenwich Mean Time (GMT), even if my mail system won't do this for me in the header. I am also going to try to get my mail system operator to change this as well. This is really the only time that makes sense for CSGnet on which the sun never sets. It bothers me when Marken answers Harnden's question before it has been asked (I know Rick is sharp, but not THAT sharp).

Rick Marken (920922.1000) mentioned parenthetically:

> (actually, I think Piaget got close
>to doing a version of the test).

Could you elaborate a bit on this. I suspect that this may be a way of introducing PCT to the considerable number of developmental psychologists out there still influenced by Piaget (I ran into quite a few of them when I spent a sabbatical year at the University of Geneva--I wonder why?).--Gary

Date: Tue Sep 22, 1992 2:37 pm PST
Subject: What mailers control for

Bill Cunningham for Gary Cziko. Sorry to inconvenience others.

Gary, if gets through as expected, `vmd.cso.uiuc.edu' is still recognized as valid host by my mailer. Can't speak for John.

Date: Tue Sep 22, 1992 3:03 pm PST
Subject: Re: questions

From Tom Bourbon [920922 -- 16:45]

Re: *really* naive questions; (from: eric harnden 920922.0940)

>mr marken, in his letter to estes, reasserts the essential
>argument that study should be directed not toward behavioral
>outputs but toward controlled inputs. further, he indicates that
>the reaction to this and other assertions is 'yeah, so what...
>that's obvious', and goes on to say that it doesn't seem like its
>obvious, as demonstrated by the experimental modes of mainstream
>behavioral researchers.

The reaction Rick described occurs very often. We have seen it numerous times in the rejections of manuscripts in which we report the results of modeling control behavior with PCT. The most common form it takes is, "But we already know that." Next most common is, "But *everyone* knows that." I guess psychologists who use the second form are being generous and suggesting that even non-psychologists understand what is going on. Truth is, most "common folks" DO understand this -- but hardly any psychologists do. The research methods used in behavioral science make that fact perfectly clear.

>now, i was under the impression that behavioral studies often
>involved the presentation of stimulus to an organism, and an
>observation of response.

Your impression is correct. There is your proof that they don't understand the phenomenon of control. They believe there are "stimuli," acting as discrete independent antecedent variables, followed by discrete dependent consequences called "responses." Do you know of places in the behavioral literature where scientists clearly, consistently, and unambiguously state that what most people call "stimuli" are really disturbances (by any name they care to call them) that affect variables controlled by the individual who is the object of study? And do those same scientists clearly, consistently, and unambiguously state that what most people call "responses" are really actions that nearly perfectly eliminate the effects of disturbances to controlled variables? My questions are not meant to be trivial, nor are they "put downs" of your post. I do not know of statements like those I mentioned -- except in the PCT literature -- and I am willing to admit that I might have missed something.

>i was further under the impression that these stimuli were
>presented under a not entirely unfounded set of assumptions of
>what, in the local terminology, the organism is controlling for.
>please explain to me why it is invalid for a mainstream
>behaviorist to say 'yes, well of course the organism is going to
>avoid pain, and prefer food. in this supposedly new jargon, i am

>being told what i already know: that it is controlling for lack of
>pain and hunger.

Here, your impression is far from what I read in the behavioristic literature. I know of no place where traditional behavioral scientists, especially those who call themselves behaviorists, who speak of their animals "controlling for" anything. The entire language of the field is dominated by claims that behavior IS CONTROLLED BY environmental stimuli. Many people who read the behavioral literature ASSUME, as you seem to have done and as I often want to do, that scientists MUST think of their animals as controlling -- as seeking and avoiding. But those terms are explicitly proscribed in the more radical forms of behaviorism. Skinner, to the night before he died, rejected the use of "states" like pain or hunger as explanations for behavior. YOU already know organisms control for their experiences -- don't be too generous in attributing the same level of understanding to "real" behavioral scientists. Make them prove it.

My experiences (bad experiences at the hands of reviewers and editors) have taught me that unless and until a person demonstrates that her or his theory or model of behavior explicitly addresses the issues I just mentioned, I should not assume the person "knows" about the phenomenon of control and about control systems.

>what interests me is the characterization of its outputs, since it
>is the nature of these that ultimately affects itself, its
>environment, and other organisms.' now, i understand that this
>demonstrates, on my part, a basic ignorance of both classical
>psychology, and (so far at least) pct. but i am hoping that, by
>framing the question so baldly, i can elicit a response strong
>enough to clear a few of my cobwebs.

There is nothing wrong with characterizing outputs. It is just that when you discover they are part of the phenomenon of control, you realize the way they are characterized in most of the behavioral sciences is wrong. W - R - O - N - G. If you look at them as caused by antecedents, and ignore the fact that they control inputs, you miss the whole point of what most actions are about. Of course hardly anyone involved in traditional behavioral science recognizes that fact, or else they wouldn't stay involved. When you find outputs that control perceived inputs, you discover that none of the traditional "laws" apply. That claim does not go down well with reviewers and editors, even when you back it up with solid data.

Believe me, when one manuscript goes through repeated rejections, all accompanied by claims that, "We already know that," but all revealing ignorance of the phenomenon of control, it is easy -- very easy -- to think, "Why bother?"

Tom Bourbon

Date: Tue Sep 22, 1992 3:07 pm PST
Subject: Influence; let's change gears

[From Bill Powers (920922.0900)] Cliff Joslyn (920921) --

>So, a request: as a thread is coming to a close, SUMMARIZE! Just
>something simple: Cziko said X, Bill said no, X', Marken retorted
>Y, and they all came around to see that Bill was right (or
>whatever).

A fine idea, provided we can figure out what we said. A lot of the interchanges represent a process, not a conclusion. Sometimes most of what's passed back and forth is junked because only the final form is even coherent. If this were a journal, all the backtracking and false starts would be hidden. I'm glad it isn't, because with this looser format it's OK to change your mind any number of times before settling down to a point of view. We don't have to give the impression that we speak nothing but Truth.

But your point is a good one and I will occasionally see what I can sum up.

Greg Williams (920922) --

>The changes in the kids' physical environment which are ... due to
>Pat then influence (or "disturb," the technical term in PCT, which
>has the unfortunate lay connotation of "harm," whereas "influence"
>is more neutral) the operation of the kids' control systems such
>that, given their unchanging reference signals at high levels
>(here, perceiving their eating "tasty" food), some of their lower-
>level reference signals will be altered so as to enable actions
>which satisfy the high-level reference signals.

But the instance we're talking about -- controlling the nutritional value of the food -- doesn't involve any disturbances technically or otherwise; just the opposite. Pat is taking care NOT to disturb what the kids are controlling for: eating, and taste. What she is influencing (without evoking resistance) is a variable of no perceptual significance to the kids: what is in the food she is giving them. She hopes that this variable, the way she controls it, will have some beneficial effect on their innards. Not being able to monitor that effect, she can't have any systematic influence on it, or control it. "It", to be perfectly clear, refers to the effect inside the kids of ingesting one kind of food rather than another. All she can do is manipulate the external variables that ought, theoretically, to have an internal effect.

One reason we're having so much trouble on this simple point is that you keep wanting to get "influence" into it when this is a very poor example of influencing anything but Pat's own perceptions. When I try to bear down on you to be explicit about the exact relationships involved, you shoot out sideways like a watermelon seed and land somewhere else.

The case you've actually landed on in the above paragraph is one of directly applying a disturbance or an influence to a controlled variable, thus eliciting a change of actions aimed at opposing the effect of the disturbance and keeping the controlled variable undisturbed. Rather than avoiding any effect that would be resisted, you deliberately invite resistance by affecting something that's under control.

To do this, Pat would have to, for example, present the kids with less food than she knows they want, leaving their desire to eat less than satisfied. Or she would have to present them with food with a taste she knows they dislike. They would then probably take some kind of action to get some more to eat or

something better-tasting. If that was what she intended, then she could be said to be influencing, on purpose, their attempts to get more or better-tasting food, doing so by creating a disturbance that could only be corrected by such attempts. I don't know why she'd do that, but if she had that goal she could control for it.

It's just too general to say that "changes in the environment" influence actions. They may or they may not, depending on whether those changes are sensed by the other, and whether the sensed changes disturb some perception that's under control. If the changes are designed deliberately to avoid disturbing any controlled variable, then they will not affect any actions.

And it's too general to say that such changes, even if they do amount to disturbances, "influence the operation of the .. control systems." Influence WHAT about the operation of the control systems? All that can be influenced is the action used to correct any induced error, and that can be influenced only in the sense of making more or less of it necessary. The nature of the action is already determined by the organization of the control system, which takes into account the environmental link through which action affects perception. This mode of influence can have no effect unless there is a variable already under active control by the system. That means that the organism is already using a particular action or set of actions to maintain a specific controlled variable at a reference level. By applying your influence to the state of the controlled variable, you simply call on the control system to increase or decrease the actions it is already using for control. You can claim credit for those increases and decreases, but not for the nature of the action. That was established before you arrived on the scene.

From the standpoint of the control system whose actions are being influenced by disturbances in this way, it doesn't matter whether those influences are being used for a purpose or simply occur accidentally. Similar influences probably occur all the time, naturally. There's no difference in the action used to counteract an accidental influence or an intentional influence. In either case, the organism prevents any changes that matter to it, by varying that which does not matter to it: its actions. It is of no concern to the control system that another person wanted to see those actions, for private reasons, either to provide some imagined help or some imagined harm.

As I understand your goal in this argument, it is to prove that the environment has influences on the organization of the organism's control systems, shaping the organism in some causal way. This goal includes proving an effect of both non-living and living influences on the organism, independent of its heredity.

I have no objection to reaching such a conclusion. I don't think, however, that it can be reached by the route you've taken so far. All the specific examples of the shaping of behavior that have come up turn out to be examples in which an already-existing control system prevents environmental effects from influencing what matters to it, its structure of goals. Even when a disturbance results in a shift of lower-level reference signals, when the disturbance is removed those reference signals go right back to where they were before. Disturbances that are successfully resisted have no effects that outlast the disturbances. At best, they illustrate how a control system operates. They don't illustrate how its organization can be changed so it will

control something else or control by a different means after the disturbance disappears.

If you want to find ways in which the environment can actually alter the organization of a control system, rather than just inducing transient changes in its actions, you will have to turn elsewhere. One place you can look, and have already mentioned, is in areas where reorganization becomes involved. Nothing important about a control hierarchy is going to change unless there is reorganization. We've been going around and around talking about changes in what the organization is DOING. That's only changing the content of behavior, not the form.

Mary brought this up the other day, and I'm just getting around to seeing the significance in what she said. All that stuff about con-men and advertisers and feeding kids healthful food is beside the point. None of that stuff changes how people are organized. It just says that they will do one thing instead of another, maybe. Whatever the people do, they will do it with a particular organization of control systems, and exercising those control systems or bumping them up against each other is not going to change anything.

Over to you.

Best, Bill P.

Date: Tue Sep 22, 1992 4:08 pm PST
Subject: "Naive questions"

[From Bill Powers (920922.1700)] Eric Harnden (920922.0940) --

Eric, if you hadn't asked those naive questions, I might have asked you to do it. You've expressed exactly the problems that most psychologists seem to have with control theory. I keep hoping that there will be some magical way to answer such questions, one that will show control theory in so clear a light that it will all become plain in a flash, as it has done for all the people who now call themselves control theorists of the PCT sort. But that illumination has to come from within; there are no magic words for those who haven't prepared themselves to hear them. And when the hearer is prepared, the words are banal and simple. It's all perception. All action is control. All organisms are organized to control what happens to themselves, not to react blindly. Hearing those words, people who already understand nod and say yes, that's it all right. That's how we work. Those who haven't seen the simple truth yet ask what the hell we're talking about.

As you say, it's no surprise for a psychologist to learn that behavior is motivated by desires for food, water, sex, companionship, and other such things we require to live. This is the first dim view of what is really going on, and it would take a really stupid or really stubborn person not to admit that behavior has such things as its ends -- even as its purposes. But there is a huge gap between recognizing these abstract sorts of "requirements" and realizing that all behavior is "motivated," that most motives are not such vitally important or philosophically significant things, but simple and ordinary things, the stuff of ordinary life as it goes on from moment to moment.

The littlest words are probably the oldest, and point to the most important things in life. One of the littlest words is "do." We speak of behavior as what organisms "do." Control theory unpacks this little word into its details. It asks, and answers, the question "How do you do what you do?"

In conventional psychology, behavior is taken for granted as a given. Organisms "respond." Their responses can be simple or complex, but they are treated as simple acts, occurrences, units. Behind this concept is a rather vague notion that somehow the nervous system makes the muscles produce those responses; a pattern in neural impulses is transformed into a pattern in space and time outside the organism. That is how "doing" happens in the traditional view.

But control theory shows us that this can't be the right story. It calls attention to all the variations that slip in between the commands of the nervous system and the consequences that we recognize as behavior, and shows that the assumed simple causal chain does not in fact exist as it's been assumed to exist. And it shows that the only way to explain any regular act, however trivial and simple it looks, is to trace out the closed loop of cause and effect in which the act is embedded.

So sometimes, when I hear psychologists and others talking so assuredly about the world of organisms as they see them, and when doubts creep into my mind, I will perform some simple act and look most carefully at it and feel it with close attention to what is happening. I will hold out a hand, and say "There, I see a hand. I see the fingers stretched out. And now I will make the fingers look and feel as if they are curling up, more and more, until I feel the fingers against the palm, see the shape of a fist, feel the sensation of effort in the forearm muscles." I look and feel very carefully, and I confirm that that is all there is. What my hand is doing is known to me only through sight, kinesthetic and tactile sensation, and the sense of effort. There is nothing else. It is truly all perception.

By going through this little exercise over and over, and looking even more carefully, I can detect just in advance of the changes in perception something almost too familiar to name: the sense of willing that the act should occur. There is a foreknowledge of what I will see and feel, a foreknowledge that is like imagining the sight and feeling. And in fact, to create the act in my perceptions, all I need to do is conjure up that imagined sight and feeling in a particular familiar way, and the actual sight and feeling obediently follow. I will that my hand look and feel open, and so it looks and feels. I will that it look and feel closed, and that happens, too, as swiftly as the imagination can change.

That's what a baby is doing. The baby isn't just waving its arms and legs around; it isn't just fascinated by the sight of its own hands. It's learning how to make those perceptions obey its will. This is difficult; there are connections to be made that grow into place only slowly. When a baby in the midst of cooing and wiggling its fingers and waving its arms suddenly breaks into anguished crying, I know what's wrong. It's trying to recreate a perception of its own hands and arms in the space around it, and the damned things won't do what it wants. I remember that, I think. Gaining control over these visual and tactile and kinesthetic perceptions takes endless trying for hours every day. There is no need for any external stimuli to cause this behavior. The baby is learning to be human, learning to control, and nothing

else is more important. The need to control, which results in what we call "doing," resides in the baby, as do the mechanisms for acquiring that control.

When I really need to check out the principles of control theory from the ground up, I don't go to a servomechanism textbook or reread what I have written. I go back to the source. I hold out a hand before my eyes and move the fingers, watching them and feeling them and sensing the play of the muscles moving them. It's all perception. All action is control. All organisms are organized to control what happens to themselves, not to react blindly. Action is the matching of perceptions to desired states of those perceptions.

It isn't control theory that tells me about these experiences. It's these experiences that tell me about control theory. Try wiggling your fingers. Maybe you will suddenly see what this is all about.

Best, Bill P.

Date: Tue Sep 22, 1992 5:42 pm PST
Subject: Re: Influence; let's change gears

> [From Bill Powers (920922.0900)]

>

> A fine idea, provided we can figure out what we said. A lot of the
> interchanges represent a process, not a conclusion. Sometimes most of
> what's passed back and forth is junked because only the final form is
> even coherent.

Exactly, and I wouldn't change any of that. But if in fact the final form IS coherent, then please let us know! "Here comes a good final summarizing place..."

Thanks.

Cliff Joslyn,

Date: Tue Sep 22, 1992 6:08 pm PST
Subject: Re: turing test

(penni sibun 920922.1400)

Date: Tue, 15 Sep 1992 08:38:36 -0700
From: "(Cliff Joslyn)" <cjoslyn@bingsuns.cc.binghamton.edu>

The fascination with the Turing test has always intruiged me. To me, the following are obvious. Perhaps I need to get educated.

1) At first consideration, X passing some test must only be a NECESSARY, and not a SUFFICIENT condition for X to have the property tested for. Thus X passing the TT is not determining that X is intelligent: surely it could be the case that something that is not "really" intelligent can be made to look like it to any finite approximation; indeed, that seems to be very common. Thus passing the

TT is at best non-determining EVIDENCE for the INDUCTIVE INFERENCE that X is intelligent.

but what is the criterion for intelligence? the turing test presupposes that the only criterion for intelligence is being judged intelligent. if there is exactly one condition, then meeting it is both necessary and sufficient. (the test can make mistakes, generating both false positives and false negatives, but that's an orthogonal issue.)

2) Point (1) depends on a PRIOR DEFINITION of the property being tested for, i.e. intelligence, so that the TT is performed relative to it.

yes, your point 1 depends on that. for the tt, there is no a priori definition of intelligence, though certainly the judges' prior experience is terribly important.

Therefore it would be possible in principle to distinguish between the "really" intelligent and the simulations. But that just begs the question: this "new" test just replaces the TT.

right. turing was not begging the question.

In the end, we must KNOW, or at least SAY, what intelligence is, and then construct the TT to measure X's satisfaction of the definition.

nope. your formulation requires such a definition; however, i'd bet you a lot (if i had it, which i don't!) that you will never ever discover such a definition that works in the real world. i don't think you should be spending your time looking for it.

I'm glad that Control Theory seems to be founded on less shaky methodological ground than traditional AI.

the turing test is not a methodology. it's more a philosophical basis.

=====

[From Bill Powers (920915.1000)] Penni Sibun (920914.2100) --

>i've been trying to establish that ``what they [sexes] REALLY are'' has
>two answers, and that these answers are distinct (sex-by-chromosomes
>doesn't always match sex-by-judgement). i've further insinuated that
>the answer that is used 99.99% of the time is the one by judgement.
>and i've left the conclusion to be drawn that the judgement answer is
>in fact the really real one (unless you're a doctor or scientist and
>have some reason to muck about in somebody's chromosomes).

Take the category "sex." If you base your notions of sex strictly on external appearances -- habits of speech or dress, mannerisms, length of hair, use of makeup, pitch of voice, attitudes toward children, roles in relationship to people in a disjoint category, and so on -- you will base your expectations and wants on those appearances. If, for scientific or other reasons, you decide to base your concepts of

sex on chromosomes, skeletal conformity, organ configuration, and the like, you will base your expectations and wants on these appearances instead.

just for the record, my distinction is between everything you can see w/o a microscope (including shape of genitalia) and chromosomes.

Which one of these ways of categorization is "right?" To me it's obvious that this is a non-question.

i agree.

So NONE of the ways of perceiving sex is the "really real" way.

no, i think one is always the really real way (except, as you point out, to weirdos like us who insist on sitting around and thinking about this stuff too hard). organisms just don't go around doubting things most of the time.

=====

[From Bill Powers (920915.1200)] Cliff Joslyn (920915) --

Right on the button, as usual. Your second point, about defining intelligence, is the most fundamental in my opinion. If I were constructing an intelligence test, I would make sure that everything I know know to do well would score high, and the things I've never mastered would contribute nothing. Intelligence, of course, means understanding control theory, so Dizzy Gillespie would score low, allowing me to get even for being a lousy jazz musician.

You'll notice that when I discussed the Turing test, I expressed it as trying to find out whether the entity at the other end was a human being or a program.

how does finding out whether the entity is a human or a program have anything to do with intelligence as you suggest defining it immediately above??

=====

[From Rick Marken (920918.1330)]

penni sibun (920914.1300) says:

>no. the test doesn't work backwards; it doesn't have anything to work >backwards to.

If this is true, then the Turing test doesn't seem to "test" much.

i don't see how this follows.

As

you describe it, the test involves judging whether the behavior I see seems intelligent or mindlike or whatever.

the test involves judging whatever based on anything you can perceive.

How can that judgement be wrong?

insufficient data, inattention, perversity (see below)...you can tell a judgement was wrong is if it is revised at a later point. (though that's not really sufficient, since the revision may be warranted.)

If this is really the Turing test then
it would probably be better described as the Turing experience.

i won't even ask you how you distinguish test and experience....

I just
recently had this experience in an airport where there were some Coke
cans dancing to the music coming out of a radio. Those cans seemed
like they were really bright; they passed the Turing test. Do you have
any reason to suspect that they didn't?

yes i do: everything i know about your views and style of argumentation.

I received and read Preston's paper. It was VERY long.

nonsense. it is typical of the length of journal papers, at least in the
journals i'm used to reading.

Thus, a measure of the success of the mechanism you build
(or postulate) for producing intelligent behavior is whether it produces
intelligent behavior.

a measure of success for building a bridge that supports cars is whether it
supports cars. so?

And, since the only standard for intelligent behavior
is that which is produced by humans (and described as such),

listen hard: the only standard for intelligence (i'm not going to play your
linguistic games of calling intelligence ``intelligent behavior'') is that it
be *judged* as such, by a human, who we stipulate is indeed intelligent. the
turing test, which compares programs to humans, tests to see whether programs
can judged as intelligent as humans.

then the
way to test the mechanism is to compare it's behavior (blindly) to that
of a human.

you like that word ``blind'' don't you? i suppose you are referring to the
fact that the original test proposed a teletype link. *you* started this
thread by making fun of a paper suggesting that the turing test be made more
comprehensive by being applied in a much less constrained setting, eg,
performing daily routines.

If the behavior of the mechanism is indistinguishable from
that of the human, then the mechanism is "artificially
intelligent".

you don't have to say it's *artificially* intelligent. that's not the important part. the test does however suggest it's intelligent, or ``can think.''

The Turing test (as I understand it) would be just fine if behavior were what Turing (and everyone else) imagined it to be -- actions and consequences of those actions.

the turing test is not about behavior. it lies in the interactionist realm, in my opinion. (i don't speak for turing here: i actually know very little about his philosophy beyond the one paper.)

The Turing
test focused on actions and their consequences

no. it focuses on an *interaction* between a judge and an entity.

the mindframe that you believe I should break. It's hard for me to imagine how I can break it -- even with your help -- because in the small domains in which I have used this mindframe, it has been most successful. But

anything works in a small domain. if your domain is small enough, you can take care of most or all of the contingencies. that something works in a small domain is fine and encouraging, but hardly definitive.

please keep trying to convince me. Maybe we could get more concrete and discuss your perspective in terms of experiments and models.

maybe we can get more concrete and talk in terms of phenomena in the real world.

=====

[From Bill Powers (920918.1600)]

Of course this isn't what the AI types would do. I don't think they really know what they would be looking for.

they know exactly what they are looking for. as whichever justice said, ``i know it when i see it.''

this applies to intelligence and porn and just about everything else in the world.

Didn't Penni say that the
best program fooled about half the people, at the last big contest?

the judges were regular, relatively computer-naive people off the street, just in case you thought something else.

That's about chance, isn't it?

right. if something is distinguished at chance, then by definition, it's indistinguishable. (i've been talking w/ a lot of japanese lately. when they say an ``l'' or an ``r'' i hear it right about half the time (unless the

context happens to make it very clear); their ``l/r'' pronunciation is indistinguishable.)

In other words, the best the judges could do with a machine that didn't give obvious clues about its machine origins was to guess whether it was "intelligent," yes or no.

they weren't doing anything w/ a machine, cause they never *that* they were dealing w/ a machine, only that it was a possibility.

i find it interesting that you are imputing so much to the judges. the individual judges weren't guessing, or at least you have no reason to suppose that's what they were doing. for all you know, each one felt 100% certain. and we also have no idea how they were making their judgements (the article i read didn't really go into that).

And they guessed at random, which means they thought a human being was a machine half the time, too.

they *didn't* guess at random.

They couldn't even identify intelligence for sure when they were interacting with a human being.

again, they weren't ``interacting w/ a human being''--they were interacting w/ something they knew was either a human or a program.

the structure of the test is not ``is x intelligent'' but ``only one of these is human (intelligent)--which?''

How did they expect to tell when it was a machine behind door number 2?

what?

Maybe that wasn't how it came out, though. Penni, how many people were misjudged as machines?

there were 6 programs, 2 humans, 10 judges, and 87 people in the audience (who got to compare notes w/ each other). the 8 terminals (6 progs, 2 hums) were ranked (rankers knew that 2 of the 8 were humans). the 66 aud. members who did a complete ranking ranked the hums first and second, the winning program 3rd, and the others indistinguishable. only 5 aud. members called the winning program a human; and each human was identified as a computer about 5 times. the aud. did ``better'' than the judges, probably cause they could confer (and weren't on the spot!). half the judges mistook the winning prog for a human. one other program was mistaken for a human 2ce, and another once. one human was judged a computer by 3 judges.

i don't think these statistics are complete, but they are all i have here.

And so on. No matter how you define intelligence, it comes down to doing something well or fast or both.

i think that's only a part of it. well, no, maybe that *is* what it is, if you are going to *define* it. but if you're just deciding whether it's there or not, who knows what it's coming down to.

cheers. --penni

Date: Tue Sep 22, 1992 6:30 pm PST
 From: Control Systems Group Network
 EMS: INTERNET / MCI ID: 376-5414
 MBX: CSG-L&UIUCVMD.BITNET@pucc.princeton.edu

TO: * Dag Forssell / MCI ID: 474-2580
 TO: Robert K. Clark / MCI ID: 491-2499
 TO: Hortideas Publishing / MCI ID: 497-2767
 TO: Henry James Bicycles Inc / MCI ID: 509-6370
 TO: Multiple recipients of list CSG-L
 EMS: INTERNET / MCI ID: 376-5414
 MBX: CSG-L&UIUCVMD.BITNET@pucc.princeton.edu

Subject: random verbiage

hello ladies and e. coli, this is the family man. i had a wonderful time meeting you guys in person and would like to give a special hello to the PLAYER. i have just started on a research project for my tech writing class and i have decided to analyze post-slave revolt Sparta in view of PCT. i think its a wonderful example on how on a macrolevel to contol x is to oblige oneself to restrictions. ex., if i decide to keep my car on the road i am obligated to act against disturbances that would change the relationship ipso facto limiting my action-parameters (possibly poor word choice) meaning i can't sleep, play baseball or build a house. i do have to keep my eyes open looking at the road ahead of me and use the steering wheel to counter disturbances. now on a macrolevel, if a group is going to oppress another in order to use them for slave labor that group of oppressers is limited in what they can do (this is blatant when the slave population greatly outnumbers the "master" population) [they cannot allow civil liberties to the oppressed (freedom of movement, of congregation, of bearing arms)] and they are obligated to do certain things - establish a powerful military, become isolationist to "keep their eyes on the road" and numerous other actions if the oppressed group is an extreme threat: establishment of military state with rigid hierarchy, disinterest in nonutilitarian endeavors (art, music, dance), etc. i hope this does not come off with any group mind overtones, that is the opposite of my intention. i do wish to show that on a macrolevel,also, control necessarily limits actions and obligates others. in fact that is the sin qua non of control.

Date: Tue Sep 22, 1992 10:25 pm PST
 Subject: random verbiage

hello ladies and e. coli, this is isaac the family man. i have already sent a message but i no idea whether it sent because this machine is the nexus of hell, AAAAAAAAAUUUUUUUUUGGGGGGGGGHHHHHHHHHHH!!!!!!!!!!!!!!!!!!!!!! if the original did send then ignore this rewrite. first i'm glad i met you guys and would like to give a big hello to THE PLAYER and a thankyou

you guys and would like to give a big hello to THE PLAYER and a thankyou ignore those last two lines i AM going to put my foot through this piece of *%#@##!!!! ... and a thankyou to the PLAYER and Runkel for sending their papers. i have just started on a reseach paper for my tech writing class. i've decided to analyze post-slave revolt Sparta in view of PCT. i basically going to attemt to show that control entails restriction and obligation on a macrolevel also. to keep my car on the road i am obligated to stay awake, keep my sight on the road ahead of me, and use the steering wheel to act against disturbances to that relationship. i am also restricted from sleeping , taking a shower, building a house, etc. this was shown clearly in Bourbon's demonstration where one person could "control" another's handle movement if 1) that person was able to maintain their desired perception 2) the "controller" obligated himself to certain actions and restricted himself from others [obligation-move in any way in order to produce desired action in the "controlled" person, restriction-ability to control own

cursor, close eyes, close other person eyes, leave, kill other person, etc.]. this is the same on a macrolevel;if a group is going to oppress another (assuming that the oppressed population doesn't particularly enjoy their lot and the oppressed population is as large or larger in numbers than the "masters") then one will see the same restriction and obligation of action-parameters. The oppressors CANNOT allow the slaves any civil liberties(freedom of movement,congregation,speech,bear arms). The oppressors are obligated to become isolationistic (to "keep their eyes on the road"), and establish or strengthen the military class. if the slave population is large and volitile enough then to keep them down more drastic measures are NECESSARY: establishment of military state and rigid hierarchy,abandonment or lack of concern for nonutilitarian endeavors (art,music,dance), etc. this was the case in Sparta where extreme social,political, and cultural changes were necessary to keep the slave population checked and working. to control them they were obligated to certain action (not in the least keeping the slaves alive) and restricted from others. it is NOT my intent that this has any group mind (fie!) overtones. obligation and restriction is the sine qua non of control.

Date: Wed Sep 23, 1992 3:33 am PST
Subject: Purposive influence

From Pat Williams (920923) >Bill Powers (920922.0900)

Hi Bill. I'm sick of all this talk of me (Pat) and the feeding of "healthy food" to my kids, so I'd like to try to clear up the confusion about this particular instance of what Greg has been calling purposive influence, myself.

>But the instance we're talking about -- controlling the nutritional
>value of the food -- doesn't involve any disturbances technically or
>otherwise; just the opposite. Pat is taking care NOT to disturb what
>the kids are controlling for: eating, and taste. What she is
>influencing (without evoking resistance) is a variable of no
>perceptual significance to the kids: what is in the food she is giving
>them.

Yes, I am taking care NOT to disturb what the kids are controlling for! That's the whole point. And I am also achieving my goal: they are eating (not just being served) "healthy" food. Everyone is happy! It is purposive because I'm doing it on purpose; and it is influence because the kids are acting differently (eating "healthy" food) than they would be if I served them just any old tasty food or let them scrounge for themselves (then they would be eating tasty food, most of which would not meet my definition of "healthy"). Yes, I am influencing a variable of no perceptual significance to the kids! But that doesn't MATTER. It is of significance to ME -- I would feel badly if they were consuming what I consider unhealthy food. I DON'T CARE whether it is of perceptual significance to them! I'm getting what I want by changing their ACTIONS (not their GOALS), and they are getting what they want.

>She hopes that this variable, the way she controls it, will have
>some beneficial effect on their innards. Not being able to monitor
>that effect, she can't have any systematic influence on it, or control
>it. "It", to be perfectly clear, refers to the effect inside the kids
>of ingesting one kind of food rather than another. All she can do is
>manipulate the external variables that ought, theoretically, to have
>an internal effect.

That's correct; I do hope that feeding the kids my definition of "healthy" food does have a beneficial effect on their innards (and I do have some evidence that it does -- for example, they have NO cavities), but I realize that because I have no direct monitoring of how their innards deal with the food, I can't very effectively control that. All I am controlling for is my perception that they are eating "healthy" food. My definition of "healthy" food doesn't really matter as long as I know exactly what it is -- it could be food that contained absolutely no animal products, or it could be food that was almost entirely derived from animals. It doesn't really matter if their health is improved by this food or not (unless it is greatly affected), because, as you say, I can't monitor the changes effectively. What matters to me in this instance of purposive influence is that they are eating what I define to be healthy food. That's all there is to it.

I'm not succeeding in getting my kids to do something they don't want to do. I'm only succeeding in getting them to act the way I want them to through my understanding of what they want! That's ALL Greg has been saying that purposive influence is. I can think of hundreds of examples of this definition of purposive influence -- I just have to look around me in families, schools, workplaces, stores, etc. The "agreements" and "bargaining" and "bed carrying" and "passing the salt" you have been talking about are all examples of purposive influence (both parties are getting what each wants by understanding and taking into account the other person's wants). Adults pay to take classes (painting, welding, auto repair, gymnastics, etc.) to be purposively influenced all of the time. You seem to dismiss all these interactions as "unimportant," since the influencee is not being made to do something he or she doesn't want to do. I think most people would consider all of these interactions to be IMPORTANT -- they are what large portions of peoples' lives consist of.

All that Greg wants to do is to look at the interactions between people that "work" (without force or threat of force) from the point of view of PCT and see how they work, why they work, and what the limitations are. Why are you resisting this so much?

Pat

Date: Wed Sep 23, 1992 4:44 am PST
Subject: CLOSED LOOP AND OTHER TOPICS

CHUCK TUCKER [920923]

TOPICS: War on Drugs; Suggestions for "Elephant" paper; What Innards?; Influence and Control, Again and Again; CLOSED LOOP

CLOSED LOOP GW 920919

My vote is for an issue on "Influence and Control" (seconding Hugh's vote actually) but if I can influence you Greg (which I can't prove that I have even after you produce the issue !?!) I would suggest that it be arranged as a conversation between Bill and you with whatever "comments" by others that seem to be relevant to the matters mentioned in the conversation. The conversational format would show the "linkages". You may want to delete some of the repetition and select those statements which most clearly (to you, of course) make the point. Perhaps a preface to the issue would give a reader an explanation of the style and format used and point out that the original posts are available on CSGNet. Finally, a summary statement by you and Bill about the conversation and the major points each of you have been try to make might end the issue. I think it would be a nice monograph to show others the central ideas of PCT. Since many others have been silent during this conversation the next issue of CLOSED LOOP could be composed of comments by others on the net and elsewhere about this conversation. These two issues of CLOSED LOOP could be then turned into another CSG publication in book form. Sounds like a great idea to me!

INFLUENCE AND CONTROL WTP 920918.0800; 18.1600; 19.0900; 20.0999 AGAIN AND AGAIN GW 890918-3; 19.0800; 19-2; 20

I have reached the tentative judgement that there is no way for one human being to influence (in Greg's terms) another human being and still maintain the principles and processes outlined in the HPCT or PCT model as proposed by Bill Powers. Bill has been extremely consistent and even precise about his model even though he states quite often that it is just a model and still has to be tested; I can find no recent significant modifications (except the addition of level names, which I usually collapse) in his model. The latest attempt by Greg to discuss "linkages" (along with comments by Bill about his cat) convinced me that influence was NOT possible in the model (by the way, I think my statements about "taking" as in "taking the role or attitude of another", "taking another into account", and such notions as "socialization", "acculturation," "learning", are very similar to Greg's "linkages").

Greg keeps insisting that THE TEST will demonstrate that person A has influenced person B to control for X. But that is not so, THE TEST only shows (if done repeatedly and systematically with careful records) what B is controlling for BUT IT CAN'T SHOW THAT B TOOK X FROM A. A can only say that the X that B is controlling for seems to be exactly like the X I use and asked (told, demanded, etc.) B to use but since B is now using X it belongs to the "control system" of B and no longer to A even though A might use it also. Now this seems like a very "picky" point to make but it is crucial to PCT: my

perceptions are mine even though they may look exactly like the perceptions that you say are yours. Runkel writes of "borrowing" (I would say "taking" again) reference signals from another but still when I am using it it is mine and it may become yours (again?) to use when I "return" it or you "borrow" it back.

The PCT model has to insist on this autonomy or it will very quickly lose its differences from all other theories (as Rick indicates in his "elephant" paper). Thus, anything that looks like a "collective social act" is actually two or more persons controlling their own individual conduct in what appears to each and every party to action as influencing the others.

NB: This is a "higher level" judgement by the parties not a "lower level" one since a "lower level" analysis, precisely done, should show a tremendous variation in actions of the parties. I also believe that as we associate with others we are told (either from a S-R or another perspective) that we SHOULD AND DO INFLUENCE EACH OTHER.

The "moving the bed" example strikes home (no pun intended) since I just helped my daughter move furniture from Columbia to Baltimore (yes, others did look at us funny as we carried the bed up I 95). But, in most instances of moving a piece of furniture I have established a "division of labor" where one person "carries" the major load while the others "guide" or at least "don't resist". But, even in these cases I think that careful analysis would show that each party to the "collective act" is acting on their own.

If my tentative judgment (i.e., that influence [A affects or determines what B controls for] is not possible between human beings if they are negative feedback control systems) is turned into a "warranted assertion" with testing and evidence the question arises: Is the PCT (or HPCT) model useful for collective social action? That is an important question to answer for those of us who have a concern for human group life, social life, collective behavior or even Society. Maybe a new model has to be constructed !?!

WHAT INNARDS? GW 920918-2

Thanks Greg for the statements from Skinner which indicate to me that he (and perhaps other behaviorists) did not consider the "innards". I guess I just find it strange that someone would not wonder more than he did about the contribution of the "innards" to behavior but people I pay attention to (especially myself) firmly believe that the CNS is important to say nothing of the liver and bladder.

SUGGESTIONS FOR "ELEPHANT" RM 920918.1330

You are disturbing the readers of your paper that don't already firmly believe in PCT and you know this; the evidence is overwhelming (and has been for years). Now you can either find out what they are "controlling for" (usually the S-R model) and "control" them indirectly by "rubber-banding" them or you can find a more direct way to say (again) that all other models (lets pretend they have models for the moment) of behavior presume PCT but none of them clearly recognize it and if they did they would have available to them a whole host of empirical studies which would make their efforts all they have wished for: a true science of human behavior.

I would suggest that you get rid of the "elephant" metaphor (as I mentioned before - I think it is distracting) and use the same approach that John Dewey used in his 1896 "Reflex Arc" paper to show (with the vocabulary and technology of our day) the same thing he was trying to show: the human being is a negative feedback control system!!!

WAR ON DRUGS RM 920918.1330

I agree - the legalization of drugs (like alcohol is now) still is a method of forcing others to control their own drug using actions in a particular way (by decreasing them). The hope is that with legalization there would be a reduction of conflict in the "drug war" (raids by DEA agents and gang members killing each other) and that it would make it possible for "controllees" (see above note on influence) to reorganize themselves to reduce their drug use. But it does not eliminate all conflict or violence.

What do you think of my other suggestion - suspend all laws and law enforcement with regard to these special drugs. Perhaps, we could extend it to all drugs by eliminating the FDA and allowing anyone to self any substance they wished w/o any regulations or laws at all (the classic libertarian [actually this is Szasz's position on drugs and suicide] position)!?!

If these are not adequate ways to solve the problem (which I admit they are not), then what would you suggest from a PCT perspective? NB: this might be the sort of thing your readers want from you in your "elephant" paper: practical solutions to problems.

Regards, Chuck

Date: Wed Sep 23, 1992 4:47 am PST
Subject: PLEASE SEND YOUR OFFERINGS REAL SOON!!!

%%&&&%%&&&%%&&&%%&&&%%

A N N O U N C E M E N T

Please send title and abstract of a paper by October 7, 1992 for:

A PERCEPTUAL CONTROL THEORY OF MACRO STRUCTURES

to

Charles W. Tucker
Department of Sociology
University of South Carolina
Columbia SC 29208
V- (803) 777-3123
F- (803) 777-5251
H- (803) 254-0136
N050024 @ UNIVSCVM

Other important events at this meeting are:

A session on Kuhn' Self Theory organized by Bob Stewart with presentations by Carl Couch, Tom McPartland, Clark McPhail and Chuck Tucker

A Presidential address by Clark McPhail

MIDWEST SOCIOLOGICAL SOCIETY MEETINGS
APRIL 7-10, 1993
HYATT REGENCY
CHICAGO, ILLINOIS

%%&&&%%&&&%%&&&%%&&&%%

Date: Wed Sep 23, 1992 4:52 am PST
Subject: TUTSIM info; influencing ideology

From Greg Williams (920923)

Information about the TUTSIM simulation program for IBM-compatibles is available from TUTSIM Products, 200 California Ave., Suite 212, Palo Alto, CA 94306; phone 415-325-4800. Ask for the "Classic Edition" of TUT'S BLOCKS, a tabloid newsletter which contains several examples of systems simulated using TUTSIM.

>Bill Powers (920922.0900)

>It's just too general to say that "changes in the environment"
>influence actions.

...

>And it's too general to say that such changes, even if they do amount
>to disturbances, "influence the operation of the .. control systems."
>Influence WHAT about the operation of the control systems?

WHICH ACTIONS ARE PERFORMED BY THE CONTROL SYSTEM.

>All that can be influenced is the action used to correct any induced error,
>and that can be influenced only in the sense of making more or less of it
>necessary.

Yes. Well, actually, "more or less" in a not strictly real-number-line way, but possibly very complexly "qualitative," being multidimensional. Even in plain old rubber-banding, the actions needed to compensate for the applied disturbances are not simply "more or less" in one dimension, but simultaneously in TWO dimensions.

>The nature of the action is already determined by the organization of the
>control system, which takes into account the environmental link through which
>action affects perception.

Take another look at the generic PCT diagram. The action is determined jointly by BOTH the reference signal AND the disturbance, via subtraction, one from the other, at the comparator. To say that the reference signal is the SOLE determiner of the action is simply false. To say that the reference signal is "more important" than the disturbance in determining the action can be considered true for at least some ideologies -- but not for other ideologies.

>This mode of influence can have no effect unless there is a variable already
>under active control by the system.

Correct.

>That means that the organism is already using a particular action or set of
>actions to maintain a specific controlled variable at a reference level. By
>applying your influence to the state of the controlled variable, you simply
>call on the control system to increase or decrease the actions it is already
>using for control.

Yes, with "increase/decrease" interpreted broadly, as above.

>You can claim credit for those increases and decreases, but not for the
>nature of the action. That was established before you arrived on the scene.

"You" cannot claim FULL credit. Only SHARED credit. SHARED with the organism. The nature of possible actions, given possible disturbances, WAS established by the organism, but YOU select WHICH of the possible disturbances ACTUALLY is applied. The ACTUAL resulting action is determined by the organism's establishment of pairs of possible actions and possible disturbances TOGETHER WITH your selection of the actual disturbance. To say that one of these factors is "more important" than the other is ideology, not fact. To say that your contribution is "all important" is radical environmentalism; to say that the organism's contribution is "all important" is radical organismism; To say that the contributions are "both important" is a middle-way ideology which does not contradict PCT-science.

>From the standpoint of the control system whose actions are being
>influenced by disturbances in this way, it doesn't matter whether
>those influences are being used for a purpose or simply occur
>accidentally. Similar influences probably occur all the time,
>naturally. There's no difference in the action used to counteract an
>accidental influence or an intentional influence.

Correct. But that is beside the point. Purposive influence matters to the influencer ALWAYS.

I think I now am beginning to understand how your tenacious clinging to the organism's point-of-view, which serves you so well scientifically (one can simply get farther in psychology with models of "innards" than without, contra Skinner) seduces you into an ideological position at odds with PCT-science. The organism role IS more important than the environment's role to the

psychologist who wants to see order (invariable goals) behind the chaos (variable actions). But it is a mistake -- because it contradicts PCT-science -- to carry that importance over to ideology. Skinner made the opposite mistake, but because he didn't have a model of the "innards," he didn't have an opportunity to see the contradiction. You have that opportunity, but steadfastly refuse to take advantage of it. And, as I've said before, my ideology is that your refusal is bad for the more widespread adoption of PCT ideas.

>As I understand your goal in this argument, it is to prove that the
>environment has influences on the organization of the organism's
>control systems, shaping the organism in some causal way. This goal
>includes proving an effect of both non-living and living influences on
>the organism, independent of its heredity.

Yes, that's my aim in the second of our two arguments. The first argument involves purposive influence only, not environmental influence in general. Of course, they have been intertwined somewhat.

>All the specific examples of the shaping of behavior that have come up turn
>out to be examples in which an already-existing control system prevents
>environmental effects from influencing what matters to it, its structure of
>goals.

Not all. I've mentioned education, for example.

>If you want to find ways in which the environment can actually alter
>the organization of a control system, rather than just inducing
>transient changes in its actions, you will have to turn elsewhere.

"Transient changes" (I called them "short-term") are really all I need, but I think there can also be what I have termed "long-term" changes, requiring influencing the outcome of reorganization.

>We've been going around and around talking about changes in what the
>organization is DOING. That's only changing the content of behavior, not the
>form.

You think the "only" is "important"; I don't. But it doesn't really matter vis a vis PCT-science -- or to purposive influencers. The content of the influencee's behavior changes, and the influencer gets what he wants. Later, the influencee might regret or appreciate enormously that changed content. That it was not changed form, too, is a trivial consideration for the influencee.

>Mary brought this up the other day, and I'm just getting around to
>seeing the significance in what she said. All that stuff about con-men
>and advertisers and feeding kids healthful food is beside the point.
>None of that stuff changes how people are organized. It just says that
>they will do one thing instead of another, maybe.

Precisely THE point, not beside it.

>Whatever the people do, they will do it with a particular organization of
>control systems, and exercising those control systems or bumping them up
>against each other is not going to change anything.

Look around. Peoples' lives are changed in both trivial and profound ways (as judged by the influencees themselves) everyday, all over the globe, by acts of purposive influence. Ask THEM whether "anything" has changed.

Best, Greg

Date: Wed Sep 23, 1992 7:44 am PST
Subject: THE model for sociology

From Greg Williams (920923 - 2)

>CHUCK TUCKER [920923]

>CLOSED LOOP GW 920919

>My vote is for an issue on "Influence and Control" (seconding
>Hugh's vote actually)...

Perhaps for the January 1993 CLOSED LOOP. I'm already well into the next (October) issue, which will be on "Standards" and closely related topics. I like your suggestions (not re-posted here) for structuring an "Influence/Control" issue (and perhaps a follow-up "comments by others" issue, to which I would want to see added "final remarks" by Bill and me. But I wouldn't feel particularly comfortable with editing either issue, since I am a principal in the argument. Might someone else (or more than one person) agree to be guest editor(s)? If nobody will volunteer, then I might be willing to go through with it, providing that Bill and everybody involved in the "comments" issue have an opportunity to pass on the final version. So, let's discuss the possibility further -- we have a bit of time to make decisions (work would start in late November/early December).

>INFLUENCE AND CONTROL

>I have reached the tentative judgement that there is no way for
>one human being to influence (in Greg's terms) another human
>being and still maintain the principles and processes outlined
>in the HPCT or PCT model as proposed by Bill Powers.

I'd like to hear more about the basis for your judgement. Do you mean to say that no one actually will sign his/her name in rubber-banding if presented with certain kinds of disturbances by another? This is an instance of what I call "purposive influencing." It is a simple instance, to be sure, because the influencee's (relevant) controlled perception (keeping the knot over the mark) is "obvious" to the influencer (who simply asked for the influencee's agreement that that is the influencee's controlled perception), and because there is no deception involved. But that is the sort of interaction which I call "purposive influencing": the influencer makes a model of what the influencee wants, then manipulates the influencee's environment so that the

influencee gets what he/she wants in such a manner (different from the manner which would be employed by the influencee in the absence of the influencer) that the influencer also gets what he/she wants. Or did you think I meant something else by "purposive influence"? If so, sorry; I didn't.

>Greg keeps insisting that THE TEST will demonstrate that person
>A has influenced person B to control for X.

No. What I have been insisting is that a person who wants to purposively influence someone else would be well advised to use some form of The Test to make a model of what the prospective influencee wants, rather than just guessing about what the influencee wants. The Test DOES NOT demonstrate that person A has influenced person B to control for X. The Test FACILITATES making a model of X, which A DOESN'T WANT TO ALTER (that would require FORCE or influencing the outcome of reorganization by B)! A USES his/her model of B's (and solely B's) want, X (which B is controlling for THROUGHOUT the entire episode of purposive influence), to INFLUENCE B to successfully control for X (assuming A is correctly modeling it) by ACTING so as that A will get what he/she wants.

>But that is not so,
>THE TEST only shows (if done repeatedly and systematically with
>careful records) what B is controlling for BUT IT CAN'T SHOW
>THAT B TOOK X FROM A. A can only say that the X that B is
>controlling for seems to be exactly like the X I use and asked
>(told, demanded, etc.) B to use but since B is now using X it
>belongs to the "control system" of B and no longer to A even
>though A might use it also.

What is this "taking"??? I never said that A takes ANYTHING from B. A needs to make an accurate guess ("model") of what B is controlling for if A is to succeed in purposively influencing B. A DOESN'T have to TAKE ANYTHING from B! A's model of B's X certainly "belongs to" A, not B. But, so what?

>Now this seems like a very "picky"
>point to make but it is crucial to PCT: my perceptions are mine
>even though they may look exactly like the perceptions that you
>say are yours.

I have no problems with this. That's why I said that A makes (his/her OWN) MODEL of B's want X; A doesn't "TAKE" X "FROM" B. Using A's model of B's X, A tries to arrange B's environment so B performs actions to achieve X so that A gets what he wants. A could be wrong about his/her model being correct, in which case A would try to make a better model -- and The Test is the key to making models of other peoples' wants (sez PCT-science).

>The PCT model has to insist on this autonomy or it will very
>quickly lose its differences from all other theories (as Rick
>indicates in his "elephant" paper). Thus, anything that looks
>like a "collective social act" is actually two or more persons
>controlling their own individual conduct in what appears to each
>and every party to action as influencing the others.

I agree wholeheartedly. This consideration in no way obviates my notion of "purposive influence," which "IS ACTUALLY TWO OR MORE PERSONS CONTROLLING THEIR OWN INDIVIDUAL CONDUCT[S] IN WHAT APPEARS TO EACH AND EVERY PARTY TO [THE] ACTION AS INFLUENCING THE OTHERS." What is crucial to recognize is that the purposive influencer is arranging things SO THAT he/she gets something he/she wants WHICH DEPENDS ON THE ACTIONS OF THE INFLUENCEE AND WOULD NOT OTHERWISE OCCUR. The non-purposive influencer does NOT so arrange things, but influences the other(s) in ways which he/she DOES NOT CONTROL FOR ("accidental influence").

>careful analysis would show that each party to the "collective
>act" is acting on their own.

Each is WANTING on their own, but each is not ACTING on their own. Given the wants, the actions of each depend on what influences (both purposive and accidental) the other has. The influences might be symmetric or asymmetric.

>If my tentative judgment (i.e., that influence [A affects or
>determines what B controls for] is not possible between human
>beings if they are negative feedback control systems) is turned
>into a "warranted assertion" with testing and evidence the
>question arises: Is the PCT (or HPCT) model useful for
>collective social action? That is an important question to
>answer for those of us who have a concern for human group life,
>social life, collective behavior or even Society. Maybe a new
>model has to be constructed !?!

I AGREE with your judgement/warranted assertion. That is what PCT science says (barring threats/force/guided reorganization). I also think that "purposive influence," as based on PCT science, is all we need to explain "collective action." The key to "collective action" in general is that one person can make a model of what another person wants, and use that model to purposively influence the other person. No new model is needed for sociology; the PCT model is The Model!

>WHAT INNARDS?

>Thanks Greg for the statements from Skinner which indicate to me
>that he (and perhaps other behaviorists) did not consider the
>"innards". I guess I just find it strange that someone would
>not wonder more than he did about the contribution of the
>"innards" to behavior...

He believed he could solve the problems he wanted to solve sufficiently well without modeling "innards." It's as simple as that: he thought he was controlling well (in PCT terms), and so he thought he didn't need to act differently.

>but people I pay attention to (especially
>myself) firmly believe that the CNS is important to say nothing
>of the liver and bladder.

I, too, think that Skinner was deluding -- handicapping -- himself. Yet his argument that modeling "innards" is PREMATURE was more reasonable in the 1950s and earlier, when much less was understood about the CNS than now.

Best wishes, Greg

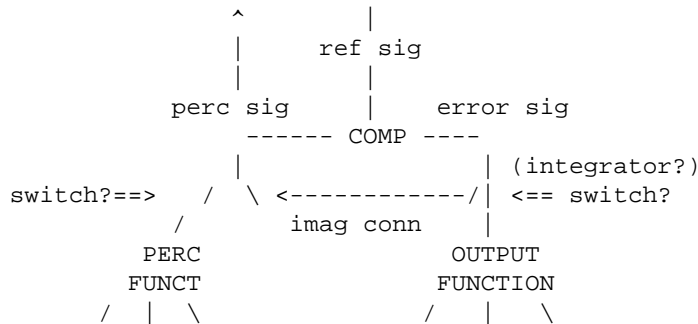
Date: Wed Sep 23, 1992 8:22 am PST
Subject: Re: Imagination

[Martin Taylor 920923] just back and catching up (Bill Powers 920906.0600)

Bill wrote, way back when I had just left ...

=====

It's the error signal that has to be looped back to become the perceptual signal, like this



Now it's very simple, because both the error signal and the perceptual signal are one-dimensional. No weights, no keeping track of how many input signals to the perceptual function there are, no multiple-pole switch. Maybe one integrator is needed in the error signal line before the output function. But that still leaves only one signal to be fed back.

=====

I don't understand this at all. You seem to have a feedback loop with no gain. The error "becomes" the perceptual signal? Then what? The error is the difference between thge perceptual signal and the reference signal, nicht wahr? So the error now becomes the difference between itself and the reference signal, meaning that the reference signal is FORCED to zero. Can an ECS do that? I think not.

If you remember the earlier discussion, we decided that imagination consisted of the assumption that ONE lower-level ECS had provided the perceptual signal that its reference from this ECS demanded. That meant that one of the output signals was short-circuited back to the corresponding perceptual input signal, but the feedback loop dynamics and the effect of real (data-driven) perceptual input on the imagined one were both unaffected. All that happened was a possible speed-up of the feedback dynamics (imagination works faster than the real world) and a loss of conflict that might occur in real, non-imagined, operation.

There is indeed a problem of how and when "a switch gets thrown" to set something into an imagination mode, and whether a very high-level ECS can require imagination at a specifiabile much lower level. If I imagine a golf shot, I can imagine just the ball soaring away beautifully, or I can imagine

the muscular events involved in the swing. The choice seems mine as to which level I use for imagination. I think there is a difficult problem here.

There is another problem in the integration of imagination with real sensory input. Imagination has a close link with planning, and it seems reasonable to suppose that at levels above sequence (if not at lower levels) the imagined sequence is in some way compared to the ongoing sequence, much as if imagination supplied a kind of reference signal as well as a perceptual signal. Again, how this could fit in the model is not clear.

It will probably be a few days before I get through the 200 messages awaiting me, and only then will I see (probably) responses to this.

Still jet-lagged. Martin

Date: Wed Sep 23, 1992 11:03 am PST
Subject: Re: turing test

[Ray Allis 920923.0845]

> (penni sibun 920922.1400)

> right. turing was not begging the question.

Indeed, he was arguing that it was not worthwhile answering the question.

> In the end, we must
> KNOW, or at least SAY, what intelligence is, and then construct the TT
> to measure X's satisfaction of the definition.
>
> nope. your formulation requires such a definition; however, i'd bet
> you a lot (if i had it, which i don't!) that you will never ever
> discover such a definition that works in the real world. i don't
> think you should be spending your time looking for it.

Unquote, Turing and most "Artificial Intelligence Researchers". Doesn't it seem odd that people who are purportedly searching for intelligence don't want to be pinned to a definition of what they're looking for? Does this have any relationship to the way money is handed out?

> I'm glad that Control Theory seems to be founded on less shaky
> methodological ground than traditional AI.
>
> the turing test is not a methodology. it's more a philosophical
> basis.

Basis for ...?

> listen hard: the only standard for intelligence (i'm not going to
> play your linguistic games of calling intelligence ``intelligent
> behavior'') is that it be *judged* as such, by a human, who we
> stipulate is indeed intelligent. the turing test, which compares
> programs to humans, tests to see whether programs can judged as
> intelligent as humans.

The "linguistic game" here is use of the word, intelligence, as if it denoted some real thing in the world, like blood or meat, and pretending that since the name exists, the thing named exists. This scam is called reification. "Intelligent" is an adjective modifying "behavior", a noun denoting an observable phenomenon. Turing (and most "AIers", e.g. Marvin Minski who defines intelligent behavior as "whatever we would ordinarily call intelligent if a human were to do it") are just ducking the hard problem of defining what it is you are talking about. Seems to me this is one of the first steps in any really scientific endeavor.

Respectfully, Ray Allis

p.s. It isn't necessary to discover some "real" definition of ANY category: you *create* the definitions.

Date: Wed Sep 23, 1992 12:54 pm PST
Subject: Rick: new ideas in psyc. info

[From Francisco Arocha, 1304;920923] For RicK Marken:

This is a copy of a message I sent to you yesterday, but I did not get any answer (I don't know if you received it) so I'm resending it through the csg-l.

Rick:

I just sent you the information about the journal New Ideas in Psychology. Since it first has to go through Canada Post, it may take longer than usual to get to you. Judging from the journal's aims and scope, it sounds promising, given the quality of your paper (I copy textually from the aims and scope, although any typos are mine):

There is growing dissatisfaction with fragmentation in theoretical psychology and without doubt there is a felt need for a journal containing articles which aim to integrate these fragmented ideas and theories. Since the usual methods and perspectives have failed to achieve such integration, common sense suggest that we try new and more daring hypotheses.

We place particular emphasis on the fact that informed discussions of psychological ideas needs to be interdisciplinary. We are intrigued by the possibility of facilitating a rapprochement between traditions of thought that currently ignore each other. We also anticipate that fresh ideas may come from fields indirectly related to psychology.

They publish papers with peer commentary and rebuttals from the author. The journal publishes 5 types of articles:

1) How ideas develop. These deal with conceptual issues in theory construction, scientific creativity and growth of knowledge (philosophical papers, I guess).

2) Exploratory theorizing. Unusually significant theoretical articles integrating subfields of psychology. Your paper would be in this category.

It also publishes 3) new, sketchy ideas; 4) book reviews; and 5) interviews.

Editor:

Type 2 papers only: Pierre Moessinger, 46 Maraichers, CH-1205, Geneva, Switzerland.

Send 3 copies, 25 pages, double spaced, generous margins. First page should contain author's name and affiliation, and full current address. The rest of the paper should not identify the author in any way. Reference styles is APA.

Saludos, Francisco

Date: Wed Sep 23, 1992 1:15 pm PST
Subject: post: dennet & powers

[From francisco Arocha, 01:22, 920923]

Found in NSP-L list (Noble Savage Philosophers):

> I've heard of people describing how our perceptual experience of the world
> contrasts in ways from the way the world is: objects become smaller as we
> see them as being more distant, though in reality the objects don't shrink;
> part of an object disappears from vision when that object moves behind
another,
> whereas in reality the first object remains whole and intact, etc. Does
> anyone know some references concerning people who write on such subjects?
>
> Thanks,
> Irwin Goldstein

Richard Gregory from Bristol University, UK, has written much on this subject, mainly from a psychological point of view, but he is unusually thoughtful for a psychologist (which means that I agree with him, I guess!) Try <The Intelligent Eye>,1970, which discusses The Muller-Lyer illusion as well as some of the others you mention.He also wrote <The Eye and Brain>and edits the journal Perception.

More interesting philosophically, though about the same vintage, is Norwood Russell Hanson,< Perception and Discovery> from Northwestern Uni.

For a more modern, though difficult, book using systems theory, W.T. Powers "Behavior: The Control of Perception" will give you a non-behaviorist view compatible with Dennett's intentionality.

Felicitas end of post

Date: Wed Sep 23, 1992 2:10 pm PST
Subject: turing test; Spartans; Healthy thinking; control is control

[From Bill Powers (920923.0900)]

What a fine lot of posts this morning! Penni Sibun (920922.1400) --

I'm not as far from your point of view as I may sometimes seem. I think that direct private experience is the final arbiter. Any theoretical idea has to explain EXPERIENCE. This is true even in physics. The more points of contact there are between theories and experiences, the more we can trust the theories to carry us over the gaps in experience. So yes, intelligence, in the final analysis, is what a human being experiences as intelligence. You can define it objectively only through a theory that others, you hope, may accept.

That means, of course, that there are as many brands of intelligence as there are observers. This can lead to pointless arguments, or it can lead to an effort at communication, in which each observer tries to say what she is looking at and how she is interpreting what is seen, always with the humble awareness that one's point of view is private and basically incommunicable. I suppose that's why we talk so much.

Your criterion for determining sex, you say, is anything that can be seen without a microscope. That's not quite saying "I know it when I see it," because the criterion implies enumerable characteristics, no one of which by itself is necessarily sufficient. And it doesn't really handle behavioral aspects of sex, on which considerable segments of the population don't agree. What is the sex of a nominally male homosexual who is playing the cultural role of a female with respect to his-her partner? Pat Buchanan, I'm sure, wouldn't bend his definition of "male" to accomodate such a situation, but clearly the male partner does, and those who have homosexuals for friends accept the other person as the other person wishes to be defined. I don't think that "I know it when I see it" is really sufficient to cover the territory.

You say "organisms just don't go around doubting things most of the time." I'll go further: Organisms don't generally realize that there's anything to doubt. The curse of the theoretician is the discovery that all is not as it seems, which puts an end to living a simple life in which appearances are accepted at face value. It isn't that appearances themselves are doubted, but once you begin to wonder why things appear as they do, you discover the gaping holes in our understanding and begin to see that the whole structure is a lot more rickety than you would like. For Pat Buchanan, there are no holes; for him life is simple. Why don't I envy him?

>how does finding out whether the entity is a human or a program
>have anything to do with intelligence as you suggest defining it
>immediately above??

I was being a little sarcastic, suggesting that people generally define intelligence to their own advantage and as a way of putting down others, illustrating by showing how I would do that.

RE: the prickly correspondence between you and Rick. Why don't you both lower the barriers and write down exactly the system concept you have formed of the other person? Seeing it in print in a public place might make it look a little different.

RE: turing contest

>half the judges mistook the winning prog for a human.

Since there were only 2 humans, this meant that half the judges mistook at least one human for a machine. To some people, although possibly not you, the Turing test is a test of a program, to see whether it exhibits an essential human characteristic labelled "intelligence." The assumption is that if the entity at the other end of the teletype DOES exhibit that characteristic, a human being will know it by seeing it. The concept behind the test is greatly weakened if it turns out that the human judge does NOT recognize this characteristic in an interaction with an entity that the judge would later agree is a certified human being. Certainly negative judgements -- that a program has failed to capture this characteristic -- are invalidated. And by implication, positive judgements, too, are cast into doubt.

If the judges were AI experts and the audience was not, the superior performance of the audience shows that the criteria the judges were using were irrelevant to the characteristic being identified. So I stand by my statement that AI types don't know what they're looking for: they DON'T know it when they see it, even in the private terms they would use when the nature of the entity is revealed to them.

Isaac Kurtzer (920922) --

Give your machine a kick for me. I think your analysis of Sparta is a brilliant idea. Now that Tom's gone, I visualize you with a chair in one hand and a whip in the other, surrounded by slavering toadies of the psychological establishment snarling and swiping at you from their pedestals while you force them one at a time through their hoops. When you take your bow, watch your back.

Pat Williams (920922.0900) --

>Hi Bill. I'm sick of all this talk of me (Pat) and the feeding of
>"healthy food" to my kids, so I'd like to try to clear up the
>confusion about this particular instance of what Greg has been
>calling purposive influence, myself.

I wondered how long you could stand it.

>Yes, I am taking care NOT to disturb what the kids are controlling
>for! That's the whole point. And I am also achieving my goal: they
>are eating (not just being served) "healthy" food. Everyone is
>happy! It is purposive because I'm doing it on purpose; and it is
>influence because the kids are acting differently (eating "healthy"
>food) than they would be if I served them just any old tasty food
>or let them scrounge for themselves (then they would be eating
>tasty food, most of which would not meet my definition of
>"healthy").

We agree that you're controlling your own perceptions, on purpose. We disagree that your kids are acting differently because of your controlling for nutritional value of the food. What they're getting is different because of that, but what they're controlling for isn't. Their actions would remain the same no matter what you served them, if they ate it. What we're trying to sort out here is not whether you're doing the right thing by your kids but who is influencing and controlling what.

You would be frustrated if your kids didn't eat what you served, but the eating is still up to them, not you, even when they eat it. You can't operate their arms and mouths and swallowing muscles. You may be gratified at seeing that good food disappear irrevocably into their mouths, but you can't influence that process or control it: matters have passed outside the sphere of your influence at that point. Your influence has reached its limit when you put the food down in front of them.

What I'm talking about is much simpler than the arguments we're going through. It has nothing to do with your feelings about what the kids are eating, your reasons for giving them the food you give them, or what the kids would do if you behaved differently. It has to do with who is in charge of which perceptions and which means of controlling them. The kids are controlling their own perceptions. You are controlling yours. As a result, there are certain interactions between you and the kids -- disturbances, or if you like, influences. These interactions come about because when you do what's necessary to control your own perceptions, your actions disturb the kids' worlds, and they adjust their actions to prevent those disturbances from having any significant effect on the variables THEY are controlling, if they physically can. Just try helping them put the food in their mouths: reach out, grab a wrist, and push the fork toward the mouth. They will resist. Even helping is a disturbance if it's applied to something that's already under control.

I get the impression, possibly unjustified, that you and Greg are looking for something more than that. You don't want to see this interaction as just "disturbances," or as just neutral "influences," but as a way in which one organism can have an intentional effect on the interior life of another organism. In other words, you want to think you're doing some good for your kids, not just going through the motions and having the kids' control systems counteract everything you do. This has got to be a sensitive issue; you go to great trouble to do the right thing for your kids, even to the time-consuming labor of schooling them at home. And here comes Bill Powers saying that all you're doing is disturbing them, and that all they're doing is making sure your disturbances have no effect on them. Makes it all look pretty futile, doesn't it?

But I don't think we can come up with a correct understanding of human relationships unless we're willing to face the (seemingly) worst possibilities and put them aside. Perhaps it will turn out that your efforts with your kids aren't futile. That would be nice. But I think it would be better to know the true relationships than to prove that you've been right all along. It's very difficult to come up with a valid explanation or analysis if there are possible outcomes that are ruled out from the start, and others that MUST be supported. This is why I have tried to point out to Greg that he has a GOAL in this argument; that he's trying to make it come out a particular way. It's trying to make an argument come out a particular way that leads to switching

definitions of words, generalizing upward by one set of rules and particularizing downward by a different set, and all the other slippery ways that human beings have of directing their own reasoning toward a conclusion that they want to be true. If you want to know the truth, you can't care how the argument comes out. You just have to follow it through and see where it leads, trying not to make any mistakes.

I think there are ways in which we can intentionally influence the interior lives of children. I don't yet have a clear picture of what they are, and when the answer gets closer it may turn out to be the opposite of what I think. But at the moment I think that because children demand instruction, we can give it to them. We can't do the reorganizing for them; if they don't do it, it won't get done. We can't control how this reorganizing will come out. But we can make sure that the environment contains everything we know about that might be useful to them, and that what we tell them about is as near to true as we can approach. We can invite their attention to things, relationships, problems that they might not come across by themselves. We can look for difficulties they are having, and demonstrate how another person would deal with them. We can make sure that when they're ready to perceive in a new way, there are things to perceive at the new level that will be generally useful to perceive -- although of course we can't say just how they will end up perceiving the situations we perceive.

Chuck Tucker has the right view, in my opinion. We don't give perceptions or reference levels to others; they TAKE them. They suck them in, when they're children. Whatever they've taken in, they work over, rearrange, test, try out, think about, and convert into something they, not the external world, find useful.

Chuck Tucker (920922) --

I agree with you so much that I have nothing to say to you. Hello. Hope you're feeling fine.

Greg Williams (920923) --

>>All that can be influenced is the action used to correct any
>>induced error, and that can be influenced only in the sense of
>>making more or less of it necessary.

>Yes. Well, actually, "more or less" in a not strictly real-number-
>line way, but possibly very complexly "qualitative," being
>multidimensional. Even in plain old rubber-banding, the actions
>needed to compensate for the applied disturbances are not simply
>"more or less" in one dimension, but simultaneously in TWO
>dimensions.

"More or less" in each dimension being controlled and disturbed, you're right. In the rubber-band experiment on a tabletop, you can disturb independently in two dimensions. The other person controls independently in two dimensions (possibly two different dimensions, but two). The "more or less" applies in each dimension.

>>The nature of the action is already determined by the organization
>>of the control system, which takes into account the environmental

>>link through which action affects perception.

>Take another look at the generic PCT diagram. The action is
>determined conjointly by BOTH the reference signal AND the
>disturbance, via subtraction, one from the other, at the
>comparator.

Take another look at the statement. The NATURE of the action means what KIND of action has to be produced, not how much of it. The KIND of action is not determined either by the reference signal or by the disturbance. It's determined by the physical links that exist between the system's output and the controlled variable, by the perceptual function defining that controlled variable, and by the effector design of the control system. The reference signal and the disturbance determine, conjointly (if that means anything) only HOW MUCH action will occur, and IN WHAT DIRECTION. They do not determine WHICH KIND of action will occur.

>>You can claim credit for those increases and decreases, but not
>>for the nature of the action. That was established before you
>>arrived on the scene.

>"You" cannot claim FULL credit. Only SHARED credit. SHARED with the
>organism. The nature of possible actions, given possible
>disturbances, WAS established by the organism, but YOU select WHICH
>of the possible disturbances ACTUALLY is applied.

The organism determines WHICH controlled variables it will be controlling at a given time. You can then pick any of them to disturb. When you do, the organism will vary its action on that controlled variable, cancelling your disturbance. It is, of course, already varying its actions on ALL the controlled variables, or is ready to do so, at all times, if they're under active control. You're not determining which variables will in fact be under control or which actions they will use. That's been decided already.

Our image of "rubber-banding" is unfortunate in one respect, because this demonstration has deliberately been made very simple, to illustrate principles. A more realistic example of rubber-banding would give the control system one rubber band attached to the knot, and twenty different people twenty rubber bands attached to the same knot. The control system won't have any difficulty in controlling the knot (unless the combined disturbance results in breaking of the control system's rubber band) because only the vector sum of disturbances matters. Control might actually be easier because independent random disturbances will sum to a net disturbance having much less variability than any one of them has.

But any one person acting as a disturbance, trying to influence the control system's hand position, is going to have great difficulties because of all the other random disturbances that are present. While control still remains possible, it's no longer possible for the disturber to estimate the best direction to move his/her own hand to achieve a correction of the other's hand position, because there is no longer any best direction. And it becomes difficult for the putative disturber to know what disturbance is actually being applied; perceiving one's own rubber-band tension is no longer indicative of the net disturbance on the other's controlled variable. The only way to make sure of applying a known disturbance is to isolate the control system from all those other influences.

>>There's no difference in the action used to counteract an
>>accidental influence or an intentional influence.

>Correct. But that is beside the point. Purposive influence matters
>to the influencer ALWAYS.

Agree. Control system control their own perceptions. Are you making some new
point here? Why the "but?"

>I think I now am beginning to understand how your tenacious
>clinging to the organism's point-of-view, which serves you so well
>scientifically ... seduces you into an ideological position at
>odds with PCT-science.
>The organism role IS more important than the environment's role to
>the psychologist who wants to see order (invariable goals) behind
>the chaos (variable actions). But it is a mistake -- because it
>contradicts PCT-science -- to carry that importance over to
>ideology.

Are you saying that a psychologist has a point of view that is not that of an
organism? Good trick. I'd like to learn how anyone can have a point of view
different from that of an organism. Go ahead -- I'm easily seduced.

>>As I understand your goal in this argument, it is to prove that
>>the environment has influences on the organization of the
>>organism's control systems, shaping the organism in some causal
>>way. This goal includes proving an effect of both non-living and
>>living influences on the organism, independent of its heredity.

>Yes, that's my aim in the second of our two arguments. The first
>argument involves purposive influence only, not environmental
>influence in general. Of course, they have been intertwined
>somewhat.

Well, that means we should be able to save a lot of time. If proving this
conclusion has been your goal from the start, you must have had other reasons
for believing it is the correct goal, reasons that go beyond, predate, or
supersede PCT. This would, of course, explain why you don't develop an
irrefutable argument based only on PCT principles: there are other principles
involved. It would also explain why, when you misinterpret a term like
"nature" (above), you always do so in the direction that leads to your goal,
and never accidentally the other way. When you are convinced that there is
only one valid outcome of an argument, which you already know, the details of
the argument cease to be important. If one argument bombs out, there's bound
to be another that will get you there.

In my lexicon, knowing the answer and coming up with whatever argument it
takes to get there is called "ideology."

>>Whatever the people do, they will do it with a particular
>>organization of control systems, and exercising those control
>>systems or bumping them up against each other is not going to
>>change anything.

>Look around. Peoples' lives are changed in both trivial and
>profound ways (as judged by the influencees themselves) everyday,
>all over the globe, by acts of purposive influence. Ask THEM
>whether "anything" has changed.

If they say anything has changed, they are ignorant of history. What you speak of is not change; it's just coping. Political intrigue is the same today as it was in Rome or Egypt. People try to push other people around just as they have always done. The interactions and the influences fluctuate and shift around in complex patterns that achieve no more today than they have ever achieved, with the exception that now there is enough power available to achieve extinction. Without a correct understanding of human organization, human beings simply cope with whatever immediate problem presents itself, only dimly aware that there are others doing the same thing, and under the false impression that they're accomplishing something of lasting worth instead of causing most of their own problems. It may well be that there is a slow drift toward understanding. But it's too slow to suit me.

What you present as "purposive influence" is simply control of others. You may say it isn't, but you still insist that purposive influence is intended to have some effect other than just disturbing what others are controlling for. Look around. See what it has accomplished. Perhaps you would get a better view from Chicago, or Los Angeles, or Albuquerque. What I see is not just a natural progression of interactions. I see the influence, if you will, of a widely-held misconception, a disastrously wrong conception, of what a human being is and how a human being works. That's what has to be changed -- not what we "do."

Best to all, Bill P.

Date: Wed Sep 23, 1992 3:51 pm PST
Subject: Re: turing test

(ps 920923.1300)

[Ray Allis 920923.0845]

Unquote, Turing and most "Artificial Intelligence Researchers". Doesn't it seem odd that people who are purportedly searching for intelligence don't want to be pinned to a definition of what they're looking for?

doesn't it seem odd that some people think they can't study something w/o writing down a definition of it first?

Does this

have any relationship to the way money is handed out?

not much. all the ai contracts i've worked on have had quite concrete and fairly mundane objectives. the funders are even less interested in pinning down intelligence than the ai'ers are.

> I'm glad that Control Theory seems to be founded on less shaky
> methodological ground than traditional AI.
>

> the turing test is not a methodology. it's more a philosophical
> basis.

Basis for ...?

ai.

> listen hard: the only standard for intelligence (i'm not going to
> play your linguistic games of calling intelligence ``intelligent
> behavior'') is that it be *judged* as such, by a human, who we
> stipulate is indeed intelligent. the turing test, which compares
> programs to humans, tests to see whether programs can judged as
> intelligent as humans.

The "linguistic game" here is use of the word, intelligence,
sigh. i knew i'd get caught if i went half-way. my original message,
following turing was framed in terms of ``ability to think.''

as if it
denoted some real thing in the world, like blood or meat, and
pretending that since the name exists, the thing named exists.

speak for yourself. i don't believe words denote.

This
scam is called reification.

if reification is such a scam, why do you employ it in your next sentence?

"Intelligent" is an adjective modifying
"behavior", a noun denoting an observable phenomenon.

``intelligent'' is an adjective (most of the time) and can modify whatever
word or words it feels like. i refer you to charles dodgeson.

Turing (and most
"AIers", e.g. Marvin Minski who defines intelligent behavior as
marvin's last name is minsky.

"whatever we would ordinarily call intelligent if a human were to do
it") are just ducking the hard problem of defining what it is you are
talking about. Seems to me this is one of the first steps in any
really scientific endeavor.

``defining'' what you're talking about is an intractable problem, and
intractable problems are best avoided. i realize some people think science is
a process of writing down immutable definitions; i find that depressing.

p.s. It isn't necessary to discover some "real" definition of ANY category:
you *create* the definitions.

why yes, you do, which is an important reason that definitions are tools at
best and not sacrosanct and why i've been arguing against them.

cheers.

--penni

Date: Wed Sep 23, 1992 4:27 pm PST
 Subject: Re: Gary's reply, Neural imaging, Turing

[Martin Taylor 920923 18:00]

(Rick Marken 920910.0900)

> But I do think that the Turing
 >Test is an EXCELLENT example of the behavioristic basis of AI (and
 >cognitive science) etc. Harnad's contribution makes it even clearer -- he
 >suggests a Total Turing Test meaning its not enough to get a simulation
 >to answer questions like a real person -- you must also get it to
 >behave in all ways like a real person -- ie -- brush teeth,
 >play soccer, build model airplanes, etc. Of course, what is interesting
 >is that behavior is defined completely superficially -- it is what you
 >SEE. There is no notion that behavior (including the conversation of the
 >original Turing Test) is a controlled consequence of simultaneous influences
 >produced by the actor and the environment.
 >...
 >
 >I think the Turing Test is an excellent model of the basic misconception
 >about behavior embraced by all the life sciences. According to the
 >Turing Test (Total version), the only problem with clockwork simulations
 >of behavior like those done in the 1700s was that they didn't mimic enough
 >behavior and, perhaps, didn't do it smoothly enough. The Turing Test
 >shows very clearly that the goal of AI and like minded "sciences" is
 >truly superficial -- they want to simulate how things "behave: not how
 >they control.

It was my impression that the Turing Test involved an interaction between the tester and the testee. That implies that the tester can apply The Test as understood by PCT. Turing, as I understand him, implied nothing about how the testee was to arrive at the observed behaviour, or whether the tester should, could, or would, observe control. So I think the quoted comments are misguided.

Martin

Date: Wed Sep 23, 1992 4:30 pm PST
 Subject: misc comments

[From Rick Marken (920923.1430)] Just some quick replies:

Gary Cziko (920922.1835 GMT) says:

>Rick Marken (920922.1000) mentioned parenthetically:

>> (actually, I think Piaget got close to doing a version of the test).

>Could you elaborate a bit on this.

Right off the top of my head, how about "conservation of number"? Young kids apparently control for "amount of goodies" in terms of linear extent instead of number. The one I love is this: put two rows of M&Ms in front of a 4 year old thus:

(1) m m m m m

(2) m m m m m

Ask which row the kid wants; the kid always goes for row 2, even when they can COUNT the M&Ms. It looks to me like the kid is controlling length in order to (presumably) get the MOST goodies. More Piagetian variations could be done to test this; eg. make row 2 shorter than row 1, add an m to row 1, stuff like that.

The point is that Piaget would vary (disturb) various aspects of a situation to see if kids "corrected" for the disturbance. If you have two equal size glasses of lemonade and pour lemonade from one into a wider, shorter glass then the kid will try to get the taller glass -- the change in height was a disturbance that needed to be corrected (avoided) in order to get the "right" amount of lemonade. I think Piaget was on to something like the PCT concept of development of levels of control. Maybe that's why PCT gets along so much better (apparently) in Europe than in the US.

Francisco Arocha (1304;920923)

Thanks -- I did get the info on the "new ideas" journal. Looks perfect.

I just sent a revised version of the "Blind men" (no, Chuck, I just won't stop controlling for that metaphor) paper to Psycholoquy -- the electronic psychological journal. I'll keep CSGNet posted on its progress (or lack thereof). If it is accepted as a target article it will be fun to hear what psychologists think of it. If Psycholoquy fails, I'll ship it off to "New ideas". What fun. Back into the fray.

penni sibun (920922.1400)

>listen hard: the only standard for intelligence (i'm not going to
>play your linguistic games of calling intelligence ``intelligent
>behavior'') is that it be *judged* as such, by a human, who we
>stipulate is indeed intelligent.

I agree. But judgement implies a comparison of some kind, I think. I think your statement implies that an intelligent person is one who can look at something (I won't say behavior -- just some aspect of their perception,OK?) and judge it to be intelligent or not. I think that, in order to do this, the one doing the judging must have some comparative basis (an image in their mind, perhaps) for deciding that they are perceiving intelligence. I think this is a fair representation of the Turing test (though, I would call the aspect of perception that Turing was asking us to judge as intelligent or not "behavior").

> the turing test, which compares
>programs to humans, tests to see whether programs can judged as
>intelligent as humans.

Yes, I never said anything different (though I did use the "b" word - "behavior" - to refer to what it was about the program and the human that was being compared).

>the turing test is not about behavior. it lies in the interactionist
>realm, in my opinion.

But don't you look at behavior to judge whether or not it is intelligent? Maybe I should make it clear that my (informal) use of the word "behavior" includes written or spoken responses to questions (which are also "behaviors")

>> The Turing
>> test focused on actions and their consequences

>no. it focuses on an *interaction* between a judge and an entity.

One kind of interaction goes like this: I ask a question, the machine makes a reply (which might be the null reply -- silence). The reply leads to another question or statement by me, then a reply by the machine, etc. If what I do is x and what it does is y (x and y being variables) then an interaction is x--y--x--y--x--y... Now I think you are saying that the intelligence of the system generating the y's is judged from the interaction x--y--x--y--x--y... Some sets of x--y--x--y--x--y... can be judged as intelligent, some not. Is that something like what you mean? If so, I certainly would agree that that is a legitimate way to interpret the Turing test. I still think it is "behavioral" but I won't even try to deal with that until we can reach some agreement about what the Turing test IS.

Best regards Rick

Date: Wed Sep 23, 1992 4:58 pm PST
Subject: Re: turing test

Penni Sibun (920923.1600)

[From Bill Powers (920923.0900)]

I'm not as far from your point of view as I may sometimes seem.

oh, i know. ;-} and where we differ, i almost always find your point of view extremely thought-provoking.

think that direct private experience is the final arbiter.

yes, i should think a pct theorist would have to think that way, which is why i'm puzzled when some don't talk as though they believe that.

Your criterion for determining sex, you say, is anything that can be seen without a microscope. That's not quite saying "I know it when I see it," because the criterion implies enumerable characteristics, no one of which by itself is necessarily sufficient. And it doesn't really handle behavioral aspects of sex, on which considerable segments of the population don't agree.

well, i don't think the criterion *does* imply enumerable characteristics; to say that ``i know it when i see'' isn't the same as saying ``these are the reasons i know it.'' and i certainly don't see why it doesn't handle behavioral aspects--don't we see those too?

RE: the prickly correspondence between you and Rick. Why don't you both lower the barriers and write down exactly the system concept you have formed of the other person? Seeing it in print in a public place might make it look a little different.

that sounds like a good idea, except i'm not sure what you mean by a system concept. can you explain?

RE: turing contest

If the judges were AI experts and the audience was not, the superior performance of the audience shows that the criteria the judges were using were irrelevant to the characteristic being identified. So I stand by my statement that AI types don't know what they're looking for: they DON'T know it when they see it, even in the private terms they would use when the nature of the entity is revealed to them.

i thought i had said that the judges were naive people off the street. the audience was made up of people who could get tickets. the article says that there were no systematic differences in the judgements of the computer and non-computer members of the audience.

cheers. --penni

Date: Wed Sep 23, 1992 5:02 pm PST
Subject: Scribbling chimps

[From Bill Powers (920923.1500)] Greg Williams (920923-2) --

In your post to Chuck Tucker you said something that brought this subject into focus:

>Do you mean to say that no one actually will sign his/her name in
>rubber-banding if presented with certain kinds of disturbances by
>another? This is an instance of what I call "purposive >influencing."

Clearly you can cause another -- even a chimpanzee -- to write any name in the rubber-band situation by applying a suitable pattern of disturbances. I had thought at one time how much fun it would be to stand before a plenary session at the annual meeting of the AAAS, invite a chimpanzee onto the stage, and say "Will the mystery guest sign in, please?"

What a furor! Is this chimpanzee actually writing his name? Can a chimpanzee write? A behaviorist would have a hard time explaining why this chimpanzee does NOT know how to write his name. But a control theorist could do so easily. The agent writing the name is not the chimpanzee, but the human being operating the other rubber band.

Does this represent an "influence on the chimpanzee's behavior?" In once sense, certainly. But it is an influence on how the chimpanzee behaves? No, it is not. The chimpanzee has acquired no new way of behaving. All it knew how to do before, and all it knows now, is how to keep the knot in the rubber bands over a target, as a way of getting an M&M later. All right, 100 M&Ms; who am I to argue with a chimpanzee's lawyer?

When I talk about "changing" people's behavior, I automatically think in terms of doing something so that a person will then know how to do something new, or do something differently from then on. I don't count manipulating people's actions by using disturbances as changing their behavior because as soon as I'm not there to supply the carefully- chosen and patterned disturbance, that behavior will simply stop. It never had any reason to exist in the first place except to protect some controlled variable from being changed. If the behavior were continued after the disturbance ceased, it would CAUSE an error. There is absolutely no reason for the "manipulee" to continue that behavior, and every reason not to do so, when the disturbance is removed.

This is what I mean by "coping." Sure, people adjust their actions because of effects from other people and from the nonliving environment. But that just prevents their hierarchies from being disturbed at the level where the disturbance has its first effect on a controlled perception. They can interact with others 10,000 times a day, and they will have learned nothing new. They will still know only how to control their perceptions in the same ways they knew yesterday. When all the disturbances go away, the people will be exactly as they were before. Nothing of any importance has changed.

Why has nothing changed? Because the disturbance-resisting behaviors have been successful. The whole reason for the existence of these control actions is to prevent intrinsic error; the current organization is the latest product of a series of reorganizations driven by intrinsic error. As long as these control systems continue to work successfully, protecting all controlled perceptions from significant deviations from their reference levels, and as long as maintaining those perceptions at those reference levels is sufficient to prevent intrinsic error, there will be no further reorganization. So nothing will change.

This is why I dismiss all the "purposive influences" you bring up as being trivial. All they illustrate is how control systems with fixed organizations will act when their actions affect other people and other people's actions affect them. Out of these fixed organizations there can arise detailed behavior of immense complexity, in the same way that all those "people" in the crowd program exhibit complex patterns of interaction -- while behaving in a simple and unchanged way.

So it seems clear to me that if people are to be changed in any meaningful way by interactions with each other, those interactions have to amount to something more than simply mutually disturbing each other, or messing about with each other's environments to alter the "schedule of reinforcements."

Without reorganization, all you will ever see in human interactions is variations on a theme as people with static organizations play out all the possible ways of disturbing each other and trying to maintain control at the same time. Those variations lost their fascination for me a long time ago, when I began to see them repeating. For most people, their lives are a replay

of their parents' lives, either symmetrically or antisymmetrically. You hate your parents and want to be totally different from them? You love your parents and want to be just like them? What a bore.

Let's talk about reorganization. That's the only area where anything interesting happens in human behavior.

Best, Bill P.

Date: Wed Sep 23, 1992 5:14 pm PST
Subject: blind men, piaget, PCT promotion

[Avery Andrews 920924.1058]

It seems to me that the problem with the blind men paper is this: academics just aren't interested in being told how concepts they are used to dealing with can be reduced to others that they aren't, and abstract arguments that some other way of looking at things would be better likewise make no impression (I know this from personal experience of trying to explain to people why formal semantics is seriously wedged, and what to do about it). What there has to be is some kind of empirically-based 'hook' that connects the new ideas to things that they're already interested in. Which brings me to the next topic: Rick's paragraph on Piagetian conservation experiments. Something that at least some people actually are interested in these days is that of what is explicitly represented, 'much less than one might expect', being the answer proposed by both Bill Powers and the interactive AI crowd. But it seems to me that the PCT Test gives you quite a powerful tool for attacking this question: anything that is controlled is almost certainly represented. And, by Occam's razor, something that is not subject to control is best assumed to not be represented.

So if the kiddies go for the longest row of M&Ms, even if it has fewer candies than a shorter row, this is evidence that that length of row is represented, but that cardinality isn't (or, if it is, somewhere, its representation is not yet well integrated into the system - my conjecture would be that a cardinality representation might be manufactured by explicit counting, but then instantly forgotten, or at least not made available as a perceptual input to the ECS's in charge of getting goodies).

So I suggest that it might be useful to present PCT as a methodology for finding out what aspects of their environment organisms are actually representing. E.g. procede as if the phenomenon of control were taken for granted, & just present it as a tool for finding out something that people already think they're interested in.

Avery.Andrews@anu.edu.au

b.t.w. anyone who wants to look at a bit of hard-core generative grammar being done, send me their snail address & I'll send them a paper on Inside-Out Unification & Semantic Case-Stacking (in Australian languages).

Date: Wed Sep 23, 1992 6:52 pm PST
Subject: I-N-F-L-U

From Greg & Pat Williams (920923) >Bill Powers (920923.0900)

>In reply to Pat

>We agree that you're controlling your own perceptions, on purpose. We
>disagree that your kids are acting differently because of your
>controlling for nutritional value of the food.

OK, now we understand your point that, from the KIDS' standpoint, they are acting the same regardless of Pat. It's only from PAT'S standpoint that they are acting differently because of her controlling. We would have understood immediately had you said that the kids would control the same (their) perceptions with or without Pat's controlling for "healthy" (her definition) food. But don't you think it is just a bit strange for Pat to think that the kids are acting IDENTICALLY whether they are eating rice or popsicles? It DOES look different! But now that we know what you REALLY mean by "the same actions," we still have a slight problem. With Pat controlling for seeing them eat "healthy" (to her) food, the kids are controlling for (some) DIFFERENT lower-level perceptions than if she wasn't controlling thusly. Aren't they? Surely the same lower-level controlled perceptions aren't involved when the kids eat what is set in front of them, rather than going to the store to buy food. Some of OUR controlled perceptions SEEM different when we eat at a restaurant than when we eat at home. But maybe we're just fooling ourselves. Still, we don't even HAVE salad spoons at home....

>What they're getting is different because of that, but what they're
>controlling for isn't.

Yep. Analogous to somebody subjected to disturbances so he/she writes his/her name in the rubber-banding demo, right?

>Their actions would remain the same no matter what you served them, if they
>ate it.

So the actions of the rubber-bander would remain the same whether he/she was signing his/her name, or John Hancock's, or writing nonsense syllables? Certainly, in all cases, the controlled perception of keeping the knot over the dot is the same. But it sure looks like different hand motions when the person writes a "JUC" instead of an "FUC" -- and in some places, we bet some onlookers might begin to be shocked by what they think of as the writer's actions in the latter case, but not the former. Well, they need to be better educated, don't they? Presumably they would be equally shocked or unshocked by seeing "JUC" OR "FUC" if they only realized that, from the WRITER'S point of view, there is NO difference in the actions producing "JUC" or "FUC." But if the letter "K" is then appended by the writer (still rubber-banding in interaction with a "disturber") to the latter nonsense syllable, somebody still might be in hot water. If the incensed onlookers are sufficiently educated, of course, they will not bother with the disturber, but immediately will denounce the writer, since the disturber obviously had no influence on the functioning of the writer's control system, including its actions. Right?

>What we're trying to sort out here is not whether you're doing the right
>thing by your kids but who is influencing and controlling what.

Agreed.

>You would be frustrated if your kids didn't eat what you served, but
>the eating is still up to them, not you, even when they eat it. You
>can't operate their arms and mouths and swallowing muscles. You may be
>gratified at seeing that good food disappear irrevocably into their
>mouths, but you can't influence that process or control it: matters
>have passed outside the sphere of your influence at that point. Your
>influence has reached its limit when you put the food down in front of
>them.

Agreed.

>What I'm talking about is much simpler than the arguments we're going
>through. It has nothing to do with your feelings about what the kids
>are eating, your reasons for giving them the food you give them, or
>what the kids would do if you behaved differently. It has to do with
>who is in charge of which perceptions and which means of controlling
>them.

Agreed.

>The kids are controlling their own perceptions. You are controlling yours. As
>a result, there are certain interactions between you and the kids --
>disturbances, or if you like, influences. These interactions come about
>because when you do what's necessary to control your own perceptions, your
>actions disturb the kids' worlds, and they adjust their actions to prevent
>those disturbances from having any significant effect on the variables THEY
>are controlling, if they physically can.

Now wait a minute. What's this "they adjust their actions" business? Didn't
you just say above that their actions were the same with or without a
disturbance? Now we're confused. You mean the disturber in the rubber-banding
experiment above (with "FUC") is responsible, even in some small degree, for
that final "K"? Really? Even though the writer was controlling that knot
equally well when writing "FUC" or "JUC" (just as the kids were controlling
equally well for eating tasty food with or without Pat's disturbance)? But the
writer DID it!! He/she WROTE OUT THE OFFENDING STRING OF LETTERS!!! The
disturber didn't even have chalk in hand!!!! Please explain.

>I get the impression, possibly unjustified, that you and Greg are
>looking for something more than that.

Nope, say both Pat and Greg. The purposive influencer doesn't want to change
others' wants, only their actions to control those wants.

>You don't want to see this interaction as just "disturbances," or as just
>neutral "influences," but as a way in which one organism can have an
>intentional effect on the interior life of another organism.

Wrong. An intentional effect on the EXTERIOR life of another organism. We are
claiming that PCT explains the nature and limits of interactions wherein one
party has an intentional effect on his/her perceptions by influencing (if you
like, disturbing) some controlled perceptions. Now, before you object that one
CAN'T disturb another's controlled perceptions, because (assuming good

control) the other will CORRECT for the disturbance, we'll rephrase the last part of the previous sentence and let YOU fill in the technical language which YOU approve of: "... by doing something analogous to what one participant does in two-person rubber-banding when the other person writes his/her name." The intention of the first party is to see certain results of actions of the second party, NOT to change the second person's "interior life."

>In other words, you want to think you're doing some good for your kids, not
>just going through the motions and having the kids' control systems
counteract
>everything you do.

No, Pat wants to SEE what she wants to see. Her motivations are beside the point, other than the fact that she does WANT to see something in particular, which depends on the kids' acting in a certain way.

>This has got to be a sensitive issue; you go to great trouble to do the right
>thing for your kids, even to the time-consuming labor of schooling them at
>home. And here comes Bill Powers saying that all you're doing is disturbing
>them, and that all they're doing is making sure your disturbances have no
>effect on them. Makes it all look pretty futile, doesn't it?

No. With the exception of the underqualified "having no effect" and "looking pretty futile" parts, we have no problems with this statement. And, supposing you mean "no effect on what they are controlling for," we have no problems with that, except for LOW-LEVEL perceptual variables. And, supposing you mean "pretty futile to attempt to affect what the kids are controlling for," we have no problems with that, either. Quite simply put, Pat affects the kids' actions in a non-futile way -- she gets what she wants, analogously to the way the first party in rubber-banding gets what he/she wants (seeing the second party sign his/her name). Tell us what you call that sort of interaction, please. The "rubber-banding interaction" seems a bit opaque to us.

>... I think that because children demand instruction, we can give it to them.
>We can't do the reorganizing for them; if they don't do it, it won't get
>done.

No argument there.

>We can't control how this reorganizing will come out.

Maybe we can influence how it comes out. That's an argument yet to get started.

>But we can make sure that the environment contains everything we know about
>that might be useful to them, and that what we tell them about is as near to
>>true as we can approach. We can invite their attention to things,
>relationships, problems that they might not come across by themselves. We can
>look for difficulties they are having, and demonstrate how another person
>would deal with them. We can make sure that when they're ready to perceive in
>a new way, there are things to perceive at the new level that will be
>generally useful to perceive -- although of course we can't say just how they
>will end up perceiving the situations we perceive.

Sounds like purposive influencing to us.

[in reply to Greg]

BP>>>The nature of the action is already determined by the organization
BP>>>of the control system, which takes into account the environmental
BP>>>link through which action affects perception.

GW>>Take another look at the generic PCT diagram. The action is
GW>>determined conjointly by BOTH the reference signal AND the
GW>>disturbance, via subtraction, one from the other, at the
GW>>comparator.

>Take another look at the statement. The NATURE of the action means
>what KIND of action has to be produced, not how much of it. The KIND
>of action is not determined either by the reference signal or by the
>disturbance. It's determined by the physical links that exist between
>the system's output and the controlled variable, by the perceptual
>function defining that controlled variable, and by the effector design
>of the control system. The reference signal and the disturbance
>determine, conjointly (if that means anything) only HOW MUCH action
>will occur, and IN WHAT DIRECTION. They do not determine WHICH KIND of
>action will occur.

I stand corrected. HOW MUCH action, and IN WHAT DIRECTION, is determined by
the reference signal and the disturbance. HOW MUCH action, and IN WHAT
DIRECTION, of the influencee is what is crucial for the purposive influencer
to see what he/she wants to see -- that's why the purposive influencer
influences the influencee's environment in PARTICULAR ways. WHICH KIND of
action is crucial, too, and because that is determined by the physical links,
perceptual function, and effector design (as you say above), the influencee
attempts to MODEL (by using some form of The Test) the influencee's
controlling.

>The organism determines WHICH controlled variables it will be
>controlling at a given time. You can then pick any of them to disturb.
>When you do, the organism will vary its action on that controlled
>variable, cancelling your disturbance. It is, of course, already
>varying its actions on ALL the controlled variables, or is ready to do
>so, at all times, if they're under active control. You're not
>determining which variables will in fact be under control or which
>actions they will use. That's been decided already.

Agreed, except that the next to last sentence needs additional qualification.
It should read: "You're not determining which variables will in fact be under
control or which POTENTIAL action, given a particular disturbance, they will
use." What HASN'T been decided already, solely by the organism, is WHICH
action of the POTENTIAL actions will be used. The disturbance influences WHICH
action is in fact used to retain control. And this influence is what the
purposive influencer is being purposive about.

>The only way to make sure of applying a known disturbance is to isolate the
>control system from all those other influences.

There is a vast range of "degree of isolation" -- from Skinner boxes to chaotic parties -- where, the empirical evidence indicates, the "rubber-banding interaction" is deemed successful by the first party (that is, he/she sees what he/she wants to see, which depends on certain actions being performed by the second party).

BP>>>There's no difference in the action used to counteract an
BP>>>accidental influence or an intentional influence.

GW>>Correct. But that is beside the point. Purposive influence matters
GW>>to the influencer ALWAYS.

>Agree. Control system control their own perceptions. Are you making some new
>point here? Why the "but?"

No, an OLD point: the purposive influencer is controlling for his/her perception, which depends on the influencee's action, and the accidental influencer is not so controlling.

>Are you saying that a psychologist has a point of view that is not
>that of an organism?

No. I'm saying that many non-PCT psychologists, as you have noted several times, take the point of view of the experimenter (an organism, to be sure!) rather than their experimental organisms (i.e., rats).

>If proving this conclusion has been your goal from the start, you must have
>had other reasons for believing it is the correct goal, reasons that go
>beyond, predate, or supersede PCT.

Well, to a degree, I suppose. I mean that if you want to model organisms mechanically (that is, lawfully), using differential equations, then I WAS taught that ALL terms in the equations DO have influences on the trajectories of the equations' variables. But basically my goal came from seeing (I think) what PCT-science says, and then seeing contradictions between that and the apparent ideology of some PCTers.

>This would, of course, explain why you don't develop an irrefutable argument
>based only on PCT principles: there are other principles involved.

I HAVE been trying to develop such an argument, based on PCT principles and empirical evidence that purposive influence is common and often successful and persistent over time and characterized as sometimes "good" and sometimes "bad" by the influencees. My ideology comes in when I say I think the ideology of radical organismism is detrimental to greater acceptance of PCT by non-PCTers.

Seeing as how you've worked some constructive criticism of me into your post, I'll reciprocate. My major difficulty with your end of the argument is that you simply don't answer many of my explicit questions, I suppose because paying attention to them wouldn't suit your own purposes. I have tried to be careful to answer virtually all of your explicit questions. It would be nice if you started to do the same. (I have a fair backlog I'd still like answered, too.)

BP>>>Whatever the people do, they will do it with a particular
BP>>>organization of control systems, and exercising those control
BP>>>systems or bumping them up against each other is not going to
BP>>>change anything.

GW>>Look around. Peoples' lives are changed in both trivial and
GW>>profound ways (as judged by the influencees themselves) everyday,
GW>>all over the globe, by acts of purposive influence. Ask THEM
GW>>whether "anything" has changed.

>If they say anything has changed, they are ignorant of history. What
>you speak of is not change; it's just coping.

That's a big "just," isn't it? I thought "coping" (that is, attempting to
control one's perceptions) meant EVERYTHING to you (or at least all else
besides "coping" was trivial). The kind of change I meant is to be found in
the enriched or shattered lives of those who have been purposively influenced
by others, some for better and some for worse -- in THEIR OWN opinions.

>What you present as "purposive influence" is simply control of others.
>You may say it isn't, but you still insist that purposive influence is
>intended to have some effect other than just disturbing what others
>are controlling for.

Guided reorganization aside (which we have yet to discuss in detail, and I
WON'T prejudge the outcome of THAT aspect of the argument!), I claim that
purposive influence is intended (by the influencer) to have ONE effect, namely
enabling the influencer to control some of his/her perceptions which depend on
some of the influencee's actions. The influencer controls by disturbing what
the influencee is controlling for (or however you want to name the "rubber-
banding interaction" discussed above), without force or threat of force, and
without trying to intentionally change what the influencee is controlling for.
Is that "control of others"? Is it, whether "control of others" or not, always
"bad" in some sense? If so, in what sense? Why isn't teaching somebody to do
something a kind of purposive influence? Why isn't Ed Ford's counselling a
kind of purposive influence? Why isn't your arguing a kind of purposive
influence? Do you think PCT can say ANYTHING about interactions among people
other than that they are (a) nonexistent, (b) virtually impossible, (c)
trivial, (d) unimportant, and/or (e) always bad? Should sociologists look
elsewhere for explanations of the nature and limits of such interactions, as
Chuck Tucker suggests? I'd appreciate some answers.

Greg & Pat

Date: Thu Sep 24, 1992 4:56 am PST
Subject: Better Late than Never

From Greg Williams (920924)

>Bill Powers (920923.1500)

If this post of yours had appeared on the net just a bit earlier, Pat and I would certainly have acted differently last night! A slightly late attempt at purposive influencing....

>Clearly you can cause another -- even a chimpanzee -- to write any
>name in the rubber-band situation by applying a suitable pattern of
>disturbances.

See, Chuck?

>A behaviorist would have a hard time explaining why
>this chimpanzee does NOT know how to write his name.

No problemo. Skinner would have said that your "responses" had "come under the control of certain environmental stimuli through a history of your being reinforced" and that your "responses served as stimuli controlling the chimpanzee's responses due to a history of reinforcement." You and I think that the offered explanation is, to be charitable, incomplete, but Skinner wouldn't have had a hard time either spewing it off or accepting it himself.

>But a control theorist could do so easily.

Yeah, it only took a month for you to agree with me that purposive influence is as I've been saying it is. Chuck is still thinking about it.

>The agent writing the name is not the chimpanzee, but the human being
>operating the other rubber band.

HOORAY!!!!!! Now we both are in agreement with PCT science.

>Does this represent an "influence on the chimpanzee's behavior?" In
>once sense, certainly.

I prefer that you say "influence on the chimpanzee's ACTION (or OUTPUTS)." That properly (in PCT terms) distinguishes influencing another organism's action/outputs from influencing another organism's behavior/outcomes. The former can occur without physical force, threats of same, or reorganization by the organism; the latter can't.

>But it is an influence on how the chimpanzee behaves? No, it is not. The
>chimpanzee has acquired no new way of behaving. All it knew how to do before,
>and all it knows now, is how to keep the knot in the rubber bands over a
>target, as a way of getting an M&M later. All right, 100 M&Ms; who am I to
>argue with a chimpanzee's lawyer? When I talk about "changing" people's
>behavior, I automatically think in terms of doing something so that a person
>will then know how to do something new, or do something differently from then
>on. I don't count manipulating people's actions by using disturbances as
>changing their behavior because as soon as I'm not there to supply the
>carefully-chosen and patterned disturbance, that behavior will simply stop.
>It never had any reason to exist in the first place except to protect some
>controlled variable from being changed. If the behavior were continued after
>the disturbance ceased, it would CAUSE an error. There is absolutely no
reason
>for the "manipulee" to continue that behavior, and every reason not to do so,
>when the disturbance is removed.

Agreed. Now suppose I purposively influence you by telling you a fact which you didn't know before, then walk away. Believing the "new" fact to be true, you realize that that belief allows you to control something you were controlling before much more efficiently (in some way). Before you knew about the EXISTENCE of the "new" fact, you weren't worried about the efficiency with which you were controlling. Now you decide to "use" your newfound belief and control more efficiently -- after all, why not? HAVE YOU REORGANIZED? If so, I take it that such reorganizations are going on all the time in our control systems. Is that what you think? Or can there be "true" (to you) changes in an organism's behavior (that is, in its outcomes) following something short of reorganization? Regardless, the EFFECTS of my disturbance on you didn't die out when I walked away. You continue to "do something new, or do something differently from then on," so the results meet your criterion of "true" behavior changes. I want to hear whether you think it's ALWAYS reorganization -> behavior changes, and NEVER something-other-than-reorganization -> behavior changes.

>This is what I mean by "coping." Sure, people adjust their actions
>because of effects from other people and from the nonliving
>environment. But that just prevents their hierarchies from being
>disturbed at the level where the disturbance has its first effect on a
>controlled perception. They can interact with others 10,000 times a
>day, and they will have learned nothing new. They will still know only
>how to control their perceptions in the same ways they knew yesterday.
>When all the disturbances go away, the people will be exactly as they
>were before. Nothing of any importance has changed.

Nothing of importance to YOU. But much is happening in the myriad instances of purposive influencing WITHOUT reorganization (or perhaps something else resulting in truly changed behavior) which IS of importance to many sociologists, therapists, educators, family members, workers, politicians, advertisers, etc., etc. Many of them would welcome ideas on how to influence others' actions without generating conflict. Perhaps PCT might actually be useful beyond the ivory tower?!?! You can ask them if you don't believe me.

>This is why I dismiss all the "purposive influences" you bring up as
>being trivial. All they illustrate is how control systems with fixed
>organizations will act when their actions affect other people and
>other people's actions affect them. Out of these fixed organizations
>there can arise detailed behavior of immense complexity, in the same
>way that all those "people" in the crowd program exhibit complex
>patterns of interaction -- while behaving in a simple and unchanged
>way.

You admit that there IS purposive influencing and that it doesn't contradict PCT science. I have no problems with how you personally judge the importance of the phenomenon -- just don't tell others that THEY should consider the phenomenon "trivial," too. That would be anything BUT putting yourself in the other organism's shoes. I myself think the details of South African politics are uninteresting. But it would reach the heights of arrogance for me to advise people who live in the Homelands that those details are "trivial." They would likely dismiss me as the "trivial" thinker! And this has been a theme in my argument all along: ideology can have seriously negative effects on the dissemination of PCT science.

>Let's talk about reorganization.

Let's. You can start be answering the questions I posed above.

>That's the only area where anything interesting happens in human behavior.

You mean: the only area where YOU THINK anything interesting happens. I would welcome further discussion on "short-term" purposive influence with others on the net who ARE interested in this ubiquitous phenomenon.

Looking forward to a "new" argument,

Greg

Date: Thu Sep 24, 1992 9:17 am PST
Subject: PPCT Promotion Idea

Hi! I'm a student at NEIU who was introduced to PCT via Professor Robertson's Introduction to Modern Psychology - The Control Theory View initial class of 1990. I was very impressed with the volume and depth of material - and with how many answers I got to why and how people behave. That was my first university (and psychology) class having just retired from a 31-year federal (management in human rights law enforcement) career. I doubted my judgement about the value of CT because I had virtually no other theories to compare it to. Since then, I have completed my major in psych, and no longer have doubts about the value of CT (PCT). It is as amazing as I originally thought it was. I did two co-teaching classes, both with Professor Robertson along my route to accumulating credits. I began a glossary for the CT class which I am now planning to finish (hopefully by December). Professor Robertson has accepted my chapter submissions, edited them, and then we used them in the class as hand-outs. I'd be happy to share that when it is done. I've been reading posts for about a month now with an objective of understanding PCT better and learning how to communicate it effectively. I've tried to do this (in a friend/family individual setting) and, as it seems happens to a lot of PCT promotion/teaching, there was no significant effect. It occurred to me that Closed Loop (or perhaps here) might be used to gather information about the sources that actually "hooked" those of us who recognize the value of PCT to the field of psych and then, perhaps focus efforts to promote/teach it on those sources instead of the ones who do not seem to want to know about it. It would probably also be helpful to identify specific examples or experiences that convinced PCTers of its value. In reading many of the posts, I find those that relate real life experiences the most impressive and convincing. I feel like I understand PCT well as it is contained in my CT text, but I seem to need a lot of practice using the technical language to describe real situations. I am getting many exposures to that in reading the posts.

Carol Kiran NEIU Student E-Mail: Kiran @NEIU

Date: Thu Sep 24, 1992 10:37 am PST
Subject: Turing, Blind men, Hi Isaac

[From Rick Marken (920924.1000)] Martin Taylor (920923 18:00)--

>(Rick Marken 920910.0900)

>> But I do think that the Turing
>>Test is an EXCELLENT example of the behavioristic basis of AI (and
>>cognitive science) etc. Harnad's contribution makes it even clearer -- he
>>suggests a Total Turing Test meaning its not enough to get a simulation
>>to answer questions like a real person -- you must also get it to
>>behave in all ways like a real person -- ie -- brush teeth,
>>play soccer, build model airplanes, etc. Of course, what is interesting
>>is that behavior is defined completely superficially -- it is what you
>>SEE. There is no notion that behavior (including the conversation of the
>>original Turing Test) is a controlled consequence of simultaneous influences
>>produced by the actor and the environment.

>>...

>>

>>I think the Turing Test is an excellent model of the basic misconception
>>about behavior embraced by all the life sciences. According to the
>>Turing Test (Total version), the only problem with clockwork simluations
>>of behavior like those done in the 1700s was that they didn't mimic enough
>>behavior and, perhaps, didn't do it smoothly enough. The Turing Test
>>shows very clearly that the goal of AI and like minded "sciences" is
>>truly superficial -- they want to simulate how things "behave: not how
>>they control.

>It was my impression that the Turing Test involved an interaction between the
>tester and the testee.

I guess I gave the wrong impression somehow, but, really, I know that the Turing test is interactive (a dialog) -- note the fifth line of the first quote above. My point is only that the judgement of "intelligence" is based on the answers to the tester's queries. These answers are the behaviors to which I refer. I presume that the tester will take into account the fact that s/he has asked a question or made a statement and that the answer relates to it in some way (maybe).

> That implies that the tester can apply The Test as
>understood by PCT.

Absolutely. The tester can, but doesn't. Similarly, psychologists can view their manipulations of independent variables as disturbances to hypothetical controlled variables and watch for lack of effect and/or resistive actions. But they don't either.

> Turing, as I understand him, implied nothing about how
>the testee was to arrive at the observed behaviour, or whether the tester
>should, could, or would, observe control. So I think the quoted comments
>are misguided.

Maybe not clear, but not misguided. Of course the Turing test can be done as a test for the controlled variable; I've said that over and over again. My point is that it is NOT done that way (nor is psychological research done that way). I'm not criticizing Turing. My point is that the Turing Test (as it is typically carried out) is behavioristic in the sense that evaluations of intelligence are based on observed characteristics of answers to questions

(or, in the Total Turing Test, responses to stimuli) rather than on an understanding that what real intelligent systems (humans) do is control. I would be thrilled to find evidence that someone, somewhere had been doing the Turing Test as a test for controlled variables. But I am not aware that this has been done. Turing described his test in behavioral terms -- he had no idea that an "intelligent" system might be a control system and this is certainly understandable; no one else did either. What I mean by "behavioral" is that the judgment of "intelligence" was based on whether replies (outputs) to questions seemed like the kind that would be generated by an intelligent system. All that is missing from this is the notion that these outputs might be part of an effort to control perceptual input relative to a reference spec. This may seem like a small thing to miss -- one which is easily remedied -- and in a sense it surely is. But it's the whole ball of wax from a PCT perspective; and in practice, this "easily remedied" little omission has been actively defended by practicing students of psychology. That is, psychologists have insisted on keeping this omission (the test for the controlled variable) omitted from psychological research practice.

But, if you know of examples of the Turing Test or Psychological Research that was done as a test for controlled perceptual input (other than that done by PCT aficionados) I would really love to hear about it.

Avery Andrews (920924.1058)--

>It seems to me that the problem with the blind men paper is this: academics
>just aren't interested in being told how concepts they are used to dealing
>with can be reduced to others that they aren't, and abstract arguments that
>some other way of looking at things would be better likewise make no
>impression

I agree. Some of the best points are sort of hidden away -- like the main point which is:

When you understand that you are dealing with a control system rather than an SR, reinforcement or cognitive system then conventional approaches to research tell you nothing about how that system works; you have to change the goal and methods of research. The goal is to discover controlled variables and the method is "the test".

This message is hidden in the last paragraph of the paper.

I just like the paper so damn much (I feel like Pygmalion -- maybe I should call the paper "Galetea"). I think the "Blind men and the elephant" image is so beautifully appropriate -- and, of all things, it turns out that there are THREE ways to describe control (corresponding to the three major approaches in psychology -- I'd throw psychoanalysis in with the cognitive approach) just as there were three blind men and the elephant. So, I admit, the paper's an ego trip -- I really am not interested in converting people to PCT with it (and, as you note, I will probably succeed at that with flying colors).

Isaac Kurtzer -- Welcome to the net.

The Player Rick

Date: Thu Sep 24, 1992 12:18 pm PST
Subject: Re: Education as Manipulation

[Martin Taylor 920924 14:45] (Rick Marken 920924.1000)

Rick, it's just a fluke that I saw your response so quickly. I'm still some 150 messages behind in my catching up. But since I did...

I agree that, as far as I know, the Turing Test has not been used in the context of the PCT Test. I read your posting as saying that it intrinsically would not be so used.

If you think about it in the light of applying the PCT Test, there is still a problem, because in any human hierarchy there will be thousands of controlled or controllable variables, and only at the top (my view, not Bill P's) there will be the intrinsic variables. The computer has an entirely different set of "intrinsic variables" that have been designed into its software. One could conceive of distinguishing reliably between a silicon and a carbon intelligence if one could do things that might affect these intrinsic variables (one would observe the actions that related to their control, not changes in them, of course, and these actions would be exactly what you were talking about--the behavioural manifestations, not the CEV, which are unobservable in the case of intrinsic variables).

Also, in the case of a linguistically communicating machine, the only percepts that it could control are those it gets from the Turing Tester, so if it is to control anything, it must do so by eliciting what it needs from the Tester. As I understand the Harnad TTT, the subject has access to outside sensory and motor connections (outside the dialogue, that is). Such a machine could be observed to control variables that the Tester could disturb, so it would be more likely to be able to show behaviour (i.e. different levels of control) that would be hard to distinguish from a biological intelligence.

Martin

Date: Thu Sep 24, 1992 1:29 pm PST
Subject: What gloating gets you

[From Bill Powers (920924.0900)] Greg & Pat Williams (920923) --

>>We agree that you're controlling your own perceptions, on purpose.
>>We disagree that your kids are acting differently because of your
>>controlling for nutritional value of the food.

>OK, now we understand your point that, from the KIDS' standpoint,
>they are acting the same regardless of Pat. It's only from PAT'S
>standpoint that they are acting differently because of her
>controlling.

You don't understand yet. From Pat's point of view they are also not acting differently. They are still putting their spoons into the mush and lifting it into their mouths and swallowing it. They are still controlling for good taste and sufficient quantity. From Pat's point of view, something different is

HAPPENING because she has changed the composition of the food. But that has nothing to do with the kids' actions. "Action" means what they are doing with their muscles. She hasn't changed that. She's only changed an effect of the actions in a dimension that leaves the control systems undisturbed, doing exactly what they were doing before. If Pat believes she can see something different in their actions, she is mistaken. She may see a difference in consequences of their actions and in the appearance of the environment, but that is not a difference in the kids' actions. I hope you are keeping in mind that we use "action" to mean "output," to distinguish it from its consequences after passing through environmental functions.

>With Pat controlling for seeing them eat "healthy" (to her) food,
>the kids are controlling for (some) DIFFERENT lower-level
>perceptions than if she wasn't controlling thusly. Aren't they?

They are not. They are controlling for taste and quantity just as before. They can "control for" only what they perceive and have reference levels for. They do not control for the nutritional value of the food. You're assuming that because they're GETTING something different by way of nutrition that they're CONTROLLING for it. That's like saying that when you swerve your car to avoid hitting a dog, and run into a lamppost as a result, you're controlling for running into the lamppost.

>Surely the same lower-level controlled perceptions aren't involved
>when the kids eat what is set in front of them, rather than going
>to the store to buy food.

You're introducing a whole batch of new controlled perceptions when you talk about Pat not setting ANY food down, so the kids have forage elsewhere for oral input. The kids were controlling for being seated at a table in order to eat; now they have to control for a lot of other things before they can even think of eating, such as how to get to the store and how to get the money and how to get Mom to stop acting so weird, and so on. When they finally get all that straightened out, they will get to the store and control for (a) taste, and (b) quantity. They might end up ordering oatmeal porridge, because evidently they think that tastes fine. Or they might order popsicles, if not having popsicles has been an error signal for them. The popsicles, however, would pose a quantity problem as well as a cost problem. Easier to eat at home.

My point was that the kids are NOT, according to hypothesis at least, controlling for nutritional content, so Pat can vary that as she pleases without disturbing them, as long as bad tastes are avoided.

>>What they're getting is different because of that [change in
>>nutritional content], but what they're controlling for isn't.

>Yep. Analogous to somebody subjected to disturbances so he/she
>writes his/her name in the rubber-banding demo, right?

No, not analogous at all. Varying the nutritional content of the food disturbs NOTHING that the kids are controlling for, and so elicits NO change in action. Rubber-banding works ONLY when you are explicitly and directly disturbing the variable that the other person IS perceiving and controlling, the position of

the knot. In the case of the kids, you would have to disturb the taste and quantity of the food.

>>Their actions would remain the same no matter what you served
>>them, if they ate it.

>So the actions of the rubber-bander would remain the same whether
>he/she was signing his/her name, or John Hancock's, or writing
>nonsense syllables?

No, you're following a false trail from the false analogy. By serving any kind of food that the kids will eat, you leave their controlled variables UNDISTURBED, so there is no way you can influence the actions by which they control those variables. Rubber-banding requires applying a disturbance to the controlled variable that would change it if no corrective action occurred.

>Certainly, in all cases, the controlled perception of keeping the
>knot over the dot is the same. But it sure looks like different
>hand motions when the person writes a "JUC" instead of an "FUC"

But you can only get the person to move the hand by deliberately disturbing the controlled variable, not by avoiding disturbing it. We don't need to follow the rest out, do we? Your argument is based on a sloppy use of PCT principles. I wish you weren't so eager to prove you're right; it makes you careless. Your conclusions are running way ahead of your argument.

>Now wait a minute. What's this "they adjust their actions"
>business? Didn't you just say above that their actions were the
>same with or without a disturbance?

I most certainly did not. If you disturb a controlled variable, the person's actions that are used for affecting it will change, or they will have no control. You're talking about disturbing a variable that is NOT controlled. Of course if you mean that their actions in controlling one variable will remain the same when you disturb a different variable, then you're right, that's what I mean. You can disturb nutritional content all you like, without altering the kids' actions, as long as the result is not a disturbance of taste or quantity, the variables the kids are controlling.

I think I'll skip the rest, because I want to see your response to my last post. I hope you now agree that Pat's varying the nutritional content of the food while not disturbing the kids' controlled variables is NOT an example of purposive influence in the rubber-banding sense, precisely because the controlled variables are NOT disturbed.

Greg Williams (920924) --

>>Clearly you can cause another -- even a chimpanzee -- to write any
>>name in the rubber-band situation by applying a suitable pattern
>>of disturbances.

>See, Chuck?

>>A behaviorist would have a hard time explaining why this
>>chimpanzee does NOT know how to write his name.

>You and I think that the offered explanation is, to be charitable,
>incomplete, but Skinner wouldn't have had a hard time either
>spewing it off or accepting it himself.

>>But a control theorist could do so easily.

>Yeah, it only took a month for you to agree with me that purposive
>influence is as I've been saying it is. Chuck is still thinking
>about it.

Your glee at concluding that you've been right all along and that you've finally persuaded me of it (and are WAY ahead of Chuck) is noted. Evidently some extremely important reference level, having little to do with PCT and a lot to do with winning a competition, has been satisfied. Unfortunately, you have satisfied it the way Skinner would have done, not as I would have done. You haven't had much trouble spewing off explanations of why you are right, or accepting them yourself, despite the fact that your argument has been shot through with errors, some of which are discussed above.

Let's look at the prize that you've awarded yourself. You have shown that it's possible for one person to control the actions of another person by disturbing the variable that the action is involved in controlling. You have cited, as a way of overriding my objections to your argument, the rubber-band experiment and the particular use of it that shows how one person can control another person's finger position -- both of which demonstrations I invented and used to illustrate exactly the same point.

Unfortunately, you chose as an example a case in which this demonstration doesn't apply. This suggests that you believed in the conclusion, but didn't understand how it was reached. Of course I know that you DO understand how it was reached. I can only conclude that you thought it more important to prove that you were right than to do so by a rigorous argument.

Practically all of my objections to your arguments in the past month or two have been aimed at showing errors in the way you were applying PCT. You took those objections as arguments against your conclusion, and simply shifted, each time, to a different argument instead of examining what I said and realizing that your previous argument was invalid in PCT terms. The outcome was that you have slipped and shifted and misstated, flailing around with the sole apparent purpose of convincing me that situations analogous to the rubber-band experiment work exactly as I say they do. I agree. You win. These relationships work exactly as I believe, and have believed for a long time, that they do.

In my post last night, I finally got around to saying why I didn't consider the rubber-band type of relationship to be of great importance in human interactions. When it's used successfully, it does nothing to frustrate or aid the control processes of the other person, it leaves the other person unchanged, and its only purpose is to satisfy a desire or need on the part of one person to involve another's behavior in control of the first person's perceptions. It is therefore a trivial interaction, hardly amounting even to an interaction because it's important only to one of the parties. To the other person, it's just another disturbance, easily and even automatically counteracted.

Interactions that matter far more can be seen only when we ask WHY a person might specifically want to involve another person's actions in his or her own control processes. This consideration leads me to see your discussions of Pat's feeding the kids as disingenuous (even if you think it has something to do with rubber-banding, which it doesn't). The way the development turned out, Pat was just controlling for her own perceptions of a certain type of food going into the kids' mouths. She was not trying to control anything internal to the kids.

>Nope, say both Pat and Greg. The purposive influencer doesn't want
>to change others' wants, only their actions to control those wants.

Aside from the sloppiness of saying that actions control wants (instead of perceptions), this statement says that Pat has no higher-level motives. If the kids shovel her food into themselves, that's all she wants. She isn't trying to influence anything inside the kids, like their state of health or their food preferences. For all she cares, they can eat Big Macs and popsicles now and for the rest of their lives, as long as she can see them putting the food she prefers into their mouths any time she wants to see it there.

Balderdash. Pat wants to keep her kids healthy and she hopes that they will develop a preference for healthy food that will last outside the home and after they have left home. She wants to control their bodies and even their preferences and beliefs about food. I would conclude something similar if Pat were a mad killer, explaining "Oh, I just like to perceive arsenic sprinkled on people's food; I'm not trying to poison anybody." Pat is a mad child-nurturer. You can't fool me by trying to keep my attention on the lower levels of control.

>>If they say anything has changed, they are ignorant of history.
>>What you speak of is not change; it's just coping.

>That's a big "just," isn't it? I thought "coping" (that is,
>attempting to control one's perceptions) meant EVERYTHING to you >(or
at least all else besides "coping" was trivial).

When have I said or hinted that reorganizing is trivial? I use "coping" to mean the way people interact when their control systems are fixed in organization. I'm willing to discuss other forms of learning (mentioned in BCP), but I define coping as how systems act when they're not learning.

>The kind of change I meant is to be found in the enriched or
>shattered lives of those who have been purposively influenced by
>others, some for better and some for worse -- in THEIR OWN
>opinions.

If those lives are enriched or shattered without any learning being involved, it's just the luck of the draw. Next time, they'll get enriched or shattered again in just the same way. Play the record again. HOW many abortions did you say you've had, little girl?

>Guided reorganization aside ...

Yes, far aside. You'd better explain how you can guide a random process first.

> ... I claim that purposive influence is intended (by the
>influencer) to have ONE effect, namely enabling the influencer to
>control some of his/her perceptions which depend on some of the
>influencee's actions.

And again I say balderdash. If that's all that purposive influence amounted to, we wouldn't even need a name for it. In fact, purposive influence even in the rubber band experiment is specifically aimed at controlling _a perception of the other person's action_. The point is not just to control some other perception that's dependent on those actions; the point is to make those actions, as perceived, be exactly what you want them to be. I tried to point out the difference in your passing-the-salt example. If all I want is to control a perception that depends on some of your actions -- in other words, if all I want is to perceive the salt shaker in my hand -- then I don't care who passes me the salt or how. I am not controlling the action by which you pass me the salt. It can vary all over the place and I will do nothing to restore it to any particular form of passing the salt. What I'm controlling for is the salt, and perception of your action in passing it to me is not part of my controlled variable.

This is a BIG difference. If I say "Please pass me the salt," you not only are free to comply or not, but you can achieve the result of passing the salt in any way that's convenient to you, including asking someone closer to me to do it instead. But if I want to have a PURPOSEFUL influence on your ACTION, I will say "Please move your hand to the salt shaker, grasp it, lift it six inches, move it in a straight line to a point over my hand, lower it, and let go." That is purposeful influence of your ACTION, if it works. If it's an upperclassman talking to a plebe at a military academy, it might well work, because of the TREMENDOUS threat of physical force that lies behind it. In any normal situation, the other person would look at you as if you were crazy and pass you the salt as he or she pleased to do it -- or tell you to buzz off.

>The influencer controls by disturbing what the influencee is
>controlling for (or however you want to name the "rubber- banding
>interaction" discussed above), without force or threat of force,
>and without trying to intentionally change what the influencee is
>controlling for. Is that "control of others"?

Yes. It's control of an action by another. No controlled variable in the other person is disturbed significantly, meaning to an extent that the other can't easily oppose without inconvenient effort.

>Is it, whether "control of others" or not, always "bad" in some sense?

It's not "bad" or "good" when it works as intended. It has no effect on any controlled variable in the other.

>If so, in what sense?

When the action you're making the other perform has a side-effect of disturbing some other variable that's under control, the other will resist the side-effect. If that resistance is successful, there's no problem. If, however, a very large effort is required in order to counteract the side-effect (for example, performing the action may result in fatigue), then the purposive influence will most likely fail. If the purposive influencer

doesn't change the purpose, he or she will simply try harder. This will lead to conflict. If the purposive influencer considers conflict "bad", the only solution is to change the purpose.

>Why isn't teaching somebody to do something a kind of purposive influence?

The "teaching" part is: you have complete control over your methods of teaching. What you have no control over is their effects -- that is, on whether learning takes place, and if it does, what will in fact be learned. If the learner intends to learn what you have to teach, the learner can try things, ask questions, ask you to show it again, and do all sorts of things that can result in learning. This may please you as a teacher; it may be what you wanted to happen. But you have no control over its happening. You have an influence, to be sure, but it can't be purposive. Purposive influence requires a complete control loop.

>Why isn't Ed Ford's counselling a kind of purposive influence?

Ed will tell you. You're drifting away from your original definition of purposive influence again. I remind you of it:

>The influencer controls by disturbing what the influencee is
>controlling for without force or threat of force,
>and without trying to intentionally change what the influencee is
>controlling for.

To this should be added WHAT the influencer controls: the influencee's ACTION. That isn't what Ed does.

Skinner thought that teaching and psychotherapy amounted to nothing more than getting organisms to produce particular ACTIONS.

>Why isn't your arguing a kind of purposive influence?

Because I have no control over your actions, your perceptions, or your reorganizing processes. My aim now and for the past month has been to correct what I perceive as errors in the application of PCT principles. I state my corrections. What you do with them is up to you. The nearest I have come to a direct attempt to control your actions was in saying "I wish you weren't so eager to prove you're right." And even then, I marked it as a wish, not something over which I have any control.

>Do you think PCT can say ANYTHING about interactions among people
>other than that they are (a) nonexistent, (b) virtually impossible,
>(c) trivial, (d) unimportant, and/or (e) always bad?

Yes.

>Should sociologists look elsewhere for explanations of the nature
>and limits of such interactions, as Chuck Tucker suggests?

Yes, in addition to using the principles of PCT. Interactions among organisms do not follow the rules of PCT because the elements interacting are separate organisms. You can have negative feedback, positive feedback, and no feedback in relationships among organisms. The individual organisms continue to obey

the rules of HPCT, but their interactions bring in new considerations. The arcs and rings that develop in the crowd program can't be deduced from examination of any one organism in the crowd; those are outcomes, but not intended or controlled outcomes. In themselves, they have no significance to the individuals. But they may have consequences that influence variables that the individuals are controlling.

 I am very tired of this argument. I'm beginning to understand why people will confess to crimes they didn't commit.

Best, Bill P.

Date: Thu Sep 24, 1992 2:19 pm PST
 Subject: Re: control/influence discussion

[From Jeff Hunter (920923 - 1)]

In this discussion of "control of control systems" there is an important point that does not appear to have been explicitly stated:

People (sometimes) have a goal of recognising when they are being successfully controlled, and breaking this control.

For example:

> [From Bill Powers (920906.0700)]

> The simplest example of this that I know of is also the first one I
 > fell for. Someone says "Hey, your sister just fell off the swing!" I
 > spin around and look, and see nothing of the kind. The other kid
 > screams out "Made-ya-look, made-ya-look!" and collapses laughing.

> Once this has happened, it creates a conflict. ...

Why should it? All it has cost is one spin, a minuscule portion of a kid's daily energy expenditure.

The other kid tricked you solely to prove that he could. In other words he was not controlling for you to turn your head, but controlling for the sake of control.

A more extreme example of this is The Conqueror forcing The Helpless Captive to lick his boots. This is not done because it is an effective method of shoe shining. It is done to show the captive, the conqueror, and the audience that the conqueror is higher in the pecking order than the captive (and presumably the audience).

This gives us two common meanings for the word "control"

- a) I can make something do what I want it to
- b) I can make something do what it doesn't want to
 (or I can make it choose the lesser of two evils)

I think the reason Greg is concerned about the discussion so far is that the PCT term for control seems to follow meaning (a) for inanimate objects, but suddenly jumps to meaning (b) for people. I PCT-control a rock if I wish to perceive it holding open a door and I maintain this percept against environmental disturbances, but I don't PCT-control a person holding the door unless I maintain this percept **against*the*will** of the person (i.e. with force or trickery).

I don't think this usage is set in stone yet. Rick and Bill seem to hold to it, while Greg doesn't.

Oh well, hope this introduces some middle ground into the discussion.

... Jeff

> That kid, by the way, stopped doing that when someone else beat him up
> for doing it.

Control-by-trickery countered by control-by-force. The subject of innumerable tales.

Date: Thu Sep 24, 1992 2:48 pm PST
Subject: Re: misc comments

(ps 920924.1600)

[From Rick Marken (920923.1430)]

Right off the top of my head, how about "conservation of number"? Young kids apparently control for "amount of goodies" in terms of linear extent instead of number. The one I love is this: put two rows of M&Ms in front of a 4 year old thus:

(1) m m m m m

(2) m m m m m

this example is cool.

penni sibun (920922.1400)

I agree. But judgement implies a comparison of some kind, I think.

yes, i'm sure there's something against which a comparison is made; i'm just not sure we can (usefully) characterize it.

>the turing test is not about behavior. it lies in the interactionist
>realm, in my opinion.

But don't you look at behavior to judge whether or not it is intelligent? Maybe I should make it clear that my (informal) use of the word "behavior" includes written or spoken responses to questions (which are also

"behaviors")

oh, i agree verbal behaviour is behavior (too bad my spelling's not consistent!). my contention is w/ yr (apparent) assertion that it's ``just'' behaviour, or, worse, just ``the other's'' behavior. what *i* do in my interaction w/ an entity is inseparable from what it does.

>> The Turing
>> test focused on actions and their consequences

>no. it focuses on an *interaction* between a judge and an entity.

One kind of interaction goes like this: I ask a question, the machine makes a reply (which might be the null reply -- silence). The reply leads to another question or statement by me, then a reply by the machine, etc. If what I do is x and what it does is y (x and y being variables) then an interation is x--y--x--y--x--y.... Now I think you are saying that the intelligence of the system generating the y's is judged from the interaction x--y--x--y--x--y... Some sets of x--y--x--y--x--y... can be judged as intelligent, some not. Is that something like what you mean?

something like it. we can assume that a teletype forces linearity in the interaction, so x and y constitute turns. however, by x, you are simply specifying a stretch of typed text bet. two ys. these demarcations may not be particularly significant ones. (eg, a y can be composed of an answer, a long pause, and a question.)

If so, I certainly would agree that that is a legitimate way to interpret the Turing test. I still think it is "behavioral" but I won't even try to deal with that until we can reach some agreement about what the Turing test IS.

ok. it's behavioral; we agree behavior is involved. i'm not sure it's Behavioral, in that it constitutes a behaviorist explanation.

cheers. --penni

Date: Thu Sep 24, 1992 3:44 pm PST
Subject: Re: Turing

[From Jeff Hunter (920924 - 2)]

> (Rick Marken 920910.0900)
> But I do think that the Turing
> Test is an EXCELLENT example of the behavioristic basis of AI (and
> cognitive science) etc. Harnad's contribution makes it even clearer -- he
> suggests a Total Turning Test meaning its not enough to get a simulation
> to answer questions like a real person -- you must also get it to
> behave in all ways like a real person -- ie -- brush teeth,
> play soccer, build model airplanes, etc.

Harnad's addition to the Turing test is spurious. Steven Hawking (A Brief History of Time) cannot brush teeth, play soccer, or build model

airplanes. He only talks through a device very much like a teletype. Yet he is generally considered intelligent.

> [Rick Marken (920924.1000)]

> Maybe not clear, but not misguided. Of course the Turing test can be
> done as a test for the controlled variable; I've said that over and over
> again. My point is that it is NOT done that way (nor is psychological
> research done that way).

Well here's an excerpt from Turing (via Hofstadter).

Interrogator: In the first line of your sonnet which reads "Shall I
compare thee to a summer's day", would not a "spring day" do
as well or better?

Witness: It wouldn't scan.

Interrogator: How about a "winter's day"? That would scan all right.

Witness: Yes, but nobody wants to be compared to a winter's day.

The interrogator is clearly trying to find which variables in the sonnet the witness is trying to maintain against disturbance. Turing probably would have phrased it as trying to discover what internal concepts and goals were held by the witness. PCT talk would describe these as CEV's and references.

> ... I would be thrilled to find evidence that someone, somewhere
> had been doing the Turing Test as a test for controlled variables.

A slight frission running down your spine? :-)

... Jeff

--

De apibus semper dubitandum est - Winni Ille Pu

Date: Thu Sep 24, 1992 4:40 pm PST

Subject: Re: Interactionism

[Martin Taylor 920924 19:30] (Wayne Hershberger 920915)

>Martin Taylor: (re: Martin Taylor 920831 1330)
> I believe that the question you are raising about whether
>the controlled variable is an environmental variable or a
>perception of an environmental variable is ambiguous for precisely
>the same reason that Agre's "leaning on the environment" is
>problematic. That is, what I said above applies here as well.
>The answer to your question depends upon the realization you
>are talking about. When you conceive of my driving (steering an
>auto) as the control of one of your conceptual realizations (i.e.,
>distinguishing between physiology and physics), then the

>controlled variable is physiological rather than physical (i.e.,
>it is not an environmental variable). However, since my driving
>necessarily involves the control of one of my perceptual
>realizations, and since that controlled variable is in the
>environmental part of our common perceptual realization (I trust
>that you do not perceive the car as being in my body any more than
>I do) the answer to your question is: in the environment. (This
>ambiguity is what accounts for "the 3 stages of Satori.") If two
>individuals share a common perceptual realization of the world, it
>should be possible for them to control the same environmental
>thing--at least in the perceptual sense of the term. Is
>that part of what you are driving at?

I don't think so. If I remember, the issue is that no-one can perceive or control anything in the environment, but, given that there exists a common (unknowable) environment, the percept I am controlling is unknowable to you. God may know if we "share a common perceptual realization of the world" but we can't know it.

There is no way that you can be sure, when applying the PCT Test, that you are perceiving the same function of the environment that I am. You disturb some environmental variable perceived by you. I do something that results in that environmental variable returning to its undisturbed position as perceived by you. You say I am controlling the environmental variable that you perceive, but that may not be true at all. The environmental variable that relates to the percept I am controlling may correlate with the one you perceive, or may include or be included in the one you perceive. For example, suppose you think I am controlling for the sum of A, B, and C, which are sensory values that you, by some magic, can determine I am receiving. So you manipulate the value of A, and find I compensate by changing B and/or C. You do various such disturbances, and find I always compensate. But I am controlling for $1.1A + 0.9B + C$, or perhaps for $A + B + C + D$, where D is something you haven't perceived, and which remains stable during your tests or which varies only slightly. You probably will never know that your percept is of a different environmental variable than mine. For one thing, your sensory input cannot discover exactly what mine is, so you would not know that my A is slightly different from yours. For all that Rick keeps saying that all PCT measurements are Ivory Soap, 99 44/100% pure, this cannot hold for percepts of more complex environmental variables.

A completely different aspect of the question is that even within me, I can know only my percepts. I may identify them with some CEV that I may try to describe to you, but I cannot know that my description is accurate, even forgetting that I cannot know whether your interpretation agrees with my intention.

So the 3 stages of Satori don't enter into the issue. When we talk loosely, we can talk about controlling aspects of the environment, and for the most part we won't go too far wrong. But sometime, that loose language is misleading, as in much of the argument about "controlling other people." We can't and don't control other people. They are in our environment, and we control only our perceptions of that environment. If our actions cause our perceptions of other people to approach our references for those perceptions then we can say (loosely) that we control the other people. But we control only our perceptions, not the people.

This is probably no clearer than before. My mind is still a bit jet-lagged. But I thought I would answer this before it passed too far by.

Martin

Date: Thu Sep 24, 1992 6:51 pm PST
Subject: harnad on intelligence

[Avery Andrews 920925.1236]

Here's a Harnad posting from comp.ai.philosophy that I think is rather a propos wrt the Turing Test:

"Intelligence" is an arbitrary word denoting clever performance that ordinarily requires "X" (where "X" is the mental state that allows humans and animals to do those same clever things). This simply shifts the burden from "intelligence" (now trivially "defined") to "X," which is not only still undefined, but clearly something that is so completely un-understood that our "defining" it now would be as sensible as Aristotle's "defining" gravity, 2000 years before Newton. (Even less sensible, because of the extra dimension of perplexity added by the mind/body problem.)

So don't bother trying to define intelligence; it's an empty, arbitrary exercise. Focus instead on what it takes to generate the clever performance that ordinarily requires X.

Harnad, S. (1992) The Turing test is not a trick: Turing indistinguishability is a scientific criterion.
SIGART Bulletin 3(4) October
(Retrievable by anonymous ftp from host: princeton.edu
directory: pub/harnad filename harnad92.turing]

--

Stevan Harnad Department of Psychology Princeton University
& Lab Cognition et Mouvement URA CNRS 1166 Universite d'Aix Marseille II
harnad@clarity.princeton.edu / harnad@pucc.bitnet / srh@flash.bellcore.com
harnad@learning.siemens.com / harnad@gandalf.rutgers.edu / (609)-921-7771

Date: Fri Sep 25, 1992 3:51 am PST
Subject: Attn. P. Sibun; repentance (might the end be near?!?)

From Greg Williams (920925)

Penni Sibun: please post your USPS address so Ed Ford can send you a sample copy of CLOSED LOOP.

>Bill Powers (920924.0900)

[in reply to Greg & Pat]

>You don't understand yet. From Pat's point of view they are also not
>acting differently.... "Action" means what they are doing with their muscles.
>She hasn't changed that. She's only changed an effect of the actions in a
>dimension that leaves the control systems undisturbed, doing exactly what
they
>were doing before.

You're right. I was wrong. This example of interaction is NOT analogous to rubber-banding. Pat is being careful NOT to disturb the kids' controlled perceptions (of relevance here) in order to control certain of her own perceptions which depend on certain of the kids' actions. So, there are at least TWO kinds of "short-term" purposive influence:

1. By disturbing the second party's control of some perception in such a way that the second party performs actions (to compensate for the disturbance) which enable control by the first party. ("Rubber-banding")

2. By altering the environment of the second party in such a way that the second party's ongoing control of perceptions is not disturbed, and that the second party's ongoing control (given the altered environment) enables control by the first party. ("Food-spiking"?)

For either of these types of interactions to be successful from the point of view of the first party (that is, his/her control is good), it is important that the first party have a good model of what the second party is controlling for/not controlling for -- perhaps more to the point, a good PREDICTIVE model of what the second party WILL control for/not control for WHEN the first party disturbs (type 1 interaction) or alters the environment so as not to disturb (type 2 interaction). Applying The Test (in some form, perhaps repeatedly) is key to developing and refining such a model. Empirically, I see both types of interactions occurring frequently, with widespread (but not universal) success from the standpoint of the influencers, and (afterwards) characterized as ranging from "good" to "bad" for the influencees, as judged by the influencees. Your earlier comment that the second kind of "short-term" purposive influence is barely an interaction at all appears to discount the model-building aspect (applying The Test), which is usually HIGHLY interactive.

All right, so (despite my density!) we've made some progress. It appears that we've discovered a kind of interaction wherein party A can (in theory, at least) control certain of his/her perceptions which depend on another organism's controlling and which is not of the rubber-banding type or of the (hypothetical) "guiding" learning/reorganization type. (You might recall that, a few weeks back, you thought it best to lump deceptive "short-term" purposive influence with rubber-banding, because both involve applying disturbances, even though in somewhat different ways.)

Whether or not you are interested in such "short-term" purposive influencing, I'd appreciate your remarks on the above conclusion. I'd hate to set out on the wrong foot again!

[in reply to Greg]

>You have shown that it's possible for one person to control the actions of
>another person by disturbing the variable that the action is involved in
>controlling.

Well, I THOUGHT the Pat-feeding-"healthy"-food example showed that. But, as you say, I was wrong. What I really showed was that it's possible for one person to control his/her perceptions by being careful to NOT disturb another's controlled variables, while altering the other's environment in certain ways. YOU showed, long ago, by inventing the rubber-banding demo, that "it's possible for one person to control the actions of another person by disturbing the variable that the action is involved in controlling."

>You have cited, as a way of overriding my objections to your argument, the
>rubber-band experiment and the particular use of it that shows how one person
>can control another person's finger position -- both of which demonstrations
I
>invented and used to illustrate exactly the same point.

Yes I did, wrongly. The Pat... example is NOT analogous to rubber-banding, which is your brain-child. Actually, the Pat... example is of a completely different sort than rubber-banding.

>Unfortunately, you chose as an example a case in which this demonstration
>doesn't apply.

Yes, I hear you. At least my mistake was not all for nought.

>This suggests that you believed in the conclusion, but didn't understand how
>it was reached. Of course I know that you DO understand how it was reached.

Honestly, I was confused by the "low-level mechanics" of the kids' actions being different if Pat was there vs. if she wasn't.

>I can only conclude that you thought it more important to prove that you were
>right than to do so by a rigorous argument.

I thought it was rigorous at the time. I did not and do not think it more important to prove that I am right than to argue (as best I am able) rigorously. Perhaps there is no way to convince you of that. Nevertheless, I will say explicitly that I am telling the truth.

>The outcome was that you have slipped and shifted and misstated, flailing
>around with the sole apparent purpose of convincing me that situations
>analogous to the rubber-band experiment work exactly as I say they do. I
>agree. You win. These relationships work exactly as I believe, and have
>believed for a long time, that they do.

I greatly admire your many accomplishments, including inventing the rubber-banding demo, of course. Far be it from me to not give credit where it is due. And now you have also participated in the discovery that PCT shows there is the theoretical possibility of DIFFERENT kind of controlling-dependent-on-others' -controlling IN ADDITION to your rubber-banding kind. As I said above, I think that's real progress in using PCT to understand the nature and limits of

social interactions, and I appreciate your patience, especially since say you aren't very interested in (as I call it) "short-term" purposive influencing.

I hope to extend this work, further considering the constraints on social interactions which follow from PCT. I know that you are interested in talking about what I've been calling "long-term" purposive influencing. I'd also welcome discussion with others who are interested in social interactions in general, whether they involve learning/reorganization or not. Wherever we go from here, thanks, Bill, for helping me get this far!

GW>>Guided reorganization aside ...

>You'd better explain how you can guide a random process first.

There are random processes and there are random processes. Some kinds of "stochastic" random processes combine random aspects with non-random aspects so that you can predict where they are headed, but not the trajectory they will take or how long they will take to get there.

As a metaphor, suppose that a ball with several motors distributed over its surface (not unlike the hairs on microorganisms) is in water in a swimming pool. Each motor propels the ball along a particular direction with respect to a ball-centered coordinate system (the ball doesn't float), and each motor turns on (and, simultaneously, all others turn off) randomly in time, for random-length time intervals. At one end of the pool is a hole through which the ball could pass, say, to ball-nirvana. The ball will either eventually move through the hole and reach ball-nirvana, or never (in infinite time) move through the hole. In this simple case, there is only one hole to move through, no others.

If the ball's random motion is analogous to learning/reorganizing and going through the hole (conveniently provided by the God of Balls -- an outside influencer) is analogous to learning/reorganization ceasing (due to finding the one solution to the problem set by the "guider"), there's your answer. IF the solution is found, it is the "guided" one. Of course, the solution might never be found -- possibly because the control system doesn't know the Pythagorean Theorem. Solution: More "guiding," through intermediate stages of learning/reorganization. Please don't presume I am implying that there necessarily must be ONLY ONE solution, or even a FEW solutions to the problem, set by the influencer. The influencer might be happy to see any of a HUGE CLASS of solutions. And please don't presume that the influencer must want to seeing the solution reached in a highly particular way; the influencer might care not at all about all of the trial reorganization attempts, and value only the eventual success -- perhaps just seeing the number "6," regardless of the manipulations the influencee went through to write it.

In such a case, the obvious strategy for the "guider" who wants "tight" control of his/her perceptions of where learning/reorganization is "taking" the influencee's control system is to break up the whole process into little "nearly obvious" steps. "Nearly obvious" means that the (class of) "new" outputs of the influencee which the influencer wants to see after a bout of learning/reorganization has finished are not very "far" is modeled by the influencer as requiring little alteration in the influencee's control system. For example, the influencer doesn't expect to successfully "guide" the

influencee, in one bout of learning/reorganization, from competency in counting to competency in long division.

Influencing so as to "guide" learning/reorganization is, no doubt, an art which is not guaranteed to succeed. But it doesn't differ in principle from control of inanimate entities, which also requires applying a model (granted, very crude models are sometimes sufficient here -- like "this rock will do the same thing as it did last time under these sorts of conditions"); the difference is that good control of one's perceptions which depend on someone else's outputs after learning/reorganization requires, in general, fairly sophisticated modeling of the other's controlling before and after learning/reorganization, in order to alter the other's environment so each step in the process is "nearly obvious" to the influencee. Another metaphor: the "guider" casts a sufficiently large net that the random floundering of the fish won't carry the fish outside the net, and then the "guider" makes sure (for each "nearly obvious" step in the learning/reorganization process) that the net has one outlet into a similarly arranged next net. Modeling is key to deciding how large the net must be.

It seems to me that "guiding" learning/reorganization would be made easier if the learning/reorganization process does NOT proceed by making LARGE "jumps" in particular reference signals, but by making SMALL deviations. If, during learning/reorganization, the values of reference signals resulting after one "trial" alteration of reference signals had absolutely NO correlation with the values before, then it might be EXTREMELY difficult for "guiding" (in particular, through "nearly obvious" stages) to work. But if such a lack of correlation INVARIABLY characterized learning/reorganization, it seems to me that successful learning/reorganization would almost always never happen. So I suspect that much -- though not necessarily ALL -- of the time, learning/reorganization is NOT characterized by such a lack of correlation. Perhaps unless success has been a VERY long time in coming via a series of successive small deviations in (sets of) reference signals, large jumps in reference signals aren't made.

All this is highly hypothetical, of course. But, so far, I see no internal inconsistencies. But I hope you will enlighten me about where, if any place, I am going wrong.

GW>> ... I claim that purposive influence is intended (by the
GW>>influencer) to have ONE effect, namely enabling the influencer to
GW>>control some of his/her perceptions which depend on some of the
GW>>influencee's actions.

>... I say balderdash. If that's all that purposive influence amounted to, we
>wouldn't even need a name for it.

I don't understand why you say this. To me, it seems like a pretty big "all"; one that sociologists, among many others besides the would-be purposive influencers and influencees themselves, are greatly interested in.

>In fact, purposive influence even in the rubber band experiment is
>specifically aimed at controlling _a perception of the other person's
action_.

>The point is not just to control some other perception that's dependent on
>those actions; the point is to make those actions, as perceived, be exactly

>what you want them to be.

I didn't say "some other perception" than the perception of the other person's action. I said "some... perceptions which depend on some of the influencee's actions." Isn't the perception of the other person's action a member of the class of perceptions which depend on the other person's actions? My intent was to not limit the possible controlled perceptions of the purposive influencer to only the perception of the influencee's actions. My definition INCLUDES what you say "amounts to" MORE than my definition! I think what you really mean is that you want to RESTRICT the definition to controlling ONLY perceptions of influencees' actions. Right? Assuming so, then consider the following.

>I tried to point out the difference in your passing-the-salt example. If all
>I want is to control a perception that depends on some of your actions -- in
>other words, if all I want is to perceive the salt shaker in my hand -- then
>I don't care who passes me the salt or how. I am not controlling the action
>by which you pass me the salt. It can vary all over the place and I will do
>nothing to restore it to any particular form of passing the salt. What I'm
>controlling for is the salt, and perception of your action in passing it to
>me is not part of my controlled variable.

>This is a BIG difference. If I say "Please pass me the salt," you not
>only are free to comply or not, but you can achieve the result of
>passing the salt in any way that's convenient to you, including asking
>someone closer to me to do it instead. But if I want to have a
>PURPOSEFUL influence on your ACTION, I will say "Please move your hand
>to the salt shaker, grasp it, lift it six inches, move it in a
>straight line to a point over my hand, lower it, and let go." That is
>purposeful influence of your ACTION, if it works.

It appears that these two types of relationships between the influencer's controlled perceptions and the influencee's actions correspond to the two kinds of purposeful influence (1 and 2). In rubber-banding, as you say, "purposive influence... is specifically aimed [by disturbing the influencee's controlled perception in certain ways] at controlling _a perception of the person's action_." But in food-spiking, purposive influence is aimed at controlling a perception which depends on some of the influencee's actions, but does not involve having a purposeful influence on the influencee's actions [it involves altering the influencee's environment, taking into account (a model of) the influencee's controlling so as to NOT alter the influencee's actions].

Do you have any problems with this reformulation?

GW>>The influencer controls by disturbing what the influencee is
GW>>controlling for (or however you want to name the "rubber- banding
GW>>interaction" discussed above), without force or threat of force,
GW>>and without trying to intentionally change what the influencee is
GW>>controlling for. Is that "control of others"?

>Yes. It's control of an action by another. No controlled variable in
>the other person is disturbed significantly, meaning to an extent that
>the other can't easily oppose without inconvenient effort.

Rubber-banding-type interactions are "control of others' actions," but food-spiking-type interactions are not. Still, of course, both afford party A a means of controlling certain of A's perceptions which depend on B's controlling. But only in the former case are disturbances applied deliberately by party A (those disturbances being effectively counteracted by party B's good controlling, so that party A sees the actions of B which party A desires to see). In the latter case, non-disturbing changes in B's environment are made by A so that when B controls in the absence of disturbances, A sees the perception A wants to see, which depends on B's actions, which are NOT controlled by A. Whew! I think I've got it at last.

>It's not "bad" or "good" when it works as intended. It has no effect
>on any controlled variable in the other.

Of course, rubber-banding does have SLIGHT effects on some controlled perceptions, unless control is PERFECT. "Food-spiking," if done well from the standpoint of the influencer, has NO effect. Right?

>When the action you're making the other perform has a side-effect of
>disturbing some other variable that's under control, the other will resist
>the side-effect. If that resistance is successful, there's no problem. If,
>however, a very large effort is required in order to counteract the side-
>effect (for example, performing the action may result in fatigue), then the
>purposive influence will most likely fail. If the purposive influencer
>doesn't
>change the purpose, he or she will simply try harder. This will lead to
>conflict. If the purposive influencer considers conflict "bad", the only
>solution is to change the purpose.

Of course, that goes for rubber-banding-type interactions only; not for food-spiking-type interactions, in which you DON'T "make the other perform" a certain action. Right?

>Skinner thought that teaching and psychotherapy amounted to nothing more than
>getting organisms to produce particular ACTIONS.

Yes. That is what he called "controlling" the organisms. At least I'm not THAT sloppy!

>My aim now and for the past month has been to correct what I perceive as
>errors in the application of PCT principles. I state my corrections. What you
>do with them is up to you.

I hope I am doing justice to your aim. One thing is for sure, whatever you might make of it: I am treating your ideas seriously.

Thanks again for your patience, and especially for providing explicit answers to many of my questions in my last post. I promise not to post the back log!!

I'm off to the Great University of Kentucky Book Sale, surpassed locally only by the Great Dayton Planned Parenthood Book Sale. Greg-nirvana! See you tomorrow.

Greg

P.S. Is the arm program all fixed? How's the paper rewrite coming?

Date: Fri Sep 25, 1992 5:47 am PST
Subject: INFLUENCE AND CONTROL AND OTHERS

CHUCK TUCKER [920925]

TOPICS: Future issues of CLOSED LOOP; Influence and Control;
Bothering with PCT

CLOSED LOOP GW 920923-2

I know it would be a great amount of work but I think that you (Greg) can edit the transcripts of the conversation so as to represent the major matters mentioned about "Influence and Control". I do think that it would be useful to have a "summary" of the discussion (as Cliff suggests). The other part of the project I suggest, i.e., comments by other on the Net and a future book, can wait to see if there is any interest in it. I think that one way to get people interested in these ideas is to have our books in their libraries in the country and ask (in some instances FORCE) people to read them.

INFLUENCE AND CONTROL WTP 920922.0900; 23.0900 GW 920922; 23; 23-2
PW 920923

Unfortunately, I still have a problem with your specification of what you mean and how you determine that A has influenced B. My dictionary has a definition similar to yours except it has some additional words: Origin of "influence" = to flow in - "the flowing of ethereal fluid or power of the stars affecting the character or actions of people"; "the power to produce an effect by indirect means"; "the power or capacity to produce a desired effect" SYN: clout, leverage, pull, sway, authority, power, control. My dictionary of word origins states: (see: affluent) "INFLUENCE first was a term in astrology, the forces determined by the flowing in of the stars; so accident first meant the falling into place (accidere) of the stars - only unbelievers deemed it chance"!

The problem with these definitions that I find is that they are stated in terms not consistent with PCT since they contain attributes, i.e., "the power or capacity," or notions which are S- R, i.e., "to produce an effect" or "to flow in". But I don't think you are using INFLUENCE in that way (at least most of the time) but rather as a description of interaction where in A figures out what B wants (call it "X") and provides it for A in such a way that their interaction continues w/o force on the part of either A or disturbances on the part of B.

But it seems to me that in PCT terms for A to influence B to "produce a certain effect" that B would have to take and use a purpose (reference signal "X") "provided" by A since "a certain effect" in PCT terms is action taken by B consistent with B's reference signal X. Thus, with this specification, Pat (please excuse me if this is offensive to either of you) has influenced your childrens' food preferences to the extent they have adopted similar notions of "healthy" and /or "tasty" food and it seems that there is no other source for

these preferences. But even if this is the case, Pat's actions have served only as an "occasion" for the children to develop these preferences from a PCT perspective. I say this not as a criticism of Pat (or anyone else, including myself) but to simply point out that in PCT terms all that we DO serves as an "occasion" for what others DO; we do not determine what they DO.

This discussion has influenced me to think about THE TEST for cooperative acts performed by 2, 3, 4, or more persons where the disturbances do not come from the persons engaged in the cooperative acts but from others. The simplest one we have already discussed: ask two or more people to use rubber bands to keep one knot over a target and then apply a disturbance to the knot and observe what they do to maintain the knot over the target. Now what about other cooperative acts: sawing wood together, a crew rowing a boat, moving a bed together, singing a song or playing music together, walking as a "with", preparing a meal together, an "assembly line" operation to make a product (think of the Charlie Chaplin movie)? What do people do when disturbances are introduced to maintain the social act to its completion? Sounds like some ideas for student projects in a "collective behavior" course!!

A NOTE ON BOTHERING WITH PCT

Using PCT, it seems to me that a person would not adopt the ideas of PCT unless it was consistent with their current notions and PCT offered some assistance or improvements in the development of said ideas OR that the disturbances were so great that definite reorganization was in process and PCT was useful in that process. Now if most professional psychologists (sociologists, social workers, educators . . .) find no problems (I have written most of my papers as critiques of conventional theories and I have never found anyone that told me that what I wrote disturbed them so much that it converted them) with their current views (see the delightful post of Francisco Arocha 920921 13:07) then we should not expect them to find PCT useful. Then add to this that there is very little in the way of "support" and "encouragement" to be found in the relationships that people find themselves for these ideas (which are so different and strange and "non-common sensical"). I think it is a wonder that anyone ever adopts even one little part of the PCT formulation (I also wonder when I find people that use the ideas of Dewey, James, Pierce, Kantor, Mead, and Bentley). So, keep writing and telling ourselves about the formulation and maybe 100 years from now there will be a revolution of the "behavioral sciences" when someone with a problem discovers the logs of CSGNet on a tape somewhere.

Regards, Chuck

Date: Fri Sep 25, 1992 5:52 am PST
Subject: LASTEST VERSION OF RUBBER BAND EXPERIMENT

CHUCK TUCKER 920925B - CSG-RBND

INSTRUCTIONS FOR A RUBBER BAND EXPERIMENT

1. Review the "rubber band experiment" described by Runkel in his statement (Chapter 10) on "control theory." You will need a rubber band knotted in the middle (called RB below), a target diagram (three examples included and others suggested) and a

flat surface.

2. Select a person, P with which to carry out a modification of the RB experiment. Instead of having P guess what is reasonable for your behavior, you will ask P to adopt a particular reference signal and perform accordingly.

3. Place the target diagram (see below) on the table and ask P to keep the knot in the middle of the RB "over" the center of the target by saying: "Please put your finger in this loop of the rubber band and keep the knot in the center of the rubber band above the "X" in the middle of the target. I will put my finger in the other end of the rubber band and move my end but you should keep the knot over the "X" until I say STOP."
(PAUSE UNTIL BOTH OF YOU HAVE YOUR FINGERS IN THE ENDS OF THE RB AND HAVE PLACED THE RB ON THE TARGET DIAGRAM WITH THER KNOT OVER THE "X") Say "STOP" when you have accomplished one of the purposes described below.

4. Your assignment is to move your loop of the RB such that you can place P's finger (which is inside the other RB loop) over the letters on Diagram A to spell out the word 'CONTROL' by having placed his/her finger on these letters in sequence: C - O - N - T - R - O - L. When you have done that say "STOP". Then ask P: "Do you recall what word you spelled when your finger touched the letters?" WHAT EVER THE ANSWER TELL P: "The word you spelled was 'CONTROL'." THEN SAY TO P: "I want you to keep the knot over the "X" again I and will tell you to stop for each of the letters". DO AS ABOVE AND HAVE P SPELL THE WORD 'CONTROL' STOPING WHEN EACH LETTER IS TOUCHED BY HIS/HER FINGER.

HINT: If P maintains her/his RS (knot over the "X"), you should be able to place P's finger over the letters. P cannot control both her/his finger and the knot: they are connected by P's maintenance of the RS and P's resistance to your disturbance. If P wants to control the knot over the "X", P must resist your disturbance. With ingenuity you can therefore get P to place her/his loop finger over all of the letters.

5. You can do this assignment with Diagram B (colors) by have P's finger touch a sequence of colors. You can make a diagram of your own with figures (triangles, circles, squares, rectangles) at both ends of it; another with numerals like 1, 2, 3, 4, 5, 6, 7, 8, 9 at both ends of it, as long as you can get a sequence of moves, remember them and have P touch them.

6. Repeat this exercise with a total of three persons.

7. Your report is due {see syllabus}. It should be typed (no more that 3 pages, double spaced) and include:

(a) a description of your procedures, P's verbal and non- verbal actions, and the outcome for all three persons.

(b) based on the observations reported in (a), write a one paragraph analysis of the hypothesis that you can influence another person only if it serves some purpose which is important to that person.

8. You could add an Appendix to your report by describing what happens when you ask P to see if he/she can get you to spell a word, touch colors, figures, symbols or numerals.

DIAGRAM C

@	%	&	*	\$!	?
?	!	\$	#	&	%	@
*	&	@	*	?	\$	%
@	%	&	*	\$!	?
?	!	\$	#	&	%	@
*	&	@	*	?	\$	%

```
*****  
|  
| X |  
|  
*****
```

@	%	&	*	\$!	?
?	!	\$	#	&	%	@
*	&	@	*	?	\$	%
@	%	&	*	\$!	?
?	!	\$	#	&	%	@
*	&	@	*	?	\$	%

Date: Fri Sep 25, 1992 5:52 am PST

Subject: LATEST VERSION OF COIN TEST

CHUCK TUCKER 920925C - CSG-COIN

INSTRUCTIONS FOR OBSERVING AND RECORDING "THE COIN TEST"

1. Have a person, P, arrange four "coins" (DON'T USE ACTUAL MONEY BUT RATHER USE EITHER CIRCULAR PAPER DISKS EACH THE SIZE OF A QUARTER OR SAME COLOR POKER CHIPS), on a table, making a specified pattern which P writes down in advance but does not show you. You might do this by asking a classmate or friend to help you with a class project which will take about twenty minutes and can be done without going to any other place [this assumes you have 4 "coins", a writing instrument and some paper with you and a table top, flat board or desk]. If s/he agrees then say something like this: "I will give you four [NAME OBJECTS YOU ARE USING] and a piece of paper. What I want you to do is to think of a pattern in which you can arrange these (name of objects) on the flat surface of this table (board) (desk) but first I want you to draw and name that pattern on this paper but DO NOT SHOW IT TO ME. Next, arrange the (name of objects) in the pattern that you have on your sheet of paper. Then, without having any more conversation between us (except for the announcement "No change in pattern" when my moving of the (object)(s) does not change the pattern) until I say "Game is over," I will change the pattern and you are to put the coins back in the pattern after each time I change it. Do you understand this game?" IF NOT - repeat the instructions as written above again and clarify where necessary.

2. Your task is to discover what P has in mind without asking any questions or using any verbal communication at all. Your discovery procedure is to change or disturb the pattern among the objects by moving the objects around in some ways and noting what you have done to change the coins. So before you make a change, draw and name the pattern of the objects on a piece of paper without showing it to P and then after you make the change draw and name the pattern that you made with the changes.

a. If your change alters the pattern P has specified, P must correct the error by re-arranging the objects to re-make the original pattern.

b. If your disturbance does not alter the pattern that P has specified, P must announce "No change in pattern."

3. Repeat this process until you are certain that you can:

a. specify and demonstrate three disturbances that will call for P to re-arrange the objects and correct the pattern; this you should be able to discern from what you have drawn on your pad. [YOU WILL HAVE TO CHANGE THE OBJECTS MANY TIMES TO OBSERVE THREE DISTURBANCES AND THREE "NO CHANGE IN PATTERN" INSTANCES]

b. specify and demonstrate three disturbances that resulted in P announcing "No change in pattern"; this you should be able to discern from what you have drawn on your pad.

4. Compare your drawings with P's drawing of the pattern and

a. report the extent of agreement, including whether you identified the pattern but named it something other than what P named the pattern.

b. report at least one example of failing to see the pattern to which P returned the objects following a change that you introduced.

5. Your report is due {see syllabus}. It should be typed (no more than 3 pages, double spaced) and include (use the list below as outline):

- a. your report of agreement between you and P (4a)
- b. your answer to 4b
- c. a brief statement on what you learned from "The 'Coin' Test."
- d. a copy of P's written specification of the pattern
- e. your drawings of three change (3a) and no change (3b) patterns

Date: Fri Sep 25, 1992 6:51 am PST
Subject: Re: Why not influence?

[Martin Taylor 920925 10:00] (Bill Powers to Greg Williams 920920.0900)

I've been reading through the long discussion between Bill and Greg, and I guess I have got about halfway through what went on while I was away. But this comment pulled me up short. Perhaps Greg answered it explicitly. Bill answered it in earlier parts of the dialogue, but seems not to have recognized his own answer. Here's the comment, and my answer to it...

=====

This is part of Bill's response to Greg's discussion of Pat giving foods that she believes likely to enhance the kids' health, "linking" her choice to what she believes the kids will like.

>I have a problem with terms like "linkable wants." What good are such
>terms? They just make us try to interpret them in PCT terms. Why not
>simply express what you mean in PCT terms in the first place? Is
>"linking" a new phenomenon that requires a new explanation? To me, it
>seems a step backward into vagueness and ambiguity. The general term
>"linking" can apply in a lot of specific situations. You could "link"
>my eating what you put in front of me to treating me nicely; i.e.,
>make your nice treatment of me contingent on my wanting what you want.
>If you didn't mean THAT kind of linking, then you shouldn't use
>"linking."

The word that comes to mind is "side-effect." Everything you do has side-effects that do not contribute to the controlled percepts you have. Bill pointed this out in the earlier part of the discussion about rubber-banding. If I want a gorilla to write its name, I train it to keep the rubber-band knot stationary, and then with my end of the rubber band, I write its name. The side-effect of its keeping the knot stationary is that it moves its hand in the pattern of its name. It doesn't care about that, but I have "linked" the thing it is controlling for with the thing I am controlling for.

Pat allows the children to control for the perception of getting tasty food, but links that with what to them is a side effect of satisfying that reference. She controls for them getting what she believes to be healthy food. To her, the tastiness of the food is a side-effect, useful only insofar as its side-effects when the children control for it allows her to satisfy her perception of them getting healthy food.

If you don't like "linking" for this phenomenon, why not? Controller A sets up the environment so that a controlled (by assumption) perception of B results in actions that have side-effects that reduce the error of some reference held by A. These side effects would not in general occur, but A's acts cause them to be a consequence of B's control for something else. That's "linking" to me, and it is the prototype of much that we do when we claim to "control" other people. We set up situations so that the anticipated side effects of their actions that form part of their control loops are main effects of our control loops, whereas what they actually control is of no interest to our control systems. ===== Secondly, I have to disagree strongly with Bill about whether Pat controls for her perception of the children's health (or rather, whether she could do so). Bill says this is impossible because she can only observe statistically whether they seem to be healthy, and if they are she does nothing. But Bill pointed out to me that he thinks every muscle in our body is continuously controlled, even if it is relaxed and need do nothing because its reference is satisfied. If Pat is satisfied that the children are healthy, how does that mean she is not controlling for her perception of their health? And if, as is necessarily the case, her judgment of their current state of health is somewhat uncertain, how does that mean she is not controlling for their health being perceived as good? If she is not sure whether providing certain food is benefiting their health, and other influences make its contribution both time-varying and uncertain, how does that mean that using the food is not an aspect of control? Not all micro-loops in a functioning control system are always known to have the correct feedback sign. If they did, there would be no reorganization. We would all be perfect, always.

Statistics is all-important when it comes to assessing the appropriateness of the signs in our output functions. To dismiss an observation because it is statistical is as shortsighted as to dismiss the base premise of PCT because a lot of behaviour can be discussed without it.

Two beefs in one posting. Linking (side-effects) and statistics. Maybe I shouldn't have linked them. If I were a politician, I wouldn't.

Martin

Date: Fri Sep 25, 1992 7:38 am PST
Subject: THE TURING TEST?????

CHUCK TUCKER 920925D

To whom it is relevant:

I have not paid careful attention to the discussion of the "Turing Test" so I probably missed the basic idea. Would someone please describe this "test" to me (if possible on Net) or point me to a source so I can read about it? By

the way, I got a tape of Noam Chomsky's Killian Lecture that he gave last year and he has a line which is quite critical of the Turing Test in which he said it was not an advancement over the thinking of Decarte; sounded very damning to me.

Thanks in advance to whomever, Chuck

Date: Fri Sep 25, 1992 9:19 am PST
Subject: Re: THE TURING TEST?????

[Ray Allis 920925.0900]

> CHUCK TUCKER 920925D
>
> To whom it is relevant:
>
> I have not paid careful attention to the discussion of the "Turing
> Test" so I probably missed the basic idea. Would someone please
> describe this "test" to me (if possible on Net) or point me to a
> source so I can read about it? By the way, I got a tape of Noam
> Chomsky's Killian Lecture that he gave last year and he has a line
> which is quite critical of the Turing Test in which he said it was
> not an advancement over the thinking of Decarte; sounded very damning
> to me.

In short form: The question was "Can machines think?" Turing asserted that a definitive answer to the question was meaningless; not worth any effort. He proposed that a parlor game, in which the object was to decide whether an entity with which you could only communicate by written words was male or female, be adapted to give people a way to decide whether machines could think for themselves. If you could not distinguish whether the entity, with which you would communicate by typewriter, was human or not then it must be able to think. Sort of an early version of "If it walks like a duck, and quacks like a duck..."

As you can see, this obviates the need to say what is meant by "think".

Ray Allis

Date: Fri Sep 25, 1992 9:22 am PST
Subject: Re: Influence; let's change gears

[Martin Taylor 920925 11:00 (getting impatient and shooting from the lip)]
(Bill Powers 920922.0900 to Greg Williams)

>And it's too general to say that such changes, even if they do amount
>to disturbances, "influence the operation of the .. control systems."
>Influence WHAT about the operation of the control systems? All that
>can be influenced is the action used to correct any induced error, and
>that can be influenced only in the sense of making more or less of it
>necessary. The nature of the action is already determined by the
>organization of the control system, which takes into account the
>environmental link through which action affects perception.

Whoa, there. One of the consequences of the PCT approach is the idea that the same end can be achieved by many different means. If one possibility is blocked, another can be used. And often the various means are mutually exclusive in any case. I can't get to my destination by car, bicycle, and train simultaneously, though any one of them would get me there. So, A can definitely "influence the operation of the control system" of B, just by making one way of B's control relatively easier than another. Take away B's car, and you influence B to take a bike. If that isn't affecting the operation of B's control system, I don't know what is. And I think this is the sense in which Greg was using the phrase, not in the sense that A changes the linkages or functions in B's control systems.

Martin

Date: Fri Sep 25, 1992 9:32 am PST
Subject: Intell.;influence; demos; Larsonian physics

[From Bill Powers (920925.0900)]

On Sunday morning (920927), Mary and I are taking off for a camping trip into Arizona and New Mexico. We're going to look at the Very Large Array and one radio telescope in the Very Long Baseline Array (west of Socorro on the Plains of Agustin), then go sideways to another campground east of Socorro, then zoom up to Carrizoso and buy some extraordinary apples they grow there, then maybe (if we're not too bushed) down to Las Cruces for a sentimental look at the observatory I designed way back when, and home. Four or five days, proolly. When I disappear from the net, don't worry.

Avery Andrews (920925.1236) --

Thanks for Harnad's note on intelligence. Nothing for me to argue with there!

Apropos of that, how many people have read Krishnamurti? This was a wise fellow. To give you an idea (as near as I remember the facts): he was raised by the Theosophists to be the Messiah because of his extraordinary intelligence and spiritual development. When he was 17 years old, he was brought before a great international meeting to be declared the saviour of the world or something, and he stood up and say "Sorry, folks, I think there's been some misunderstanding -- thank you very much, but no thanks." He spent most of the rest of his life talking with people who wanted to ask him about things. He never wrote anything, but his conversations have been recorded in books. My copies are still packed, so no refs.

In one conversation, he was being asked about intelligence -- what makes some people more intelligent than others, how can I be intelligent, and so on. He said, these questions have no meaning. Intelligence is not a quantity. It's unique. To be human is to be an intelligent being. So relax.

I thought Penni would like that. I put this in a reply to you because you said "intelligence" and that stimulated me to respond in this way.

Greg and Pat Williams (920925) --

Good. Handsomely said. Now let's give up the idea that anything at all must have purposeful influence on organisms, and begin working out whether it does or not, without caring how the investigation turns out. Without staking out a position and defending it.

I'm starting a summary of basic principles and their application to the shaping of behavior. When it gets far enough along I'll post it for revisions, additions, corrections, and deletions. This is strictly an attempt to lay out what we can derive from HPCT as it stands, not an attempt to justify HPCT. I'm not ruling out conjecture, but am trying to eliminate anything that's obviously false as seen from in here. The objective is to arrive at a scenario for human development that a True Believer in HPCT would find consistent with the theory. As we find out what's wrong with this scenario, we may gain some insights that will improve the model. All this may not get done until after the camping trip.

The Arm is all fixed. I'm waiting for the permissions to put in the day or so that will be needed to write up the detailed changes and ship them to you. This is only because of busy-ness with the net and other things, and the fact that I need a break.

Chuck Tucker (920925) --

Thanks for the writeups on the RB and Coin experiments. Clark McPhail sent me, some time ago, copies of his students' lab reports on similar assignments, and they were fascinating. Maybe you could post exerpts from the results you get. I think this is a beautiful example of how people make sense of PCT when they're presented with a problem to solve by using it.

Maybe in our attempts to publish in the conventional literature we should focus on describing phenomena that only PCT can explain, and then explaining them. As several people have been pointing out, it doesn't seem to work when we use PCT to explain things that others think they can already explain. Maybe we should start presenting them with things they can't explain.

A thought about the coin game. An extremely difficult pattern to guess would "All the coins exactly where they are and oriented as they are." The experimenter would never get a "no error" response! It seems to me that the pattern must contain at least one dimension in which it's NOT controlled in order to give the experimenter enough hints.

Oded Maler (920925) --

I think, I hope, that your posting of the description of Larsonian Physics was intended as an admonition: DON'T LET THIS HAPPEN TO CONTROL THEORY!

The observatory where I worked in the 60s was a magnet for papers like this; for some reason, people seem to think that astronomers have a special interest in revolutionary ideas, perhaps because of misinterpreting the title of Copernicus' _De Revolutionibus_. So we got every nutty new idea in the world sent to us. For years, being a revolutionary myself, I conscientiously studied the materials and wrote back what I thought of them, because who knows? Maybe someone really had something new. But it was all sad garbage in the end, a result of a sort of cargo cult approach to science (imitate the outward

trappings while missing the point altogether). I think that I probably saw some of Larson's stuff, probably indirectly. Don't remember now.

Reading that stuff at the observatory was good for me, because it made me realize how easily I could cross the line, too, and start ranting and raving and sounding like a madman. It made me sympathize more with people in the conventional branches of psychology who came across my work and found it just as strange as I found all the new theories of electrical vortex gravitation, tachyon generators, hollow earths, space visitors in the White House, particle physics, and health radiation resonators. If you want to know how a new idea looks to others, insert it in the middle of a list like that and see what impression it gives.

What I've found in common among all the writings like this is an utter lack of self-criticism. It never occurs to the author of such a work that the insights which thrilled the soul at three o'clock in the morning might collapse the next day, given any sort of judicial review. Everything that sounds plausible gets written down and elaborated on -- "But that means and that means ... holy smokes, and that means ..." The person gets caught up in a great manic wave of optimism that simply brushes aside all mundane considerations like experimental proof or internal consistency.

Even when experiments are done, they're done in tremendous haste and with absolute certainty that they will bear out the theory. So of course if there are difficulties, as there always are, they are attributed to everything BUT the theory -- we need better bearings, we need a more powerful Van de Graf generator, we need to balance the weights more exactly, we need to build the whole thing LOTS bigger, we need more subjects, and so on. In a lot of cases, when gadgets are actually built, the inventor is tempted into just a little cheating and misrepresenting and abracadabra, because he sincerely believes that if he can just raise enough money to improve the apparatus, it really will work, so he's really not cheating or lying, not in the long run. All those investors will get their money back BILLIONS AND BILLIONS of times over. So the manic rush leads to fraud, without any evil intent at all. Just ignorance of how the search for knowledge really has to work, with a healthy dose of greed and a longing for fame.

I should think that any honest scientist trying to promote a new idea would feel a little uncomfortable reading that last paragraph.

Even now, I get guilty twinges at stuffing things like Larson's theory into the wastebasket. What if he's really on to something? What a horrible injustice, to dismiss him just because he and his supporters write like opinionated nuts and can't spell Reagan! What if people did that to my ideas?

But there it goes, plop. I can't be that open-minded any more. At least I can try not to do the same things when I get brilliant ideas at three o'clock in the morning.

Best to all, Bill P.

Date: Fri Sep 25, 1992 9:55 am PST
Subject: Who was Turing?

[Ray Allis 920925.0930]

This just arrived from another net I subscribe to. Maybe it will help you place Alan Turing in the scheme of things.

> Sender: History of Computing Issues
<SHOTHCL-SIVM.BITNET@pucc.Princeton.EDU>

> The following quote is attributed to Frankel in:
> Randell, B.,
> The Colossus in a History of Computing in the Twentieth Century,
> (Eds. N. Metropolis, J. Howlett and G-C. Rota), Academic Press 1980,
> pp 47-92.
> I am wondering if anyone can tell me who Frankel was and how valid
> the quote might be. Is there any direct quote of Von Neumann attributing
> the fundamental concept of the (stored program) computer to Turing?
> .in 5
> "Many people have acclaimed von Neumann as the 'father of the computer'
> (in a modern sense of the term) but I am sure that he would never have
> made that mistake himself. He might well be called the midwife, perhaps,
> but he firmly emphasized to me, and to others I am sure, that the
> fundamental conception is owing to Turing--insofar as not anticipated by
> Babbage, Lovelace and others. In my view von Neumann's essential role
> was in making the world aware of these fundamental concepts introduced by
> Turing and of the development work carried out in the Moore School and
> elsewhere. Certainly I am indebted to him for my introduction to these
> ideas and actions. Both von Neumann and Turing, of course, also made
> substantial contributions to the 'reduction to practice' of these
> concepts but I would not regard these as comparable in importance with
> the introduction and explication of the concept of a computer able to
> store in its memory its program of activities and of modifying that
> program in the course of these activities."

Date: Fri Sep 25, 1992 10:51 am PST
Subject: Linking; statistics

[From Bill Powers (920925.1100)] Martin Taylor (920925.1000) --

My objection to invented terms like "linking" is that it leads too easily to logical errors, particularly in verbal discourse (as opposed to mathematical). What happens is that one phenomenon is given the name linking; call it linking-A. Then another phenomenon is called linking-B. Aha, says the brain, they're both "linking." Therefore the first phenomenon must be just like the second. I gave Greg an example of a different kind of linking that was not at all like the kind he was thinking of.

This is what I've been calling the "category error." It's my main objection to verbal generalizing as a mode of theorizing. By categorizing, we strip away differences, so we can say that phenomena that are actually quite different "can be seen as" examples of the same thing -- that is, members of the same category. But this works only one way; you can't say that because A belongs to category C and B belongs to category C that A and B can be treated interchangeably from then on. Particularly when only words are involved, this

is really a form of punning. My cat is attached to mice and my computer is attached to mice. Therefore cats are like computers.

This error is committed ALL THE TIME in psychology. It's compounded by saying that people do certain similar things that can be seen as category A, under various circumstances that can all be seen as belonging to category B; therefore, circumstances in category B result in doing things in category A. Now we have categories causing categories. This is one of the main reasons that theories in psychology predict so poorly; this sort of category error generates huge numbers of counterexamples when you try to reduce the generalization to specific cases again. As long as you STAY at the category level, the generalizations continue to hold more or less true. But any attempt to apply them in real cases will succeed only by luck. That's because generalizations have no causal force in nature; they're only conclusions by an observer. By the way, they do a lot of this in physics, too, although mathematics helps to limit the damage.

When I said Pat couldn't control for her kids' health because her knowledge of her effect on their health was statistical, I didn't mean that her perception of the effects was statistical. Generally, there isn't any perception of the effects at all. You can't see the effect of feeding good food to the kids unless you sometimes feed them bad food, too, and observe the difference. The statistics I was referring to was the "findings" of nutritional science, which deal with populations, not Pat's kids. We (to take the onus off Pat) feed our kids food that we understand to be good for them, for personal theoretical reasons and because of what other people have found through nutritional studies of populations. This has to be mostly an open-loop process. We don't know if our own kids are typical, and we don't have the means of monitoring the effects that are available to a research laboratory. Mainly, though, we don't have control over their health because we don't experiment with it; we don't really know whether varying their nutrition would vary their health. If the kids look poorly, we'll look for other causes, not at what we're feeding them (unless we decide that it wasn't nutritional after all -- but then when they get well we'll stick to the new diet, not go back to the old one). Pat's not going to put her kids on a diet of Big Macs and fries for a month to vindicate her views on nutrition, and if they get sick and recover, she's not going to decide that their former diet made them sick (healthy food doesn't make people sick).

There are lots of things we have to do open loop (or think we have to) because we really don't have any direct control over them. This is what makes medicine, psychotherapy, playing the stock market, and cleaning the chimney every fall such grey areas -- did we have an effect, or was it something else? Or would the same thing have happened no matter what we did? We can't just let the world happen without trying to do SOMETHING about it, can we? So we go through the motions anyway, hoping to have an effect even if we'll never know what it was.

Ray Allis (920925.0900)-- RE: turing test

>If you could not distinguish whether the entity, with which you
>would communicate by typewriter, was human or not then it must be
>able to think. Sort of an early version of "If it walks like a

>duck, and quacks like a duck..."

>As you can see, this obviates the need to say what is meant by
>"think".

Right on. All the Turing test can really ask is whether the other entity is a human being or a program (or, I suppose, a duck).

Best, Bill P.

Date: Fri Sep 25, 1992 10:56 am PST
Subject: control/influence, Turing

[From Rick Marken (920925.1000)]

Jeff Hunter (920923 - 1) --

> This gives us two common meanings for the word "control"

> a) I can make something do what I want it to
> b) I can make something do what it doesn't want to
> (or I can make it choose the lesser of two evils)

> I think the reason Greg is concerned about the discussion
>so far is that the PCT term for control seems to follow meaning (a) for
>inanimate objects, but suddenly jumps to meaning (b) for people.

I like the idea of having just one technical meaning associated with the word control. I suggest: acting to keep a perceptual variable equal to a reference variable. Your meaning (a) captures this technical meaning when "I" is the controlling system, "something" is the perceptual variable, "make" is the actions of the control system, and "do" is keeping the perceptual variable equal to the reference variable. If the perceptual variable is also controlled by another control system (relative to a reference variable that is generally uncorrelated, over time, with the reference variable of the original system) then we have the situation described by your meaning (b). I didn't mean to use two different meanings of "control". The fact of the matter, however, is that if you try to control a perception that is controlled by another control system then you get CONFLICT. Conflict occurs BECAUSE systems control (in my technical sense). But when there is conflict, one or both of the systems involved will not be able to control successfully! Maybe this is one the problems in this whole discussion of controlling behavior. My technical definition of control assumes that control is working (perception matches reference; there is virtually no error). But it is possible to be organized as a control system with respect to some variable and to be unable to control that variable; this is what happens when there is conflict. When there is conflict (which occurs when control systems try to control the same perceptual variable relative to difference references) there are control systems involved -- that's the only way that conflict can occur. But one or both of the systems will not be controlling SUCCESSFULLY (if at least one of the systems has sufficiently high gain -- so that it would be able to control the variable successfully when there is no conflict). If one system has MUCH higher gain and output range than the other then it will successfully control the variable while the other system experiences massive error (while producing maximum

output). The latter system is definitely NOT in control -- though it IS controlling FOR a particular perception.

So maybe we need to distinguish between the PROCESS of control and the RESULT of control. My technical definition focuses on the result; keeping perception matching reference. This result does not occur when there is conflict; but the process of control DOES (and must) occur when there is conflict. The process of control refers to a high gain negative feedback relationship between action and perception. Nominally, this process results in perception matching a reference. When there is conflict, the process may or may not produce this result -- depending on the strength of the control system.

Re: The Turing Test

Jeff Hunter (920924 - 2)]

> The interrogator is clearly trying to find which variables
>in the sonnet the witness is trying to maintain against disturbance.

Yes. What you described is part of "the test". I found another example of "test-like" interrogation described in Hofstadter's "Metamagical Themes". On p 501-502 there is a discussion of how you might tell if the machine had "emotional responses". It is suggested that, when asked its reactions to certain pieces of music or literature, the machine might "steer clear of emotional topics". Hofstadter's discussion gives the impression that he has some notion that a living system might be trying to control perceptions -- like "degree to which I discuss emotional topics". But in fact, when you read the entire exchange it is clear that Hofstadter has no idea that the evasion he mentions might be a sign of a controlled variable; indeed, on p. 502 he clearly misses this point completely. He has one of his conversationalists say: "This would certainly make me suspicious. Any consistent pattern of avoiding certain issues would raise serious doubts in my mind as to whether I was dealing with a thinking being". A PCTer would immediately see "avoiding" a symptom of control. And this quote might suggest that Hofstadter might notice this too. But NO. It turns out that the reason that the speaker would be suspicious of answers that avoid emotional topics is because "I simply can't believe emotions and though can be divorced." So these Turing tester's are looking at the CONTENT of the machines answers, rather than what the answers suggest about what the machine might be trying to accomplish. Thus, if, when asked about it's reaction to a certain piece of music the machine says "I don't know that piece" this is taken as evidence that the machine has no emotions (because it is trying avoid a discussion of emotional topics) -- the "avoiding" that MIGHT be happening is considered important only in terms of it being an indication of the existence of emotions (or lack thereof) on the part of the respondent.

So I am not getting "frisson" over your example of "the test" being done as part of the Turing Test. What I am impressed by is how people can come SO CLOSE to the notion of testing for controlled variables -- and still miss the point completely. If, in your sonnet example, there was an explicit hypothesis about the variable being controlled and efforts to redefine this variable based on constant questioning aimed at disturbing the hypothesized variable, THEN I would be "frisson"-ing. But, as it sits, I would say that the

conventional Turing testers are as far from understanding the notion of controlled variables as are your basic Skinnerian operant conditioner (who is unwittingly doing a version of the test when s/he tries to figure out what constitutes a reinforcer).

Best regards Rick

Date: Fri Sep 25, 1992 12:32 pm PST
Subject: Re: Linking; statistics

[Martin Taylor 920925 14:30 --finally all caught up!!!]
(Bill Powers 920925.1100)

> We (to take the onus off Pat) feed our
>kids food that we understand to be good for them, for personal
>theoretical reasons and because of what other people have found
>through nutritional studies of populations. This has to be mostly an
>open-loop process. We don't know if our own kids are typical, and we
>don't have the means of monitoring the effects that are available to a
>research laboratory. Mainly, though, we don't have control over their
>health because we don't experiment with it; we don't really know
>whether varying their nutrition would vary their health. If the
>kids look poorly, we'll look for other causes, not at what we're
>feeding them (unless we decide that it wasn't nutritional after all --
>but then when they get well we'll stick to the new diet, not go back
>to the old one).

One doesn't reorganize when there isn't error. Sure. But for a lot of parents with unhealthy children, changes in diet often are among the changes that they try, until they find what the kid was (say) allergic to. The kid becomes healthier, and they no longer try to change the diet. The effect may have been what Skinner called "superstition"--something else caused the good result by coincidence, but nevertheless, the error has gone away, and reorganization stops. No doubt diet change was not the only action involved in the reorganization process, but it was the perception of the kid's health that was the subject of control, and the perception of the relation between the parent's action and the child's health is itself statistical. There is, to help, the population-statistical information about the probabilities that people with certain diets more often have certain symptoms than do people with other diets. That information can help the parent to choose which diet changes to try, and to that degree, perhaps the reorganization is not random. Maybe this relates to Greg's "guided reorganization?" We are getting into a murky area here, I think.

>My objection to invented terms like "linking" is that it leads too
>easily to logical errors, particularly in verbal discourse (as opposed
>to mathematical). What happens is that one phenomenon is given the
>name linking; call it linking-A. Then another phenomenon is called
>linking-B. Aha, says the brain, they're both "linking." Therefore the
>first phenomenon must be just like the second. I gave Greg an example
>of a different kind of linking that was not at all like the kind he
>was thinking of.

"I have to use words when I talk to you" at least by e-mail. I don't think our mathematical descriptions of the feedback loops are either accurate or would be intelligible if they were complete. This is especially true since the feedback loop incorporates the disturbed world.

We use "reference" and "percept" and "error" with technical meaning. Why not "linking" which seems to describe a phenomenon that is important in interpersonal interaction. I'm not sure it occurs when only one hierarchic control system interacts with the inanimate world. It's an extra phenomenon and need a word. "Linking" seems to fill the bill.

I agree that verbal puns can lead one's theorizing astray. It happens a lot. But often the verbal mode is the best we can do. I doubt you could mathematize Ed Ford's juvenile delinquents very effectively.

Say "OWA TANA SIAM" and see the six-year-old laugh.

Martin

Date: Fri Sep 25, 1992 1:21 pm PST
Subject: Re: Attn. P. Sibun; repentance (might the end be near?!?!)

penelope sibun
xerox parc
3333 coyote hill road
palo alto, ca 94304

Date: Fri Sep 25, 1992 2:40 pm PST
[From: Dennis Delprato (920925)]

TO: CSG-L RE: Postulates

Some Fundamentals of B. F. Skinner's Behaviorism
(From article with above title, American Psychologist, 1992, 47,
November, in press)

Dennis J. Delprato
Eastern Michigan University
&
Bryan D. Midgley
University of Kansas

Despite claims of scientific superiority, neither B. F. Skinner nor any adherents of his behaviorism have seen fit to explicitly identify in one place its fundamental postulates. Indicative of their ties to positivism, several votaries of Skinner's brand of psychology served as referees of the paper cited above, and all rejected the idea that Skinner had any postulates or assumptions at all! Thus, the paper was only accepted for publication after the offending term, postulates, and, subsequently, the term, assumptions, were struck from its title.

Why would one who has rejected Skinner's psychology complete a laborious undertaking that entailed scrutiny of Skinner's corpus with the aim of

extracting and, to some extent, systematizing his approach to psychology? The most basic answer is because I am of the opinion that it is seriously flawed in a most fundamental way--its postulates (not necessarily all). Yet extant attempts to bury it are unconvincing. For example, perhaps the most-cited attack, Chomsky's, does not even stick with confronting Skinner's position, attributing Hull's drive reduction theory to Skinner's views, for example. I do not deny that Skinner's psychology is a step forward. Indeed, frequently those rejecting it move us backward as far as what it means to take a scientific (i.e., naturalistic) approach to psychological events. My reasoning is that in laying out Skinner's bedrock assumptions and generalizations in the most faithful way possible I have provided a legitimate target for scientific and scholarly debate.

My idea for a fair, indeed very sympathetic, presentation of Skinner's science came after years of hearing that it represented the hope of the future and that those of us who were putting forth field and system views were offering nothing new. My decision to develop an organized presentation of Skinner's views in such a way that it met the approval of prominent Skinnerians derived from what I saw as a need to get in print an authoritative version of his most fundamental assumptions and generalizations so critics had *something to address.*

Note that the motivation behind my systematization of Skinner's psychology is in no way included in the paper cited above. The paper was first submitted to the American Psychologist long before someone decided to put out a special issue devoted to Skinner. However, while it was in the review process, Skinner died and plans for the special issue evolved. What happened was that the submission was then placed into the hands of the special issue editor who gave it further scrutiny. This all greatly delayed editorial action but turned out to give it more legitimacy as a fair representation of Skinner's basic views.

Below I have simply listed each of 11 points (note conflicting position on reductionism--VIII. A. & VIII. B.). The points (assumptions/generalizations) we derived are "data based" in that each is supported by at least two quotations from Skinner's professional writing (presented under each in the paper as written). Actually, all were supported by considerably more than two quotations; space considerations limited presentation of supporting data. In the actual paper, the quotations are followed by clarifying text with the major aim being to address them from Skinner's point of view. In several cases, we brought up different sub-issues (such as position on theory, "private events," and fundamental datum) under particular points.

Why submit these to CSG-L?

1. Skinner's psychology, it seems, more than any alternative to PCT comes up. To me, this is appropriate. Despite the potential for feedback control in operant theory, Skinner's steadfast adherence to a sophisticated brand of environmental determinism makes his views a challenging foil for PCT. It may be useful to have Skinner's fundamental assumptions handy.

2. From what I have observed, those commenting about Skinner's psychology here *are* informed. In fact, along with a few who follow the interbehavioral literature (such as Noel Smith, Bill Verplanck, and Linda (Parrott) Hayes), commentators/critics here are the most informed I know of. Thus, I want to

emphasize that I am *not* offering this posting here as a corrective to any of what I identify as misrepresentations of Skinner's views.

3. For different reasons, commentators on CSG-L address behaviorism. At times, there is a tendency to equate Skinner's brand of behaviorism with generic or historical behaviorism. One of the most common mistakes critics make is failing to distinguish between Skinner's behaviorism and others of the *numerous* versions of behaviorism we have had, beginning around the end of the 19th century. The unadorned listing of Skinner's fundamental position below may not communicate this, but his was an atypical form of behaviorism. As Greg Williams previously noted, Skinner was no simpleton. He was progressive, in contrast, e.g., to so-called cognitive-behavioral types who often would have us back at Cartesian interactionism. At different times, CSG-L commentators have cited Skinner as coming close to a feedback-control approach. In my view, he could not get beyond a mere glimmer of a truly radically different conception of behavior because his assumptive base remained unexamined, hence unaltered, with the result of making it impossible to move beyond the operant as he viewed it.

4. Perhaps most importantly, I would like to encourage a thoughtful effort to lay out the fundamental postulates of PCT/HCPT. I know that, who else, W. T. Powers, has come close to doing this. More than anything, the recent inter-changes between Greg Williams and especially Bill over interpersonal control, influence, manipulation, ... led me to conclude that it might be useful for their purpose, as well as in general, if the fundamental postulates were to be set forth and frequently reiterated. I also presume they would not be set in stone. [After I prepared all this, a posting from Bill today (920925) suggests he may be working on a presentation of the fundamental position. Grand idea.]

So, what follows are some fundamentals of Skinner's approach to matters psychological. I realize that they may not be particularly interesting or useful without accompanying text, but I'll bet they still function as a disturbance. I hope the adjusting behavior (behavior-1) will not be pressing SELECT followed by pressing DELETE.

I. Purpose of Science:

The primary purpose of science is prediction and control.

II. Methodology:

The methodology is functional analysis which relates environmental independent variables to behavioral dependent variables.

III. Determinism:

Behavior is determined; it is lawful.

IV. Locus of Behavioral Control:

The causes of behavior are localized in the environment.

V. Consequential Causality:

Selection by consequences is the primary causal mode by which environment determines outcomes in living systems.

VI. Materialism:

Dualism is false. The only world is a physical world.

VII. Behavior as Subject Matter:

The subject matter of psychological science is behavior and behavior only.

VIII. A. Reeducations:

The subject matter of psychology is reducible (at least to biology).

VIII. B. Nonreductionism:

Behavior cannot be completely explained in terms of biology or any other "lower-level" discipline.

IX. Organism as the Locus of Biological Change:

It is the organism that changes through evolutionary and environmental histories, and the changes are biological.

X. Classification of Behavior into Respondent and Operant:

There are two major classes of behavior, or more completely, functional relations: respondent and operant.

XI. Stimulus Control of Operant Behavior:

Operant behavior can be brought under the control of antecedent stimuli, and description of operant behavior usually requires three elementary terms and their functional interrelationships.

XII. On The Generality of Behavioral Principles:

The full complexities of human activity--including language, thinking, consciousness, and science--are behavior to which all the above apply.

Date: Fri Sep 25, 1992 7:09 pm PST

Subject: Unifying theories of psychology

[from Gary Cziko 920926.0235 GMT]

I have had some disk problems over the last few days and so have not been able to keep up with CSGnet. If anyone had directed queries to me in CSGnet traffic, I should be able to get to it in the next couple of days.

Rick Marken:

Here's the kind of argument I'd like you to use in your Blind Men paper. These quotes come from chapters in a book edited by P. C. W. Davies and J. Brown (Eds.) (1988), Superstrings: A theory of everything? Cambridge University Press.

"The history of physics is the history of successive stages of unification. . . . Maxwell unified the laws of electricity and magnetism, and in addition established a link between electromagnetic field theory and optics by showing that light consists of electromagnetic waves. Einstein found a connection between space and time, and energy and mass, and then went on to link spacetime to gravitation." (Davies & Brown, 1988, p. 8)

"[one] might remark that history shows us that reconciling inconsistent physical theories is a very good way of making fundamental progress. If we

look at some of the advances in the twentieth century we see that Einstein's theory of special relativity came from a wish to reconcile two outstanding theories of the day, namely, Maxwell's theory of electricity and Newtonian mechanics. Einstein's theory of general relativity likewise came from an effort to reconcile his own special relativity with Newtonian gravity. Finally, quantum field theory came from an effort to reconcile nonrelativistic quantum mechanics with special relativity. So many of the most far reaching advances in the twentieth century have come about because previous theories weren't compatible with one another. History teaches us that reconciling incompatibilities between theories is a good way to make really fundamental progress." (Witten, 1988, p. 97)

So you see, Rick, you're in good company in your attempt to reconcile S-R, reinforcement, and cognitive theory by showing them to be special cases in PCT. This is why I think that this is such an important paper.

By the way, I got these quotes from a Lev Goldfarb's critique of late Allen Newell's Unified theories of cognition which appeared in the September Brain and Behavioral Sciences. It seems to me that PCT has it all over Newell's SOAR approach to unified psychology, but then again I just can't make much sense of SOAR (I think it would have made more sense to me if I had read about it before discovering PCT).--Gary

Date: Fri Sep 25, 1992 7:18 pm PST
Subject: Chjoosing actions; "make"; controlling symptoms; nice, dennis

[From Bill Powers (920925.2000)]

Martin Taylor (920925.1100) --

>Whoa, there. One of the consequences of the PCT approach is the
>idea that the same end can be achieved by many different means. If
>one possibility is blocked, another can be used. And often the
>various means are mutually exclusive in any case.

This is a "levels" problem, which I was avoiding during the argument to keep from getting more complexities into the act.

A control system that acts by choosing one action from a selection of others can't also choose the amount and direction of that action that ensues from disturbances. That takes two levels (more than one, anyway). Achieving the same end by variable means (which is required by disturbances of various kinds) is the result of a control system's reacting to error by altering its output. If the system is an action-choosing system, then the disturbance would amount to a change in the environment of a rather complex sort, that requires switching from one KIND of action to another KIND (through sending reference signals to THIS system instead of THAT at a lower level). In order for this level of control to work well enough to achieve real control, this switching must be quick and smooth: automatic. If it isn't -- if there are false starts and wrong choices -- then we would probably guess that reorganization is going on.

My point is that even when we speak of switching KINDS of control, the situation is quite analogous to what happens at a lower level where the

dimension of change of output is strictly HOW MUCH of the perceptual variable is to be maintained. I'm using this HOW MUCH idea in a pretty metaphorical way, because I just want to mean a shift along whatever dimension of change is appropriate at the level under discussion. Obviously, at higher levels the dimension along which change of output takes place isn't the simple sort of continuous thing we see at the level of intensities. But it can still be ordered to create a dimension of stepwise change.

Taking away one of the possible choices is directly analogous to disturbing a lower-level system by tying a rope to one arm. It alters the feedback function connecting the output (which selects choices) to the input. The analogy would be to altering the schedule of reinforcements, not to directly adding or subtracting reinforcements. This kind of parametric disturbance shows up only when the control system tries to act; then it discovers that the effect on the controlled variable is different or missing. In a choosing-type system, the result of the increased error would be a further change in the choice of action -- picking the next one in order of ease, handiness, effectiveness, or whatever the ordering scheme is. If the bicycle is missing take the car; if the car is missing take the bus; if the bus doesn't come, take a taxi.

If there IS no natural "next" choice, then the error will not be corrected. The perception in question won't be controlled, and the next higher level will be disturbed. Instead of an automatic next choice of a category of action, we might see the person reasoning out some alternative, or dropping the goal that called for making the choice (if both the bicycle and the car are missing, call in sick). Failing that, the system might be thrown into a state of reorganization.

>If that isn't affecting the operation of B's control system, I don't know what
>is. And I think this is the sense in which Greg was using the phrase, not in
>the sense that A changes the linkages or functions in B's control systems.

I have never argued that control systems can't be affected, even on purpose. But so far I don't see what the term "linkage" describes that isn't already described in any basic analysis of relationships that uses all the resources of HPCT.

Jeff Hunter (920923 - 1) --

>This gives us two common meanings for the word "control"

>a) I can make something do what I want it to
>b) I can make something do what it doesn't want to
> (or I can make it choose the lesser of two evils)

I agree with Rick Marken on this. This is really just control. The word that's changing its meaning between these two definitions is "make." If you ask me for the salt, I can "make" you reach out a foot (that is, 12 inches) to receive it if I hold it one foot from where your hand is. If you don't want to let go of the salt, I can still "make" you give it to me, by prying your hand open using a force too great for you to resist. In both cases, I'm controlling; it's just that in the second case I have to use a little more effort to overcome environmental resistances.

Of course in the second case, a conflict is created, while in the first case there is none (as Rick said).

Martin Taylor (920925.1100) --

>But for a lot of parents with unhealthy children, changes in diet
>often are among the changes that they try, until they find what the
>kid was (say) allergic to. The kid becomes healthier, and they no
>longer try to change the diet.

Sure. In this case there's a symptom that can be perceived, so to control that perception one varies the diet to try to make it go away. If you have a stock of remedies at hand, it's an ordinary control process. If you don't know what to do, you just try different things until you hit on one that makes the symptom go away -- reorganization.

Feeding healthy food as a matter of principle isn't quite the same thing. If the kids develop a symptom, you can't feed them still healthier food -- if you could you'd be doing that all along. When you don't actually have any control, all you can do is imagine that you have control, or give up.

I know people using herbal remedies and megavitamins. I asked recently how one of them was doing. "Wonderful," the answer came back. "He's vomiting and coughing up a lot of black stuff, really getting rid of the poisons." I said "That's nice." He's a sensible fellow; I'm pretty sure he will soon tell his helpers that he feels a whole lot better now and they can stop giving him that stuff.

>We use "reference" and "percept" and "error" with technical
>meaning. Why not "linking" which seems to describe a phenomenon
>that is important in interpersonal interaction.

Because I don't think it describes anything we can't already handle, and because I like to use "link" in other ways, as in "the feedback function links the output to the controlled variable" and "properties of the environment link the disturbance to the controlled variable." I don't want to have to keep explaining that I mean one kind of linkage, not some other restricted and made-up sense of the term. I have to do enough of that already. I think that creating a lot of new terms just makes for jargon, and conceals simple relationships by lumping special cases into categories of their own. Maybe later.

Dennis Delprato (920925) --
RE: Analysis of Skinner's World.

That's a really grown-up job. It's the sort of analysis we need done on PCT. Why not give it a whack yourself? Maybe you'd be a good person to edit the forthcoming issue of Closed Loop on influence and control, with editorial and historical comments!

Mary and I both think that we aren't done with this thread yet, and that the Closed Loop on this subject shouldn't be put out prematurely. I guess Greg reached the same conclusion. We seem to be headed slowly in some direction; let's wait until an obvious pausing place appears.

Congratulations on the "in press." Somebody around here knows how to get published.

Best to all, Bill P.

Date: Sat Sep 26, 1992 12:39 pm PST
Subject: Fundamentals of HPCT, Version 0

[From Bill Powers (920926.2200)]

RE: Organismism. This is my first attempt to lay out basic principles of HPCT in a brief and systematic way. It undoubtedly needs thorough examination for non-sequiturs, unwarranted (by HPCT) assumptions, and gratuitous conclusions. I've tried to avoid saying things that are obviously false in terms of the basic theory. I haven't defined everything as if for a newcomer. While I haven't restricted my imagination, I've tried to remain noncontroversial in my assumptions. Have at it.

A. The basic architecture of human organization, according to HPCT.

A-1. The behavioral outputs of thge human organism consist of forces applied by effectors to the environment (which includes the physics of the body). All actions of the CNS on the environment take place through these effectors.

A-2. The environment can affect the CNS through the sensory nerves.

A-3. The environment can affect the human organism physically and chemically, through the food, water, and air it consumes and through direct impingement of forces and energies on its body.

A-4. The basic principle of organization of the organism is that of a control system. This organization is repeated at all levels from DNA to the highest levels of the CNS.

A-5. In addition to being composed of control systems, the organism can control through the acquisition of new control systems. The process of acquisition is called, broadly, reorganization. The adult CNS consists primarily of acquired control systems, brought into being through reorganization.

A-6. Reorganization of the behavioral systems (the neuromotor systems) is driven by a control system, perhaps of complex nature, that is concerned only with maintaining the physical organism in a viable state. This is called the reorganizing system.

A-7. This "viable state" is defined in terms of a set of critical variables and a corresponding set of reference signals. The critical variables are a sampling of biochemical (and possibly other) states in the organism. The reference signals are genetically set -- quite possibly, in many instances, simply as fixed thresholds of response in chemical or neural sensory mechanisms. Such an arrangement could combine sensing and comparing in a single neuron.

A-8. Changes of behavioral organization, including the acquisition of the first behavioral control systems and all subsequent

systems, are driven by the difference between the values of the critical variables and their respective reference signals.

A-9. A minimal model of reorganization adds the absolute values of all these difference-signals together, the sum determining the frequency with which random reorganizations are applied to the CNS.

A-10. The process of building new control systems works entirely without a "teacher." The sole criterion for starting and stopping reorganizations is the state of critical variables relative to their reference levels. The sensors that detect critical variables and the built-in reference signals and comparison process are inherited, designed and built by evolutionary processes expressed through DNA. The state of the critical variables in which reorganization will cease is thus defined independently of the present-time environment.

A-11. There is no direct connection between behavioral control systems and the reorganizing system. The acquired control systems become organized to vary their actions on the environment in a way that keeps perceptual signals matching CNS reference signals. For each controlled perceptual signal, there is a state of the environment corresponding to the state of the signal. Reorganization will cease only when this variable state of the environment is in a state that keeps the critical variables at their respective reference levels or prevents disturbing them.

A-12. This method of acquiring behavioral organization requires no a priori knowledge of physical laws, entities, forces, relationships, and so on in the environment. It makes no assumptions about the environment other than the minimal ones concerning continuity and consistency required to allow random reorganization to work.

A-13. It is probable, however, that the CNS is preorganized to make the building of control systems efficient. There are strong suggestions that the raw nervous system has a modular organization. Perceptual computations occur in sensory nuclei and specialized layers of the brain; output computations occur in motor nuclei and other specialized layers. Comparison functions appear to occur within restricted physical regions. There are, furthermore, layers of organization in which different types of neural components exist, presumably adapted to make possible certain type of computations. This arrangement into layers may reflect evolutionary "knowledge" of the environment, in that certain types of functions are important in controlling all real environments regardless of their details (and independent of our ways of describing physical laws).

B. The growth of behavioral organization.

B-1. The first aspect of behavior that must become organized after birth is concerned with the use of muscles to control the perceptions critical to survival. The neonate is helpless; it does not have enough organization to take care of the needs represented by its critical variables. It has just enough organization so that when a critical error occurs, the reorganizing system can activate the motor equipment to produce movement and noise in a global and unorganized way.

B-2. If caretakers did not interpret these output effects as a signal to provide food or other aid, the baby would simply die of starvation, dehydration, cold, heat, injury, illness, or suffocation. Its critical variables would depart too far from their built-in reference levels to allow life to continue.

B-3. This initial random-output process quickly becomes organized. Sensory signals indicating various states of the environment (and body) become connected to motor outputs. The sense of the connection and the level of sensory signal at which an output will cease are arrived at through random reorganization.

B-4. Further reorganization will be prevented when a sensory input indicates a condition before it has effects on critical variables, and when this sensed condition leads to action that removes the condition before reorganization can begin. If the initially effective action amounts to a signal to caretakers, the caretakers will correct the condition, completing the control loop. This has the side-effect of preventing critical errors, so reorganization does not further change the neuromuscular connection.

B-5. Thus the CNS acquires the first control systems for controlling the sensory representations of body states related to the critical variables. The particular motor actions that are used to control these sensory signals depend on the caretakers -- what they regard as a signal of distress, and how they interpret the signal in terms of the remedial action they take.

B-6. These first control systems are crude and unreliable. They depend on the continuous presence of caretakers, and on their ability to interpret basically meaningless signals as indicating specific problems. While major problems having to do with food and comfort are usually handled properly, the critical variables inside the baby are connected only loosely to events in the external world, and represent far more detailed aspects of the organism than can be controlled through a few crude gestures.

B-7. The result is that while major critical errors are prevented, there must be a continual state of critical error that produces continual reorganization. For example, even though the baby can signal hunger sensations before they indicate a severe deviation from a reference state, and even though the mother responds by picking up the baby and nursing it, the baby still has to find the nipple and suck. The baby may be put down in an uncomfortable position; it may have gas bubbles that are painful; there may be a light shining too brightly in its eyes; it may find its limbs cramped or immobilized, or its skin itchy or inflamed. The mother can't be present every instant, and the mother does not know how to interpret the random-looking struggles of the baby in any but the broadest terms.

B-8. The baby therefore is continually required to learn to control aspects of its experiences known only to it, as a means of preventing critical errors that only its own reorganizing system can detect. It must continue to refine its signals to the caretaker to make the consequences more specific to the need, but it must also learn to control other parts of the sensory world by moving its limbs, its body, its eyes in ways that lessen the critical errors without the help of the caretaker.

B-9. Through processes like these, the baby learns how to communicate its needs to the caretakers in terms the caretakers are prepared to (or learn to) recognize, and at the same time how to affect other sensed aspects of its world that directly or indirectly bear on the states of critical variables inside it.

B-10. Thus the foundations are laid for a process of ever-increasing control of perceptions of the outside world and the body, with the underlying reorganizing system always providing the ultimate judgment of whether the resulting organization achieves results that are good for the organism or bad for it. That judgment rests finally on evolutionary experience.

C: Relationships of the behavioral systems to the environment.

C-1. The outside world can affect the organism's body directly, or affect it through its sensory organs.

C-2. Direct effects on the body that are not represented as sensory information (interoceptive or exteroceptive) are not controlled by the acquired behavioral control systems. Their only possible influence on behavior is through direct effects on critical variables that can result in reorganization. Other effects on the body that do not affect critical variables are either resisted by the biochemical systems or alter the state of the organism without resistance and without affecting behavior.

C-3. Environmental effects can influence the body directly at the same time that they influence sensory information. Ingested food, for example, affects the biochemical processes in the body, and simultaneously affects taste, smell, tactile, and visual sensory organs. What makes a substance into "food" is not its sensory effects, but its effects on critical variables. Those effects have first priority; as a consequence, the sensory aspects of the substance become controlled (through reorganization of the behavioral systems) in favor of ingesting the substances or in favor of doing something else with them -- spitting them out or avoiding them. Sensory aspects of the food do not determine which result will occur; only effects on the critical variables can make that difference.

C-4. The basic capacity to sense the environment and body is built-in in the form of sensory nerve-endings. Higher levels of perception, derived from the basic level, are not built in but are acquired through reorganization. All higher perceptions are acquired as a component of a control system. The only purpose for acquiring any control system is to maintain critical variables near their reference levels. Therefore all higher perceptions are organized in order to control them, and thus, indirectly, to control critical variables.

C-5. The sensory effects of the environment at levels higher than the first thus always affect perceptions that are, or sometimes are, under control.

C-6. If the effect of the environment is to alter a controlled perception, the amount of alteration is normally kept small by the actions of the control system. The sensitivity to error and the capacity of the outputs of a control system are such that any remaining effect of a normal external disturbance will not cause critical error. Reorganization simply continues until that is the case. Abnormal disturbances cause reorganization to begin again. If the organism survives, these disturbances then become "normal," in that the

control systems change until the new disturbances can be sufficiently resisted. Perceptions that are controlled well enough to prevent critical error are said to be "unaffected" by disturbances.

C-7. Actions that control perceptions have side-effects on the body and on other perceptions. Side-effects on the body are unimportant unless they cause critical errors, in which case reorganization occurs and persists until those side-effects are removed.

C-8. Side-effects of control actions on perceptions not directly related to the control action are important only if they disturb other controlled perceptions. If the disturbances are successfully countered by other control systems, without inducing further disturbing side-effects that can't be controlled, there is no important effect on the behaving system as a whole.

C-9. Perceptions that are not currently under control can be affected by the environment without any action that resists the disturbance. These perceptions then simply change. If that change does not disturb any other control systems to which those perceptions are inputs, there is no important effect on the behaving systems. If the external circumstance producing the disturbance does not also impinge on the body so as to cause critical errors, there is no important effect on the organism as a whole. If such errors are caused, reorganization will commence and will not end until the perception being affected (as well as the circumstance doing the affecting) is brought under active control. So a perception that was not controlled under those circumstances becomes controlled under the same circumstances. The result is that uncontrolled perceptions remain uncontrolled only as long as disturbances of them make no difference to the organism.

C-10. The external world can disturb perceptions in two ways: by disturbing variables, and by changing parameters.

C-11. Disturbing a variable perceived by a control system means applying the same kind of influence to the variable that the control system's own outputs have. The aspect of the environment sensed by the control system determines whether an affect in its environment is or is not a disturbance. This also determines the kind of output that the control system will have; the output must have an effect that can change the sensed aspect of the environment. So disturbances are always expressed in units commensurable with the units of output.

C-12. Parametric disturbances cause a change in the form of the environmental function through which a control system's outputs affect its own perceptions. Such changes have an effect on behavior only in the dimension of the original effect of the output on the perception. If the parametric change results in the output having new side-effects, those new side-effects are irrelevant unless they result in an autogenic disturbance of some other control system in the organism -- if the output begins to have effects on other perceptions that it did not have before -- or if they affect the body in a way that induces reorganization. The only relevant effect on the original control system that a parametric disturbance can have (without involving reorganization) is to change the amount of effect a given output has on the perception -- changing the loop gain of the control system. The greatest effect a parametric disturbance can have is to reduce the effect of output on input to zero, or to increase the amount of that effect to the point where the control system

becomes unstable. Between those extremes, parametric disturbances change only the amount of effect the system's output has on its input.

C-13. Parametric disturbances change a system's loop gain. They can also increase or reduce the range of effective control and affect the stability of control. They do not directly alter perceptions, as disturbances of variables tend to do. Their effects are visible only during active control behavior.

C-14. These special effects of parametric disturbances are otherwise like those of disturbances of variables, with similar consequences when they are not successfully resisted by the organism.

D. Relationships among organisms

D-1. The control actions of one organism can affect other organisms. These can be direct effects on the bodies of other organisms, or they can be disturbances of controlled and uncontrolled variables or parametric disturbances of other control systems.

D-2. All considerations under C above hold when the effects of the environment on an organism are produced by the actions of another organism.

D-3. There are two added considerations when the source of an effect is the action of an independent organism. The stability of control in both systems can be affected because of feedback effects; this is mutual parametric disturbance. Also conflict can develop; this is mutual disturbance of variables. The two cases go together; unless the two systems are tightly coupled to each other, neither conflict nor instability will result.

D-4. Tight coupling between independent organisms exists when either the input or the output variables are so nearly identical that one system cannot act on its own perceptions without having an equal effect on the other's perceptions of the same kind.

D-5. When the input variables are tightly coupled, one system's input variable must necessarily be the same as the other system's input variable. But the other system's action must be tightly coupled to its own perception if control is to exist, and the two systems will not necessarily control the input variable relative to the same reference level. Therefore tight coupling of input variables leads to conflict.

D-6. When the output variables are tightly coupled, one system's action must necessarily be the same as the other system's action. As the same action can't in general affect both inputs in the ways each system requires, again conflict occurs.

D-8. The effects of conflict depend on the loop gains in the two control systems.

D-9. If both loop gains are low, the systems will produce outputs within their normal range of output, opposing each other. Both will experience large error, but with low loop gain will do little to correct it. This implies that the errors are not important -- that is, that critical error does not depend much on keeping those perceptual errors small. If such errors are important in that sense, both systems will begin to reorganize and increase their loop gains.

D-10. If one loop gain is high and the other low, the system with the high loop gain will keep its error small while the other system experiences large error -- in the same range it would experience if had no control. The system with lower loop gain will begin to reorganize if that amount of error results in critical error.

D-11. If both loop gains are high, both systems will produce maximum output and experience maximum error. Both systems will lose control of their perceptions. These perceptions became controlled as a way of preventing critical error. Loss of control therefore implies that reorganization will start.

D-12. Conflict, if important in the sense of affecting the critical variables in the systems that have lost control, will lead to reorganizations that escalate the conflict, bringing both loop gains into the high region and increasing the outputs to their maximum. This will continue until one or both organisms dies, or until one organism reorganizes in a way that eliminates the conflict altogether.

D-13. Therefore the normal steady-state interaction between organisms will be one that involves no conflict.

D-14. Interactions without conflict require that independent organisms in close contact control perceptual variables that are linearly independent of each other. This is the solution to conflict that will occur, through reorganization, if both organisms survive. It is the only solution, other than breaking contact, that can persist over time.

D-15. When there are no conflicts, each organism controls variables that are independent of the other's controlled variables. The actions of each organism can amount to disturbances of variables or parametric disturbance of the other's control systems. These disturbances, however, remain within the range where each system can continue to control all its perceptions by varying its own output actions. Each organism continues to maintain its own critical variables near their reference levels, independently of the other.

D-16. The shared environment will then come to a state in which it simultaneously satisfies all the reference conditions in both organisms. This can involve a close intertwining of physical relationships, because for the behaving systems, "environment" is all that is not CNS or reorganizing system. It includes, therefore, the bodies of both systems, not just the world between the systems.

Best, Bill P.

Date: Sat Sep 26, 1992 3:28 pm PST
Subject: Behavior of Perception

[From Dag Forssell (920926)]

The first day of three went well. The customer is apparently enthused. Now, I am preparing the last details for the third, technical day.

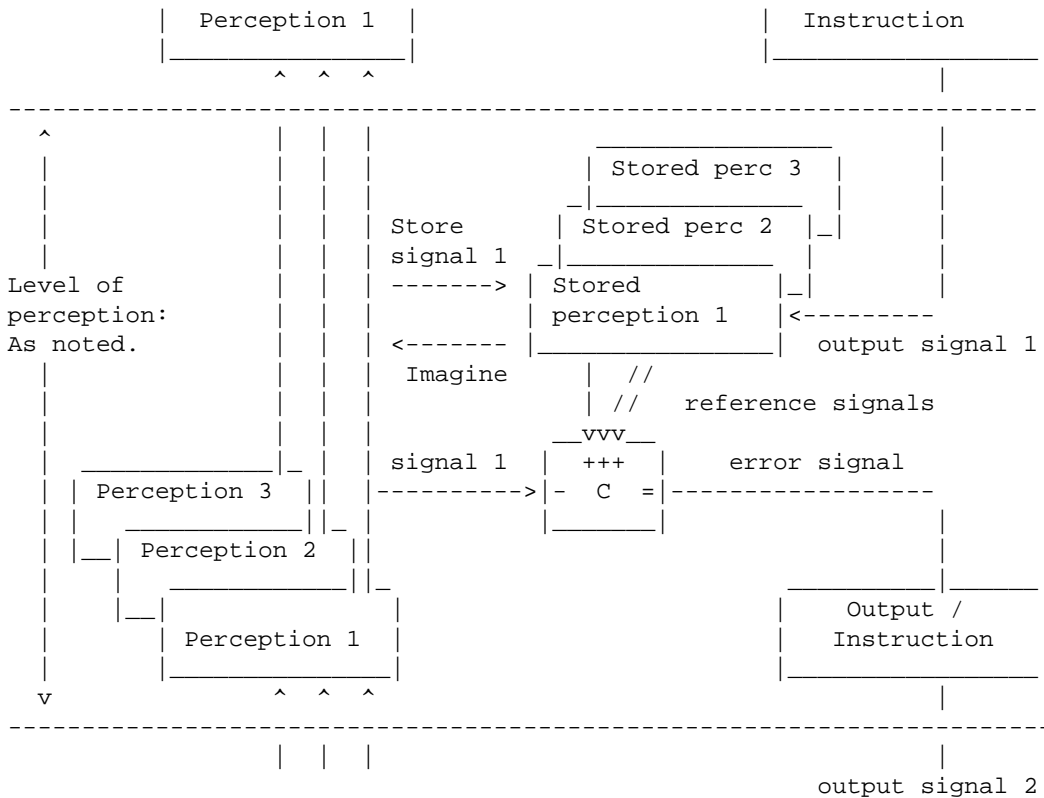
Behavior of Perception is the most comprehensive summary of HPCT and a challenge both to portray and understand. It shows the integration of perceptions and reference signals. You can see how "Perceptions" "Behave".

I would very much like some help with suggestions on how to word and illustrate this. Some of you may remember my presentation in Durango on this. The time has come for me to produce (version one).

I think that a version focusing on vision would be of interest. Would you like to help, Wayne? Here I will attempt to portray Kent's example in his paper, relating to the cruise control. (As I do, I find I want to take a stab at vision. This is much more detailed than Kent's narrative).

A year ago, the sequence of the levels was questioned. In this muscle/ motor example, I will show Intensity, Sensation, Transition and Configuration. I believe that is appropriate. In the vision example, (see Bill's post on levels of perception (920324.0300)), the regular sequence Intensity, Sensation, Configuration and Transition makes more sense. Is there any problem suggesting that there have evolved different level functions for different purposes? Different numbers of levels? I would think not.

The following diagram is my basic module:



I: Motor Control example.

LEVEL OF PERCEPTION: 11; Systems Concept (I MOTOR CONTROL)
 Output signal 1: Not applicable.

Stored perception 1: I am a successful entrepreneur
Stored perception 2: Customers benefit from my product
Stored perception 3: My work requires frequent travel
Error signal: N/A
Perception 1: N/A
Perception 2: N/A
Perception 3: N/A
Signal 1: N/A
Output: Provide memory address?
Output signal 2: Memory address?

LEVEL OF PERCEPTION: 10 Principle (I MOTOR CONTROL)
Output signal 1: See output signal 2 above.
Stored perception 1: Face to face meetings create orders
Stored perception 2: I have informed this customer about product
Stored perception 3:
Error signal: Go to San Francisco
Perception 1: A customer wants to meet me in San Francisco
Perception 2: I am in Los Angeles
Perception 3: This can be a large order
Signal 1: See perception 1.
Output: Consider ways to get to S.F.
Output signal 2: Memory address

LEVEL OF PERCEPTION: 9; Program (I MOTOR CONTROL)
Output signal 1: See output signal 2 above.
Stored perception 1: Get to S.F. by train/bus/walk: 12 hours
Stored perception 2: Get to S.F. by car. Walk from parking lot: 6 hours
Stored perception 3: Get to S.F. by Air / rental car / walk: 3 hours
Error signal: Not yet in San Francisco
Perception 1: Time is 12 noon
Perception 2: Customer wants to meet 4 pm
Perception 3: A plane is available
Signal 1: ?.
Output: Decide to Fly / drive / walk
Output signal 2: Memory address

LEVEL OF PERCEPTION: 8; Sequence (I MOTOR CONTROL)
Output signal 1: See output signal 2 above.
Stored perception 1: Driving
Stored perception 2: Flying
Stored perception 3: Walking
Error signal: Driving on schedule?
Perception 1: Driving in progress
Perception 2: Flying completed
Perception 3: Walking ahead
Signal 1: See perception 1.
Output: Convert error to memory address
Output signal 2: Memory address

LEVEL OF PERCEPTION: 7; Category (I MOTOR CONTROL)

Output signal 1: See output signal 2 above.
Stored perception 1: Passing
Stored perception 2: Being stuck behind car
Stored perception 3: Open road ahead
Error signal: To pass / pass completed
Perception 1: Passing
Perception 2: Road clear
Perception 3: Safe again See: II VISION 7;
Signal 1: See perception 1.
Output: Convert error to memory address
Output signal 2: Memory address

LEVEL OF PERCEPTION: 6; Relationship (I MOTOR CONTROL)

Output signal 1: See output signal 2 above.
Stored perception 1: Passing car dynamics
Stored perception 2: Sounds / engine
Stored perception 3: Margins at this speed
Error signal: Passing progress
Perception 1: Pressing car forward
Perception 2: Passing car on right See: II VISION, 6; relationship
Perception 3: Clearances on all sides
Signal 1: See perception 1.
Output: Convert error to memory address
Output signal 2: Memory address

LEVEL OF PERCEPTION: 5; Event (I MOTOR CONTROL)

Output signal 1: See output signal 2 above.
Stored perception 1: What driving posture feels like
Stored perception 2: What accelerating feels like
Stored perception 3: What looking alertly feels like
Error signal: Posture error
Perception 1: Pattern of leg posture
Perception 2: Pattern of body posture
Perception 3: Pattern of head posture
Signal 1: See perception 1.
Output: Convert error to memory address
Output signal 2: Memory address

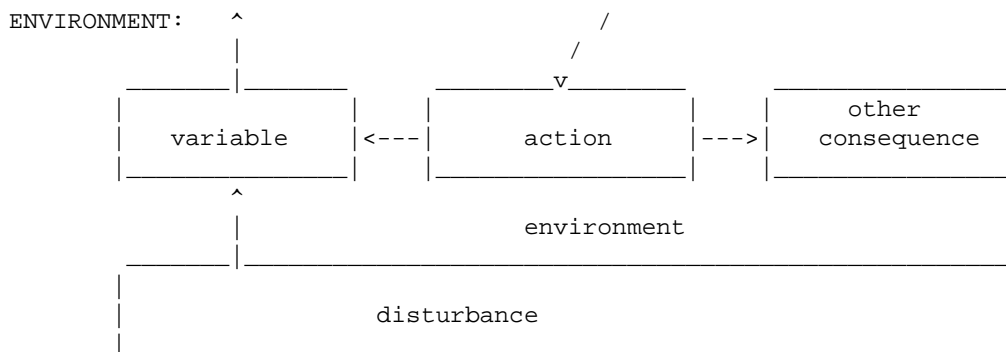
LEVEL OF PERCEPTION: 4; Configuration Note: Motor control only.

Output signal 1: See output signal 2 above.
Stored perception 1: Position
Stored perception 2: Posture
Stored perception 3:
Error signal: Position error
Perception 1: Position and orientation of this body part
Perception 2: Position and orientation of other body part
Perception 3: Position and orientation of other body part
Signal 1: See perception 1.
Output: Conversion / multiplier
Output signal 2: Velocity reference

LEVEL OF PERCEPTION: 3; Transition Note: Motor control only.
 Output signal 1: See output signal 2 above.
 Stored perception 1: N/A
 Stored perception 2: N/A
 Stored perception 3: N/A
 Error signal: Velocity error
 Perception 1: Velocity and direction in this body part
 Perception 2: Velocity and direction in adjacent part
 Perception 3: Squeeze on sole
 Signal 1: See perception 1.
 Output: Conversion / multiplier
 Output signal 2: Acceleration reference

LEVEL OF PERCEPTION: 2; Sensation (I MOTOR CONTROL)
 Output signal 1: See output signal 2 above.
 Stored perception 1: N/A
 Stored perception 2: N/A
 Stored perception 3: N/A
 Error signal: Acceleration error
 Perception 1: Sense of force / acceleration from small muscle group
 Perception 2: Signals from opposing muscles, working on same joint
 Perception 3: Pressure on foot sole
 Signal 1: See perception 1.
 Output: Conversion / multiplier
 Output signal 2: Tension reference signal

LEVEL OF PERCEPTION: 1; Intensity (I MOTOR CONTROL)
 Output signal 1: See output signal 2 above.
 Stored perception 1: N/A
 Stored perception 2: N/A
 Stored perception 3: N/A
 Error signal: Tension error
 Perception 1: Muscle tension receptor: Tension
 Perception 2: Muscle stretch receptor: Stretch
 Perception 3: Adjacent muscle receptors: Tension & Stretch
 Signal 1: Tension
 Output: Conversion / multiplier
 Output signal 2: Signal to muscle fiber: Contract.



Variable: Muscle tension
Action: Contract muscle
Other consequence: Buildup of waste products, fatigue
Environment: Blood flow, tissues
Disturbance: Inertia, Foot pedal

II: Vision example. (I top out at level 7).

LEVEL OF PERCEPTION: 11; Systems Concept (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1:
Perception 2:
Perception 3:
Signal 1: N/A
Output:
Output signal 2: N/A

LEVEL OF PERCEPTION: 10 Principle (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1:
Perception 2:
Perception 3:
Signal 1: N/A
Output:
Output signal 2: N/A

LEVEL OF PERCEPTION: 9; Program (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1:
Perception 2:
Perception 3:
Signal 1: N/A
Output:
Output signal 2: N/A

LEVEL OF PERCEPTION: 8; Sequence (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1:
Perception 2:
Perception 3:

Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 7; Category (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Safe passing
Perception 2:
Perception 3:
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 6; Relationship (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Passing car on your right
Perception 2: Oncoming car far away
Perception 3: In passing lane
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 5; Event (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Passing car
Perception 2: Trees zooming by
Perception 3: Cars approaching in distance
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 4; Transition (II VISION)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Movement of needle
Perception 2: Sound of engine
Perception 3: Closing in on car in front
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 3; Configuration (II VISION)

Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Speedometer needle
Perception 2: Number
Perception 3: Red line
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 2+; Picture elements (II VISION) - Needed level?

Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Edge
Perception 2: Curve
Perception 3: Shading
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 2; Sensation (II VISION)

Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Brown pixel
Perception 2: Black pixel
Perception 3: White pixel
Signal 1: N/A
Output: N/A
Output signal 2: N/A

LEVEL OF PERCEPTION: 1; Intensity (II VISION)

Output signal 1: N/A (perception only)
Stored perception 1: N/A
Stored perception 2: N/A
Stored perception 3: N/A
Error signal: N/A
Perception 1: Red sensitive cone: Red intensity
Perception 2: Green sensitive cone: Green intensity
Perception 3: Blue sensitive cone: Blue intensity
Perception 4: B/W sensitive rod: light intensity
Signal 1: N/A
Output: N/A
Output signal 2: N/A

I didn't mean to introduce a levels problem at all. What I had in mind was much simpler and more general. For one thing, I haven't yet begun to think seriously about ECSs that can direct their output to different supporting ECSs. I am still thinking of "blind" ECSs that simply produce more or less output depending on the error. But an outside observer can see that prior reorganization has given this ECS several different parallel paths (which I have called micro-loops or something similar) by which this output passes through the outer world and affects the ECS's percept. Normally, the path with the lowest resistance (I like the word "impedance") in the outer world provides most of the effect. If the real-world situation makes the micro-loops mutually incompatible, then only the one with the lowest impedance has any effect at all--we choose the path of least resistance. If the car is there, we use it. If not, we bicycle if that's available. Otherwise, we walk or take the bus. We can't do all at once, but all have at some time or other been by reorganization made available as expressions of the output of the ECS.

If a low-impedance loop is not available, perhaps the ECS has to produce more output to have any effect at all on its percept (these things are not linear when we get to this kind of level). But it may not have to do so. The alternate means may be used without any change in output level.

I don't think you need any "choosing" systems when you are actually acting in the world. I think you need them when you are imagining (planning), because then there are no blocks to the execution of any path (they all have essentially zero impedance). And the whole issue of imagination needs a discussion like the one you have been having with Greg. I don't think it is clear (at least not to me) yet. Martin

Date: Sat Sep 26, 1992 6:33 pm PST
Subject: Dag's diagrams; impedance

[From Bill Powers (920926.1900)] Dag Forssell (920926) --

An ambitious undertaking. It makes me nervous, because it takes the 11 levels as Gospel about 20 years prematurely. Also, it doesn't (can't, really) capture the parallel nature of systems at these levels, and the branching networks that underly each system at each level.

One reason this is difficult to do is that's hard to get a true set of examples that show what the levels are about. At the system concept level, for example, "I am a successful entrepreneur" isn't really a reference signal at this level, because "successful" implies that the dimension of variation is in degrees of success. Actually the reference level should be "I am an entrepreneur." Success is defined by lack of error -- that is, how closely your perception of yourself matches the reference perception. simply picking an image of an entrepreneur DEFINES what success means.

I don't see how the other perceptions at level 11 amount to system concepts: "Customers benefit from my product" sounds to me like a principle, or even a logical deduction. And how is a relationship between your work and frequency of travel a system concept?

I've never seriously tried to fill in all 11 levels like this. The only way I could think of doing it realistically would be to head for the analyst's couch and start taking apart my own goal structure. This sort of example isn't too difficult to construct at the lower levels, but it's damned hard to think up arbitrary examples that really show the difference between the levels. At the higher levels I think you would have to study the organization in a real setting with a real person.

Try starting with a system concept like "I am an entrepreneur." Then try to answer the question, "What makes you think you're an entrepreneur?" The answers (for some person) might be things like "Well, I get out and hustle, I take risks, I hope to get rich, I organize things..." from which a listener would derive "Hustle; take risks; get rich; organize" and so on. In other words, you should get a series of generalizations which are the elements that for this person make up being an entrepreneur. Then you have to take each of those principles, and ask for examples, which will get you specific program-level perceptions.

I don't think you can just make up diagrams like this, not if you want them to have any persuasive force. They have to be based on actual personal experience, yours or someone else's, or they just won't be convincing. In effect, you have to start doing the research that we need, with real people, to find out how these higher levels actually hang together. This can't be done at arm's length, in an uninvolved and abstract way. If you do, it will SOUND and BE uninvolved and abstract.

If you just want to get the general idea of the levels across, I think it's best to avoid drawing such specific diagrams, unless they contain actual data from a live person.

Martin Taylor (9209025.2000) --

The "impedance" concept is sort of ingenious, but I can't see how to model it so that a particular output from a control system would be spread out among all the different possibilities -- what would keep all of them from trying to happen at once? It is really harder to drive a car than to ride a bicycle? When I try to reduce your clever idea to a specific case, I can't think of a way to make it work.

Best, and goodbye until about Thursday ... Bill P.

Date: Sun Sep 27, 1992 8:32 am PST
Subject: Behavior of Perception

[From Dag Forssell (920927-1)]

Bill Powers (920926.1900) In quick response to my diagrams:

>...it takes the 11 levels as Gospel about 20 years prematurely.

>I don't think you can just make up diagrams like this, not if you want
>them to have any persuasive force. They have to be based on actual
>personal experience, yours or someone else's, or they just won't be
>convincing.

I really appreciate your quick feedback, before you went camping.

My wording for the upper levels was hasty (and sloppy), pressing against time to post something for the sake of feedback. I had spent more time and felt more comfortable with the lower levels.

My thinking now is to introduce the analysis of levels with guidance of your post on levels of perception (920324.0300). (Any thoughts on my description of the bottom levels of vision. We have never mentioned the three color sensitivities of cones, but they are well documented).

I grant you that

>Also, it doesn't (can't, really) capture the parallel nature of systems
>at these levels, and the branching networks that underlie each system
>at each level.

But this diagram can be supplemented by others (drawn to a smaller scale, (such as LCS I: page 278) designed to convey the parallel and branching nature of the model. I am working on that too.

>This sort of example isn't too difficult to construct at the lower levels..

I would like to do just this, then leave the upper reaches blank. My post yesterday will provide a start and a vehicle for communication on the net. The net result for now is simply to scale back my ambition.

Any contributions, corrections and comments from other netters are solicited. I will share the graphic result by snail mail on request, which will at least look good and be suggestive of further research.

I agree that it is a good idea to leave the upper reaches blank, leaving them to individual introspection.

For purposes of Vision/Mission statements, which are a leadership application of PCT, I will be the first to suggest that people use the general concept of layering, but use labels that make sense to themselves as they see fit.

As I see it now, suggested wording in the upper reaches may include:

System concept: Understanding
 Belief
 Vision

Principles: Principles
 Values
 Priorities
 Mission
 Standards

Programs: Action Plan

Sequences: Methods

By the time a person has created a statement for a business, using some of this terminology, it will likely appear to have more layers than we talk about in the model.

Leaving the upper reaches of the model blank invites people to think this through and create layers of terminology that make sense to them.

Again, thanks for fast feedback, Bill.

Dag

Date: Sun Sep 27, 1992 6:57 pm PST
Subject: turing test]

(ps 920927)

the following was meant to go to the list. i've included it verbatim below, and followed it w/ my comments.

From: (Cliff Joslyn)

> (penni sibun 920922.1400)
>
> but what is the criterion for intelligence? the turing test
> presupposes that the only criterion for intelligence is being judged
> intelligent.

That's fine, but then the TT is not really a test at all. If all it says is "if someone (you, me, we all) says it's intelligent, then it's intelligent", then the TT is not even a heuristic. The whole purpose of a test for X is to distinguish between what merely APPEARS to satisfy X and what REALLY satisfies X. We don't talk about a test for redness; under normal circumstances, the perception of redness is its own test. The TT helps us understand when we're PERCEIVING intelligence. But like redness, that perception is intuitive, and automatic; and frequently faulty.

As you've described it, the TT is not a TEST for intelligence, but rather a DEFINITION of intelligence. That's OK, but let's not get confused. Of course all definitions are necessary and sufficient; that's what a definition IS.

>(the test can make mistakes,
> generating both false positives and false negatives, but that's an
> orthogonal issue.)

No, that's not orthogonal; that's exactly my point. The concept of a false positive or negative assumes that there's some criteria for Xness independent of the test for Xness. The test provides imperfect evidence for Xness, perhaps accurate, perhaps not. When the test is USED, then it provides us all the information we have; but when it's DESIGNED, then we have the "higher" definitinoal level as well. Is the TT the criteria, or the test? If it's the test, then what's the criteria (the prior definition?) If it's the criteria (as I say it is), then what's the REAL test?

> i'd bet

> you a lot (if i had it, which i don't!) that you will never ever
 > discover such a definition that works in the real world. i don't
 > think you should be spending your time looking for it.

On the contrary, I think the TT IS such a definition, and that it will work in the real world.

Cliff Joslyn

=====

> (penni sibun 920922.1400)
 >
 > but what is the criterion for intelligence? the turing test
 > presupposes that the only criterion for intelligence is being judged
 > intelligent.

That's fine, but then the TT is not really a test at all.

random house has a concise definition of ``test``:

the means by which the presence, quality, or genuineness of anything is determined.

i don't see how the tt fails at this.

As you've described it, the TT is not a TEST for intelligence, but rather a DEFINITION of intelligence. That's OK, but let's not get confused.

i'm not confused. i don't think you're confused either, though i'm pretty sure you don't understand what i've been saying.

>(the test can make mistakes,
 > generating both false positives and false negatives, but that's an
 > orthogonal issue.)

No, that's not orthogonal; that's exactly my point. The concept of a false positive or negative assumes that there's some criteria for Xness independent of the test for Xness.

maybe i shouldn't have used the terminology of false positives and negatives. what i meant was that over time the tester might revise her judgement, eg, based on new evidence. determining intelligence is a process of finding out, and you are not going to find out everything about an entity by the time you have to make a judgement.

you're actually saying tis right here:

The test provides imperfect evidence for Xness, perhaps accurate, perhaps not. When the test is USED, then it provides us all the information we have;

but i don't know what you're talking about here. the tt doesn't have a definition embedded in it.

but when it's
DESIGNED, then we have the "higher" definitinoal level as well.

Is the TT the criteria, or the test?

`criteria' is plural; `criterion' is singular. your criterion/test distinction seems to add an extra level of abstraction to how you talk about things w/o having anything more to say about the phenomena in question (eg, intelligence). you've also set yourself up for infinite regress, i think.

cheers. --penni

Date: Mon Sep 28, 1992 5:22 am PST
From: Martin Taylor
EMS: INTERNET / MCI ID: 376-5414
MBX: mmt@ben.dciem.dnd.ca

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Behavior of Perception

Dag, some quotes from July by me, Bill P, and me...
I could get the specific date references if you want.

=====

Me...

>>If the division into levels of control is required to be explicit, it must be
>>localised in a single ECS (that is, one ECS for each variable under control at
>>that level). That's one level of the hierarchy, right? To require this to be
>>"explicit" sounds a lot like the symbolic AI approach. In a distributed
>>connectionist system, a single node can participate in the (non-localised)
>>representation of more than one concept, depending on the global dynamical
>>activation of the network.

>This kind of distribution of function exists in the hierarchical control
>model. For example, the reference setting for a perception at level N is
>often the sum of several higher level outputs, and is set to satisfy the
>goals of all these higher level systems. Similarly, several systems at level
>N may be involved in satisfying the goal of a single system at level N+1.
>Perceptual functions are also distributed in this way -- the hierarchy is a
NN.

Well, yes and no. I suspect Allan was getting more at the question of whether a group of ECSs at a level can act as a population vector, though I admit his wording is ambiguous. I look on the question more as "given that we have sensors only for THIS red, THIS green, and THIS blue, how is it that we can control for a pretty close match to a wide range of blends." We presumably do not have perceptual input functions that match all different possibilities. Indeed, if we did, our absolute colour recognition would be as good as our colour matching. We control for simultaneous values of the levels of three different references, an emergent population behaviour, I think.

=====

Bill...

>... "given that we have sensors only for THIS red, THIS green, and THIS
>blue, how is it that we can control for a pretty close match to a wide
>range of blends."

The obvious answer to this penetrating question is not very satisfactory to me. In simplistic terms, the target color has to be remembered (all three components, in different control systems), and then each remembered value must be reproduced by varying the individual color dimensions. My biggest problem with this explanation is that I don't experience colors in terms of their trichromatic components, but simply as whatever colors they are. Another problem is that people can compare two colors and pronounce them similar or different (hearking back to a previous conversation).

A somewhat more pregnant solution, if you will allow degrees of pregnancy, involves imagination, or model-based control. A control system might ask, "How would I have to change THIS color to make it the same as THAT color?" I think we do something like this when trying to match colors: that white is too bluish, the other one is too yellowish. This could be a report on the way a color sensation has to be changed to achieve a match; the implied meaning of the statements is "too bluish TO BE THE REFERENCE COLOR." Experiments with shape matching seem to involve manipulations of perceptions in imagination, such as rotations. The object being compared isn't compared as it is, but after mentally shifting it, changing its scale, flipping it, or rotating it -- all operations that we seem to be able to do in imagination. The amount of difference between the original and the target perception could then be sensed as the amount of mental action required to eliminate the difference.

It's hard to get out of the habit of using the old model. One day soon I'll have to try out the model-based scheme in a serious way. An appropriate experiment would be one in which the controlled variable disappears occasionally, moving behind an obstacle or being blanked out in the middle of a control run. Another might be a high-noise situation in which the controlled variable can't ever be perceived very accurately. It might even be possible to deduce the characteristics of the mental model by seeing how control behavior continues during times when the controlled variable isn't visible or becomes ambiguous.

=====

Me...

>>... "given that we have sensors only for THIS red, THIS green, and THIS
>>blue, how is it that we can control for a pretty close match to a wide
>>range of blends."

>

>The obvious answer to this penetrating question is not very satisfactory to
>me. In simplistic terms, the target color has to be remembered (all three
>components, in different control systems), and then each remembered value
>must be reproduced by varying the individual color dimensions. My biggest
>problem with this explanation is that I don't experience colors in terms of
>their trichromatic components, but simply as whatever colors they are.
>Another problem is that people can compare two colors and pronounce them
>similar or different (hearking back to a previous conversation).

(I don't quote your "more pregnant" solution, because it goes off in a different

direction that is unaffected by the following comment.)

The core of my question, and of my interpretation of Allan's, has to do with the concept of "coarse coding." Coarse coding is a very important principle that we seem to use a lot, all the way from the retina (which uses it both in space and in colour) to the motor systems. The colour system is probably the simplest to describe, but the concept applies in multidimensional spaces equally well. We have in our retinas three types of cone, all of which respond

to light over a wide range of frequencies. Call them R, G, and B. Forget B for the moment. The sensitivities of R and G vary across frequency, R being most sensitive to a colour we might call orange, and G to a yellow-green.

But we can make exquisite discriminations over a wide range of frequencies because anywhere between green and red the ratio of the R cone outputs to the nearby G cone outputs depends strongly on the frequency. We perceive some function of this ratio as a colour. For bluer colours, and for non-spectral hues, the B cones come into play, and we use the ratios among all three--or more correctly, between the B and the sum of R and G.

In general, coarse coding depends on there being a set of detectors each of which is most sensitive to some specific values of some features that can have a continuum of values. These would be the perceptual outputs of some ECSs (i.e. the outputs of their perceptual input functions). The ranges of values to which these detectors are appreciably sensitive overlap substantially, and a subsequent system can obtain precise information about the location of some entity in the feature space by deriving ratios among the outputs of the detectors. The resolution that can be obtained by this procedure can be much greater than the grain of the detector centres in the feature space.
=====

So we have had at least one interchange on RGB coding of vision. Would you like more?

Martin

Date: Mon Sep 28, 1992 7:44 am PST
Subject: Re: Dag's diagrams; impedance

[Martin Taylor 920927 18:30] (Bill Powers 920926.1900)

>The "impedance" concept is sort of ingenious, but I can't see how to
>model it so that a particular output from a control system would be
>spread out among all the different possibilities -- what would keep
>all of them from trying to happen at once? It is really harder to
>drive a car than to ride a bicycle? When I try to reduce your clever
>idea to a specific case, I can't think of a way to make it work.

Well, I must have misinterpreted you some months or more ago, because I thought it was your idea. It's "always" been a part of my concept of PCT. Anyway, if it wasn't your idea, here's how I see it. They all DO try to happen at once, but they can't. They inhibit one another. It is the world that stops them all happening at once. If they could, they would. Outputs are going in all directions, but the world prohibits some percepts from actually being

controlled. So the "taking bicycle" percept/reference cannot be satisfied if you are sitting in the car. I guess in the background there is another point I meant to discuss at some time--"giving up." When there is persistent error, one possible and often used response of an ECS is simply to reduce its gain to zero, to give up on a hopeless situation. If that happens, the "taking bicycle" reference is simply relegated to the scrapheap of unfulfilled goals (don't take that seriously) when you start the car.

In general, if the world is such that the taking of some action precludes the taking of another, only one will happen (or neither). You have no choice in the matter, and generally the one that happens is the one that seems easier at the moment (i.e. involves the most readily satisfied lower-level reference levels).

I don't see choices happening at the time of the event. Only in the imagination or preplanning operations, and in them you can imagine at the same time bicycling, driving, and walking. You could think of choices being made then, possibly as a consequence of the overall error resulting from the imagination of each (aching muscles and wet clothes from bicycling, for example). But imagination/planning is a whole can of worms perhaps best left unopened for now.

Martin

Date: Mon Sep 28, 1992 7:55 am PST
Subject: Arm (to B.P.); Influence Summary

From Greg Williams (920928)

Attn. Bill Powers: Getting photos made sounds OK. I just hope they're in the \$20 range, not the \$100 range! I'll work on the paper corrections today and get them to you this week. I assume that you'll print out the figure captions yourself.

In the belief that what Gregory Bateson has termed "double description" might be usefully applied to our discussion on "influence," I offer the following TENTATIVE summary, to be compared and contrasted with Bill's summary (which I have not yet read). MULTIPLE description might be even more useful. I hope others will provide comments on, corrections to, and elaborations of both the following and Bill's summary.

Because labeling/terminology has proven problematic in this discussion -- unnecessarily so, I believe -- in the following, I have tried to avoid all interpretive wording which relates PCT ideas to folk and sociological terms (such as "influence," "manipulation," "exploitation," "facilitation," "domination," and "aggression"). This doesn't mean that I don't think that defining (at least some of the) conventional terms in PCT terms is an important undertaking. But that remains for the future.

An aside: What's going on in HYPNOSIS in PCT terms???

ACCORDING TO PCT, HOW ONE PERSON (A) MIGHT CONTROL FOR HIS/HER PERCEPTIONS THAT DEPEND ON ANOTHER PERSON (B) ACTING IN CERTAIN WAYS

(The "might" is significant. PCT says that the following methods are possible in theory, but extra-PCT considerations might make any one of them difficult or even impossible in certain cases.)

1. A disturbs particular perceptions being controlled by B so that B compensates for the disturbances with actions which A wants to perceive.
2. A arranges B's environment so that when B controls for particular perceptions, A perceives what he/she wants to perceive.
3. A arranges B's environment so as to trigger learning/reorganization in B's control system resulting in actions which A wants to perceive.
4. A applies physical constraints or threatens to apply physical constraints to B so that B's actions are as A wants to perceive.
5. Combinations of the above.

No other methods are possible.

In methods 1 and 2, B's control system functioning is not conflicted; in methods 3 and 4, it is conflicted.

ACCORDING TO PCT, WHAT AN ORGANISM'S CURRENT CONTROLLING DEPENDS ON

1. Current operation of an organism's control system (not undergoing learning/reorganization) is a function of current reference signals, current input/output functions, and current environmental disturbances.
2. Current reference signals and current input/output functions are functions of the path of the most recent bout of learning/reorganization.
3. Onset of learning/reorganization at a particular time is a function of reference signals, input/output functions, and environmental disturbances at that time.
4. The path of learning/reorganization is a function of (possibly randomly generated) successive sets of changes in reference signals and/or input/output functions, each successive set of changes being made to the result of the previous set of changes, and another set of changes being made only if certain criteria are not met for ceasing learning/reorganization.
5. Whether or not the criteria for ceasing learning/reorganization are met by the reference signals and/or input/output functions at any point on the path of learning/reorganization is a function of reference signals, input/output functions, and environmental disturbances at that point.

(Note: Here it is assumed that memories of environmental disturbances count as environmental disturbances, so that, for example, the disturbance of being called a "pig" would, as a memory, continue to act as a disturbance perhaps for a long time after the sound waves had dissipated, with similar results as if "pig" were being repeated over and over again by the disturber.)

6. At any time, the criteria for ceasing learning/reorganization are functions of reference signals and input/output functions at that time.

Best wishes, Greg

Date: Mon Sep 28, 1992 7:57 am PST
Subject: REVISED DIAGRAMS FOR RUBBERBAND EXPERIMENTS

CHUCK TUCKER 920928A

I am posting three revised diagrams for the rubber band experiment. These diagrams have the symbols (colors) arranged in a mirrored reverse arrangement. This arrangement makes it possible for the E to place his/her finger on a "spot" on his/her side of the diagram and the P will have his/her finger on that same "spot" be it symbol, letter, color or whatever. The E does not have to watch the finger of the P but simply his/her own finger. This is a way to show that E can control for a particular symbol and if P is controlling for the knot over the "X" then P's finger will be on the same symbol as E's. Now, what do you want to call this: influence, exploitation, manipulation or ordinary social transaction??? Oh no, not this whole discussion again!!!!

Regards, Chuck

Date: Mon Sep 28, 1992 7:57 am PST
Subject: PREV POST

DIAGRAM A

A	R	C	U	L	O	T
N	L	O	C	B	P	F
S	I	Y	O	N	C	V
Q	W	R	T	O	N	D
A	G	E	I	U	Y	S

Date: Mon Sep 28, 1992 8:06 am PST
Subject: "Influence"

From Greg Williams (920928 - 2)

>CHUCK TUCKER [920925]

>CLOSED LOOP GW 920923-2

>I know it would be a great amount of work but I think that you
>(Greg) can edit the transcripts of the conversation so as to
>represent the major matters mentioned about "Influence and
>Control". I do think that it would be useful to have a "summary"
>of the discussion (as Cliff suggests). The other part of the
>project I suggest, i.e., comments by other on the Net and a
>future book, can wait to see if there is any interest in it.

Sounds OK to me, as long as ALL contributors get to OK the final version(s?).

> I

>think that one way to get people interested in these ideas is to
>have our books in their libraries in the country and ask (in some
>instances FORCE) people to read them.

FORCE?!?! What an interesting suggestion for promoting PCT. How about burning
all non-PCT psychology books, for good measure? Do you teach a course in
ideology, by any chance?!?!

>INFLUENCE AND CONTROL WTP 920922.0900; 23.0900
> GW 920922; 23; 23-2
> PW 920923

>Unfortunately, I still have a problem with your specification of
>what you mean and how you determine that A has influenced B.
>...
>The problem with these definitions that I find is that they are
>stated in terms not consistent with PCT since they contain
>attributes, i.e., "the power or capacity," or notions which are S-
>R, i.e., "to produce an effect" or "to flow in". But I don't
>think you are using INFLUENCE in that way (at least most of the
>time) but rather as a description of interaction where in A
>figures out what B wants (call it "X") and provides it for A in
>such a way that their interaction continues w/o force on the part
>of either A or disturbances on the part of B.

That's right. I hope my just-posted summary helps to clear up the problem.

>But it seems to me that in PCT terms for A to influence B to
>"produce a certain effect" that B would have to take and use a
>purpose (reference signal "X") "provided" by A since "a certain
>effect" in PCT terms is action taken by B consistent with B's
>reference signal X.

It's OK with me if you want to use "influence" solely in this way. I'm not
really very concerned about clinging to PARTICULAR matches between PCT

descriptions and PARTICULAR everyday/sociological terms; to me, the important thing is to identify SOME matches between certain PCT descriptions and everyday/sociological terms. If "influence" is best construed (because of some sorts of historical precedents, fine) what PCT talks about as B's learning/reorganization which depends (in part) on A, so be it. We're all trying simply to make sense, first and foremost, and we don't make sense de novo (except maybe to ourselves only)!

>Thus, with this specification, Pat (please
>excuse me if this is offensive to either of you) has influenced
>your childrens' food preferences to the extent they have adopted
>similar notions of "healthy" and /or "tasty" food and it seems
>that there is no other source for these preferences. But even if
>this is the case, Pat's actions have served only as an "occasion"
>for the children to develop these preferences from a PCT
>perspective.

That is correct. Pat would then be A contributor to the result, not the ONLY contributor. She would be necessary (assuming the kids wouldn't adopt the notions without her "influence") but NOT sufficient for the outcome.

>I say this not as a criticism of Pat (or anyone
>else, including myself) but to simply point out that in PCT terms
>all that we DO serves as an "occasion" for what others DO; we do
>not determine what they DO.

Correct, unless we apply overwhelming physical force. This is what I've been saying all along. Notwithstanding your feelings expressed above about the proper use of "influence," I've been at pains to use that term because it (in my opinion!) connotes PARTIAL cause and so stands in distinction to "determination," which connotes COMPLETE cause.

Best, Greg

Date: Mon Sep 28, 1992 10:03 am PST
From: Hortideas Publishing / MCI ID: 497-2767

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Diagrams

From Greg Williams (920928 - direct)

Sorry I've been too busy to comment on your diagramming before now. I'll be gone tonight and tomorrow, at home Wednesday on.

>Dag Forssell (920926)

>I'll be happy to mail good looking charts of the final result to those
>deserving souls who help out.

And for the CSG archives, too, I trust!

>As I labor on this, I get uncertain about the levels where I have placed the
>different happenings of passing a car. Have I stretched this too much?

My overall suggestion, considering your motivation for doing the diagramming and the intended audience, is that you should not be worrying too much about whether every detail is perfectly worked out and whether there are any inconsistencies. I think your main point should be that this is the level of analysis necessary for a detailed description of the underlying phenomena involved in any behavioral scenario. Your managerial/engineering audience will get the message that the required analysis is anything but trivial, but that PCT is up (or is getting to be up) to the task. Remember, this whole exercise isn't to convince non-PCT psychologists that PCT is useful for understanding the genesis of behavioral events, but has the primarily heuristic function of relating the theory to "real life" events. At least that's the way I see it. If I'm approximately correct in my assumptions, the demonstration of how such an analysis would look is FAR more important than whether all the details are exactly correct. Besides, how would you verify (oops, sorry -- falsify) the theoretical details at this stage of our knowledge about actual peoples' CNSs?

>Bill Powers (920926.1900)

>An ambitious undertaking. It makes me nervous, because it takes the 11
>levels as Gospel about 20 years prematurely.

Maybe 100 years or even more, with respect to falsification possibilities!

>Also, it doesn't (can't, really) capture the parallel nature of systems at
>these levels, and the branching networks that underly each system at each
>level.

Yet it gives the flavor of the depth of analysis required even WITHOUT getting into the parallel business.

>I've never seriously tried to fill in all 11 levels like this.

But see "A Cybernetic Model for Research in Human Development," in LCS [I].

>Dag Forssell (920927-1)

>My wording for the upper levels was hasty (and sloppy), pressing against time
>to post something for the sake of feedback. I had spent more time and felt
>more comfortable with the lower levels.

As I said above, I don't think that is a big problem.

>But this diagram can be supplemented by others (drawn to a smaller scale,
>(such as LCS I: page 278) designed to convey the parallel and branching
>nature of the model. I am working on that too.

Sounds good.

BP>>This sort of example isn't too difficult to construct at the lower levels..

>I would like to do just this, then leave the upper reaches blank. My post
>yesterday will provide a start and a vehicle for communication on the net.
>The net result for now is simply to scale back my ambition.

How about the arm demo (new version almost ready) and/or the gatherings demo as explicit examples of lower levels operating in real time? There, you KNOW how the levels are organized!

For purposes of Vision/Mission statements, which are a leadership application of PCT, I will be the first to suggest that people use the general concept of layering, but use labels that make sense to themselves as they see fit.

>By the time a person has created a statement for a business, using some of >this terminology, it will likely appear to have more layers than we talk >about in the model.

Quite frankly -- I've told Bill, of course -- I have never bought into Bill's hierarchy AT ALL except as a (not very well fleshed-out) existence proof for AN instantiation of PCT details. At this stage, I think PCT has a lot to learn from living organisms before pushing HPCT as THE instantiation. Your leaders are providing data which shouldn't be ignored by the theoreticians.

>Leaving the upper reaches of the model blank invites people to think this >through and create layers of terminology that make sense to them.

Right on!

Best wishes, Greg

Date: Mon Sep 28, 1992 12:16 pm PST
Subject: Misc replies

[From Rick Marken (920928.1030)] Gary Cziko (920926.0235 GMT)--

Thanks for the quotes about unifying observations in psychology. I promise that my next submission of the paper to a journal will present the ideas in the "Blind men" paper in terms of PCT as a unifying model -- and I think this change merits addition of your name as co-author (if you don't mind). When I make the revisions I'll send you a copy for your approval.

Dennis Delprato (920925) -- Congrats to you and Brian Midgley on publication of your paper. I agree with Bill; it's nice to know that some PCT people know how to get published.

Greg Williams (920928) --

>An aside: What's going on in HYPNOSIS in PCT terms???

Great question. I think it's a consciousness phenomenon, which means I haven't got a clue, other than that it has to do with systems we (I) don't understand that have to do with "willing" arbitrary changes in reference signals (ie. making changes that have nothing to do with achieving higher order goals). I think hypnosis has to do with that aspect of ourselves that let's us "wiggle our index finger" for no other reason than to do it.

>ACCORDING TO PCT, HOW ONE PERSON (A) MIGHT CONTROL FOR HIS/HER PERCEPTIONS
>THAT DEPEND ON ANOTHER PERSON (B) ACTING IN CERTAIN WAYS

Your list of control methods is excellent; very helpful.

By the way, Greg, Pat, Chuck, Bill and others involved in the "Purposeful influence" discussion: the idea of an issue of "Closed loop" on this topic is a very good one. And the idea of eventually turning this into a book is also a great. I have been trying to write a book on this subject for four years (and avoiding it by writing nonsense like the "Blind men" paper); I think my book would have a slightly different emphasis than the "purposeful influence" book. Actually, I think the "purposeful influence" book could serve as a nice supplement to my book (if I ever get it written) or vice versa. So don't feel bad that you are beating me to the punch with a better book by smarter people and depriving me of a possible source of income that would have allowed me to send my kids to college and buy my wife a new coat (well, I guess you can see through that BS; who needs a coat in California. I'd get her a skimpy bathing suit; she still looks great in one).

Suggested new topic: PSYCHODYNAMIC PSYCHOLOGY

As the proud papa of TWO currently active threads on CSGNet -- the ones on the Turing test and "purposeful influence" -- (yes, folks, I unintentionally started the "purposeful influence" thread on august 26 when I asked Greg Williams if he knew of any reference to behaviorist comments that suggested an awareness that their model of people (as controlled agents) might be a problem for their belief in their own ability to control).

The new topic is: psychodynamics -- models of human nature that emerge from Freudian "theories" of the mind; theories like those of Jung, Adler, Maslow?, etc etc. You know, the fun stuff. We spend a lot of time here explaining why PCT is better than behaviorism (pick your flavor) and cognitivism (pick your branch). But most lay people, I think, see the psychological theory to beat as some form of psychodynamic theory. Most pop psych stuff has a lot of psychodynamic influences in it (when I was young it was people like Fromm, Byrne, etc). I get the impression that psychiatry is still heavily influenced by psychodynamic thinking. I think a lot of people think that stuff like Jung is really deep whereas more scientific approaches (like PCT) are shallow and not getting at the real nitty gritty, slimy stuff that's really going on in our minds (especially our unconscious). These people are particularly interested in dreams, literature and art as sources of information about our essential selves. I have a feeling about where they are coming from; stories like oedipus rex, greek, roman and even abrahamic mythologies have such rich imagery and seem to "touch" some of our deepest concerns in ways that are not captured by saying things like "family is a system concept".

My question is: what do you think PCT could say to the psychodynamic types that might speak to their interests/concerns? Are there any Freudians or Jungians on the net? What is the Freudian or Jungian impression of PCT? Do we care?

Best regards Rick

Date: Mon Sep 28, 1992 1:09 pm PST
Subject: dennet & powers

September 28, 1992 [From Pat Alfano] Francisco Arocha

For an ecological perspective of "our perceptual experience" try James J. Gibson's 1966 book, *The Senses Considered as Perceptual Systems*. Also, his 1979 book; the title of which is something like *An Ecological Approach to Visual Perception*. He also has written many articles. His ideas are quite different than those of Richard Gregory and the mainstream perceptionists.

Pat

Date: Mon Sep 28, 1992 2:50 pm PST
Subject: Giving up (The Test in practice)

[Martin Taylor 920928 16:45]

The following question was posed to me by a colleague, and I thought it of sufficiently general interest to communicate it and my answer to CSG-L.

=====

>On the first day, David moves the garbage can into the middle of the aisle.
>Martin moves it back.
>On the second day, David moves the garbage can into the middle of the aisle.
>Martin moves it back.
>On the third day David moves it again.
>Martin leaves it.
>
>Is Martin controlling for the position of the garbage can?
>
>No, he stopped.
>Yes, he has determined that if he ignores the disturbance, it will
>go away, eventually.
>
>How does David test whether Martin is controlling for the position of the
>garbage can?
>

You are talking about a phenomenon I label, with great originality, "giving up." When two ECSs are in conflict, it is impossible for both to bring their errors to zero simultaneously. That's almost a definition of conflict. Typically, both control systems are inside one body, and in many real-world situations there is more benefit to the body from allowing one of the control systems to satisfy its reference, than from leaving the real-world variable at a compromise level. Sometimes, the world is such that there exists no compromise. If one ECS, for example, has a reference to perceive the body as bicycling to work, and the other has a reference to perceive it as driving, there is no compromise intermediate position (possibly a motorcycle, but let's suppose none is at hand). The mechanism is that one ECS gives up attempts to control. Either the person cycles, or the person drives, but not part of each.

(There are other reactions to persistent error, including reorganization, which could lead to the possibility of both references being satisfied, the problem being seen in a new way that eliminates the conflict).

The "giving up" mechanism is needed in a related situation, a shift of attentional focus in response to an alerting signal. The alert causes some ECS that was not controlling to attempt to exert active control. This means that some ECS that had been controlling must "give up." If that did not happen, the body would not survive very well--imagine a primitive man eating berries when an alerting system showed him a tiger about to pounce. If the berry-eating percept HAD to be satisfied, or even if it caused error in the perceptual control for escape, there would be a nice dinner for the tiger. And this might happen, if the berries were particularly tasty so that the man hesitated about the escape (conflict situation).

So, we must have a built-in mechanism for giving up in the face of difficulty. It can't be too strong, of course, and I suspect that in humans it is too weak for our own good (leading to war and similar bad things). But it is there. If attempts to control a percept are persistently thwarted, there are two possible reactions--try harder or give up, fight or flight, for example.

Now to your original question, how to apply The Test. I don't think you can, unambiguously. The very first time you move the wastebasket, and I move it back, you know that I am controlling for some percept that includes (at that moment) having the wastebasket in its original place. But you don't know just what it is that is really being controlled for. Perhaps the wastebasket is blocking a draft that makes my feet cold, and you could have replaced it with something else that would have kept me happy. Perhaps I want to see a rigidly aligned set of objects in the office, and you could have satisfied my reference by moving two or three other things into the aisle. All you know, and this you do know, is that I am controlling now for a percept that includes the wastebasket where it was.

My claim is that The Test is always ambiguous. P can tell that Q is controlling for a percept that incorporates some CEV that P has disturbed, but P can never tell that the percept Q is controlling for corresponds exactly to the CEV that P's percept corresponds to. The CEV is in the world, but only percepts can be controlled. The percepts are private, so in a sense the CEVs are, too. But the correspondences of CEVs can be tested, and that is what you do when you try variants on moving the wastebasket, such as putting some other obstruction in the aisle, or moving the wastebasket to some other unobtrusive place. Such experiments can never give exact answers (pace Rick Marken and Bill Powers) but can only be evaluated statistically.

The fact that on the third day I leave the wastebasket where you moved it is irrelevant. You got your information on the first trial. On the third day, I obviously am not controlling for a percept that includes the wastebasket in its original position, or if I am, there is an internal conflict in me that causes some other controlled percept to override that one. For example, I may perceive that by persisting in moving it back I will cause you to get angry, and this would create error. My reference for seeing you satisfied has overridden my reference for having the wastebasket where it was.

Reference levels change. What we control for at one moment may not be what we control for at another moment. Why not? Because some other percept changed at a higher level, altering the error in this higher ECS, and this changed the reference levels in all sorts of ECSs at levels below it. Another reason might be reorganization, which naturally changes reference levels by virtue of the changed interconnections of ECSs. The changes in reference level are an

independent reason why experimental tests can only be statistical, and why one can never be sure to what degree one's percepts coincide with the percepts of other people.

Martin

Date: Tue Sep 29, 1992 2:23 am PST
Subject: Change of address

To whom it may or may not concern. My new address as of 1.9.92 is:

Oded Maler, LGI-IMAG (Campus), B.P. 53x, 38041 Grenoble, France
Phone: 76635846 Fax: 76446675 e-mail: maler@vercors.imag.fr
Home address: 7 chemin de vence, 38700 Corenc, phone: 76880711

Date: Tue Sep 29, 1992 10:27 am PST
Subject: The Test in practice

[From Rick Marken (920929.1000)] Martin Taylor (920928 16:45) --

Your example of the test is, indeed, interesting.

In your example, it looks like Martin is controlling the position of the waste basket, then stops. I don't see this as giving up -- it could be just a change of reference. There really wasn't much testing going on in the example; only one level and type of disturbance, no revision of the hypothesis about the CEV when there was lack of disturbance resistance, etc.

Your comments on conflict seem reasonable. I would just say that I think the test should not involve setting up a conflict with the testee. The tester should just be disturbing variables, not trying to prevent the testee from controlling the variable. The test should cause no inconvenience for the testee -- or higher order systems are likely to become involved and, ultimately, reorganization. This would mess up the clarity of the results of "the test".

>My claim is that The Test is always ambiguous. P can tell that Q is control
>ling for a percept that incorporates some CEV that P has disturbed, but P can
>never tell that the percept Q is controlling for corresponds exactly to the
>CEV that P's percept corresponds to. The CEV is in the world, but only
>percepts can be controlled. The percepts are private, so in a sense the
>CEVs are, too.

I suppose this is ultimately true. But when your guesses about a controlled (perceptual) variable are allowing you to predict behavior variations with 99% accuracy, I think worrying about whether you have identified the absolutely true CEV is rather academic. For example, in my "area" vs "perimeter" control study, I did the test for the controlled variable to determine whether the subject is controlling $x+y$ vs $x*y$ (where x and y are height and width of a quadrilateral figure). Using $x+y$ as the hypothesized controlled variable the error in predicting responses was, I think, about 2%. With $x*y$ as the hypothesized CEV, the error was halved, to 1%. You could probably do slightly better with some other hypothesized CEV -- maybe $\sqrt{x*x+y*y}$ -- but clearly you are on the right track with $x*y$. I guess you reach a point where you have

to determine how much trouble it's worth to try to improve your model of the behaving system; it seems to me that, when you are predicting over 99% of the variance, you are so far ahead of the conventional psychology game that such improvements are really mute (until PCT science develops to a level at which such tiny improvements are significant).

> But the correspondences of CEVs can be tested, and that is
> what you do when you try variants on moving the wastebasket, such as putting
> some other obstruction in the aisle, or moving the wastebasket to some other
> unobtrusive place. Such experiments can never give exact answers (pace
> Rick Marken and Bill Powers) but can only be evaluated statistically.

> The changes in reference level are
> an independent reason why experimental tests can only be statistical, and
> why one can never be sure to what degree one's percepts coincide with the
> percepts of other people.

I agree that there are many difficulties inherent in doing "the test". One only needs to actually try the "coin game" to see how difficult and frustrating it can be in practice. But I think that, in this early stage of the development of PCT science, the attitude that the test can only be evaluated statistically, is just too defeatist for me. I say this because 1) we already have some examples of very precise results of the test -- where there is no need for statistical evaluation at all. My "area/perimeter" study mentioned above is one; Bill's analysis of the "shock avoidance" task is another; my mindreading demo should count as another. These are examples of the test where the quantitative results are so precise that the use of inferential statistics (I presume this is the kind we're taking about) would be nothing more than pompous posturing (the results are obviously "significant") 2) we know that people are controlling perceptual variables and there is no reason to suspect that we cannot get a pretty accurate description of them -- even if we can't actually perceive what people are perceiving. So the goal of the test should always be high quality data -- we should work to improve the test (changing disturbance techniques, revising hypotheses about CEVs, whatever) until the results are nearly perfect (as in the examples we have) before throwing in the towel and starting to turn the statistical crank again; statistics should be the last resort; but given my understanding of the PCT model I can't see how they could tell you much about the system you are controlling. I don't rule out statistics completely -- I have used a statistical measure of performance in my mind reading demo (the stability factor) and done an inferential statistical procedure to decide whether or not an object on the screen was or was not under control. But this was a crude approach to doing a version of the test -- if I really wanted to know WHAT variable was being controlled, I would have done things that obviated the need for the statistical approach.

I think that this is an important topic and, perhaps by discussing it we will get some people out there to actually try the test and see what problems they run into.

I don't want to be "ideologically" anti-statistical, by the way. I am willing to believe that I don't fully understand your position on statistics, Martin. I consider it entirely possible that statistical methods be could useful tools in PCT -- they already have been. I just object to using statistics the way they are used in conventional psychology -- to see whether or not anything

happened in a study. Maybe it would help if you described exactly how you see statistics fitting into methodology of "the test". Perhaps give a concrete example of your idea of using statistics as part of "the test".

Thanks Rick

Date: Tue Sep 29, 1992 12:12 pm PST
Subject: Blind Men; Touting CSGnet

[from Gary Cziko 920929.0250 GMT] Rick Marken (920928.1030) says:

>Thanks for the quotes about unifying observations in psychology. I promise
>that my next submission of the paper to a journal will present the ideas in
>the "Blind men" paper in terms of PCT as a unifying model -- and I think
>this change merits addition of your name as co-author (if you don't mind).
>When I make the revisions I'll send you a copy for your approval.

While I would love to have you add my name to a revision of this paper, I think it would be a bit misleading if you did so just because you added some quotes from others that I provided and wrote a bit about PCT as a unifying model for psychology. I think I would will need to put some of my own words into the paper (like _the_, _of_, and _or_) for it to properly carry my name. If you wish to revise the paper on your own, an acknowledgement of my suggestion would be appropriate, along with others whose comments were helpful.

And while you're at it, why don't you try to work in a comment (or at least a footnote) about CSGnet giving it's address? You could also mention me (with my e-mail address) as a source of more information about CSGnet. I have done this in my paper on PCT that will be out in a couple of months in Educational Researcher_. I would also hope that other CSGnetters who are publishing on PCT will mention CSGnet in their work.--Gary

Date: Tue Sep 29, 1992 2:07 pm PST
Subject: Blind Men, UCSD anyone?

[From Rick Marken (920929.1330)] Gary Cziko (920929.0250 GMT) re: Blind men

> I think I would will need to put some of my own words into the paper

Of course. I think we should also change the title to reflect your suggested change in the paper's point of view. This will be after the Psychology review -- which might not be long, especially if it is rejected (THAT would be a coup -- being rejected by a journal that accepts everything).

Correction: In my post to Martin Taylor (today) I used the word "mute" when I meant "moot"> Interestingly, the mistake makes little difference to my intended meaning but I thought I'd point it out anyway.

General question: Is there anyone on CSGNet who is at UCSD? Or does anyone on the net know anyone with an e-mail address at UCSD? Is so, please let me know. Thanks.

Regards Rick

Date: Tue Sep 29, 1992 2:17 pm PST
Subject: Re: The Test in practice

[Martin Taylor 920929 16:00] (Rick Marken 920929.1000)

Let's take this backwards, in terms of answering your posting.

>I consider it entirely possible that statistical methods be could useful
>tools in PCT -- they already have been. I just object to using statistics
>the way they are used in conventional psychology -- to see whether or
>not anything happened in a study.

I also object to the way statistics are used in conventional psychology. I started this objection when I was in graduate school, and have maintained it ever since. Any time I see a significance test in a paper, my first question is "was it inserted at the demand of the editor, or did the author really not know what his/her data meant." A significance test tells you nothing except that your experiment was sensitive enough to detect an effect that you knew beforehand had to exist. To quote Ted Nelson, "everything is deeply intertwined." To quote myself in a letter refusing to incorporate significance tests into a paper that satisfied what Ward Edwards called "the Interocular Traumatic Test" (the test that Rick approves): The inclination of the rings of Saturn affects the curl of a puppy dog's tail. So what use is it to demonstrate that you have made enough measurements to show it to be so. The interesting question is how much influence the rings of Saturn have on that puppy dog's tail, and how reliable is that estimate under different circumstances.

> Maybe it would help if you described
>exactly how you see statistics fitting into methodology of "the test".
>Perhaps give a concrete example of your idea of using statistics
>as part of "the test".

That's an important question. I'd prefer to answer it in a separate posting, for two reasons. (1) It needs a lot of background presentation, some of which I have done in other postings, but which should be pulled together, and (2) it needs me to think of a mid-level experiment that is unlikely to satisfy the 99% predictability that can be obtained from lower-level tasks. My presumption is that you get the 99% prediction because the subsystems (perhaps ECSs) that are involved in the task are those that support very many different kinds of behaviour, and so are not readily disturbed by contextual differences. Since I am at the moment being asked to think of PCT-driven experiments in planning and decisions-making tasks, I may be able to satisfy condition (2) in that context. (Any help from the CSG-L readership in devising such an experiment would be most welcome, by the way.)

> These are examples of the test where the quantitative results are
>so precise that the use of inferential statistics (I presume this is the
>kind we're taking about) would be nothing more than pompous posturing (the
>results are obviously "significant")

As you see from the above, "inferential statistics" is not what I am talking about. What I am talking about would be more along the lines of parameter estimation for the perceptual functions, gain estimates with variance, and the like.

```
> we know that people are controlling
>perceptual variables and there is no reason to suspect that we cannot get
>a pretty accurate description of them -- even if we can't actually perceive
>what people are perceiving. So the goal of the test should always be high
>quality data -- we should work to improve the test (changing disturbance
>techniques, revising hypotheses about CEVs, whatever) until the results
>are nearly perfect (as in the examples we have) before throwing in the
>towel and starting to turn the statistical crank again; statistics should
>be the last resort; but given my understanding of the PCT model I can't
>see how they could tell you much about the system you are controlling.
```

...

```
>For example, in my "area" vs
>"perimeter" control study, I did the test for the controlled variable to
>determine whether the subject is controlling x+y vs x*y (where x and y
>are height and width of a quadrilateral figure). Using x+y as the
>hypothesized controlled variable the error in predicting responses was,
>I think, about 2%. With x*y as the hypothesized CEV, the error was halved,
>to 1%. You could probably do slightly better with some other hypothesized
>CEV -- maybe sqrt(x*x+y*y) -- but clearly you are on the right track
>with x*y.
```

(These statements are, as I proposed, in backwards order).

Do you see the contradiction between your two statements? Now I grant that you followed up with a comment that it isn't worthwhile to try to do better, and in this case you may well be right. But suppose your intention was to try to find the form of the perceptual function that defines the CEV. This form could be fairly critical in determining the internal linkages within the control net, because if it happened to be based on the sum of squares rather than a simple product, and you found on another occasion a different CEV that also used the squares of simple dimensions, you might well suspect that the individual squares could also be controlled if necessary. But if the better function was the product, you might ask whether actually it was the sum of logarithms instead. If that were the case, then your variance would depend differently on the magnitude than it would if the product were the actual controlled variable. (I suspect that the sum of logarithms is a more likely perceptual function than the product, by the way). Having reduced the extrinsic variance so much gives you an opportunity to determine where the rest comes from. Perhaps it is mainly due to a low gain function, perhaps to delays in the feedback loop. You can model that, to see whether it improves the statistics (and I know you have done so, in other experiments). I don't find the use of statistics "throwing in the towel" or a "last resort." I see it as fundamental to finding out what is going on.

Going on backwards...

```
> The tester
>should just be disturbing variables, not trying to prevent the testee from
>controlling the variable. The test should cause no inconvenience for the
>testee -- or higher order systems are likely to become involved and,
```

>ultimately, reorganization. This would mess up the clarity of the results
>of "the test".

Yes, ideally any observer should avoid disturbing the thing observed. But any disturbance to a controlled variable causes error in the controller, even if momentarily. The tester is inevitably controlling for perceiving a change in a variable, or a resistance to change (if the circumstances are appropriate). The tester may relinquish control very quickly ("give up") if the testee shows signs of controlling the disturbed variable, but there has been an effect by the tester on the testee, which has some (we hope vanishingly small) probability of inducing reorganization. And if the testee happened to have been on the verge of reorganizing anyway, a vanishingly small disturbance might tip the balance. When David moves the wastebasket that I had been controlling for seeing under the table, I might "finally" be tipped into deciding that I don't want to use a wastebasket at all. (Some people think that is the case when they see my office, in any case).

>In your example, it looks like Martin is controlling the position of the
>waste basket, then stops. I don't see this as giving up -- it could be just
>a change of reference. There really wasn't much testing going on in
>the example; only one level and type of disturbance, no revision of the
>hypothesis about the CEV when there was lack of disturbance resistance, etc.

Yes, sure it could be a change of reference. I hypothesised that it was a case of "giving up." And your other comments just illustrate the point I made above about the importance of statistics. There are always other disturbances going on, and you will get near-perfect results only if several conditions hold: (1) you have identified correctly the controlled CEV, (2) the controller gain is high compared to the effects of other systems or sources of disturbance that might affect that CEV, (3) the controlling reference level is fixed or changes according to a known pattern. There are probably other conditions, but I can't think of them at the moment. These three are characteristically different in their effects, but failure to meet them cannot be distinguished readily from measurement noise without careful analysis of the statistics, and I'm not sure whether there is any reliable way to do so even then.

Perhaps some of this answers your final comment:

>Perhaps give a concrete example of your idea of using statistics
>as part of "the test".

Does it? Martin

Date: Tue Sep 29, 1992 3:46 pm PST
Subject: Re: Dreams, etc.

[Martin Taylor 920929 16:50] (Rick Marken 920928.1030)

I guess no other fool has rushed into the psychodynamic gritty slimy stuff dumped into CSG-L by our resident loose canon, so maybe I'll try a first slip.

> These people are particularly interested in dreams,
>literature and art as sources of information about our essential selves.

One of my earlier reasons for getting interested in PCT is that it seemed to give a very straightforward explanation of both dream structure and dream content. Since I discovered the degrees-of-freedom argument for the coexistence of massively parallel passive processing and active sequential control processing, the relevance of PCT seems even more pronounced. Dreams, in this view, are not the same as planning processes, but they are closely related. One would expect any organism to dream that based its interaction with the world on a wide-ranging sensory system and a low-bandwidth effector system, and that had a sensory system that was ineffective part of the time.

One assumption is that the ineffectiveness of our main sensory system (vision) at night puts us at risk from more effective night predators ("we" being organisms without sonar or effective night vision). Most sensory systems that work in low light or without light are better at detecting moving objects than stationary ones. It requires either an internal energy source in the target or an energy flood from an outside source like the sun to allow easy pattern detection that allows a predator to identify stationary prey. So we presume that for the most part it is a good thing for day animals to stay still at night (and for complementary reasons for night animals like bats to stay still during the day).

If a living control system must stay still except in an emergency, it makes sense for (a) the muscular output systems to be paralyzed and (b) the sensory systems that can function under those conditions to remain functional while the sensory systems that are minimally functional are closed off. This is one of the characteristics of sleep. The different stages of sleep have different kinds of paralytic effect, but I have no theory about that (nor much knowledge of the different effects).

Even under normal daytime conditions, with all output systems functioning perfectly, most of the potentially controllable percepts are not being actively controlled, because of the resource limitations imposed by the restricted degrees of freedom. So a lot of them are presumably being pseudo-controlled through the imagination loops, in coordination with incoming sensory data (I think it likely that the imagination predicts and tracks percepts as part of the alerting system, but I have no explicit model for how that happens).

When one is asleep, there are no actively controlling systems, at least not for the external musculature, most of the time. But the imagination has no such restrictions. Only if the imagination activates some kind of alerting system (or if the sensory systems not closed off do so) will active control return. And as we all know, both kinds of awakening stimulus occur. It is not noise in itself that awakes us, but noise of prespecified patterns or of unexpected loudness or sharpness. My almost inaudible alarm clock that my wife does not hear beeps me awake, when the loud rumble of a passing truck does not. She might awake to the almost inaudible cry of her baby (many years ago), when I would not hear it even after she nudged me. And for imagination, I think few people sleep to the end of the fall off a dream cliff, unless it ends with them flying away.

What about the content of dreams? Most psychodynamic theories seem to deal with the dream as a kind of problem-solving analogue, and I think the PCT approach would reinforce this view. But in imagination, as in dreams,

impossible things are not prohibited, contradictory things do not necessarily conflict, and the temporal coherence enforced by the real world does not apply. So a dream can be locally coherent (that coherence is enforced by the references supplied from relatively high-level ECSs) but globally nonsensical, as when, for example, one walks along a sand dune into the bar of a hotel that is/was/becomes the bridge of a ship. The high-level content is supplied by high-level reference signals that are the outputs of ECSs experiencing error, perhaps even at the level of principles or system concepts. Other content is supplied in part from sensory data from sound or touch or internal sensing systems (which are not closed off in sleep), or from mid-level ECSs experiencing error.

Why is the content rich? Why do myth and art resonate so strongly? I would assume that the richness depends on the widespread interconnections that make the hierarchy a distributed system. There are lots of things going on in there at any one time, and quite possible several different high-level ECSs with error that produce output to drive them. The resonance of myth and art may signal that lots of us, perhaps by the nature of the natural and social world, experience conflicts of the same kind, and these conflicts can be expressed in many different but equally rich ways. The actual form of expression is a reflection of the particular conflict-induced error in the artist/author, but it also provides the input necessary to reduce the error in the audience, because the underlying hierarchies and "normal" error patterns are alike in artist and audience.

I think this approach could be much improved by careful thought, observation, and experiment, but it appeals to me very much. It makes dreams, imagination, myth, and art a very natural consequence of life. The first two need no explicit communication devices. They are internal, and should be shared by most animals. The second two need communication, and in the case of myth the communication must be linguistic. Myth I would assume to be strictly a human structure, but I'm not at all sure about art.

There is a town in the Black Forest called Freudenstadt. It's a very ordinary small town, and a great traffic bottleneck.

All very slimy speculative, but I like it.

Martin

Date: Tue Sep 29, 1992 4:50 pm PST
Subject: The Test & Laboratory Training

[From: Dennis Delprato (920929)]

Recent discussion of The Test reminds me this could be the subject of a great laboratory exercise. Some of you programmers out there need to work on converting demonstrations to full-fledged labs that can be conducted at perhaps intro. and advanced undergrad. levels. A lot of students at present are working on some dreary labs that are not PCT in orientation. If nothing else, PCT work can be more interesting to students than much of what is done now. A side effect may be that some elect to pursue PCT further.

Any progress on the "Goal-Seeking with Random Consequences of Responses" lab, Rick?

I don't program computers, but I do a pretty good job of systems analysis in the sense of incorporating computer applications with users--as in applications of computers to instructional labs so the entire product hangs together. This includes testing and development of a package. I would be interesting in working with anyone who would like to put together a sequence of PCT labs that could be published and offered to others (lab instructors who would not have to know anything about PCT to use the labs). They would simply be interested in a coherent package that enables them to use the computers they have available. It seems to me that more and more institutions from 2-yr. colleges on up have computer labs, but little interesting and informative software to use.

The initial "product" needn't be an entire term's worth of material. I suspect many instructors would be interested in first using a few labs.

I surmise that one of the appeals of operant psychology was that it was activity-oriented. Students and instructors who neither knew nor cared about the theory initially could actually do something (train rats, pigeons) with results at least interpretable (if erroneously) in terms of the theory. Operant psychology had/has a generally-applicable research preparation. PCT has a generally-applicable research preparation that does not even require nonhumans. Students can use themselves as subjects. One factor working in favor of PCT today is the increasing difficulties instructors are experiencing at using nonhuman subjects. Furthermore, I believe psychology finally has come around to realizing that not very much can be learned by nonhuman research.

Note the advantages of PCT for instruction I've noted in the above brief paragraph: activity-oriented, a generally-applicable research preparation, strictures against nonhuman subjects, psychology has gone human. There are more. The time has come for knowledgeable individuals to exploit advantages such as these. PCT will then be scrutinized eventually on a larger scale and will stand a chance of being more a part of biopsychosocio-science--to rejected, modified, embraced, or whatever.

Dennis Delprato
Dept. of Psychology
Eastern Mich. Univ.
Ypsilanti, MI 48197
Psy_Delprato@emunix.emich.edu

Date: Wed Sep 30, 1992 4:20 am PST
From: CHARLES W. TUCKER
MBX: N050024@univscvm.csd.scarolina.edu
Subject: On behavior of perception

Dear Dag,

I think your effort to spell out the levels is great and I encourage you to continue such an effort. I would like to have any of the charts that you have available. I would also like to have your 920926 post again so I could down

load it and examine it more carefully. I do agree with Bill that you should try to use a real person to examine the levels.

Thanks again for you work on this issue.

Regards, Chuck

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

Date: Wed Sep 30, 1992 4:24 am PST
Subject: HPCT AND OTHER TOPICS

CHUCK TUCKER 920930

TOPICS: HPCT; COIN TEST; UNEXPLAINED PHENOMENA

UNEXPLAINED PHENOMENA WTP 920925.0900

I believe that it is an excellent suggestion for us to write on phenomena that the other "behavioral scientists" can't explain. My major difficulty is finding a way for them to admit that they can't explain it since they refuse to use the notion of explanation and the strict criteria for an adequate explanation that I (should I say WE) do. Clark and I (he more than I) have been trying for nearly 20 years to convince students of "collective behavior" that their theories and explanations and data and everything else are severely inadequate and in fact (or by fact) they have no adequate (even by moderate standards) explanations of civil disorders, riots, demonstrations and the like. So the "trick" is to do this indirectly by showing how PCT explains a particular phenomena and THEN say, "Oh, by the way did you notice that this is a different explanation based on a different model then you have been using for these phenomena, golly gee what a surprise, maybe we have something different here, I wonder what it could be??" [Did you notice that your finger movements spelled CONTROL; did you notice that this model accounts for all those phenomena covered by S-R, Cognitive and Psychoanalytic theories] But unfortunately when you are asking them to publish in their journal (Editors speak of their journal not the journal of the discipline) one has to use such tactics and can't be direct about it. Such is the state of the "behavioral sciences" today. IF YOU DISAGREE AND/OR HAVE SOME OTHER SUGGESTIONS ABOUT THIS MATTER I SURE WOULD APPRECIATE THEM AS I IMAGINE OTHERS WOULD ALSO.

COIN TEST WTP 920925.0900

I tried to image what pattern would be "coins" arranged "All coins exactly where they are and oriented as they are" and the only one I can come up with is one where P is controlling for the relationship of the "coins" to the surface; as long as the "coin" is on the surface then there is "no change". We have tried to eliminate that by having P arrange the "coins" in a pattern

and name it which makes it extremely difficult for P to control for such a reference signal as "relationship to the surface". I have changed the game so that coins are not used to eliminate the "heads and tails" situation, different type coins, different color coins and other such variables. I would also suggest the one demonstrate this game with those you are asking to do it before they do it. In a class I use either an magnetic board with letter "o" or a felt board with disks (such arrangements could also be used for conferences and meetings; one of my students found a computer program [Harvard Graphics?] which draws coins and did the game on the computer) for showing the game. There are ways, in other word, to reduce the variability of patterns for this game. Of course, if you want to introduce a variety of variables that are possible for controlling then make the game more complex.

HPCT WTP 920926.2200

This is a wonderful statement. I hope that everyone will consider it carefully and we can have a serious conversation about it.

Regards,

Chuck

Date: Wed Sep 30, 1992 4:55 am PST

Subject: Postulates

From Greg Williams (920930)

>Dennis Delprato (920925)

>RE: Postulates

>So, what follows are some fundamentals of Skinner's approach to
>matters psychological. I realize that they may not be particularly
>interesting or useful without accompanying text, but I'll bet they
>still function as a disturbance.

> I. Purpose of Science:
> The primary purpose of science is prediction and control.

> II. Methodology:
> The methodology is functional analysis which relates environmental
> independent variables to behavioral dependent variables.

> III. Determinism:
> Behavior is determined; it is lawful.

> IV. Locus of Behavioral Control:
> The causes of behavior are localized in the environment.

> V. Consequential Causality:
> Selection by consequences is the primary causal mode by which
> environment determines outcomes in living systems.

VI. Materialism:

- > Dualism is false. The only world is a physical world.
- > VII. Behavior as Subject Matter:
> The subject matter of psychological science is behavior
> and behavior only.
- > VIII. A. Reeductions:
> The subject matter of psychology is reducible (at least to biology).
- > VIII. B. Nonreductionism:
> Behavior cannot be completely explained in terms of biology or any
> other "lower-level" discipline.
- > IX. Organism as the Locus of Biological Change:
> It is the organism that changes through evolutionary and environmental
> histories, and the changes are biological.
- > X. Classification of Behavior into Respondent and Operant:
> There are two major classes of behavior, or more completely,
> functional relations: respondent and operant.
- > XI. Stimulus Control of Operant Behavior:
> Operant behavior can be brought under the control of antecedent
> stimuli, and description of operant behavior usually requires three
> elementary terms and their functional interrelationships.
- > XII. On The Generality of Behavioral Principles:
> The full complexities of human activity--including language, thinking,
> consciousness, and science--are behavior to which all the above apply.

>Bill Powers (920925.2000)

>That's a really grown-up job. It's the sort of analysis we need done
>on PCT. Why not give it a whack yourself?

Since Dennis has not (so far) given it a whack, the following is a tentative presentation of PCT postulates in the same format in which Dennis presented Skinner's postulates. I'll bet they function as a disturbance.

I. Purpose of Science:

The primary purpose of science is explanation.

II. Methodology:

The methodology is modeling which relates co-determining environmental variables and organismic variables.

III. Determinism:

Behavior (controlling for an outcome) and action (output) are determined; they are lawful. The path of reorganization is not determined; it is not lawful.

IV. Locus of Behavioral Control:

The causes of behavior (controlling for outcomes) are localized in the organism. The causes of action (output) are found in the organism and the environment.

V. Consequential Causality:

Negative-feedback control is the primary mode by which organisms determine outcomes.

VI. Materialism:

Dualism is false. The only world is a modeled "physical world."

VII. Behavior as Subject Matter:

The subject matter of psychological science is behavior (whether observed "externally" or hypothesized "internally") and behavior only.

VIII. A. Reeductions:

The phenomena of psychology are explainable in terms of physiological phenomena AND particular observations at the level of psychological phenomena.

VIII. B. Nonreductionism:

Behavior cannot be completely explained in terms of biology or any other "lower-level" discipline.

IX. Organism as the Locus of Biological Change:

It is the organism that changes through evolutionary and environmental histories, and the changes are biological.

X. Classification of Behavior:

There is one major class of behavior: control of outcomes.

XI. Antecedent Influence on Behavior (Control of Outcomes):

Control of outcomes depends on, but is not fully determined by, antecedent environmental events within an organism's lifetime and in evolution, and explanation of such behavior requires modeling of organismic reference signals.

XII. On The Generality of Behavioral Principles:

The full complexities of human activity -- including language, thinking, consciousness, and science -- are behavior to which all the above apply.

Now, how about somebody else trying to put the postulates of mainstream cognitive scientists in the same format?

Best, Greg

Date: Wed Sep 30, 1992 5:48 am PST

Subject: Why 99%?

From Greg Williams (920930 - 2) >Martin Taylor 920929 16:00

>My presumption is that you get the 99% prediction because the subsystems
>(perhaps ECSs) that are involved in the task are those that support very many
>different kinds of behaviour, and so are not readily disturbed by contextual
>differences.

I believe that the highly precise predictions which have been achieved by some PCT tracking models are due to the condition of "good" control being satisfied in the experiments to which the predictions are applied. That is, the subject

is able to track close to the target (because the disturbances are "easy" to compensate for). ANY negative-feedback model with a reasonable frequency response would predict tracking close to the target, yielding highly precise predictions. But this is misleading, because the important question (if you already believe in PCT) is how to choose BETWEEN different negative-feedback models, and the only way to do this WITH HIGH PRECISION is by looking at transient (temporarily "poor") control. The "good" control case data simply don't allow sufficient sensitivity for comparing different models. But apply some square-wave disturbances, and you'll see immediately whether, say, a proportional-integral-derivative linear negative-feedback tracking model with PARTICULAR values of parameters gives more precise predictions than the same kind of model with DIFFERENT parameter values or than other kinds of models with various parameters.

Best, Greg

Date: Wed Sep 30, 1992 7:02 am PST
Subject: Postulates

[From: Dennis Delprato (920930)] (Greg Williams):

Thank you for your statement of PCT's fundamental postulates.

Your expression is such that each of Skinner's has a corresponding one in PCT. Another consideration is to identify what is most central (what points) to PCT and begin from here. The ones I extracted from Skinner's written corpus were those that in my evaluation of the data appeared most central; furthermore, some of his close followers made suggestions that I took into account.

Either way, I think it is important to "get the cards on the table" in matters scientific, thus exposing fundamental (often hidden) assumptions to closer scrutiny.

I made a typo in VIII A.: Reeducations should be Reductionism. I might note here that one outcome of my efforts was to show quite a bit of coherence in Skinner's assumptive framework (in the absence of careful scrutiny). His advocates liked this. What they did not like was his conflicting statements on the reductionism-nonreductionism question. I clearly exposed this striking incoherence (through his most recent writing). Advocates were displeased--more with him than with me, for the data were undeniable--he was inconsistent.

Date: Wed Sep 30, 1992 9:26 am PST
Subject: Dreaming as ref. sig. calibration

From Greg Williams (920930 - 3) >Martin Taylor 920929 16:50

>One of my earlier reasons for getting interested in PCT is that it seemed
>to give a very straightforward explanation of both dream structure and
>dream content.

A few years back, I had the idea that dreaming might serve to recalibrate/fine-tune our control structures. After much perusal of the (enormous!) literature on sleep (especially regarding the effects of sleep deprivation), I gave a presentation on this at the 1990 CSG meeting. Unfortunately, I made the mistake of thinking that the "drift" needing to be offset by dreaming is in

the gain of the input function -- but in fact that DOESN'T matter appreciably if control is good. More recently, Rick Marken (in posts on CSGnet) pointed out that there might be drift (i.e., due to purely local biochemical fluctuations) in reference signals which might need correcting. Dreaming might provide an opportunity for correcting such "additive" (as opposed to "multiplicative" -- gain) drifts: the output side is shut down, so the person doesn't get hurt, then "imaginary" disturbances WHICH ARE NOT DUE TO INDEPENDENT ENVIRONMENTAL EVENTS, BUT CAN BE KNOWN TO THE ORGANISM are applied to various perceptual input pathways, allowing reference signals to be checked to see whether they "correctly" offset the disturbances (and allowing correction if necessary).

There's still a problem with this idea, at least at the low levels. Some researchers have reported "no" deficits in fine motor control (i.e., playing ping-pong) even after hours and hours of sleep deprivation. But maybe ongoing correction, while awake, of these low-level processes is easily achieved, while higher-level processes need dreaming? There is a lot of evidence that long-term sleep deprivation can result in bizarre kinds of behavior. And dreams are often "about" high-level reference levels -- maybe.

>What about the content of dreams? Most psychodynamic theories seem to deal
>with the dream as a kind of problem-solving analogue, and I think the PCT
>approach would reinforce this view.

Maybe some of this is going on, as well. Bill in BCP raised this possibility.

Slipperily (on purpose, this time),

Greg

Date: Wed Sep 30, 1992 10:04 am PST

Subject: skinner and reduction

[From Francisco Arocha, 920930;13:38] (Dennis Delprato, 920925)

> VIII. A. Reeductions:
> The subject matter of psychology is reducible (at least to biology).
>
> VIII. B. Nonreductionism:
> Behavior cannot be completely explained in terms of biology or any
> other "lower-level" discipline.

and (920930)

> His advocates liked this. What they did not like was his
> conflicting statements on the reductionism-nonreductionism question.
> I clearly exposed this striking incoherence (through his most recent
> writing). Advocates were displeased--more with him than with me, for
> the data were undeniable--he was inconsistent.

These two statements are not necessarily contradictory. This is simply because there is more than one kind of reductionism. The most radical being physicalism (everything is reducible to physics), but there are more moderate forms, such as emergent reductionism, which involves partial reduction, not

total, because of the acknowledgement that matter (at any given level) has properties that lower forms do not possess. The way you phrased VIII.B: 'cannot be COMPLETELY explained...' is coherent with this more moderate version of reductionism. So you could always explain higher levels by using lower levels, with properties. But, this may not be the way Skinner thought about it, in which case, as you said, he was inconsistent.

Salut, Francisco

Date: Wed Sep 30, 1992 10:17 am PST
Subject: Re: The Test in practice, Why 99%

[From Rick Marken (920930.1100)] Martin Taylor (920929 16:00) --

>Since I am at the moment being asked to think of PCT-driven
>experiments in planning and decisions-making tasks, I may be able to satisfy
>condition (2) in that context.

I am starting to do something similar with satellite control tasks. I don't know how it's going to work out but here's my thinking:

1) identify what seem to be the (perceptual) variables that need to be controlled in the task. For example, in a satellite contact one variable that must be controlled is "degree of connectivity". This is a complicated variable whose value depends on all kinds of things; antenna availability, location, status of processing equipment, etc. Of course, during the pass this variable should be in a state very close to "connected".

2) identify factors that influence the states of these variables; distinguish external (disturbance) factors from operator produced factors.

3) test the definitions of the variables in 1 by doing the test; I hope we will be able to run simulations where, say, the operator is asked to set up for a pass and will have various ways of influencing connectivity -- while (simulated) external factors also affect connectivity (antennae go down, demuxes become unavailable, etc).

4) try to map variables in terms of being means or ends relative to each other. This can be done by asking experts "how do you do this; why do you do that" where doing is expressed in terms of the perceptual variables identified in 3.

This is just talking through the ol' hat at the moment; if the group I'm working with makes any progress I will keep you posted.

>As you see from the above, "inferential statistics" is not what I am talking
>about. What I am talking about would be more along the lines of parameter
>estimation for the perceptual functions, gain estimates with variance, and
>the like.

Then we are in furious agreement. I love it.

Dennis Delprato (920929)--

>Recent discussion of The Test reminds me this could be the
>subject of a great laboratory exercise.

Yes.

>Any progress on the "Goal-Seeking with Random Consequences
>of Responses" lab, Rick?

Nope. My kid took my computer so I can only work on it at work -- but that would be wrong (to quote that great moral philosopher, R. M. Nixon).

But I think I would like to do the following: I would like to write the lab in Basic. I would suggest that the lab be a study of "reinforcement" and discriminative stimulus theory. Have them collect data from a couple subjects; then have them try to explain the results. They could choose from several models in the computer. Then they could compare the behavior of the model to that of the subject. The goal of the lab is to help the student understand that the concept of reinforcement implies that there is a controlled perceptual variable -- in this case, what is controlled is the relative position of the moving dot and one of the stationary dots on the screen.

I think it could be fun -- especially if you have a condition which lets the experimenter manipulate the direction of the movement of the dot after each press; once the experimenter has discovered the controlled variable, he/she can then "control" the subject's bar pressing rate by making the dot tend to move away from the target dot after each press. So you can show that this kind of "control by reinforcement" is a consequence of the disturbance resistance characteristics of a control system.

Greg Williams (920930 - 2) --

>I believe that the highly precise predictions which have been achieved by
some

>PCT tracking models are due to the condition of "good" control being
>satisfied in the experiments to which the predictions are applied.

> ANY negative-feedback model with a reasonable
>frequency response would predict tracking close to the target, yielding
highly
>precise predictions. But this is misleading, because the important question
>(if you already believe in PCT) is how to choose BETWEEN different negative-
>feedback models, and the only way to do this WITH HIGH PRECISION is by
looking
>at transient (temporarily "poor") control.

I agree. It is somewhat disingenuous of me to posture about the 99% accuracy predictions when, in fact, they are virtually guaranteed by the fact that we are dealing with situations where there is clearly control and we have a damn good notion of the controlled variable. I do think that the improvements in prediction that I spoke of are good indications of an improved description of the controlled variable. But I'll try to lighten up the bragging about 99% prediction. The reason I have harped on it is because I want PCT to avoid succumbing to the statistical "cop out" which, to me, means that you do a study (in our case, test for a controlled variable, say), find noisy results (the subject resists the disturbance on x % of the trials) and INSTEAD OF

TRYING TO GET A BETTER DEFINITION OF THE CONTROLLED VARIABLE, you apply a statistical test and say "the effect of the disturbance was significant so I conclude that I have found a controlled variable". I want PCT to avoid this approach in favor of working hard (and it will be hard -- doing the test for the controlled variable correctly will be a lot harder than doing an ANOVA correctly; the test can be done "cook book" -- at least not yet) to improve one's results (with 99% accuracy as the goal) by trying to find improved definitions of the controlled variable and improved ways of testing it. I hope someone knows how to do this; I sure don't.

Best regards Rick

Date: Wed Sep 30, 1992 4:56 pm PST
Subject: Instructional Labs

[From: Dennis Delprato (920930)] (From Rick Marken (920930.1100))

>>Dennis Delprato (920929)--

>>Any progress on the "Goal-Seeking with Random Consequences
>>of Responses" lab, Rick?

>Nope. My kid took my computer so I can only work on it at work -- but that
>would be wrong (to quote that great moral philosopher, R. M. Nixon).

I take that this also means you are "not a crook." Furthermore, I'll trust you with my audiotapes any time.

>But I think I would like to do the following: I would like to write
>the lab in Basic. I would suggest that the lab be a study of
>"reinforcement" and discriminative stimulus theory. Have them collect
>data from a couple subjects; then have them try to explain the
>results. They could choose from several models in the computer. Then
>they could compare the behavior of the model to that of the subject.
>The goal of the lab is to help the student understand that the concept
>of reinforcement implies that there is a controlled perceptual variable --
>in this case, what is controlled is the relative position of the moving
>dot and one of the stationary dots on the screen.

>I think it could be fun -- especially if you have a condition which
>lets the experimenter manipulate the direction of the movement of the
>dot after each press; once the experimenter has discovered the
>controlled variable, he/she can then "control" the subject's bar
>pressing rate by making the dot tend to move away from the target
>dot after each press. So you can show that this kind of "control by rein-
>forcement" is a consequence of the disturbance resistance characteristics
>of a control system.

Go for it. But, to me this is a step above the present version of the lab based on your E. coli study. The "Goal-Seeking With Random Consequences of Responses" lab might best precede the one you have just proposed. By way, I have access to several IBM compatibles, as well as to the Macs.

I do feel that for purposes of instruction, it is a good policy to keep things as simple as possible so as not to obscure the really big points of PCT with matters that are best considered by the more advanced student. Bill's Demo 1 and Demo 2 (not currently in a form suitable for labs) nicely illustrate a gradual build up in complexity.

Dennis Delprato
Dept. of Psychology
Eastern Mich. Univ.
Ypsilanti, MI 48197
Psy_Delprato@emunix.emich.edu