

Date: Wed Jul 01, 1992 5:46 am PST  
Subject: Narcissus is us; learning

[From: Bruce Nevin (Wed 920601 08:59:24)]

The mirroring around level 0 reminds us to what extent apparent structure in the environment is a reflection of structure in the control hierarchy, projected there by the observer. We assume vice versa, but that can only be an assumption. (Right, Wayne?)

So it is worse than perhaps Martin has said. Not only may the observer (the investigator) identify the wrong environmental variable V as that controlled by the observed control system, the two parties may also have differently structured control hierarchies, and so may parse the environment differently into environmental variables.

We have to assume some commonality as a working hypothesis, in order to proceed at all. We must not forget that this is an assumption, which itself must be tested so soon as some other hypotheses have been tested and found stable in its context. Start with a temporary scaffold (I assume you and I perceive the world alike; or I assume that there is a like world for each of us, however we perceive it). In its context, propose some environmental variables (as perceived by oneself) as perceptions controlled by the other. Having attained some good results and some less good results, go back and try to get more uniformly good results by replacing parts of the scaffold. Question whether the world as one perceives it is identical in a few particulars to the world as the other perceives it.

Nagel somewhere has a famous image for science, of being in a storm at sea on a boat that we must take apart and rebuild as we are sailing.

Learning:

Seems to me introspectively that often learning and reorganization involve lowering the gain on control of perceptions that I had been controlling, and listening for the voice of the next-best candidate in pandaemonium. Control at some higher level n stays the same; presumptions about the (lower level) means for controlling at level n are relaxed. Alternative means are tried out. But the alternatives are not random. They are the next-most clamorous contenders. An analogy might be drawn to recessive traits in the gene pool.

Bruce           bn@bbn.com

Date: Wed Jul 01, 1992 7:34 am PST  
Subject: System Concept, Dictionary

[From Hank Folsom 920701]

Bill Powers (920629.0930) said:

>It isn't really necessary to educate everyone about control theory. If HPCT  
>is right, everyone has the necessary levels of control and perception.

>What's needed is a set of experiences in which they can see the basis for  
>justice and constitutional rights.

Quite true. However, You can lead a horse to water, but you can't make him drink. is a very appropriate Category 1A observation. Just because people become aware of PCT does not mean that they will control for the greater good. The best we can hope for is that in a PCT world, more people will have System Concepts that will include the concept of the world being a better place.

The reason I proposed looking at four categories of people was to make it easier to target those most likely to be open to and buy into PCT.

Category 1A people would seem ready for PCT, because the concepts should not generate large error signals for them. On the other hand, since they already intuitively are living as though they understood the positive implications of PCT, they presumably are controlling effectively, and have no need to change.

There are a lot of potential Category 3's out there who have a vested interest in the status quo. Because they are controlling for no change, PCT says there is little likelihood they will accept a new way of thinking. Psychologists are the most interesting example. They have much to gain from PCT. However, PCT says they will have the greatest error signal generated by PCT. As has been said many times, PCT, is a threat to their System Concept. The only approachable psychologists are those who have large error signals within their specialities in psychology. Does this view sound reasonable to the psychologists on the net?

There are close to 5,000,000,000 Category 1 people out there who could be taught about PCT and System Concepts. Those with large error signals are probably the most open to learning about PCT.

The best potential markets for a PCT revolution are among those groups who would see PCT as a tool, not a threat. These are people who have goals that can be more easily achieved by applying an understanding of PCT. Any suggestions?

PCT Dictionary:

We Category 2's are the most difficult bunch to work with on Systems Concepts. We can & will flop back into Category 1 in mid-statement, with no warning at all. It is very difficult to keep alert for changes in another's viewpoint, especially when we respect their knowledge and ability in PCT.

I feel a need for a PC Dictionary (Perceptually Correct?). Dictionaries have different definitions for different worlds. e.g.: Post: 1. In computer usage, an electronic message. 2. In farming, what holds up a fence. (The consequences of a threat to hit you with a post depends on who the speaker is. Interestingly, most prefer the farmer.)

Is there any interest on the net in developing a list of words that mean different things in the S-R world and the PCT world, or do not even translate? For starters, how about:

Behavior  
Control  
Purpose  
Response  
Stimulus

If the list is of workable length, and I suspect it is, the next step would be the PCT and S-R definitions. The final step would be the hardest: How to communicate without running afoul of the different meanings.

Hank Folson

Henry James Bicycles, Inc. 704 Elvira Avenue, Redondo Beach, CA90277  
310-540-1552 (Day & Eve.) MCI MAIL: 509-6370 Internet: 5096370@MCIMAIL.COM

Date: Wed Jul 01, 1992 6:22 pm PST  
Subject: Martin's D of F paper

From Bill Cunningham 920701.2020

(Martin Taylor 920701.0240)

The net sure has been silent today. Everyone must be digesting the two long and very thoughtful posts from Bill P and Martin.

I find Martin's degrees of freedom argument a very powerful answer to my doubts and fears. The paper is of fundamental importance for my application.

The similarity detection/difference discussion is right on target. It was that Kanerva algorithm was a similarity detector, but John Gabriel reminded me that it can work either way. Does that help with the notion of a switch hitting ECS?

Previous past net discussion on "subjective probability" had bothered me with by its extension into type 1 and 2 decision errors, particularly the consequence of false alarms wherein a control system would lock onto the wrong percept and not let go. A switched control system with an alert mechanism would remove the false alarm problem as a deterrent to survival in nature, allowing concurrent offense and defense. I guess track-while-scan radars sorta fit this description.

Martin doesn't state so explicitly, but he has presented a case for a control system for the control system (regardless of how implemented). I have argued for a long time that investment in the control of information flow will pay larger dividends compared to generating, passing or processing more data. The degrees of freedom argument puts some meat on those bones.

A final thought on hyperlexia. That discussion immediately conjured up my exposure to high speed Morse code operation. As a teenager, I participated in a high speed plain text practice net. Somewhere between 30 and 40 words a minute, plain text ceases to be dits and dahs or even letters. Plain text starts sounding like spoken language, to the point that typing or writing becomes wasteful relative to notetaking. I got so I could carry on a normal conversation at 40 wpm, limited more by the keyers of the day. However, I couldn't do squat with 5-letter code groups.

No redundancy for error correction, just the damned letters. However, real Morse operators hear only the letters. I've seen them continue to type for 4-5 minutes after cessation of 60 wpm code groups--with zero errors. One of the folks I try to sell HPCT is trying to find out what makes good and bad operators. Maybe there's a clue here.

Martin's paper does leave me uncomfortable with the issue of how one ECS wrests control from another. Not the idea, but the mechanism, stability and sharing issues. Any ideas?

Dunno about the rest of you, but I'm still digesting.

Bill C.

Date: Wed Jul 01, 1992 6:41 pm PST  
Subject: Re: System Concept, Dictionary

[Martin Taylor 920701 14:00]  
(Hank Folsom 920701)

A dictionary seems like a good idea as a background document to be developed and amended, and to be kept by Gary. Isn't there such a thing?

I have a similar problem with my Layered Protocol model of communication. The model provides clean distinctions among concepts that map approximately, but not precisely, onto the standard linguistic concepts "lexicon" "syntax" "semantics" "pragmatics" and the like. I want to use those words, because I think LP theory makes them useful (for decades, the distinctions among the last three have seemed very fuzzy to me). But if I use them to a linguist they will be misunderstood, and will be misunderstood differently depending on which flavour the linguist prefers.

To make a PCT dictionary presupposes that the listener/reader has the proper appreciation of PCT. Otherwise the definitions will make no sense. So the dictionary would be of more use to CSG-L contributors than to neophytes.

Martin

Date: Wed Jul 01, 1992 6:43 pm PST  
Subject: DoF posting--copyright correction

[Martin Taylor 920701]

Bill Cunningham brought to my notice that the copyright wording I added to my long posting of 2:40 this morning could be interpreted as prohibiting copying the message to interested parties outside CSG-L. That was not my intent. I added the notice because I had an idea that some refined version of the posting might one day become a paper, and some journals have strong views on prior publication. Accordingly, I should like to amend the copyright notice so as to permit personal circulation but not public dissemination outside CSG-L. I might ask that anyone who does circulate it personally let me know, and if possible that they might pass back comments they might receive on it.

Martin Taylor

Date: Wed Jul 01, 1992 7:41 pm PST  
Subject: Re: Reorganization

[Martin Taylor 920701 22:40]  
(Bill Powers 920630.2000)

Sorry to post another near 400-line posting in the same day. But here it is, anyway.

>Well, good. It was beginning to look as if we would soon run out of things  
>to disagree about.

Oh, I wouldn't worry too much about that. We can always talk about statistics. But I'd like to get as big as possible foundation of things we agree solidly about, to move on to more subtle issues. What I would like to sort out about PCT is what must be true, what cannot be true, and what may be true if we allow some other assumptions. It's what you called "truthsaying," I think. No disagreement in the first two classes should be allowed to stand, at least not in the foundational structure.

This is a long comment on your long posting about my discussion of learning. For the most part I agree with what you say. You write so clearly that it is possible also to find where and why I disagree with some bits. At the end of your posting, you say:

>The one point on which we apparently have a significant  
>difference is on the relationship of intrinsic variables to what is  
>learned. I hope I have explained more clearly what I mean by intrinsic  
>variables and just why I think that there must be NO relationship to the  
>learned control systems, save for the one imposed by the natural  
>environment that relates the state of the world to the basic intrinsic  
>state of the organism itself. It would seem that you have not understood my  
>terms here; perhaps now you do.

I don't think I had misunderstood your terms, but I did not have as clear an idea of how you see the reorganization system as I now do. I stated the intrinsic variables as being those dealing with body chemistry. That was sloppy. I think you put it much better:

>By "intrinsic," I mean that  
>these variables have nothing to do with anything going on outside the  
>organism. They are concerned with the state of the organism itself. They  
>might BE variables like blood pressure and CO2 tension, or they might be  
>signals, biochemical or even neural, standing for such variables.  
>...  
>For each intrinsic variable or signal, there is some state that is at the  
>center of a range that is acceptable for life to continue. That is a  
>\_definition\_ of an intrinsic variable and its reference level in this  
>context, that separates such a variable from all the other variables in a  
>living system.

Nevertheless, your description describes more or less what I had in mind.

You presented very clearly why you think that there should be a reorganizing system outside the sensory-motor control hierarchy, and why the intrinsic variables should be represented only in that separate system. But (so far) I don't buy the argument as truthsaying, and I am not at all sure that the end product of the argument, your separate reorganization system, is even plausible. Maybe you can convince me.

>[Intrinsic variables]

>are variables in the system that, if maintained within certain limits  
>(Ashby's way of seeing it) or close to certain reference levels (my way),  
>would assure or at least promote a viable organism.

>

>It was then only a short step to realizing that if a hierarchy of control  
>were ever to come into being, this process of reorganization had to be  
>operational from birth, and most likely from early in gestation. This meant  
>that it could not use any perceptions of intensities, sensations,  
>configurations, transitions ... system concepts, as the nervous system  
>would come to perceive such things, before the ability to perceive (and, I  
>would now add, control) such things had been developed. This immediately  
>ruled out any principle of reorganization that uses any familiar  
>perceptions or means of control; particularly, any programs, principles, or  
>system concepts. Those things would EMERGE from reorganization; they could  
>not cause it. There could be no reorganizing algorithm.

Total agreement so far, except for a possible question about what you mean by "algorithm." Surely for reorganization to occur, it must follow some discernable rules? You describe such rules later (random reconnection, for example), so you must mean something different from what that sentence seems to mean.

>Now the outlines of a reorganizing system begin to appear. When intrinsic  
>variables depart from their normal reference levels, something is seriously  
>wrong; survival is in question. This is "pain,"  
>generalized. If the organism is to survive, it must do something.

>

>But in the beginning, it doesn't know how to do ANYTHING. It has no  
>conception even of an external world. It has no idea of how that external  
>world bears on its well-being. It has no knowledge of how to affect that  
>external world even if it has the capacity to do so. Therefore, we have to  
>conclude, whatever is done about the intrinsic error, it must be done at  
>random -- without any systematic relationship to the outside world.

>

>We can debate, of course, just how much of a head-start evolution actually  
>provides for this process. In lower organisms, it's quite a large amount.  
>But I wanted to solve the worst case because that would establish an  
>important principle; any organization capable of working in the worst case  
>would naturally work even better with a head start. So we can ignore that  
>consideration.

Again, total agreement, except that I might be prepared to argue that in higher organisms, evolution provides even more of a head start. But going for the worst case is fine. We started there. I made that presumption in the posting that evoked yours, by questioning where the first ECSs came from.

>In my model of the reorganizing system, I posit intrinsic reference signals  
>and a comparator for each one. This is a metaphor; in fact, all we need  
>assume is that there is some reference state established by inheritance,  
>and that deviations of the intrinsic variables from their respective  
>reference levels lead to reorganization. We don't need to guess at the  
>mechanisms or even the locations of these processes -- that kind of  
>question can be answered only by a kind of data that nobody knows how to  
>obtain yet. We can only speak of functions, not about how they are carried  
>out. A control-system diagram illustrates the functions and how they must  
>be related.

Fair enough, but the connotations here are beginning to bring us onto  
treacherous ground. You are beginning to assume that the control systems for  
intrinsic variables are eventually going to be found to be organized in a  
system separate from the system that interacts with the outer world. You  
haven't said it yet, but the end of the paragraph is stated in a way that  
leads one's thinking in that direction.

Here, I'd like to do a little extrapolation of the discussion before rejoining  
yours. Let us assume that there develops (exists?) a control hierarchy for  
intrinsic variables. Does this not have a formal structure like that of the  
familiar sensory-motor hierarchy, in that higher-level perceptions based on  
the values of intrinsic variables are controlled through reference levels  
supplied to lower level EICSS (Elementary Intrinsic Control Systems)? If so,  
then are these higher-level EICSS not controlling percepts that correspond to  
some Complex Internal Variable (CIV), just as higher level ECSS control  
perceptions of Complex External (Environmental) Variables? If there exists  
such a hierarchy how is it reorganized? If no such hierarchy exists, how are  
errors in particular intrinsic variables going to be organized so that they  
target parts of the sensory-motor control hierarchy relevant to them? More on  
this after the next segment of your posting.

>The outputs of the reorganizing system change neural connections, both as  
>to existence and as to weights. They cause no neural signals in themselves;  
>they merely change the properties of the neural nets. In doing so, they can  
>connect sensory inputs to motor outputs, and thus, in the presence of  
>stimulation, create a motor response to a sensory stimulus. They don't  
>create any particular response; they only establish a functional connection  
>so that motor output bears a relationship to sensory input according to the  
>amount of input, should any such inputs occur.

>

>[...big chunk omitted]

>The state of the world affects intrinsic variables.

>

>Therefore if reorganization results in stabilizing certain aspects of the  
>external world against disturbance, and brings those aspects to specific  
>states, the result may be -- MAY be -- to bring some intrinsic variables  
>closer to their reference states. This is purely a side-effect of what the  
>new control system is doing. What the control system is sensing and  
>controlling may have nothing directly to do with the side-effect that is  
>changing an intrinsic variable. But if the result of sensing and  
>controlling in that way is to lessen intrinsic error, reorganization will  
>slow or even cease. And that control system will continue in existence.

I assume that the beginning of this passage refers to the creation of an ECS, since the provision of an S-R connection would normally be useless. But there is a third connection to be made if this is so--to the reference input of the new ECS. Where does that come from? Is this new ECS a top-level one, in which case the reference signal is (so far) undetermined? or is it interpolated in an existing hierarchy?

>[The intrinsic control system is] a system that may involve hundreds of  
>intrinsic variables with specific reference levels, and hundreds or  
>thousands of pathways that connect error signals (if signals they be) to  
>the target locations for reorganization.

In light of your insistence on rigorous ignorance and random operation, that word "target" is interesting. I can see that if a sensory-motor hierarchy exists, but is flailing about because all its connections are as yet randomly made, then a random connection of the outputs of the reorganizing system would (eventually) develop stable control, and the reorganizing system would then be found to be targeted at relevant locations in the sensory-motor hierarchy. But an organism that contained such a powerful sensory-motor hierarchy initially would probably be dead long before the reorganization had taken useful effect. So one must assume that infants do not have a well developed control hierarchy, or that if they do, then the gain on most ECSs is near zero. It is not by accident that the young of all species are either incompetent or already organized with an effective sensory-motor control hierarchy. If they do not have the control hierarchy, then the reorganizing system must develop new ECSs, with the issues as raised above. If they do have an existing hierarchy, then the reorganizing system cannot be randomly linked to it. So, your argument seems to lead to the situation that you don't want to allow: the reorganizing system DOES know about specific aspects of the control hierarchy, whether the two hierarchies are separate or no, or even whether the reorganizing system is a hierarchy at all.

>Now we can see the basic logic of reorganization. The reorganizing system  
>is not concerned at all with what control systems exist or what variables  
>in the outside world are brought under control. All it is concerned with is  
>keeping intrinsic variables near their reference levels. If there is  
>intrinsic error, reorganization commences, with the result that perceived  
>variables are redefined, sensitivities and calibrations change, means of  
>control change, and the external world is stabilized in a new state. The  
>only significance of that new state to the reorganizing system is that  
>intrinsic error may be corrected, putting an end, for a while, to  
>reorganization.

OK. This paragraph rings true, taken on its own. But it does not need the underlying idea that the reorganizing system is separate from the sensory-motor hierarchy. If you change the words "intrinsic variables" for "controlled percepts" in this paragraph, very little need change. Indeed, in the parts quoted before, intrinsic variables ARE controlled percepts, but not in the sensory-motor hierarchy. They are controlled through modifications of the organism's internal environment, not the external world.

>Given a reorganizing system that works this way, an organism can learn to  
>survive in environments having almost completely arbitrary properties. A  
>pigeon with such a reorganizing system can come to maintain its internal  
>nutritional state near an inherited reference level by pecking on keys



>instead of grain, or even by walking in figure-eight patterns -- it can  
>learn to control variables that have absolutely no natural, logical, or  
>rational connection to nutrition, and by doing do, can protect its own  
>nutritional state. It does not need to reason out why performing such acts  
>is so vital, or even what the connection is with getting food. It doesn't  
>even have to know that ingesting little bits of brown stuff has the effect  
>of keeping it from starving.

Agreed. Important.

>If you recall my posts on the origins of life a year or so ago, you'll  
>realize that this process of reorganization exemplifies exactly the same  
>process that I proposed as the way the first living molecules came into  
>being. The main difference is that the variability that creates new  
>organizations to be retained or evolved away comes not from external forces  
>but from an active process of random change driven by internal error  
>signals. The reorganizing system is evolution internalized and made  
>purposeful.

You mean this? (Powers 910731, quoting me in the first paragraph)

<>What I was trying to get across was that if part of a self-organized  
<>feedback system happened one day to evolve so that its self-corrective  
<>feedback was modified in response to some environmental disturbance that  
<>it could not previously survive, then it would be more likely to exist  
<>into a farther future. It would also be a rudimentary control system,  
<>with an externally settable reference. That reference would itself be  
<>part of a non-control-system stabilized structure, but could later  
<>become incorporated in a higher-level control system that might evolve.  
<>Then we would have a two-level hierarchic control system. That, in its  
<>turn, could evolve a higher layer, and each such layer would contribute  
<>tho the apparent stability of the entire system. But always at the top  
<>there would be a non-control-system feedback complex of some degree of  
<>apparent stability.

<

<This is very close to a proposition about the origins of life that I  
<posted before you got onto the net. I started a little farther back than  
<you do. Suppose you have a chemical reaction going on that is forming  
<complex molecules, and that these molecules, during breakdown, interact  
<with their substrate so as to influence concentrations of chemicals \*on  
<which formation of that kind of molecule depends\*. This is feedback.  
<Obviously, NEGATIVE feedback would be highly favored; modifications of  
<the complex molecules that resulted in negative feedback effects on the  
<replication-critical substances would lead to increased relative  
<concentrations of those molecules. Where feedback is positive, that  
<population of molecules would quickly disappear (changes in the critical  
<substances would be amplified instead of opposed). This is strictly  
<Darwinian evolution; nothing fancy.

<

<[....]

<

<There would, of course, be an unstoppable tendency for this kind of  
<negative feedback to become more and more effective and thus more and  
<more prevalent. The appearance of catalysts, enzymes, introduces  
<amplification that vastly improves the tightness of feedback control.

<Somewhere in here, before or after the enzymes appear, there must also be  
 <the first appearance of reorganization (and here, Prigogine's concepts  
 <may glancingly intersect with mine). The system, which must be complex by  
 <now, becomes capable of reacting to chronic error by \*causing\* random  
 <changes in the molecular structure, or the structure of molecular  
 <relationships. The changes are random, but the selection process is not:  
 <the rate of random change drops to zero if and only if the error is  
 <corrected by the new relationship of the molecule/structure to the  
 <substrate environment. So we have the effect of directed evolution  
 <without any telology and without any external direction. This introduces  
 <a principle of evolutionary progress that Stephen J. Gould would hate:  
 <evolution plus blind variation and selective systematic retention must  
 <tend toward tighter and tighter control, and greater and greater  
 <resistance to external events that tend to affect the accuracy of  
 <replication. This gives us an evolutionary scale on which to compare  
 <organisms.

<

<I can now tack your paragraph above onto the end of my exposition, as the  
 <story of what happens next (actually your story and mine probably  
 <overlap). Once we have negative feedback, and amplification with enzymes  
 <and later with neurons, and the capacity to create internal error-driven  
 <blind variation of organization, we have the ingredients for a system  
 <that can add levels of control whenever that is the only solution to  
 <error-correction that is left. And I think we end up with a very pretty  
 <picture of the whole sweep of evolutionary history from soup to nuts like  
 <us.

Back to the present posting:

>If there's one primary concept that must be understood to understand my  
 >theory of reorganization, it's that the variables controlled by this  
 >process are completely apart from the variables represented as perceptions  
 >in the learned hierarchy. The learned hierarchy is concerned primarily with  
 >sensory data about an outside world, and about those aspects of physiology  
 >that are represented in the sensory world. The reorganizing system is  
 >concerned about variables in the world beyond the senses -- with the actual  
 >state of the physical organism at levels inaccessible to the central  
 >nervous system.

"Beyond the senses" is perhaps misleading. I would substitute "beyond the  
 externally-directed senses. Clearly if intrinsic error is to be detected,  
 something must detect the state of the intrinsic variable.

This brings up a little point about reality perception. A few weeks ago we  
 had a little discussion (which I'm not going to seek out in my Hypercard  
 stacks) about the difference between haptic and tactile perception, in which  
 it was pointed out that when touching is being done by the subject rather than  
 to the subject, the subject perceived an object instead of touch sensations.  
 The control, and in particular the internal kinaesthetic sensations were  
 required if the subject were to perceive an external object. We also  
 discussed emotion. It seemed likely that internally generated sensations had  
 to be used in combination with sensations obtained through the external  
 sensors to create a stable emotional percept.

I do not find it plausible that internal states are unrepresented in the

controlled percepts of the sensory-motor hierarchy. In fact, I think that if they were not, we would have much more difficulty with real-world control. And I find it entirely plausible that there should be ECSs within the sensory-motor hierarchy that derive their perceptions entirely from internal states, including "intrinsic variables."

>It is not necessary to learn from experience  
>that error signals represent some degree of failure to control; large error  
>signals indicate serious problems, no matter what perception they relate  
>to. The reorganizing system could monitor error signals in general, en  
>masse, without any need to know what they mean, and their mere presence at  
>large magnitudes could be sufficient to cause reorganization to start. This  
>would satisfy the requirement that intrinsic variables be inheritable. And  
>one result would be that loss of control could lead to highly localized  
>reorganization, precisely in the system that has lost control.

Yes, to the first part of the paragraph. But doesn't the second part begin to suggest that Occam's razor is a bit blunt? You say that you have held in reserve the idea that reorganization might be a local ability of an ECS, any ECS. If sustained error is to be a localizable trigger for a separate reorganization system, isn't it simpler just to allow any ECS with sustained error to do its own reorganization by changing some signs at its output, or perhaps seeking new structural links? We know that new dendritic connections develop as a consequence of rich active experience in rats, so the generation of new links from local activity is not absurd neurologically.

Allowing this, it seems to me that the passage quoted from last July leads naturally to the presumption that the intrinsic variables do contribute to the perceptions in the sensory-motor hierarchy, that errors in them are strong determiners of the need for reorganization, and that a single mechanism rather than a duple one is adequate for the task.

The single localized mechanism has an added advantage to the system as a whole, not apparent at first glance. If, as you suggest, and as I think is thermodynamically essential, there is a threshold of error below which the probability of reorganization is essentially zero, then the fully reorganized hierarchy will be in a critical state of subdued tension, with minor conflicts at all levels. It will, in this state, be prepared for rapid reaction to disturbances in any of its controlled variables. Any greater internal conflict would lead to further reorganization, any less would lead to slower reactions to disturbances.

As a final note, we did agree some months ago that reorganization must at least be modular. This conclusion was based on an argument from degrees of freedom, and would apply to any mechanism for random restructuring. The underlying fact is that in a high-dimensional space almost all directions are nearly orthogonal to any selected direction. If the actual linkage pattern for optimal control of a particular percept represents a particular direction in the space described by the signs of the output linkages of all the ECSs contributing to that percept, then almost any random reorganization will have a negligible effect if control is poor, but a devastating effect if control is good. Only by reorganizing in a low-dimensional space can the e-coli effect be put to good use.

So, if reorganization is to be done by an ignorant system external to the

sensory-motor hierarchy, it should be done at any one time only in small localized groups of links, probably relating to a single ECS or a small group of ECSs controlling a very similar percept. Also, presumably, low-level ECSs that are able (usually) to provide percepts that match their references are better supports for the reorganization of a higher-level one than are incompetent low-level ECSs. One would expect low-level ECSs to settle down to a stable, productive, life long before the higher-level ones that depend on them. That doesn't mean they don't change over time, but they will usually change smoothly, and more probably in their perceptual functions than in their output connections.

-----  
I'm sure you will disagree with a lot of the above, and probably you will think I have missed your main point. Maybe I have, but it feels right, and it feels as if everything you wrote made sense, based on an unnecessary underlying presumption.

Is this enough disagreement to make you happy? Long as this comment is, it mostly does agree with what you wrote. But not all.

Martin

Date: Wed Jul 01, 1992 7:41 pm PST  
Subject: Levels, gain, and emotions

[From Kent McClelland 920630.1230]

I've been working on a revision of the long ms. distributed last year on the net, and I'm finding that some current issues discussed on the net match some of the concerns I've run into in revising the ms.

First, a possibly trivial detail: what is the currently "canonical" version of HPCT levels 1 to 11?

Bill Powers last year at the CSG meeting in Durango suggested a change in his thinking about the ordering of levels, and this possible reordering was subsequently discussed on the net, I believe. If my memory serves me correctly, Bill considered switching the order of the configuration and transition levels, arguing that configurations might be constructed out of transitions, rather than the other way around. Thus, transitions would be seen as level three, after intensities and sensations, and configurations as level four.

Bill, what is your current thinking on this? What do other netwatchers think? (What IS the appropriate term for those of us stare regularly at the messages flitting the screen?) I don't remember what the result of the discussion on the net was, so perhaps it was inconclusive. I realize that the order and even definition of levels is not a matter of great moment at this stage in the development of HPCT, but I'd just as soon keep up to date in things I hope to put in print.

Next, the subject of loop gain has been raised again in recent posts (Martin Taylor, 920701 0240 and Bruce Nevin, Wed 920601 08:59:24. [I guess that probably should be 960701, but, as Rick knows, it's hard keeping months

straight in the summertime!) and I have some questions about the concept.

1. How should a non-engineer think about the concept of gain? Is the gain of an ECS located only in the output function of loop, as Rick Marken seems to indicate in the definition he offers in his "elephant" paper? Or is gain better conceptualized as a product of functions involving input, output, and any environmental transformations (cf., Bill 's definition in his 1979 Byte article on the "Nature of Robots" Part 2?)

Parenthetically, I note that if gain is defined as in the second def, then. . .

Machines in the environment can improve gain. (e.g., eyeglasses, power tools.)

Attention may improve gain (by somehow jacking up the input function).

Increased "effort" (whatever that may be in a PCT definition) is also an attempt to improve gain.

Improvement in gain seems more or less synonymous with economist's concept of "efficiency."

Does all this sound plausible?

2. How might loop gains be regulated?

Martin, in the post mentioned above, suggests that "A monitoring ECS can become passive simply by reducing its output gain to a negligible level. . ." And further on he talks about an ECS that could have "a gain function that stiffens with increasing error. . ." Bruce, in his recent post suggests that "often learning and reorganization involve lowering the gain on control of perceptions. . ." If I'm recollecting correctly, Bill has suggested that reorganization may involve a process of selecting the right gain. But what exactly is doing the regulation of gain? And how does gain get switched on and off?

I'm interested in the subject of gain partly because of Rick Marken's intriguing suggestion that positive emotions are associated with increases in loop gain. As I've indicated before on the net (McClelland, 920322), I think control theorists haven't paid enough attention to the subject of emotion, and how it fits with memory, learning, and decision making. I've managed recently to read Bill's new LCS II collection of essays, and found in the essay on emotion (one of two "lost" chapters from Behavior: The Control of Perception, the other one being a good description of the therapeutic "going up a level" idea) Bill's especially interesting suggestion that feelings are in part perceptions of intrinsic error signals, messages from the body about how it's doing, and thus that learning and emotions must be closely related. Of course, as Bill has just pointed out in his recent post on learning (920630.2000), intrinsic error signals may be of many other kinds besides those perceived as feelings, but he describes his original view as focusing on "pleasure and pain." Do you still think, Bill, that the "affect" part of emotion can be equated with intrinsic error signals?

I don't know quite how it all fits together, but I think that one could make

a compellingly plausible account of the dynamics of the hierarchical system's operation by identifying negative emotions as the perception (overlaid by higher-level cognitive-social interpretations) of the intrinsic error signals arising from persistent errors in control loops, and by identifying positive emotions as the perceptions of decreasing error that come with increases in loop gain through successful reorganization or just the improvements in control that result from practice, i.e., "honing" of the reference value in memory. Memory, I think, is what has been neglected so far in this picture. I would stipulate that the memories used as reference values are often closely associated with emotional memories (memories of the perceptions felt as emotions), and that that association (a "ping" of emotion) experienced when the memory/reference signal is tried out in imagination often guides our decision making processes. We shy away from reference signals remembered as aversive and select reference signals remembered as pleasant, unless a higher-level loop "rationally" over-rides the choice.

---

Some other random comments:

I've found there's lots of good stuff in the LCS II volume. (Nice job, Bill, Greg, and Tom!). The chapter I found most pertinent to my sociological interests was a 1980 working paper called "CT Psychology and Social Organizations." This, it says, was written for, or possibly with, D.T. Campbell as "theoretical background" for his "investigations of the uses of social indicators in judging social programs." Rather than rehashing the basics of control theory, the paper lays out a few basic principals and then goes on to a very perceptive discussion of how PCT ideas can be used to analyze social structures, covering such topics as conflict and competition, alignment of goals, specialization, cooperation, coordination, and management. I think it ought to be required reading for anyone who is trying to work out the implications of PCT for sociology or social psychology. Does anybody know if Campbell actually used this paper in some sort of program evaluation or based any of his work on it? The paper in LCS II called "Control Theory for Sociology" was less meaty, mainly because it was mostly given over to another one of Bill's endless series of attempts to introduce control theory to new audiences. I can see why he's getting tired of these gigs after 20 years!

I've also been reading Hugh Petrie's book on the Meno paradox (The Dilemma of Enquiry and Learning, University of Chicago Press, 1981) which Bill referred to in one of his recent posts (920629.1600). The book presents a very nicely laid-out philosophical case for control theory and is well worth reading. Linguists on the net might find it interesting, because he has some things to say about Chomsky's views.

Bruce Nevin (Wed 920601 08:59:24) has again brought up the issue of intersubjectivity and social construction of a sense of reality. In various discussions on the net of these issues, I haven't seen any reference to the works of the little band of maverick sociologists called ethnomethodologists, who actually have some useful perspectives on these problems. The head honcho of the group is Harold Garfinkel (Studies in Ethnomethodology, Prentice Hall, 1967), but beginners might want to start with one of the following:

Mehan, Hugh, and Houston Wood. 1975. The Reality of Ethnomethodology. New York: Wiley.

Benson, Douglas, and John A. Hughes. 1983. The Perspective of Ethnomethodology. New York: Longman.

Back to my ms!

Best regards, Kent

Kent McClelland	Office: 515-269-3134
Assoc. Prof. of Sociology	Home: 515-236-7002
Grinnell College	Bitnet: mcclel@grin1
Grinnell, IA 50112-0810	Internet: mcclel@ac.grin.edu

Date: Wed Jul 01, 1992 7:49 pm PST  
Subject: Re: Martin's D of F paper

[Martin Taylor 920701 10:20]  
(Bill Cunningham 920701.2020)

Kind words. Thank you.

>Martin's paper does leave me uncomfortable with the issue of how one ECS  
>wrests control from another. Not the idea, but the mechanism, stability  
>and sharing issues. Any ideas?

I was basing my words on the ideas in what I actually presented at the Madrid AGARD meeting, and described in a posting around July 7. The idea is that if two ECSs conflict, the higher-gain one will have a smaller error. If an ECS has been controlling reasonably well, but suddenly finds that it can no longer do so because its corrective actions are resisted, it has sufficient information to determine that some other controller is trying to control the same CEV (Complex Environmental Variable). It may be a characteristic of some (all?) ECSs that they relinquish control by reducing their gain under such circumstances. An alternate response would be to increase gain, and try to beat the interloper.

Both strategies can be observed in conversation. If one person interrupts another, the first may relinquish control or may try to override the interruption. Usually, one or the other backs down. So it may be with ECSs suddenly subject to competition, and in the case of the aircraft, I suggested that such a mechanism (relinquishing, that is, not fighting back) be deliberately built in to automated subsystems.

Martin

Date: Wed Jul 01, 1992 7:49 pm PST  
Subject: learning and knowledge

[from Joel Judd 920630]

Martin (920729 & 0630) and Bill (920629 & 0701),

Sorry to assume too much, but I didn't want to write too much in the original post:

>Knowledge can certainly be represented, but interactions do not require the kind of representation >that implies regress.

Right; \*transmission\* requires the kind of representation that involves regress. Transmitting encodings requires that both receiver and sender understand what the encodings represent.

I also agree with your description of ECS environmental interaction, and what it means for an ECS to "know" what's going on. The latter part of my comments referred to how an ECS knows that it's getting "better"; how does it "learn"?

With reference to education, I guess the crude way to express it is to say that in many cases the whole is more than the sum of its parts. A house is NOT just the things you mentioned--it's an organized combination of them. The contractor knows what that means. The rookie apprentice does not. Sure, he's seen pictures and been inside other houses and watched movies of construction, etc., but he's never actually made one. Not having the details of BLC theory handy, I would say that not even point (2) can be taught. All three points require learner experience; the teacher can only evaluate evidence of learner.

I guess I should be happy that interactivism (i.e. organism interacting w/ environment) as the basis for learning is apparently assumed by PCTers--it certainly isn't by all educators. Bill mentioned the dilemma I was alluding to--Meno--which Hugh Petrie attempts to resolve in his book. Again, I agree that the solution necessitates action, blind action, but that's not what the transmission analogy implies. It implies that one can learn without making mistakes; BVSR says in effect that we learn FROM our mistakes (hence the title of Perkinson's book). However, it's interesting that many learning "errors" NEVER occur. One of the standard lines in language acquisition goes something like "But how do you account for the fact that children never say XYZ?" I suppose that a PCT response (assuming that "never" turns out to be accurate) would be along the lines of "Well, at such and such a point in development, the particular PC hierarchy is such that XYZ is not the kind of random variation a normal system will produce; there are constraints on the blind variation."

So instead of continuing to look for more molehills, can I say that the following approximates a PCT epistemology:

New organisms are born with genetically transferred intrinsic variables. These allow the organism to interact, from the start, with its environment in order to preserve itself. Disturbances to these variables provoke reorganization in an effort to reduce error and maintain reference levels necessary to life. In a given environment, however, the actions commenced by reorganization involve (of necessity) the sensory-motor capabilities and limitations of the organism. Perceptions which reduce intrinsic error sufficiently are remembered, and remain in place until or unless they fail to reduce future intrinsic error, in which case reorganization recommences. As sensory-motor systems unfold in their genetically pre-determined manner,



reorganization finds newer, more efficient and more sophisticated ways of satisfying intrinsic error--perhaps even anticipating such error through memories of past experiences, vicarious learning, and extrapolation. By the time a human being is just a few years old, such massive foundational learning has taken place in the sensory-motor systems (visual-aural discrimination, tactile development, figure-ground, conservation, etc.) that future learning often takes for granted such development, and /or assumes that it must be "hardwired." Adding to the confusion is the seeming convergence in development by members of the same culture and linguistic community--barring damage or abnormality EVERYONE born in particular group will grow up able to function as a member of that group.

But, after a certain "maturation" point, this reorganization system becomes somewhat suspect. For example, in the case of language, after about 6-7 years--certainly after puberty--the results of reorganization in a L2 become highly variable. There is no longer the convergence showed by children in an L1, now all kinds of perceptions seem to satisfy the intrinsic error of these adult systems.

Am I just preaching to the converted?

P.S. Martin--do you know something about Canada we don't? Are things that bad?

Date: Wed Jul 01, 1992 7:58 pm PST  
Subject: Similarities and differences

[From Bill Powers (920701.1300)] Martin Taylor (920701.0240) --

In trying to sort out the roles of ECSs in various modes -- passive observation, active control, and model-based "shadowing" -- I think you're exploring interesting territory. But I think your thesis concerning "similarities and differences" is a step backward.

I'm sure that these terms have been in your mind for a long time and that they've settled down to some meanings that are perfectly clear to you. It isn't clear to me, however, that either "similarity" or "difference" describes a basic feature of perception. I don't doubt that one can point out similarities and differences between different objects of perception, but I don't think that the labels describe what's going on. I just finished paying attention for 15 minutes or so as I made and ate lunch, and didn't find any occasion to notice either a similarity or a difference. Of course if you had pointed out similarities and differences to me I would probably have agreed with you that they seem to exist. But they're so arbitrary! And they simply don't seem to dominate experience as you claim they do (unless you go into a mode where you're actively looking for them).

How is a locomotive similar to a diamond ring? They're both expensive. How is one window in this room different from the window next to it? One is to my left, the other is to my right. The idea of a generalized similarity or difference detector seems to me impractical, because there are simply too many ways to decide that perceptions are either similar to or different from other perceptions --or both at the same time.

By examining any pair of perceptions, we can always find a multitude of

ways in which they're different, and another multitude of ways in which they're similar. To see how two perceptions are similar, we look for higher-level perceptions that can arise from either of them and are in fact the same perception. And to see how they are different, we simply note lower-level attributes that change as we transfer attention from one to the other. Two chairs can be similar in that both give rise to the perception "chair," a category, even though at the same time they may belong to disjoint categories such as reclining and electric. In terms of the simple category chair they are identical. But when we transfer our attention or gaze from one to the other and at a lower level of perception, we experience a change of configuration or sensation. So we can say they are different, and point to the attributes that are not identical in the two chairs.

I think, in fact, that differences are probably just transition perceptions, and that we experience differences primarily in terms of changes of configuration. It's difficult to state the difference between two events such as bouncing and exploding, or between two relationships such as above and between, or between dogs and automobiles, or between falling and rotating, or between honesty and persistence (to sample the higher levels). This is because they are above the transition level where we notice change. The best we can do is say that they're not the same perceptions (and logically, therefore, are "different" even though we can see no basis for the difference).

If you could characterize similarity detectors or difference detectors in some way beside just saying that they detect similarities and differences, perhaps I could come closer to understanding what you mean. Or to skip to the real issue, just what do you mean by similarities and differences?

I also have trouble with ECS's wresting control from each other. I think I have to ask you to diagram this process and explain how it works. I think some of the notions in your treatise are carrying us dangerously close to the modern Scholasticism that infects so much of academia (as some of my own recent pronouncements have been doing). We're drifting into a mode of discourse that is far from the spirit of modeling and experimentation, the approach in which models aren't just POSSIBLE explanations of experimental results, but are the ONLY PLAUSIBLE explanations, with no serious rivals or alternatives.

Maybe I'm just suffering a backlash from my own long spew about reorganization that kept me up past my bedtime last night. I woke up this morning feeling very dissatisfied with what's going on on the net. Your post hit me when I was already in that state. I suddenly got a picture of a bunch of old men (and young people trying to imitate old ones) sitting around a table comfortably debating about angels and pinheads, trying to solve the riddles of the universe by clever manipulations of words. Thinking about Brooks' subsumption architecture, I realized that one motive behind my criticisms is simple jealousy: I wish I were building little robots and making them actually do things that are interesting. I wish I were doing real experiments to test real hypothesis instead of just sitting here and writing and writing and writing. I have a longing to be doing something REAL. I want some meat to get my teeth into. RAW RED MEAT.

I need another vacation from retirement.

Grumpily, Bill P.

Date: Thu Jul 02, 1992 7:52 am PST  
Subject: Re: Similarities and differences

[Martin Taylor 920702 11:00]  
(Bill Powers 920701 13:00)

Two points, that I will try to make brief, for a change. Similarity and difference detection, and mode of discourse.

I think you misinterpret the similarity and difference detectors as relating CEVs currently affecting sensor systems. That's not what I meant. In pre-PCT terms, they refer to similarity to or difference from some template. I look at the "template" now as being a reference level. A similarity detector has a gain function concave upward (e.g. gain = error to a power greater than unity), and may well have zero gain for some finite level of error. A difference detector has a gain function that has some appreciable slope near zero error. There is a second kind of similarity detector, which is not an ECS so far as I can see, but a perceptual function akin to what neural network people call a "radial basis function". It emits a larger signal the nearer the incoming perceptual pattern matches its "template." It could well be in an ECS whose reference level is near zero, causing an error when something like the target is in the sensory data stream.

Does that help? As for whether it is a scholastic point: no it isn't, because regardless of your theoretical viewpoint, it nonetheless happens that tasks that depend on similarity give different results than tasks that depend on difference (precision). Failure to notice which is important in a particular task has led to some pretty sterile arguments among "scholasts" of different schools.

On modes of discourse, I thoroughly agree with your wish to do reality testing. It is so clear that working in the imagination mode permits one to entertain contradictions that would be soon exposed when tested. But we can't always "go real" because of resource limitations, as I discussed in the "similarity difference" posting. But we can work out the implications of our models to the extent possible, using what you called the truthsaying approach. What MUST be true? What CANNOT be true? We may be wrong in our analyses, but they do help to show where reality testing might be fruitful.

I have no resources to devote to reality testing, other than the contract to Chris Love, of which the Little Baby is a small part (and he finishes the contract fairly soon). So perforce I devote the time to consideration of the implications of fundamental truths applied to situations that might become real. Then I look in the natural world to see if there is anything that applies. The similarity-difference issue is a case in point. There are a bunch of strange phenomena out there, which become perfectly natural and obvious once one realizes that PCT must be true and that there are more degrees of freedom for sensor systems than for external output (joints, and shape-changing effects like facial expression). Similarity and difference are a natural consequence of resource limitation in a control system. They have been observed as puzzling phenomena. Now they are not. Is that "modern Scholasticism?"

Martin

Date: Thu Jul 02, 1992 8:20 am PST  
Subject: Similarity-difference

[From Bill Powers (920702.1000)]

Martin Taylor (920702.11:00\_ --

Tell you what. How about giving us a reference to the similarity-difference studies that we can look at? I want to see what kind of facts we're talking about.

If you're really talking about forms of comparison functions, we're into a different subject.

Best, Bill P.

Date: Thu Jul 02, 1992 8:28 am PST  
Subject: Reorganization

[From Bill Powers (920702.0800)]

Martin Taylor et. al. (920701) --

OK, lots of disagreements here to keep me happy. As I suspected, the basic concept of the reorganizing system as I propose it is VERY hard to grasp -- it's hard even to see how such a thing could work. Martin Taylor: "I am not at all sure that the end product of the argument, your separate reorganization system, is even plausible."

How could a system that is not AT ALL concerned with the form of behavior end up forming behavior? The problem is very much like trying to understand how a system concerned with outcomes could be unconcerned with producing specific outputs that produce just that outcome. Everybody on the net understands control of outcomes through variations in action that are systematically related to disturbances but not to the controlled outcome. Now the problem is to see how an outcome could be controlled by variations in action that are not related to ANYTHING.

Let's re-examine the lessons of E. coli. Inside E. coli there is a control system (this is my way to model it, anyway) that senses a time-rate of change of concentration of some substance. The perceived rate of change, now just a chemical signal inside the bacterium, is compared with a reference rate of change, another signal. The difference, the error signal, acts on an output function, as usual.

This is, however, a very peculiar output function. What it does is not to cause a systematic and appropriate change of direction, but periodically to create a tumble, a random change in orientation of the body of the bacterium. The tumbles themselves have no systematic effect on orientation. They simply create new directions of swimming, at random.

Clearly it's the direction of swimming that systematically determines whether the bacterium heads up or down a gradient of attractant. But the bacterium's error signal does not operate on direction of swimming. It operates on THE INTERVAL BETWEEN TUMBLES. Big error --> shorten the interval; small error --> lengthen the interval. The final result is that the bacterium proceeds quite efficiently up the gradient, in a series of zig-zags that has longer legs in right directions and shorter legs in wrong directions.

BUT THIS CONTROL SYSTEM KNOWS NOTHING ABOUT CONTROLLING DIRECTIONS OF MOVEMENT IN SPACE. It does not sense direction of movement. It senses only the rate at which a concentration changes. Koshland, in perfusion experiments, showed that the same effect on intervals between tumbles is obtained by varying the concentration in a flowing medium while the bacterium is tethered to a sticky substrate in a fixed orientation.

So the bacterium is sensing one variable, concentration change, by reorganizing a different effect: speed of swimming times the cosine of the angle between direction of travel and chemical gradient. It is not systematically affecting direction of travel, but only adjusting the interval between mutations in the direction of travel. At no time does it have any information about which way it is swimming in space. It is controlling a SIDE-EFFECT of direction of swimming -- time rate of change of concentration of an attractant. From the viewpoint of the human observer, and Koshland in particular, this bacterium knows how to "navigate" in three-dimensional space. But that is entirely the wrong interpretation. This bacterium knows nothing of three-dimensional space. Its world has no extension beyond its own membrane. As far as the control system in the bacterium is concerned, outputting an error signal causes the time rate of change of concentration to remain more or less at the right level, and that's all that's happening. The location of the body in space or its actual direction of movement is outside the world of this control system.

Let's take a small step upward in complexity. Suppose that we now have a microorganism that can actually steer: it can sense the direction of the gradient relative to its body. It compares this sensed direction with a reference signal of zero (i.e., there isn't any reference signal) and converts the error to a change in direction of swimming (perhaps by varying the speed of movement of ciliae on different sides of its body). If a positive error signal causes one side to speed up and the other to slow down, the body will turn toward the gradient and align to swim exactly up it; if the sign of the effect is reversed, the body will turn down the gradient and align in the direction down the gradient. If the loop gain is set too high in either direction, control will turn into oscillation, so progress will cease either up or down the gradient. If this is a 3-D situation, there will be two control systems of this kind. We need to consider only one of them.

This chemical gradient could be one of a substance that is helpful or harmful to the bacterium. In other words, PRESENCE of one substance would have some deleterious effect on the inner workings of the bacterium, while LACK of the other would have a deleterious effect on the same thing.

Let's now give this organism a simple reorganizing system. The intrinsic variable is the concentration of some substance inside the bacterium that is an indicator of health. As long as this substance is near a particular reference level, all is well. If the surroundings are conducive to health but at a level lower than the

optimum, in terms of this indicator, the organism should steer up the gradient if it knows what's good for it (the reorganizing system does).

But let's suppose that the output gain, the effect of the error signal on differential speed of the ciliae, starts at zero. Now, save for good luck, there will be either an excess of substances that are deleterious to health or a shortage of those that are conducive to health. In either case, the indicator, the intrinsic variable, will depart from its reference level. Reorganization will commence.

In this case, we assume that reorganization can affect only the loop gain. The loop gain can be increased or decreased by some amount  $\delta$  through variations in output sensitivity, one factor in loop gain. Changing enzyme concentrations could alter output sensitivity and hence loop gain. At intervals, the size of  $\delta$  is varied randomly within some small range between positive and negative limits and is added to the current loop gain. So the loop gain may increase by a small amount with each episode of reorganization, or it may decrease.

If an increase in loop gain causes even more error in the reorganizing system, or fails to decrease it, the next random increment/decrement of loop gain will occur sooner or at least no later. If the error is lessened, the next change will be postponed a while. The result will be that the loop gain will perform a biased random walk in the direction that lessens intrinsic error. Tom Bourbon has demonstrated that this does in fact work.

If the SIGN of the loop gain is positive, the organism will swim directly up the gradient. Simply reversing the sign will cause the organism to swim directly down the gradient (because the relationship between changes in direction and changes in error reverses).

The result will be not only that the SIGN of the output sensitivity will be correct for swimming up a gradient of beneficial substances or down a gradient of noxious substances, but the size of the loop gain will come to the maximum at which stable control of direction still exists.

If this directional control behavior succeeds in keeping the level of beneficial substances sufficiently high, or the level of noxious ones sufficiently low, the reorganizing system will go to sleep -- it will detect no error, and the rate of reorganization will drop to zero. From then on, the control system will automatically swim up or down the gradient and continue to avoid noxious substances or seek beneficial ones (one or the other, but not both, in this very simple case). It will follow changes in direction of the gradient in space without any further modification, resisting disturbances that might make its path deviate from the right direction.

In this example there's a direct relationship between the chemical gradient sensed by the steering control system and the effect on the intrinsic variable. The intrinsic variable that indicates state of health is affected by the same chemical substance that is used for steering. In fact, it is the reorganizing system that gives "value" to the substance being sensed and controlled by the direction-control system. The direction-control system is not concerned with the meaning of the chemical signals indicating direction errors: it will just as happily steer the organism up or down the gradient, depending on the sign of its loop gain. It's the reorganizing system that decides that this substance is

noxious or beneficial, in terms of effects on some intrinsic variable that is an indicator of the state of the organism.

But even that is misleading. The reorganizing system doesn't put a label on the substance used for steering. It simply monitors the EFFECT of that substance on the organism itself, and if there's intrinsic error it fiddles with the loop gain until the intrinsic error goes away. The result is that the organism either seeks or avoids that substance, according to whether the loop gain ended up positive or negative. We, looking at the final outcome, decide that this organism likes or dislikes the substance, judging by the fact that it seeks or avoids it.

With only a small further change we can further generalize this example. Suppose that this organism steers not by sensing chemical gradients but by sensing light intensity impinging on receptors on either side of its body. Depending on the sign of the loop gain, it will swim either toward or away from the light.

Suppose also that in the medium there are noxious substances that have a higher concentration where there is light, and that recombine to harmless forms in the dark. Now the behavioral controlled variable is differential left-right (-up-down) light intensity, but the basis for reorganization is some effect of a noxious substance on an internal indicator of health, an intrinsic variable that has nothing to do with sensed light intensity.

Now reorganization will create a negative loop gain and the organism will avoid light. The control system in charge of steering knows nothing about noxious substances, and the reorganizing system knows nothing about light. Yet the overall effect is that the control system avoids light and thereby keeps noxious substances from changing the state of the intrinsic variable that is the basis for reorganization. The organism has adapted to a photochemical phenomenon that is entirely outside its ken. From the observer's point of view, the organism has come to assign a negative value to light, judging from the fact that it avoids light.

Suppose that something in the environment now gradually changes, so that the physical situation reverses: now noxious substances form in the dark, and are dissociated into harmless components by light. The organism will find itself swimming into trouble by avoiding light. The noxious substances will cause the indicator of health, the intrinsic variable, to depart from its reference level, and reorganization of loop gain will commence. It will cease only when the loop gain has become optimally positive, for that will result in the organism's seeking instead of avoiding light. The steering control system will now remove the organism from the concentrations of noxious substances -- without in the slightest intending to do so.

I hope that this sequence heralds a dawning a little nearer in the future. Martin Taylor says:

>Fair enough, but the connotations here are beginning to bring us onto treacherous >ground. You are beginning to assume that the control systems for intrinsic >variables are eventually going to be found to be organized in a system separate >from the system that interacts with the outer world. You haven't said it yet, but >the end of the paragraph is stated in a way that leads one's thinking in that >direction.

Now I have said it much more explicitly. This is precisely what I am proposing.

-----  
I am concerned about some sloppiness that is showing up in various ideas proposed for how ECSs do things. We have ECSs that can alter their own loop gain, that that communicate with and control other ECSs at the same level depending on the situation, and that can voluntarily "relinquish control." This and other such proposals have the effect of putting a great deal of additional function into an "E"CS. This is all right with me IF IT'S BACKED UP BY A TESTABLE AND DEMONSTRABLY WORKABLE MODEL SHOWING HOW EACH PROPOSED NEW FUNCTION WORKS.

That requirement is sort of being ignored. I get sucked into doing it myself, which is among my main reasons for raising a red flag. New ideas are a dime a dozen, but new ideas that have been worked out to the level of modeling aren't. And I don't mean just a schematic diagram -- I mean something you could demonstrate in a computer (even a very simple version of it), just to show that it would actually do what you say it would do. Not everything you draw will actually behave the way you think it will.

What would be necessary in an ECS that could "relinquish control" all by itself? It would have to detect the conditions under which this is required -- what kind of detectors, sensing what? How much computation, and what kind, is needed to recognize the conditions? What kind of actuators would it need, driven by what, acting on what? Many proposals that sound simple when expressed in a few words turn out to entail far more complexities than I, at least, would like to see in an ECS -- some of them require whole hierarchies of control!

Of course our conjectures are always ahead of what we can actually support; I'm not trying to put a damper on creativity. But if we get too far ahead of ourselves, we'll fall back into that Scholastic mode in which all you have to do is think of something verbally plausible while the possibility of actually proving that your idea would work goes down the drain (in other words, standard psychology). That I am vastly uninterested in.

I'll get around to answering some other posts later -- if I ignore anyone it isn't out of disinterest, but due to shortage of time. Ask again if the point hasn't been covered.

Best, Bill P.

Date: Thu Jul 02, 1992 9:02 am PST  
Subject: learning?

Did anyone respond to my question a few days ago on learning? I have been unsuccessful in connecting with the University's news net so if there was a reply I missed it. I would appreciate any replies sent to my internet address.

Mark

Educational Psychology 210  
College of Education  
Univ of Illinois at Urbana-Champaign  
phone: (home) 351-8257  
(office) 244-8080

USmail: 405 South 6th St. #4  
Champaign, IL 61820  
e-mail: (Internet) m-olson@uiuc.edu  
(Bitnet) FREE0850@uiucvmd



Date: Thu Jul 02, 1992 11:57 am PST  
Subject: Re: Reorganization

[Martin Taylor 920702 13:30]  
(Bill Powers 920702.0800)

Bill, you have presented another clear exposition on the aspects of reorganization with which we are both (all?) familiar. But you haven't dealt with the primary objection to your proposal, which has nothing whatever to do with whether the reorganizing system/principle/method "knows" anything about what it reorganizes. Everything you write about it in this posting speaks to that non-issue.

I said that I didn't think that what you proposed was plausible for reasons that I hoped were clear. I tried to specify the area of agreement and the area of problems, but apparently it didn't get across. Sorry to be obtuse.

Rather than go through it again, let me just put the main objection very simply. The bacterium whose control has three degrees of freedom can reorganize with no trouble. In 3-space, a substantial portion of the randomly chosen directions are "near" (say, within 60 degrees of) a preselected direction. In higher-dimensional spaces, the proportion of directions near the selected direction is very small.

If we call the  $i$ 'th ECS at level  $L$   $ECS(n,i)$ , there is a vector with elements  $link(ECS(n,i),ECS(n-1,k))$  where  $link(p,q)$  is 1, 0, or -1, depending on the existence and sign of the connection between  $p$  and  $q$  (or takes a real value, in the general case).

Taking the simple case in which  $link()$  can be 1, 0 or -1, a random reorganization has a probability 0.33... of doing the right thing if there is only one dimension, 0.111... in two dimensions, and  $(1/3)^n$  in  $n$  dimensions. If there are three ECSs in each of two layers, that is roughly a one in twenty-thousand chance. All other sets of connections lead to some conflict, and even if we grant the probability that some other sets provide stability, the odds are not good that global reorganization by non-targeted random alteration of link sign will reach an optimum quickly. The problem is the same as that of molecular evolution as seen by the creationists. You can't do it that way. You have to grow stably. And that, it seems to me, means targeted reorganization. But not reorganization in which the reorganizing system "knows" what or why it is reorganizing, in the sense of "hunger means arrange to get something to eat." We agree that that is ordinarily nonsense (caveat: not nonsense under "teaching").

Now think what "quickly" means in respect of reorganization. It means survival. If a critical feedback loop has positive feedback, the organism may die if the problem (THAT problem) is not corrected quickly. But the reorganizing system (we agree) does not know what THAT problem is. All it can know is that there is error, and it can know that the error is manifest in a particular ECS. Now, the dimensionality of the output connection set for that ECS is much smaller than for the net as a whole. It seems reasonable to suppose that the positive feedback could be corrected more easily and quickly in the low-dimensional space of one ECS than in the high-dimensional space of the hierarchy. A blunt instrument to do it might be to change the sign of the comparator output (i.e.

reverse all the link signs at once). Any such reorganization changes the demands on (reference signals supplied to) some lower-level ECSs, and may result in them developing uncorrectable errors, and thus reorganizing. Or it may not. The system is blind to that. But each change is of low dimensionality, and thus feasible to achieve through random processes.

>I am concerned about some sloppiness that is showing up in various ideas  
>proposed for how ECSs do things. We have ECSs that can alter their own loop  
>gain, that that communicate with and control other ECSs at the same level  
>depending on the situation, and that can voluntarily "relinquish control."  
>This and other such proposals have the effect of putting a great deal of  
>additional function into an "E"CS. This is all right with me IF IT'S BACKED  
>UP BY A TESTABLE AND DEMONSTRABLY WORKABLE MODEL SHOWING HOW EACH PROPOSED  
>NEW FUNCTION WORKS.

>

>That requirement is sort of being ignored.

Agreed. I've worried a bit about that, too. But is there any evidence that the neural system is built of repetitive structures like an ECS at all? ECSs are very nice units for modelling, and make good predictions of actual behaviour in simple situations when coupled appropriately. They are functionally very effective, they implement a principle that has to be true, and they make much of the net behaviour simple and intuitive to understand. My hunch is that the real structures are much more distributed, and that we will rarely if ever be able to identify a functional piece of the brain and say "here's a classic scalar ECS."

In my mind, the situation is rather like that of the computational linguists who think they have sets of rules that determine legal and illegal strings in a natural language. The rules work well for a central core of language, but do not capture its fluidity. Likewise, I suspect that the simple ECS hierarchy will capture behaviour in many situations, but will be only an outline of what happens in the real world.

My worry is that it is too easy to solve problems by giving this cartoon-ish entity, the ECS, the intrinsic ability needed to solve each particular problem. In the "similarity-difference" paper (and in the AGARD extension) I proposed that control could pass between two ECSs, and tried to use non-formal language to describe it, so that the functional need for such a possibility could be separated from speculation about mechanism. (I re-posted a possible mechanism last night, which I hope will answer your concern about one ECS "wrestling" control from another--the functional but informal description of what must happen).

>Not everything you draw will actually behave the way you think it will.

How true!

>What would be necessary in an ECS that could "relinquish control" all by  
>itself? It would have to detect the conditions under which this is required  
>-- what kind of detectors, sensing what? How much computation, and what  
>kind, is needed to recognize the conditions? What kind of actuators would  
>it need, driven by what, acting on what? Many proposals that sound simple  
>when expressed in a few words turn out to entail far more complexities than  
>I, at least, would like to see in an ECS -- some of them require whole

>hierarchies of control!

Good questions, all. And the last point about hierarchies internal to an ECS does ring a bell when I think of all the hierarchies that exist within a single cell. Let's give a try at some kind of answer. Tentative.

Detectors: I can see two reasons why a control system will be unable to maintain its error near zero when it has been doing so previously: (1) there is a barrier in the environment stronger than the effectors in the ECSs control loop--we hit a wall, for example; (2) another controller has newly started to work on the same CEV or one in the control loop to the first (either in the effector or the sensor part of the loop). I'm not sure these cases can be distinguished internally to the ECS having a problem. But the result might well be the same--relinquishment of control. So the sensor senses the persistence of error over some period of time, just like the sensor needed for reorganization (no matter which view of reorganization you take).

Actuator: The actuator could (should be expected to?) act on the gain control, the "insistence" of the ECS. The initial response might be to increase the insistence, and after some time in which the error does not reduce appreciably, to reduce it. A second possibility concerns the imagination connection, about which I am not so clear. But trivially, rather than the actuator reducing the gain to near zero, it might switch the output connection into the imagination mode rather than providing references to lower ECSs that were not performing as desired. I find this a bit clumsy, and I'm not going to argue for it. But it's a possibility.

Imagination is an issue that is currently nagging me. Perhaps one day I will post on it.

(Bill Powers 920702.1000)

>Tell you what. How about giving us a reference to the similarity-difference >studies that we can look at? I want to see what kind of facts we're talking about.

The best I can do right now is to list the references I used in the 1983 book (pp 172ff). These don't seem to include the category search studies that seem to me to be the most dramatic. The name Joula, and the journal Perception and Psychophysics springs to mind for them, but I think Juola only did one. My referencing system is pretty haphazard, and it isn't easy to go back for something specific like this. And I haven't looked at the research since about 1985. But these seemed pretty convincing to me.

D.A. Taylor (no relation to me or JGT) Holistic and analytic processes in the comparison of letters. Perception and Psychophysics, 1976, 20, 187-190

Jones, B. The integrative action of the cerebral hemispheres. P&P, 1982, 32, 423-433

Cunningham, J.P., Cooper, L.A. and Reaves, C.C. Visual similarity processes: identity and similarity decisions. P&P 1982, 32, 50-60.

I note a comment at the end of this section of the book: "Throughout the next few chapters, we shall find examples of such cooperating process pairs. One is

always a fast, global process preferred by the RH (right hemisphere), the other a slower, analytic process usually performed by the LH. At higher levels of processing, the evidence for two processes is stronger and comes from more diverse sources (Chapter 11)."

For these, I suggest you get the book out of the library. Speaking of which, the owner of the copy of BCP that I was using has left for a job in Australia, so I can't normally refer to it any more.

Martin

Date: Thu Jul 02, 1992 12:06 pm PST  
Subject: Re: learning and knowledge

[Martin Taylor 920702 15:30]  
(Joel Judd 920630)

>>Knowledge can certainly be represented, but interactions do not require the  
>> kind of representation that implies regress.

>Right; \*transmission\* requires the kind of representation that involves  
>regress. Transmitting encodings requires that both receiver and sender  
>understand what the encodings represent.

No it doesn't, unless you use "coding" in the sense of a one-to-one mapping. PCT says you don't do that. Layered Protocol theory says you don't do that. We had some discussion about it last year (from Bruce Nevin, mostly). Who uses that sort of coding except spies? and if they lose the codebook, they're toast. Transmission of information does not require encodings that involve infinite regress. It requires some recursive beliefs that each hold about the other's belief about... And ultimately the belief system is grounded in the belief that each has some commonality of interactive experience with the world.

Martin

PS. If you don't know that Canada is in the middle of a constitutional crisis, your news organizations must be pretty poor. If, by July 15, the provinces and the federal government cannot come up with an agreed Constitution that is acceptable to Quebec, then in October Quebec will have a referendum on whether to separate and become an independent country. So far, there is no sign that the provinces and the feds can agree unanimously (apart from Quebec), and even if they do, it is unlikely to be an agreement that Quebec politicians will deem acceptable. It's much the same as Czechoslovakia (sorry, the Czech and Slovak Federated Republic; sorry, Czechia (or something) and Slovakia (next year)). They are discussed in your news services?

Date: Thu Jul 02, 1992 12:42 pm PST  
Subject: more on aftereffects to Bill Powers and Gary Cziko

[from Pat Alfano]  
(Bill Powers)

I had tried the aftereffects program not only on different people but on different computers and monitors; I don't think they were all color monitors. I also used different lighting on the monitor and in the room and angled the monitor differently. It is hard to believe that no one would get aftereffects, or that the program is at fault. When you stare at something that is moving, you should get aftereffects when you stop. Everybody (5 people that I know of) who tried your first program got after effects.

You remember correctly, I have fewer motion problems when driving. I think in part I have adapted. I have also learned some tricks. If the visual flow starts bothering me (most of my problems occur on the expressway) I move my eyes around more and try to block out or tune out some of the peripheral motion; things eventually settle down. I also realized that when the road curves I would tend to keep my head vertical to the earths' gravity; now I turn my head into the curve. People have speculated that the reason that passengers are more likely to get motion sickness than drivers is because they turn away from turns, whereas drivers turn into it. I can't be sure how much these things affect me, but I am better. I have discovered other, mostly visual, stimuli that causes problems in other situations. Understanding what is going on helps alleviate some of the anxiety associated with motion sickness.

I have never noticed any visual aftereffects after driving a car, only somatic ones. I have had visual aftereffects after long walks. One day, after a long walk, I was standing looking out over my lawn when I noticed that it was moving away from me. Quite suddenly my body jerked as if I had lost my balance. I tried to bring about the effect the next night. What I noticed then was that my busy body was falling backwards, so of course the lawn appeared to be moving away from me, and of course my body jerked upright. I was falling backwards.

However, the visual motion effect was strong enough to make me think that it would have happened to some degree even without my body moving backwards. I have had a hard time bringing about the effect again because paying attention to it changes it.

(Gary Cziko)

I would like to hear the specifics about your relative's motion problems and just how the doctor came to his diagnosis. I got interested in motion sickness and the vestibular system because of my own problems, although, I have discovered that my motion experiences (sickness) is not nearly as bad (nearly as bad) as it is for some people. I also experience anxiety with the physical symptoms of motion sickness; anxiety is a common symptom with vestibular dysfunction. I have learned some tricks that help alleviate the symptoms, like restricting peripheral vision, keeping my eyes moving, and just being aware of what it is in the environment that is causing my body to react. I was riding the el one day, sitting in a seat that faced the middle of the car (cuts down on motion in peripheral vision). I was lost in thought when suddenly I felt nauseous, dizzy and panicky. I realized that I had been looking off to my left so I turned back that way to see if there was anything there that may have caused my reaction. I could see, reflected in the large window of the car behind mine, the environment moving in an unusual way. As soon as I realized that it was the novel visual information that triggered something in me the anxiety went away and the physical symptoms diminished rapidly. All these tricks help but I still avoid the el if it is at all possible.

Hope to hear from you soon. Pat

Date: Thu, 2 Jul 1992 17:19:44 -0600  
Subject: Plausibility of random reorganization

[From Bill Powers (920702.1600)]

Martin Taylor (920702.0800) --

>I said that I didn't think that what you proposed was plausible for  
>reasons that I hoped were clear.

...

>Taking the simple case in which link() can be 1, 0 or -1, a random  
>reorganization has a probability 0.33... of doing the right thing if  
>there is only one dimension, 0.111... in two dimensions, and  $(1/3)^n$  in  $n$   
dimensions. If there are three ECSs in each of two layers, that is >roughly  
a one in twenty-thousand chance. All other sets of connections >lead to  
some conflict, and even if we grant the probability that some >other sets  
provide stability, the odds are not good that global >reorganization by  
non-targeted random alteration of link sign will >reach an optimum quickly.  
The problem is the same as that of molecular >evolution as seen by the  
creationists. You can't do it that way. You >have to grow stably.

I now remember this point that you brought up some time ago -- just didn't  
make the connection. I don't have a definitive answer, but I think that  
your analysis is making some assumptions that have alternatives. I'll not  
dispute that "targeted" reorganization might be necessary (although when I  
used the term "target" the other day, I was referring to the whole  
hierarchy). I have proposed a version of targeting based on the  
phenomenological idea that awareness directs reorganization to problem  
areas. But having no model of awareness or attention, I haven't pushed that  
very hard. Nor am I convinced that random reorganization won't do the  
trick.

One alternative to targeting that handles SOME of the statistical problem  
is the idea of critical phases in maturation. This is consistent with the  
idea that the growth of the hierarchy is almost entirely bottom-up. Under  
this concept, when it's time to learn hand-eye coordination, in the crib,  
that's the only level of organization susceptible to reorganization, and so  
on up the levels. This is not to say that reorganization occurs exclusively  
at the top level at a given time; only that there is a top level, that it  
gets progressively higher with time, and that reorganization has no effect  
above this level. But I'm not sure that even this idea is necessary.

I'm made a little suspicious by your way of framing the 3-D learning  
problem. To speak of "the chance of doing the right thing" makes it seem  
that the outcome of the random act is either right or wrong, and also that  
you have only one stab at it. If E. coli had to gamble everything on one

tumble, it would be in bad shape. In fact, after any tumble in either the 1-D, the 2-D, or the 3-D case, the chances of heading in a direction more favorable than unfavorable, after a single tumble, are about 50 percent. You're treating each dimension as an independent case, which would imply that the probability of all three cases being in the favorable half-region is only 1 in 8 (you calculate 1 in 27). If you think of a tumble as selecting a direction in space, however, the probability of this direction being in one hemisphere rather than in another is 1 in 2. It isn't necessary for any tumble to aim directly up the gradient; all that's required is a component in that direction. The probability is 50 percent that the component will be between 0 and 100 percent of the swimming velocity -- it would be interesting to know the actual average velocity but it's not zero.

Note that even in a hypersphere of  $n$  dimensions, the chances of  $n$  simultaneous reorganizations creating a change toward rather than away from a given point in the hyperspace is still 50 percent. Of course the average velocity toward the target point decreases with the number of dimensions -- but any bias will get you there eventually. This is one of the basic principles of methods of descent (I think -- I'm no expert). There are, of course, methods of STEEP descent, but I don't see right off how they would be implemented by a reorganizing system.

Putting the problem in terms of right versus wrong choices makes the probability of organizing even one level of control seem incredibly small. But I think this is the wrong way to set up the problem. EVERY form of a perceptual function will yield a perceptual signal that is a regular function of external events. There is not just exactly one combination of inputs that will yield the "right" perception, with all others being "wrong." There are many possible ways of perceiving a given environment that will allow control, and many ways of exerting control that will have at least some beneficial effect on intrinsic state. On the scale of individual perceptual signals at the lowest levels, the number of equally good alternatives must get astronomical. I think you're misstating the combinatorial problem.

One aspect of control, the sign of the effect of error on action, is binary in nature and has only a 50 percent probability of being chosen right by a random process. When a given control system such as a spinal reflex is being organized, however, the mostly likely feedback effect will be none at all, because there are dozens or even hundreds of parallel systems all hooked up more or less the same way. This makes a 50-50 chance of getting it right into a continuous distribution with the most likely one being neutral. All that's required to get SOME control is that there be more loops in the negative feedback mode than in the positive feedback mode. A biased random walk will work quite well to optimize the amount of negative feedback.

Another factor that has to be kept in mind is that an infant left to reorganize itself into a child will surely die. The infant is supported from outside while it gets its behavioral control systems into order. It can spend a long time making mistakes. It can go through millions and millions of reorganizing trials both overtly and in the imagination mode, 24 hours a day. Your point about reorganization being called upon to make rapid correct decisions simply doesn't hold up: that's not necessary. If we

leave it up to children to make immediately correct reorganizations of their systems for avoiding oncoming cars, there won't be many children left. In organisms that are not born with rather extensive complete control systems that control the most important variables, there is no alternative but to protect the developing young from the need to solve control problems by the slow process of reorganization.

Finally, only the simplest control systems involve a huge number of degrees of freedom (something you should consider in line with your DoF paper). Each successive level, up to a point, drastically reduces the number of degrees of freedom. The first new system at a given level allows for only 1! Even passing from heat intensity receptors to the sensation of warmth involves an immense convergence: heat detected anywhere on the skin is warmth. So the most difficult reorganizing tasks are those at the lowest levels -- where there is the greatest amount of preorganization of neural pathways and the highest rate of convergence.

I think that in considering degrees of freedom, you are doing your mental calculations as if all control systems are present and active from the beginning. My view is that in human beings at least, there are very few low-level behavioral control systems available in the beginning, and no higher-level systems at all.

All these considerations must considerably alter calculations of the chances of random reorganization being successful. But I will still not rule out some sort of targeting.

Best, Bill P.

Date: Thu Jul 02, 1992 8:38 pm PST  
From: Dag Forssell / MCI ID: 474-2580  
Subject: Direct Mail

[From Dag Forssell (920702)]

Bill wrote me the holograph in LCS II as his suggestion on how to introduce the subject of my seminar. I have been reluctant to use it, because it seems so powerful that it would just strain the credulity of the message if introduced prematurely. It has occurred to me that the flat earth / round earth "module" provides a framework where I can take advantage of Bill's statement to generate curiosity put it in perspective. So far, I have had a 1% response to my letters. 12 companies have requested my intro package.

Here is letter version 9, hot off the word processor. I will mail 300 or so come Monday. Any comments will be welcome and useful through Saturday.

\_word\_ means underline, right? Italicized titles not shown.

Copyright 1992 Dag Forssell. Permission is granted for quoting within the mailing list CSG-L, and for use in Closed Loop and other publications of CSG.

(Letterhead)

(Page 1)



William T. Powers, CEO  
CSG International  
73 Ridge Road CR 510  
Durango, CO 81301

July 2, 1992

Dear Mr. Powers:

You may be interested in the only fundamentally new perspective on people that has been proposed since 1637. Adopting it can mean improvements for your bottom line, productivity, quality and morale - particularly if you deal with knowledge workers and would like to lead them in the most effective and mutually satisfying way possible.

Costly people problems exist at all levels in American industry. Dr. W. Edwards Deming, pioneer in Quality Management, writes in "Out of the Crisis," page 85:

"In my experience, people can face almost any problem except the problems of people. They can work long hours, face declining business, face loss of jobs, but not the problems of people. Faced with problems of people (management included), management, in my experience, go into a state of paralysis, taking refuge in formation of QC-Circles and groups for EI, EP, and QWL (Employee Involvement, Employee Participation, and Quality of Work Life).... There are of course pleasing exceptions, where the management understands... participates..."

At the core of the design of any social or business organization lies some assumptions about people. If you question these assumptions, the implications for the design and function of your business organization may be large.

The basic perspective from 1637 that still dominates our science and culture is the cause-effect idea that events impinging on organisms cause them to behave as they do. The new one (which has been developed since 1957) is called Perceptual Control Theory, or PCT. The developer, William T. Powers, writes in "Living Control Systems, Vol II":

"Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains why it is so hard for groups of people to work together even on something they all agree is important. It explains what a goal is, how goals relate to behavior, how behavior affects perceptions and how perceptions define the reality in which we live and move and have our being.

Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single testable concept of how living systems work."

Over, please...

William T. Powers

July 2, 1992

Page 2

Understanding people no longer has to be complex and confusing. PCT can be taught in simple form with a comprehensive management application in one day and in more detail with leadership applications in three.

An executive gains insight that allows him or her to inform, influence, align and lead people with mutual respect. S/he can teach people to be more effective and cooperative. Employees can be more satisfied, while the company as a whole responds better to the leader's direction and becomes more productive.

This control perspective will also make it much easier to understand and teach Total Quality Management programs. For instance, when you review the 14 points of the Deming Management Philosophy with this insight, you will see that each point touches on one or more aspects of a system of control. Lack of fear, pride in workmanship, dignity... - all can be seen as manifestations of effective individual control.

I am personally convinced that PCT, once it is widely understood, will have the same kind of impact in the life sciences as Newton's theories did in the physical sciences. Besides a consuming interest in this new development, I have 25 years management experience in engineering, manufacturing, marketing and finance. My formal education includes an MBA from the University of Southern California and a Masters degree in Mechanical Engineering from Sweden.

I have developed the Purposeful Leadership™ programs to explain PCT and apply it to skillful use of diagnostic tools that give the executive the capability to work on productivity. That includes effective communication, teaching effectiveness, resolving conflict, supporting self-motivation in employees, team building, Total Quality Management, leadership insights, effective performance appraisals, effective selling concepts, and development of corporate and individual mission statements. The executive learns how to build confidence, build trust and develop caring relationships.

The basic principles can be taught in a day to any attentive person, who can also verify them. People trained in the "hard" sciences will appreciate the scientific approach and elegant simplicity of the theory, and everyone will be able to begin applying the principles as soon as they understand the underlying model and have had some instruction and practice.

I would like to describe this perspective so you get the point immediately, but this is an impossible Catch-22 challenge, because it is a different concept altogether from what predominates in our world today. Until you understand the principles, you cannot understand at all. I need a few hours in class to explain and illustrate the principles.

When you request it, I will send you a do-it-yourself concept demonstration /test. Until then, perhaps I can indicate how I believe this new perspective fits into the scientific evolution of the life sciences with the following illustrative analogy:

In an era when "everyone knew" that the earth was flat, scientific explanations were developed for navigation and astronomy. Many problems with

those explanations persisted, but people worked around them.

I cannot say what "everyone knows" about human behavior, but experts on the subject employ the 17th century perspective of cause and effect to guide their research. Any book on experimental psychology tells you that the scientific method to learn about behavior is to condition the research animals, set up an experiment, then vary the stimulus (independent variable) and watch the response (dependent variable).

Continued....

William T. Powers

July 2, 1992

Page 3

(This would be a valid scientific method if in fact animals and people were cause-effect organisms. But our demonstration will show you in a few minutes that they are not).

With this scientific method our experts have done many experiments and reported explanations which are now part of our language, culture and management practices.

With or without an awkward science, there have always been natural leaders, successful salesmen, wise parents and good communicators. But it is rare that they can explain what they do and why. Their insight and skill seems intuitive. Human behavior practitioners and many executives make an effort to master this important subject, which demands attention. They depend on a variety of experiences and interpretations, not hard science, to develop effective personal approaches for dealing with people.

(Imagine how good they will be when they get good understanding that applies every time. With PCT, the executive learns to function as well as those intuitively wise people. With practice even better, since s/he will have greater insight).

Many problems with the expert's scientific explanations persist despite all the research, but people work around them. Lack of success indicates that we lack a good model or "paradigm" to help us understand why people do what they do. In our ignorance, we often spend our energies in debilitating conflict instead of in productive cooperation.

\_When Copernicus and then Galileo introduced the fundamentally new insight that the earth is round\_ (it has always been round), \_the problems of navigation and astronomy were placed in a new light\_. The new insight did not invalidate the common sense observation that the earth appears flat locally, but science was able to progress.

Most experts on the old science could not comprehend the new paradigm, because they had already internalized the flat paradigm in all its details as their personal reality. With time the experts died off, and new ones grew up, embracing the new paradigm on its merits because it solved many of those persistent problems.

Isaac Newton's "Principia Mathematica," published fifty years after Galileo, was resisted in the same way, just like all dramatically new approaches have

been. It took fifty years for it to be fully accepted. The evolution of science is much more than a steady accumulation of knowledge! The process requires creativity. The opportunity for a revolution arises when a current paradigm fails to solve problems and competing paradigms are offered to provide better explanations. A struggle of many decades typically takes place, with the trained scientists continuing the development of the existing paradigm while outsiders and early converts champion a new one.

The 20th century understanding of Perceptual Control Theory (people always control) provides a fundamental new insight that puts the problems that result from human interactions in a new light.

Perceptual control is as incomprehensible at first glance to a person trained in cause-effect thinking (which we all are in our culture) as the idea that the earth is round was to a person trained in the details of a flat earth. The demonstration shows this clearly. Still, an understanding of PCT contains an explanation of the illusion of cause and effect in people, just like the understanding that the earth is round contains an explanation of the illusion of a flat earth.

Over, please....

1 The phenomenon and process is described in Thomas Kuhn's seminal book: "The Structure of Scientific Revolutions," which introduced the term "paradigm."

William T. Powers

July 2, 1992

Page 4

Another illustrative analogy is to say that we live in a maze where only the walls and passages are visible to us. The perspective of Perceptual Control allows us to rise above the maze and see the structure. We can then set and reach our goal much easier.

The new perspective does not invalidate any wise common sense observation or practice. It does provide an enhanced understanding of seemingly intractable problems. It provides new diagnostic tools and shows why cookbook rules for behavior (programs which tell you what to do under certain circumstances) do not always work.

Perceptual Control Theory is already well developed. But no doubt it will take time - well into the 21st century - before this breakthrough is known, understood and embraced by a majority of experts. You can take advantage of what "everyone will know" in the 21st century right now to improve your company's competitive position. But because it breaks new ground, you must be willing to think for yourself to do it. You will participate in a scientific revolution when you understand and adopt it.

Some people will think that the term "control theory" promises a new way to control other people. It is precisely the other way around. We show how people control themselves at all times. When you understand PCT, you can work with people rather than get into conflict despite the best of intentions.

Please request the free introductory 39 minute audio tape with script and

illustrations. It demonstrates the basic concept and explains the benefits, applications, background and content of our programs. The demonstration /test allows you to find out if your associates can recognize control in action. (I bet they can't).

When you receive the introduction, I think you will find the demonstration enlightening and entertaining. Please feel free to share it with your technical, operations and sales managers at any level for their evaluation. This is a win/win program to increase the understanding and effectiveness of anyone who deals with people.

Sincerely,

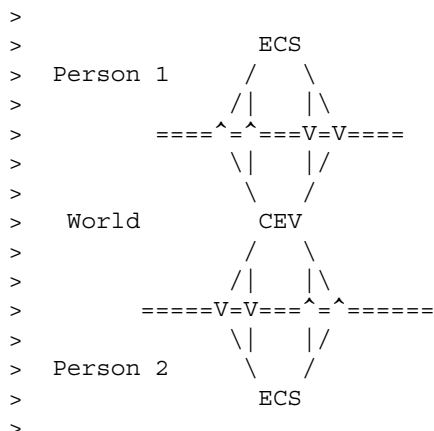
Dag Forssell

Date: Fri Jul 03, 1992 10:52 am PST  
 From: Dag Forssell / MCI ID: 474-2580  
 Subject: Taylor's diagram

[From Dag Forssell (920703)]  
 Martin Taylor 920629 1515

>>It just occurs to me that we take Martin's 23 level chart and fold it on  
 >>the mirror line, then interconnect the control systems across so we  
 >>control all the perceptions up and down. We are back to the diagram as  
 >>we know it, but with an expanded understanding of it.

>  
 >This is true, but I'd rather not do that. One of the points of making the  
 >diagram is to show a relation between a "Boss Reality" and a controlled  
 >percept. Now if that Boss Reality exists, it is accessible to another  
 >control system (e.g. an experimenter). The other control system can focus  
 >on (perceive) the same complex environmental variable (CEV) and perhaps  
 >attempt to control its perception of the CEV, disturbing the first control  
 >system's perception of it. We can diagram the interaction his way, taking  
 >the top half of this diagram as representing the whole of my earlier one,  
 >and the bottom half its mirror image:



>If you fold the original diagram about the mirror line where "the rubber  
 >meets to road", you lose this view. Much more crucial, you lose the view  
 >of somewhat off-focus disturbance, which was the core of the discussion  
 >about VOT.

Bill Powers (920629.1600)

>The 11 levels below the line through level 0 are, of course, in the observer.

Bruce Nevin (Wed 920601 08:59:24)

>The mirroring around level 0 reminds us to what extent apparent  
>structure in the environment is a reflection of structure in the control  
>hierarchy, projected there by the observer. We assume vice versa, but  
>that can only be an assumption. (Right, Wayne?)

>

>So it is worse than perhaps Martin has said. Not only may the observer  
>(the investigator) identify the wrong environmental variable V as that  
>controlled by the observed control system, the two parties may also have  
>differently structured control hierarchies, and so may parse the  
>environment differently into environmental variables.

Martin, I was gone two days, or I would have commented on your post sooner.  
I think your presentation is both exciting and dangerous.

Exciting because of the relationships you want to portray. Dangerous  
because it will invite misinterpretation.

> One of the points of making the  
>diagram is to show a relation between a "Boss Reality" and a controlled  
>percept. Now if that Boss Reality exists, it is accessible to another  
>control system (e.g. an experimenter).

Your diagram creates the illusion that the "Boss Reality" is structured in  
a hierarchy outside the control system, and is accessible as the same "Boss  
Reality" to another observer. This is patently false. That is the dangerous  
part. Your diagram may be very useful as an illustration or teaching tool,  
however. That is the exciting part.

Your description fails to make the distinction forcefully, if at all. (If  
you do, it escapes me). Bill makes the distinction - and Bruce (I think).

About a year ago, I portrayed in a diagram to Bill (as a teaching tool) the  
process of perception up all 11 levels. Separately, I tried to portray the  
process of control down all 11 levels. In a third step, I meant to marry  
the two. Bill did not let me get away with it. (I have not given up on the  
teaching tool.)

It took me time, but I now think I understand the "behavior of perception"  
reasonably well. I think it is very hard to grasp the close integration  
that HPCCT suggests of both perception and control up and down the hierarchy  
and sideways in between each level. I believe it is much easier to say: I  
understand; I agree, than to explain and demonstrate understanding.

In one of my posts over a year ago, I related the perception of room  
temperature as the difference between the rrrrrrrrr song of the sensing  
neuron and the rrrrrrrrr song of related memory reference. You may remember  
the song part. As perceptions are passed on up the hierarchy, imagination  
enters into the process as well, to flesh out the picture as needed.

While I agree that there is a Boss Reality phenomenon we call temperature (and model as molecular motion - I have never seen a molecule or its movement, but have adopted the systems concept), it seems clear that it is "accessible" to two different observers ONLY by comparison with the subjective reference in each observer. Therefore the same Boss Reality can never be accessible in the same way to two different observers.

The uninitiated reader of your diagram may be led to believe that you advocate something that is true of the "Boss Reality" and that two observers can agree on with certainty. A careful note will have to accompany the diagram each time it is used, to convey what it portrays and what it does not portray. As I understand you, the diagram may be intended only as a simplistic teaching tool for those who have not yet understood the concept of behavior of perception.

Dag

Date: Fri, 3 Jul 1992 15:36:43 EDT  
Subject: Re: Taylor's diagram

[Martin Taylor 920703 15:15]  
(Dag Forssell 920703)

>> One of the points of making the  
>>diagram is to show a relation between a "Boss Reality" and a controlled  
>>percept. Now if that Boss Reality exists, it is accessible to another  
>>control system (e.g. an experimenter).  
>  
>Your diagram creates the illusion that the "Boss Reality" is structured in  
>a hierarchy outside the control system, and is accessible as the same "Boss  
>Reality" to another observer. This is patently false. That is the dangerous  
>part. Your diagram may be very useful as an illustration or teaching tool,  
>however. That is the exciting part.

I don't think it is patently false. What IS patently false is that we can ever know whether it is true. All we have access to is our own perceptions. That we can control them suggests the existence and character of the Boss Reality (if there were nothing out there that we could perceive as a Complex Environmental Variable, we couldn't control the perception of it). I think all this goes without saying. It applies to all attempts to employ The Test. The experimenter always has to assume that there is a Boss Reality, and that the CEV being disturbed as part of The Test is one that has a corresponding percept in the subject (and that the subject exists, as part of the Boss Reality).

I don't see anything new, difficult, or dangerous in any of this. Either there is some reality out there or there isn't. If there isn't, we are (sorry, I am) just having fun in our own head(s?). If there is, we manipulate it and see others manipulating it. The only structure of the "out there" that we can deal with is the structure we give it through our perceptual/cognitive machinery. And there is neither more nor less reason to deal with the "truth" of the world as structured than to deal with it as a set of mathematical equations describing quantum chromodynamic systems, or whatever today's favoured physical substrate

might be.

Any instance of The Test depends for its success on the subject controlling the perception of the same CEV that the experimenter perceives. The Test is there to see whether this identity exists. If it does, reality has, in part, the same structure to the subject as to the experimenter. A lot of my "statistics" argument with Bill and Rick hinges on the unlikelihood of the experimenter hitting on the same CEV as the subject, or even having the same structure of reality as the subject. It becomes even harder when there is a cultural or species gap between the experimenter and the subject.

Of course it's all in the head. That doesn't mean it's any the less in the world. And that is not dangerous. What is dangerous is to assume that your way of seeing the world is the only way, and must be the right way, and that anyone who sees it differently is lying, bad, dangerous, and to be suppressed. That's what's dangerous.

-----

>About a year ago, I portrayed in a diagram to Bill (as a teaching tool) the  
>process of perception up all 11 levels. Separately, I tried to portray the  
>process of control down all 11 levels. In a third step, I meant to marry  
>the two. Bill did not let me get away with it. (I have not given up on the  
>teaching tool.)

What did Bill not allow you to get away with? If outputs derived from error signals go as references to the same lower level ECSs that the perceptual inputs come from, the upward flow of perception is intimately married with the downward flow of action. Indeed, it is this presupposition that we here are banking on, in our belief that we can train a control net much faster and more precisely than we can train an ordinary S-R neural net such as a pattern recognizer. I thought that the reciprocal connection was an aspect of HPCT that was pretty much taken for granted. Is it not?

Martin

Date: Fri, 3 Jul 1992 17:12:50 EDT  
Subject: Re: Similarities and differences

[Martin Taylor 920703 16:00]  
(Bill Powers 920703.0600)

I think I can accept most of what you wrote, and once again find myself frustrated by my evident obscurity in writing. I'll try to rephrase one or two points, and maybe we will have a common understanding. Or maybe some real disagreement lurks in the words.

>Nor do I think that "templates" are a necessary construct here. The concept  
>of a template is an alternative to the concept of a perceptual function,  
>and one that is hard to defend. The image of moving a negative around over  
>a positive image looking for a match doesn't fit the pandemonium model in  
>which all perceptions and reference signals are one-dimensional variables.

I intended "template" as a non-specific notion that included a perceptual function but was not restricted to it. Pour battre le cheval mort, if an ECS has four inputs, A, B, C, D, then if its perceptual function is  $F(3A+B-C-3D)$  it



has a template for linearity, in my terminology. I'm sorry to say that I was aware of this possible misunderstanding when I used the term, but hoped (ignoring Murphy) that it would not occur. In this case, "template" can certainly be a remembered percept. It can equally be any source of a reference signal.

-----  
 There's also a terminological problem in "similarity" and "difference", because the distinction shows up in behaviour rather than in any perceptual logic. It is a question of what one is controlling for, so far as I can see. I was describing a (to me, plausible) mechanism when I should have been standing back a little. Let me try again.

I consider three situations, with respect to the organism (not necessarily with respect only to an ECS).

- (1) some percept is being actively controlled to be as close to its reference as the other controlled percepts permit.
- (2) some percept is not being actively controlled, but if
  - (2a) it departs too far from some reference, or
  - (2b) it comes sufficiently close to some reference, then control relating to this percept must become active or bad things happen.

Condition (1) is what I have identified with the difference detection of an active ECS. I could equally have talked about "identity" detection, but that word has connotations of labelling and category. Some people use it, however. Condition (2) is what I have identified with similarity, whether the criterion for action is approach to or departure from a reference. (2a) is "I don't want to be too hot or too cold", and (2b) is "I don't want to see a tiger looking at me hungrily from too close." To complete the set, (1) is "I want to keep near the centre line of my traffic lane."

In my discussion, I put all the nonlinearity into the comparator, but as you point out, it could equally well be in the perceptual function (but see below).

In the para on "template" above, I used  $F(3A+B-C-3D)$  as the template for linearity. Now let us use this function  $F$  to illustrate a possibility. In all cases, we are dealing with a normal ECS that has a simple difference as a comparator, and a linear output gain function. We will assume that the reference level is zero. (If it isn't, the function  $F$  can be applied to the comparator output rather than the output of the perceptual function, a rather more sanitary procedure that does not affect the argument). There are at least three cases, corresponding to my three sorts of detector.

Type (1):  $F(x) = x$  ( $|x| < \text{some limit}$ )

This is a difference (identity) detector--an active zero-seeking controller. It tries to see its input as linear, and continuously controls it to maintain linearity as closely as it can.

Type (2a):  $F(x) = 0$  ( $|x| \leq t$ )  
 $= |x-t|^2$  ( $|x| > t$ )

This is a similarity detector that produces output only when  $x$  deviates from zero by a sufficient quantity. If the input is sufficiently nearly linear, it does nothing.

Type (2b):  $F(x) = t^2 - x^2$  ( $|x| \leq t$ )

$$= 0 \quad (|x| > t)$$

This is a similarity detector that produces output when  $x$  is close to  $t$ . It does not want to perceive linearity, and if the input is sufficiently far from a straight line, it does nothing. Otherwise it produces output that presumably induces actions that cause the input to deviate from linearity. This is an alerting ECS.

The thesis, from the degrees of freedom argument, is that ECSs of type 1 can exist to control any of the degrees of freedom implicit in the sensory input, but that no more than a few can be simultaneously satisfied. To affect which few degrees of freedom are controlled, many parallel ECSs of type 2 can be accepting input, but none will provide output unless the similarity condition is violated, at which time they will provide output.

I attempted in my big posting to suggest several different possibilities for what happens when one of the type 2 ECSs does start to provide output, all of which have the same functional result: a Type 1 ECS will be controlling a perceptual degree of freedom that was not previously being controlled, and that will occur at the expense of removing from control another perceptual degree of freedom.

> I anticipate that what you  
>are interpreting as reference signals or templates will also be explainable  
>as perception of a difference-relationship or a similarity-relationship  
>between two distinct percepts -- perhaps one of them being remembered. My  
>bias is always to associate objects of consciousness with perceptual  
>signals produced by input functions, not with reference signals (unless  
>imagination routes them into the perceptual channels) or error signals  
>(which in the model as it stands today are not part of the perceptual  
>system). Perception of differences should not be confused with error  
>signals.

Yes, I try to go along with that. If I say something that disagrees with it, or seems to, either I have been thinking sloppily or I have been writing sloppily. You should pull me up on such occasions. I don't think this was such an occasion, as I never conceived of error signals (or references) contributing to percepts at any level.

-----  
> The model is most easily tested when it leads  
>to flat statements allowing of no conceivable alternatives -- when it  
>stands or falls on statements of truth. And I also think that such extreme  
>statements lead very naturally to simple experimental designs. What we're  
>talking about is REAL falsifiability.

The problem with truthsaying is that language is not logic, and it is very easy sometimes to believe you are making a watertight case for the necessity of something that isn't necessary at all. I know I'm arguing out of both sides of my mouth here. The problem is that reality testing needs lots of resources, so the best truthsaying I can achieve is the most cost-effective procedure for me. But I can never be really sure that what I say MUST be true is not standing on some rickety foundation that might be swept away.

Extreme statements need to be interpreted, and as with "template," the boundaries of the intent of the statement are not always clear. It is often very difficult to say just what has been falsified by any experiment.

On falsification, my notion is that all exact theories are false, so that one gains no information in finding one to be false. Vaguer theories may include the truth somewhere (insofar a truth means something that will not be found false within a finite time), but they are harder to demonstrate to be false when they don't include the truth.

As you may have guessed some long time ago, I reject firmly the notion we are taught in school, that science consists of the rejection of falsifiable hypotheses. I take it to be more like the reorganization problem (when we agree on it), continual evolution of better descriptions of nature, by which I mean more concise descriptions that cover wider ranges of conditions.

Martin

Date: Fri, 3 Jul 1992 05:18:00 GMT  
Subject: Dictionary, Learning Disabilities

[From Hank Folson (920703)]

Martin Taylor says (920701 14:00):

>To make a PCT dictionary presupposes that the listener/reader has the  
>proper appreciation of PCT. Otherwise the definitions will make no sense.  
>So the dictionary would be of more use to CSG-L contributors than to  
neophytes.

This dictionary will have both S-R world and PCT world definitions side by side. Newcomers will see right away that there are big differences between the two worlds. We will have to complete the dictionary before we will know whether this will intrigue and attract them or simply turn them off....

---

Martin Taylor

Re: Similarity/difference processes:

My daughter told me about several students who are receiving help at her university. One major thing they do is to allow the student unlimited, or at least two times, the normal test taking time. Without the usual time constraints, these students can now get well above average grades. I don't know about you folks, but if I couldn't figure out the answers, sitting there longer would not improve things! Either I knew it or I didn't. My guess was that these kids had some difficulty in going from level to level, and the added time was of use to them for that reason. But the similarity and difference processes sound like a more logical explanation. It also suggests the possibility of training people.

Hank Folson

Date: Fri, 3 Jul 1992 07:21:55 -0600  
Subject: Similarities and differences

[From Bill Powers (920703.0600)]

Martin Taylor (920701.1100) --

RE: Similarities and differences.

>A similarity detector has a gain function concave upward (e.g. gain =  
>error to a power greater than unity), and may well have zero gain for  
>some finite level of error. A difference detector has a gain function  
>that has some appreciable slope near zero error.

Allow me to pursue the question of perceiving similarities and differences between distinct percepts. I'll get to templates at the end.

Suppose that the basic form of a perception  $p$  of a single variable  $v$  is nonlinear and approximated by  $p = k*v^2$ . This is an approximation of the low to middle range of the perceptual response. The slope at zero input is zero, increasing linearly as the variable departs from zero.

In detection of a difference relationship, I propose that the perception is derived from the difference in a single attribute between perceptions of two sets of variables (the same argument can be extended to multiple attributes). In general, the amount of one attribute in one set can be expressed as  $c + d/2$ , and in the other as  $c - d/2$ , where  $c$  is the amount common to both variables and  $d$  is the amount of difference. The difference in amount of attribute as perceived at the relationship level is  $p_1 - p_2$ . If each perception has the same nonlinear relationship to the amount of attribute, approximated as a square, we have (leaving out scaling factors)

perceived difference =  $p_1 - p_2 = (c + d/2)^2 - (c - d/2)^2$ , or

$$p_1 - p_2 = 2*d*c.$$

The variable  $d$  is the difference itself. The function  $p_1 - p_2$  is the computation that yields perception of this difference, a relationship. Note that if there is no common attribute at all ( $c = 0$ ), there can be no difference! This suggests the old "apples and oranges" observation: you can't compare things that have nothing in common. Note also that if the functions are linear, so that  $p_1 = v_1$  and  $p_2 = v_2$ , then the perception of difference is just  $p_1 - p_2 = v_1 - v_2$ , and the slope is still nonzero at zero difference. And finally, note that a nonzero slope at zero difference will still be found for other forms of positively-accelerating nonlinearity. This is easy to show graphically.

Similarity is more difficult to define because there is no simple natural definition as there is for difference. If two percepts are identical, how much similarity do they indicate? An infinite amount? A lot? According to your definition, the closer two percepts come to being identical, the higher the perceived similarity, with the slope increasing as identity is approached. If this rule applies generally, identity must correspond with maximum slope, so there is maximum slope when the amounts of the common attribute in each set of percepts reaches equality, and a cusp is found at that point, or a singularity.

A simple proposal is that similarity is perceived through a function that is the reciprocal of the difference relationship:

$$\text{sim} = k / (p1 - p2).$$

As division by zero is approached, of course, the similarity perception simply goes to maximum, not infinity, as analog dividers have finite limits. This function has the required accelerating nonlinearity as similarity shades toward identity. It changes sign abruptly at  $p1 = p2$ .

A reciprocal function would, of course, resemble a power function with an exponent greater than one -- I doubt that judgements of similarity yield data that is quantitative enough to distinguish between a best-fit reciprocal and a best-fit power function. The only reason I could see for choosing a power function would be an attempt to be consistent with Stevens' "power laws" of stimulus magnitude estimates. There's no physical reason to suppose that power laws are involved, although of course one can always fit a power-law curve to a nonlinear relationship.

I've treated similarity and difference detection here simply as perceptual functions, examples of relationship perceptions. I don't think there's any need to bring in error signals, especially as I resist including them in the world of conscious experience (although they keep getting back into it, especially with the Revised Model that we keep abandoning again).

>There is a second kind of similarity detector, which is not an ECS so >far as I can see, but a perceptual function akin to what neural network >people call a "radial basis function". It emits a larger signal the >nearer the incoming perceptual pattern matches its "template."

As I've defined a similarity detector, there doesn't seem to be any difference here, except that a perception is being compared with a template instead of another perception. I really don't think that "similarity" is an appropriate idea in relation to templates. What knows that there is a similarity? Another similarity detector?

Nor do I think that "templates" are a necessary construct here. The concept of a template is an alternative to the concept of a perceptual function, and one that is hard to defend. The image of moving a negative around over a positive image looking for a match doesn't fit the pandemonium model in which all perceptions and reference signals are one-dimensional variables.

I will try to obtain the references you cited. I anticipate that what you are interpreting as reference signals or templates will also be explainable as perception of a difference-relationship or a similarity-relationship between two distinct percepts -- perhaps one of them being remembered. My bias is always to associate objects of consciousness with perceptual signals produced by input functions, not with reference signals (unless imagination routes them into the perceptual channels) or error signals (which in the model as it stands today are not part of the perceptual system). Perception of differences should not be confused with error signals.

-----  
I agree that when reality testing is difficult, it's best to rely on truthsaying -- saying what MUST be true and what CAN'T be true, according to the model. There's little point in spending a lot of time on what MIGHT be true (as I've done above). The model is most easily tested when it leads

to flat statements allowing of no conceivable alternatives -- when it stands or falls on statements of truth. And I also think that such extreme statements lead very naturally to simple experimental designs. What we're talking about is REAL falsifiability.

-----  
Best

Bill P.

all ECSs would be in conflict about almost every percept. I assume "apple controllers" come into play only with respect to things that are sufficiently like apples to make it reasonable to try to perceive them as apples.

I'll answer the plausibility post later. I think we are even closer to agreement there, but some questions do remain.

Martin

Date: Fri, 3 Jul 1992 14:45:02 +0200  
Subject: Going south

[From Oded Maler 920703]

A-propos vacations:

Tomorrow I start to drive from Rennes to Aix in order to participate in the CACS summer school (hopefully will arrive there despite of the barriers of the truck drivers). Then I'll go to Grenoble looking for a house (reading CSG-net is not a sufficient condition to get a permanent position, as I discovered) and return here toward the end of July. But probably will have access to e-mail everywhere.

There were some very interesting postings recently especially from Martin and Bill which I hope to comment on when I return. Just a small technical comment on Bill's fair description of 'Brooks from PCT point of view': maybe in the higher levels, when you cannot add/substract abstract symbols, interaction between behaviors by enable/disable with priorities is reasonable. If you have 'approach food' and 'escape enemies' ECSs both trying to put their opposite reference signals on the same motor when 'food' and 'enemy' are close to each other, I'm not sure addition/subtraction is the best way.

Best to all and cheer up.

--Oded

Date: Fri, 3 Jul 1992 10:41:33 CDT  
Subject: Bill on Brooks

[From Dick Robertson]

(Bill Powers 920630.0800)

I read your blast on Brooks with some mischievous pleasure and amusement,

>Brooks knows nothing about my work, or at least dismisses it (he never  
>cites it).

It's frustrating, isn't it, to see the confirmation of HPCT in action and not be able to use it to make things more the way you'd like them to be. What I mean is that we know, in principle, from the theory that behavior realizes the reference signals of the highest level in one's hierarchy. So, it's really no surprise that Brooks

>did not object to anything I said or to anything in these programs -- at least  
>not to my knowledge, as he has never replied. Unless he simply dropped  
>everything in the wastebasket without looking at it, he evidently found  
>nothing of any interest in the letter or the programs.

How could he have found anything of interest? If he had, he would have had an enormous error signal somewhere in his self-system--something along the lines of: "I am a leader in my field, a scientific pioneer, a great man, an explorer opening paths that others have yet to travel...(NOT??)"

It's not hard to make these speculations about his self-system, even though I have had no opportunity to put him to The Test. The attributes that characterize a person indicate the self-image and principle level perceptions that he/she most strongly controls. So, reasoning backwards from Brooks's actions we infer that the perceptions being controlled by them requires that any input that would cause error signals in his image of himself would immediately be corrected. And we know all this already, but it's still hard to swallow, isn't it?

It's not really surprising that most people control for higher order variables involving achieving and maintaining status, income, recognition, comfort, etc. Far more interesting, I think, is the case where someone doesn't. Like you, for instance. Oh, I think you do too, to a reasonable extent, except that you have been controlling for something else even more and that puts constraints on the length to which you can go in controlling for the conventional stuff.

I'm pretty sure I have never heard you identify just what that something else is. I think it's an important question. I think it bears indirectly on the discussion that you and Martin (among others) have been having about the workings of the top level and its connection with reorganization. Yes, I know your biography fairly well, and I have heard your history of how you got started and the early days. But, while that higher priority perception might be implicit in that story it's not explicit, and I think it's worth investigating. I might well be projecting from myself, but I would guess that you made key choices at certain developmental choice points, which were both free choices and which you couldn't have made any other way without violating the self-system you already had. (And that self system was never triggered to reorganize.)

That looks like a paradox to me. Were the choices free or were they constrained by a prior requirement?

I say I might be projecting from myself because way back at the beginning of my career I heard W T Powers and his two cronies present a view of how behavior works (in hardly more than one hour) that exploded across my horizon so strongly that I felt no hesitation about working for nothing a day a week for the next couple years. In those couple of years a pair of my peers (and still close friends) layed down the foundations of their present international fame as researcgers and got positions at much more prestigious schools than little old northeastern. I'm not saying they didn't deserve it; they have done good "normal science," as Kuhn calls it.

I still have no idea as to what it was in me that turned in the opposite direction. In one way it was a conscious choice, but I have a sense that I couldn't really have done otherwise. Why is that? There is a scent of determinism in there that I distinctly don't like. For one thing it doesn't jibe with my favorite speculation that free will exists at the highest level, because, maybe, the self-system is in chronic reorganization throughout one's life. I seem to get support from that speculation from time to time when I observe in both myself and others moments of "not being oneself today." I detect some random strayings from the familiar self image. That reminds me of the spontaneous reversals of control that you and Rick, I think it was, observed while people were learning certain tracking tasks.

I've also collected what seem like observations on the other side of the paradox. While I was in Belgium I finally got to read a copy of Francois Jacob's\* The Statue Within. (\*The Nobel geneticist.) I first read a review of it in Science ten or more years ago and tried to get it, but its US publisher reneged. Anyway, what he meant by the "statue within" was his impression of his life as the unfolding of a set of implications that he felt he perceived in his earliest memories. (A not too radical idea for a geneticist I guess.)

What this boils down to is that I'm impatiently waiting for you modeling guys--you, Rick, Martin Taylor and all--to get around to modeling the highest levels, so we can see how they have to work in an autonomous organism.

Best-, Dick Robertson

Date: Fri, 3 Jul 1992 18:11:18 EDT  
Subject: Re: Plausibility of random reorganization

[Martin Taylor 920703 18:00]  
(Bill Powers 920702.1600)

Once again, we are getting close to agreement. But (luckily) we are not quite there yet, I think.

>One alternative to targeting that handles SOME of the statistical problem  
>is the idea of critical phases in maturation. This is consistent with the  
>idea that the growth of the hierarchy is almost entirely bottom-up. Under  
>this concept, when it's time to learn hand-eye coordination, in the crib,  
>that's the only level of organization susceptible to reorganization, and so  
>on up the levels. This is not to say that reorganization occurs exclusively  
>at the top level at a given time; only that there is a top level, that it



>gets progressively higher with time, and that reorganization has no effect  
>above this level. But I'm not sure that even this idea is necessary.

Yes, I was trying to push this. But rather than saying that a new top level is being developed, I still like the idea that one is inserting levels, thus redefining what was there before. The difference is one of viewpoint, I think. See later, about the baby.

>I'm made a little suspicious by your way of framing the 3-D learning  
>problem. To speak of "the chance of doing the right thing" makes it seem  
>that the outcome of the random act is either right or wrong, and also that  
>you have only one stab at it.

Right or wrong--yes. One stab, no. I never intended that implication.

>You're treating each dimension as an independent case, which would imply  
>that the probability of all three cases being in the favorable half-region  
>is only 1 in 8 (you calculate 1 in 27). If you think of a tumble as  
>selecting a direction in space, however, the probability of this direction  
>being in one hemisphere rather than in another is 1 in 2. It isn't  
>necessary for any tumble to aim directly up the gradient; all that's  
>required is a component in that direction.

That's true when there is only one degree of freedom for the controlled percept: "satisfactoriness of the environment" and three degrees of freedom for action. But the more common case is when the action degrees of freedom are fewer than the perceptual degrees of freedom. Then you have to worry about conflict, and components in those other directions do matter.

One in 27 is correct, because a link can have any of three values, not two. We are talking about reorganization that permits making and breaking links, as well as changing signs, are we not?

Even when we are ignoring the conflicts induced by components in directions orthogonal to the one causing the reorganization, my argument was not about whether the control vector would eventually point in the right direction, but about how fast it would do so. I should not think that changes induced by reorganization should ever occur faster than the Nyquist rate for the feedback loop in question, so I was interested in the probability of changes that were significant improvements. I grant that 1 in 2 will be improvements on no control, but most of those will be trivial improvements if the space has high dimensionality. When I first brought up this topic, I considered anything within 60 degrees of the optimum direction to be significant improvement.

>EVERY form of a  
>perceptual function will yield a perceptual signal that is a regular  
>function of external events. There is not just exactly one combination of  
>inputs that will yield the "right" perception, with all others being  
>"wrong." There are many possible ways of perceiving a given environment  
>that will allow control, and many ways of exerting control that will have  
>at least some beneficial effect on intrinsic state. On the scale of  
>individual perceptual signals at the lowest levels, the number of equally  
>good alternatives must get astronomical. I think you're misstating the  
>combinatorial problem.

Yes, you are quite right about that. My error is to require a particular ECS to control a particular perceptual degree of freedom. This will be valid if the perceptual input function of each ECS is prescribed in advance, but of course it cannot be. Each ECS must learn what it is perceiving as well as to control that percept. That thought carries much implication, which I haven't considered enough to pursue here. But it does carry the implication that there is an enormous amount of symmetry in the hierarchy that a random reorganization system can initially exploit. But it can't do so once the symmetry has been broken by some ECSs having learned to control their percepts. Then, with random reorganization, I think the combinatorial problem is as I stated, the more so the more ECSs have acquired control.

>Another factor that has to be kept in mind is that an infant left to  
>reorganize itself into a child will surely die. The infant is supported  
>from outside while it gets its behavioral control systems into order. It  
>can spend a long time making mistakes. It can go through millions and  
>millions of reorganizing trials both overtly and in the imagination mode,  
>24 hours a day.

Yes, that was my point about it being no accident that infants of all species being born either unable to act (very low loop gain) or with built-in control (but unable to learn new controls at the level that is inborn) like a deer or a chicken. Deer can run at birth, chicken can peck at seed (but cannot learn to adapt to prism displacements of their vision). But I think "millions and millions" may be saganesquely excessive. It's possible, especially at low levels, I grant. At higher levels, things move more slowly.

>Your point about reorganization being called upon to make  
>rapid correct decisions simply doesn't hold up: that's not necessary.

It does, if the system being reorganized is actively controlling, with reasonably high gain.

>In organisms that are not born with rather extensive complete control  
>systems that control the most important variables, there is no alternative  
>but to protect the developing young from the need to solve control problems  
>by the slow process of reorganization.

Yes, that's what I meant.

>Finally, only the simplest control systems involve a huge number of degrees  
>of freedom (something you should consider in line with your DoF paper).  
>Each successive level, up to a point, drastically reduces the number of  
>degrees of freedom. The first new system at a given level allows for only  
>1! Even passing from heat intensity receptors to the sensation of warmth  
>involves an immense convergence: heat detected anywhere on the skin is  
>warmth. So the most difficult reorganizing tasks are those at the lowest  
>levels -- where there is the greatest amount of preorganization of neural  
>pathways and the highest rate of convergence.

A really speculative point! Apart from the statistical convergence that has nothing to do with control, due to the natural redundancy of the real world, I have envisioned the possibly controllable percepts (not degrees of freedom for perception) as growing in number as we go up the levels, before reducing

at the highest levels. Sort of barrel-shaped, rather than conical.

>I think that in considering degrees of freedom, you are doing your mental  
>calculations as if all control systems are present and active from the  
>beginning. My view is that in human beings at least, there are very few  
>low-level behavioral control systems available in the beginning, and no  
>higher-level systems at all.

No, there's a misunderstanding here. I tried, as I remember, to put two alternatives into play. One was indeed the matured system, which presents a problem for random reorganization because of the likelihood that untargetted random reorganization will disrupt areas that are working very well. The other was a developing system, and for it, I suggested that we might consider the top level as representing optimum values for the intrinsic variables, even if initially there were no other levels. All other levels are inserted inbetween the top and the world-interface.

Since we are (by agreement) working from the lowest common denominator of no prior construction of ECSs, we have to include evolutionary development here. No matter what the evolutionary level, the prime concern is to maintain those intrinsic variables near optimum long enough to pass on a structure description to the next generation (whether it be by cloning, seed-spreading, or whatever). The primary control system has this function.

I think that the place where we have a disagreement is how this primary control system effects its control in an organism that can control other percepts in its environment. (Can trees?) My preference is for a single hierarchy, in which the primary control system has been elaborated to effect its control through the provision of reference signals to other ECSs. Yours is for the primary control system to be separate from another hierarchy, and to effect its control by blind modification of that second hierarchy.

I don't think there are ground other than aesthetic (Occam's razor) for choosing between these organizations, unless it can be shown that either would not work. If you allow that your primary control system can act by targetting local areas of the sensory-motor hierarchy, then I see little possibility for distinguishing them on grounds of plausibility. They come almost to mean the same thing in different words. Probably some differences do remain: I think that maintenance of error, or more particularly the uncontrolled growth of error, in an ECS seems a plausible reason for reorganizing something about that ECS, whether it be the signs of some or all of its outputs, the nature oits perceptual function, or even shutting down or inverting its gain (are not the most intense missionaries the recently converted?).

Have you modelled to reorganization of a moderately complex hierarchy? That would be a lovely demo, if you have.

Finally:

>I have proposed a version of targeting based on the  
>phenomenological idea that awareness directs reorganization to problem  
>areas. But having no model of awareness or attention, I haven't pushed that  
>very hard.

This would be the effect of "teaching" as opposed to learning, wouldn't it? Wasn't this where we came in?

Martin

Date: Sat, 4 Jul 1992 13:29:10 EDT  
Subject: Re: Plausibility of random reorganization

[Martin Taylor 920704 13:00]  
(Bill Powers 920702.1600)

This is a slight reprise in the combinatoric problem. Perhaps I should have titled the posting "combinatorics and conflict" but I like to keep the thread going with a constant title, for later reference.

>Putting the problem in terms of right versus wrong choices makes the  
>probability of organizing even one level of control seem incredibly small.  
>But I think this is the wrong way to set up the problem. EVERY form of a  
>perceptual function will yield a perceptual signal that is a regular  
>function of external events. There is not just exactly one combination of  
>inputs that will yield the "right" perception, with all others being  
>"wrong." There are many possible ways of perceiving a given environment  
>that will allow control, and many ways of exerting control that will have  
>at least some beneficial effect on intrinsic state. On the scale of  
>individual perceptual signals at the lowest levels, the number of equally  
>good alternatives must get astronomical. I think you're misstating the  
>combinatorial problem.

Yesterday I mentioned a symmetry argument that made Bill's point. I had intended to add a note about conflict that makes it more forcefully, but forgot. Here it is.

If the number of sensory degrees of freedom equal the number of action degrees of freedom, then it is possible to organize control so that different ECSs control independent percepts, and that all percepts can simultaneously be maintained at their reference levels. Under these conditions, there is a symmetry group of "right" perceptions and output links (a link, remember, in this argument has a value -1, 0, or 1). Any member of the symmetry group is an optimal control system. This reduces the combinatoric problem, but still leaves the number of optimal control systems very small in the universe of randomly connected control systems. I can't do the maths. Maybe someone else can. But that's not the end of the story. It's where I left the story yesterday.

The continuation of the story is based on there being far more sensory DoF than action DoF (degrees of freedom). Under these conditions, it is not possible for all percepts to be brought simultaneously to their reference levels. There is intrinsic conflict (I use the word advisedly, because I link it conceptually to the physico-chemical intrinsic variables that determine survival). There should be some kind of metric for the amount of conflict, that depends on the long-term average error over the control network.

If there is intrinsic conflict, there is some minimum possible level of average error greater than zero. When that minimum is exactly zero, only

the symmetry group described earlier will represent an optimum organization. But if some error is intrinsic, then a much larger group of organizations will be optimum or very close to optimum. This corresponds to Bill's words:

>There are many possible ways of perceiving a given environment  
>that will allow control, and many ways of exerting control that will have  
>at least some beneficial effect on intrinsic state. On the scale of  
>individual perceptual signals at the lowest levels, the number of equally  
>good alternatives must get astronomical.

I'm not sure how astronomical the number gets in comparison to the number that are possible and bad, but certainly this argument justifies Bill's final sentence (of the quoted paragraph:

>I think you're misstating the combinatorial problem.

I was.

The problem then becomes a practical one. To what degree was I misstating, and as a practical matter does the error affect the main thrust of the argument, that random reorganization, untargetted within a control hierarchy, is unlikely to achieve good results in a control system that is effectively interacting with the real world?

The question can, in principle, be addressed by computational experiments, but I think it would be hard in practice. To do the experiment, one would have to design a model world with controllably many degrees of freedom, in which the model hierarchy could be reorganized. There would have to be some effect of its behaviour on some simulated intrinsic variables, and so forth.

I'll leave for another day the complications that arise when the behaviour of the world is discontinuous.

Martin

Date: Sat, 4 Jul 1992 11:39:39 -0600  
Subject: Reorganization

[From Bill Powers (920704.0800)]

Martin Taylor (920703.1600) --

OK on "template." The image of a positive over a negative scene, however, is the way the term is understood by others. Let's just use "reference signal," as that's what we mean.

-----  
>I consider three situations, with respect to the organism (not  
>necessarily with respect only to an ECS).

>

>(1) some percept is being actively controlled to be as close to  
> its reference as the other controlled percepts permit.

>(2) some percept is not being actively controlled, but if  
> (2a) it departs too far from some reference, or

> (2b) it comes sufficiently close to some reference, then  
 > control relating to this percept must become active or bad things  
 > happen.

-----  
 In your explication, I had a bit of trouble with "controlling linearity"  
 until I translated into "controlling for a value of a linear function of  
 several variables." With that translation I get the picture, I think.

(1a), as you say, is an input function for a simple control system.

For (2a), you give

$$\begin{aligned} \text{Type (2a): } F(x) &= 0 \quad (|x| \leq t) \\ &= |x-t|^2 \quad (|x| > t) \end{aligned}$$

I don't think you really want that square in there, because it creates an  
 ambiguity in the error signal -- if the sign of the loop gain is right for  
 $x < t$ , it's wrong for  $x > t$ . I think what you want is more like

$$\begin{aligned} F(x) &= 0 \quad (|x| \leq t) \\ &= x - t \quad (x > t) \\ &= -(x + t) \quad (x < -t) \end{aligned}$$

This is a perceptual function with a dead zone.


For case 2b you give

$$\begin{aligned} \text{Type (2b): } F(x) &= t^2 - x^2 \quad (|x| \leq t) \\ &= 0 \quad (|x| > t) \end{aligned}$$

Again you have an ambiguity of sign due to the square -- you get the same  
 output whether  $F(x)$  is less than  $t$  or greater than  $-t$ , so there's no  
 indication of which way to apply the correction.

I think what you're trying to achieve here is a function that increases  
 from zero when  $x$  enters a region between  $-t$  and  $+t$ . But it looks to me as  
 if you're trying to get the perceptual function to accomplish the whole  
 control task. If what you want is an ECS that will avoid the region around  
 zero, and to move right or left according to whether  $x$  is above or below  
 zero, then you have to arrange for the perceptual function to indicate  
 proximity to zero (or whatever value is to be avoided) AND the sign of the  
 proximity.

I would use a function like  $F(x) = x/(a + x^2)$ , which looks sort of like  
 this:



```

|   *       *
|   * *

```

Then the appropriate control system would simply have a reference level of zero. If proximity to  $F(x) = 0$  were approached from the right, the error signal ( $r - x$ ) would be positive; if from the left, negative. A disturbance that pushed  $x$  toward zero from either direction would be resisted. If the disturbance were large enough to move  $x$  into the central region where the slope is reversed, the feedback would become positive and  $x$  would flip immediately, as fast as maximum output could move it, to the other side of the central region.

This same function could be put into the comparator, in which case we would need only  $F(x) = x$  for the input function, and any specified reference condition could be avoided (not just zero). Incidentally, if you just reverse the sign of feedback, with the above form of the comparator, the control system will seek the central region (normal control) and "give up" when the disturbance drives the error over the hump in either direction.

I don't see what we gain by calling these "similarity" or "difference" detectors. Is there some deep reason for which we need these terms, outside their normal connotation of relationships?

-----  
 RE: testing models

>The problem is that reality testing needs lots of resources, so the  
 >best truthsaying I can achieve is the most cost-effective procedure for  
 >me. But I can never be really sure that what I say MUST be true is not  
 >standing on some rickety foundation that might be swept away.

I haven't communicated clearly what I meant here. I think models should be tested against reality, of course. But before that, I think models should be tested to see if they really do what you think they will do. I've been illustrating that point above. While I didn't actually design an ECS using those input functions and simulate its behavior, I'm confident from previous experience that with my modifications the models would in fact control  $x$  relative to the specified reference level in the desired way -- whereas I couldn't design one that would actually control anything using your definitions in 2a and 2b. My version of 2a is, in fact, almost the same model I used in the "Gatherings" program for collision avoidance.

The input functions as you defined them would certainly produce a perceptual signal of the kind you were thinking of. But when you incorporate them into a complete control system, and test it via simulation, you'll find that you can only get control on one side of the reference condition; on the other side you'll have positive feedback. If you had actually carried these proposals to the point of simulation, you would have seen that there would be no point in comparing the model behavior with real behavior: the model itself wouldn't work as you envision.

That's what I'm talking about when I say that models have to be tested. They have to be tested to see if they will do what you are trying to accomplish, quite independently of whether that behavior would match real behavior. The first question that has to be answered is "What does a system

designed this way actually do?"

-----  
RE: Reorganization

>I still like the idea that one is inserting levels,

I don't, particularly. There are reasons for which I like it, but more for which I don't. How do you open up the connections from a higher to a lower system to insert a complete control system with all its connections to and from both the higher and the lower systems? This idea seems to me to entail enormous difficulties, whereas building from the bottom up eliminates those particular problems completely.

One of my main reasons for introducing the concept of levels of control, way back near the beginning, was the realization that it's impossible to go from thought to action in one jump. If I have a reference signal that says "I'm home," in words, and a verbal evaluation that says "I'm at my workplace," and an error signal that says "I'm not where I want to be," and a verbal plan of action that says "Therefore I need to get in my car and drive home," which set of muscles should I tense first? There is simply too great a gap in types of variables for this to work at all. The units of input and output in the highest system (propositions stated in words) are simply not appropriate for the units in which muscle tensions are adjusted (sensed stretch/tension). What's needed is a series of small steps in which higher-order units of error can be translated systematically into units of more detailed goals. I've tried to grasp what those steps might be, by seeing levels of perception and control that are "close" to each other in some sense, so I can at least imagine how an error of one type might be correctable by simple adjustments of reference signals for the next lower type.

In the present 11-level model, a gap leaving even one level out is hard to imagine bridging. If you leave out sequence, how can a logical error be translated into a goal for a different category? If you leave out configuration, how can a transition error be translated into a correction of sensations? Once you've seen intermediate levels, it's hard to imagine doing without any of them -- just seeing them shows why they are needed.

I don't see how higher levels of control can even exist without at least some of the lower ones existing first. Would it be possible to do logic without the underlying ability to identify sequences of actions, recognize a particular sequence? I don't think so. If the symbol string  $A \rightarrow B$  can't be distinguished from the string  $B \rightarrow A$ , how can logic be done at all?

Or consider relationships. With no objects, transitions, or events being perceivable, how could any relationships among them be perceived?

It seems to me that the very existence of any level of perception and control posits the existence of a lower level that is closely related to it, the higher system being at the same time a prerequisite for any higher level at all to form.

There's some developmental evidence that, if interpreted in a certain way, indicates that you may be right. For now, however, I don't see how you CAN be right.



-----  
ME:

>>It isn't necessary for any tumble to aim directly up the gradient; all  
>>that's required is a component in that direction.

YOU:

>That's true when there is only one degree of freedom for the controlled  
>percept: "satisfactoriness of the environment" and three degrees of  
>freedom for action. But the more common case is when the action  
>degrees of freedom are fewer than the perceptual degrees of freedom.  
>Then you have to worry about conflict, and components in those other  
>directions do matter.

I think you're looking at degrees of freedom in the wrong place -- the environment. In E. coli, there is only one output degree of freedom, the dimension being interval between tumbles. And there is only one input degree of freedom: time rate of change of concentration. If the output fans out into multiple effects in the environment, then those effects "fan in" again to an effect on the single input degree of freedom. All other effects are irrelevant to this control system. To see how this works, one only has to recall that the environment actually, according to physics, has almost an uncountable number of degrees of freedom in the quarks etc. that comprise it. If environmental degrees of freedom mattered, then even a simple stretch reflex couldn't work, because there is only a handful of motor signals reaching a given muscle, but they're affecting the states of zillions of quarks. How can just a few hundred signals adjust the degrees of freedom of zillions of quarks so as to produce a particular effect on a few hundred stretch receptors?

There's a different way to look at the problem of constructing a control system. Suppose you start with an arbitrary perceptual function. This function will create perceptual signals that can be recorded (built-in function). From the recordings you get reference signals (or in their absence, a reference signal of zero). The build-in comparison function produces an error signal. This error signal can be connected to produce many outputs. Now, which connections should be selected, and what should their signs be?

For any given connection, the resulting effect on the perceptual signal can be only to increase it, decrease it, or leave it the same. There's no other way the perceptual signal can vary -- it can't go sideways. If a particular connection has no effect on the perceptual signal, it can be undone again. If the connection causes a positive-feedback effect on the perceptual signal, it, too, can be undone again. The only connections remaining will be those that have a negative feedback effect.

It doesn't matter at all how any particular connection results in a negative feedback effect -- it could be a direct path through the environment or an indirect one. It could be a single path or 10,000 paths in parallel. It could involve one lower-level control system, many, or none (if this is the lowest level of control). The number of connections might be the bare minimum required for adequate control, or it could lead to highly redundant overkill. None of that matters, if control is sufficiently good to turn off reorganization. Of course if the arbitrary perception can't be controlled with ANY set of connections, the result won't correct intrinsic error, and the whole system will be reorganized away.

If you look at the world from inside any one control system, there is always just ONE input degree of freedom (the perceptual signal) and ONE output degree of freedom (the error signal).

-----  
 If during its formation a control system comes into conflict with other control systems, it will experience error. Reorganization will continue until it doesn't experience (too much) error any more. So each control system involved in conflict with others will be reorganized until the overall conflict is minimized. The more control systems there are that already are controlling, the more error will be induced by conflict into a new control system and the less into the existing ones, for all the existing ones are resisting disturbance together. There's really no need to keep reorganization from affecting control systems that are already working. Generally, when a new control system is added, all control systems will have to be modified somewhat, not just the new one.

Let me try to generalize this.

In a collection of control systems at a given level, there is the general requirement that all errors should be as small as possible. If, as I suggested, one intrinsic variable is simple total absolute error (perhaps at a given level), and the intrinsic reference level for total error is zero, then we have essentially a simultaneous equation in  $n$  unknowns. The perceptual signals represent the values of  $n$  functions of environmental (or lower-level) variables, the  $n$  reference signals constitute the desired values of the variables, and the coefficients are the adjustable feedback connections and perceptual weightings.

I'm assuming now that reorganization is specific to each level.

In general, for a set of arbitrary connections, there will be a difference between all perceptual signals and their respective reference signals. The total absolute (or squared) error indicates the distance in hyperspace between the actual values of signals emitted by the perceptual functions and the desired values. The problem is then that of adjusting the coefficients in the  $n$  equations for a minimum total error.

In the absence of an analytical solution (ha!), an approach to a minimum-error solution can be achieved by various methods of descent. One method is what I call reorganization -- hill-climbing, I suppose. Actually we're looking for the bottom, but it's the same idea. All the coefficients are incremented or decremented at some nominal interval of time by some small randomly-selected amount. This will either increase or decrease the total error. If the error is increased, the random change in coefficients is postponed, so the same increments continue to be added to the coefficients in each interval (bear with this plodding pace -- I haven't actually worked this out in detail before). If the error increases, the interval between random changes decreases. Thus we get a biased random walk through hyperspace which will converge toward the minimum error condition. Local minima will simply slow the process -- as long as significant total error remains, random changes will continue, and eventually a run of similar changes will pop the system out of the minimum. Eventually, of course, could be too late -- but that's a different problem.

Well, I'm going to think about this some more, and try out a method like this for solving a big set of simultaneous equations. Don't tell me it's all been worked out ages ago. I want to see for myself.

-----  
>>Your point about reorganization being called upon to make  
>>rapid correct decisions simply doesn't hold up: that's not necessary.

>It does, if the system being reorganized is actively controlling, with  
>reasonably high gain.

If we think of reorganization as instituting small changes in the parameters of control, there's no reason why a high-gain control system can't go right on controlling. Even in category control (to pick up that thread), slight changes in perceptual parameters would just move the boundaries of the category a little. I think those boundaries are fuzzy anyway, so we don't really have any discrete-variable problems here.

Look at it the other way around. If reorganization is to work as I propose, it CAN'T work fast. If rapid adaptations to changes in conditions are needed, reorganization can't do the trick. What's needed then is to learn some new mode of control. In the reversal experiments Rick and I did, we found that the human control system can reverse the sign of feedback in the tracking system within 1/2 second of a reversal of the external connection. So that is a learned control system, not a sign of reorganization. That mode of control is a PRODUCT of reorganization; once it's established, reversals no longer produce enough error to start the process of reorganization.

This is, in fact, how I think higher levels of control come about. situations arise in which existing control systems can't correct intrinsic error well enough any more, even though they're all keeping their own errors small most of the time. They're not controlling the right ASPECT of the environment. So a new level of control begins to form, setting the formerly fixed or randomly varying reference signals in a systematic way to control new perceptual variables that are functions of the old ones. This adding of levels continues as long as there is a need and as long as there are available neural components of the right kind still unorganized. You arrive at the top level when you run out of new layers of neurons that permit new types of control.

-----  
>Apart from the statistical convergence that has nothing to do with  
>control, due to the natural redundancy of the real world, I have  
>envisioned the possibly controllable percepts (not degrees of freedom  
>for perception) as growing in number as we go up the levels, before  
>reducing at the highest levels. Sort of barrel-shaped, rather than  
>conical.

That's how I visualized it in BCP -- widest in the middle levels in terms of POTENTIALLY controllable perceptions.

>Since we are (by agreement) working from the lowest common denominator  
>of no prior construction of ECSs, we have to include evolutionary  
>development here. No matter what the evolutionary level, the prime  
>concern is to maintain those intrinsic variables near optimum long  
>enough to pass on a structure description to the next generation

>(whether it be by cloning, seed-spreading, or whatever). The primary  
>control system has this function.

As I see the reorganizing system, it is concerned with controlling INTERNAL variables only. To reply to a previous comment of yours, I see these variables as including variables not available to the senses, even to proprioception. There is probably overlap. But for variables that are simultaneously intrinsic variables and sensory variables, the reorganizing system has build-in sensors and reference signals, while the sensory-based system, the growing learned hierarchy, does not. The learned systems, for example, can come to seek those conditions that lead to sensory hunger signals; the reorganizing system cannot diet.

In infancy, the human organism is kept warm, dry, fed, cuddled, and so on by external agencies. It does not need to learn behaviors that will accomplish these things; they are done for it. So its intrinsic state is maintained by the parents. Evolution provides only the biochemical control systems that maintain physiological integrity, and perhaps a few temporary behavioral systems, as for nursing.

Fortunately, parents are not able to keep all intrinsic errors exactly at zero. In fact, they probably provide just a bare minimum of support, leaving very large errors uncorrected. But these errors are not so large or of such a type that the infant can't survive.

>My preference is for a single hierarchy, in which the primary control  
>system has been elaborated to effect its control through the provision  
>of reference signals to other ECSs. Yours is for the primary control  
>system to be separate from another hierarchy, and to effect its control  
>by blind modification of that second hierarchy.

I guess that's it. I don't know which is right; yours has its points. Perhaps what we need to do, as we're not likely to resolve this conflict experimentally, is to find a way to talk about reorganization so it doesn't matter which is the right model.

However we end up talking about reorganization, we have to arrange for somatic states of the organism to have a very powerful directing effect on what control systems are acquired. Logic and principles often bow to hunger. How does that work? I can see how my model would do it, but how does yours?

-----  
>Have you modelled to reorganization of a moderately complex hierarchy?  
>That would be a lovely demo, if you have.

No, but Rick Marken is experimenting with it using his spreadsheet multilevel model.

-----  
>>I have proposed a version of targeting based on the  
>>phenomenological idea that awareness directs reorganization to problem  
>>areas. But having no model of awareness or attention, I haven't pushed  
>>that very hard.

>This would be the effect of "teaching" as opposed to learning, wouldn't

>it? Wasn't this where we came in?

Yes, and I expect we'll go around a few more times before all is said.

-----  
 RE: "millions and millions"

That was NOT saganesque. Sagan said "Billions and Billions." I am a thousand times less saganesque than Sagan.

-----  
 Best,

Bill P.

ems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>  
 From: "William T. Powers" <POWERS\_W%FLC@VAXF.COLORADO.EDU>  
 Subject: Off time

Starting at 6:30 on the 5th and continuing until some time on the 8th, my mainframe will be down for modifications or something. So I'll be off the net unless I can patch in somewhere else. Actually I hope to get a lot done. Then the following week I will also be gone for 5 days, camping and giving a talk to the International Society for Systems Science in Denver (the 15th). This will probably get me so far behind that I won't catch up until after the meeting.

While the cat's away...

Best, Bill p.

Date: Sat, 4 Jul 1992 12:25:58 -0600  
 Subject: Re: plausibility of random reorganization

[From Bill Powers (920704.1200)] Martin Taylor (920704.1300) --

RE: Plausibility of random reorganization.

I like that "symmetry groups." Maybe some day I'll understand what it means. But I get the idea.

>The continuation of the story is based on there being far more sensory >DoF than action DoF (degrees of freedom). Under these conditions, it >is not possible for all percepts to be brought simultaneously to their >reference levels.

This reminds me of a problem that comes up in speaking of controlled variables that are functions of sets of controlled variables. People have asked me "What happens when disturbances come along that alter the inputs without changing the perception?" One always tends to think of some concrete situation, and unconsciously puts into it more than the definitions require. If you're controlling  $Ax + By$ , then ONLY that quantity is controlled. If a disturbance changes  $x$  and  $y$  in the right proportions to leave  $Ax + By$  with the same value, then  $x$  and  $y$  simply change, with no resistance from the control system. ONLY what is perceived is controlled.

It's true, as you say here and have said before, that there are far more input variables than output variables, particularly at the lowest level where the ratio is (if you say so) 20,000:1. So clearly, not each individual input signal can come under independent control.

What happens is that BUNCHES of input signals are perceived as a single signal, and THAT SIGNAL is brought under control. This leaves an enormous number of ways in which individuals within the bunch can vary, but that's irrelevant to any given control system. In those dimensions, the bunch is simply not controlled.

This leaves, of course, many potential bunches that are never perceived, and thus are never controlled. In the uncontrolled degrees of freedom, those bunches simply change when they are disturbed.

Some people learn to perceive some functions of their inputs among all those that are possible. This is why individuals are different -- they perceive and control the same environment in slightly or greatly different terms. They don't need to find THE way of perceiving and controlling; all they need is A way that is sufficient to form a coherent hierarchy good enough to sustain life.

Most potential percepts never become actual percepts.

One function of higher-level systems is to employ subsets of the available lower-level systems in compatible symmetry groups.

>There is intrinsic conflict (I use the word advisedly, because I  
>link it conceptually to the physico-chemical intrinsic variables that  
> determine survival)...

Hmm .. coming closer to my concept. Goody.

>The question can, in principle, be addressed by computational  
>experiments, but I think it would be hard in practice. To do the  
>experiment, one would have to design a model world with controllably >many  
degrees of freedom, in which the model hierarchy could be >reorganized.  
There would have to be some effect of its behaviour on >some simulated  
intrinsic variables, and so forth.

This is exactly what has to be done, hard or not. Of course we don't have to start with the most complex possible physical world model. I'd settle for three environmental variables and three intrinsic variables variously dependent on the environmental ones. I've wanted to get moving on this kind of computer experiment for a long time but have always had some other project that seemed more essential for teaching purposes. Maybe, after the cockroach and the arm and the motion-illusion experiment and the formant tracker ... I seem to think that happiness is a bunch of big error signals. We really would be getting along much faster if there were more people actually working at this sort of thing.

>I'll leave for another day the complications that arise when the  
>behaviour of the world is discontinuous.

Not much of that in the physical world at the level where we experience it.

Best, Bill P.

Date: Sun Jul 05, 1992 10:26 pm PST  
Subject: Canada

[From Hank Folson (920705)]

Martin Taylor 920702 15:30 to Joel Judd:

>If you don't know that Canada is in the middle of a constitutional  
>crisis, your news organizations must be pretty poor.

As a Canadian who lives in the U.S., I have been following the American media on this. The subject comes up rarely. The U.S. media/government appear to have a goal of breaking Canada up. The articles generally speak favorably of Quebec as an independent country. They never suggest reasons why Canada should not break up, nor what problems would be created by the breakup. Perhaps the CSGnetters can suggest why from a PCT point of view the U.S. is controlling for this when ".one country, indivisible.." is part of the American System Concept (Pledge of Allegiance). If Canada ought break up on racial lines, there are several similar areas that should naturally carve their own countries from the U.S: The African southeastern area, a small Cajun/French area, and the Indian/Hispanic southwest are obvious parallel examples. These groups have even more reason to control for their own country than the people of Quebec do, as the Quebecois already have more political power than other Canadians.

Martin, what do you think the Quebec politicians, corporations and people are controlling for? What would you guess is their System Concept, and what error signals might they be trying to reduce? What do you think the rest of Canada is controlling for?

Perhaps proposing that Quebec be given back to the indigenous peoples might be the test for the controlled variable here? They have a longer and stronger claim than the French Canadians do to the land. This should meet with strong American approval, as the U.S. has controlled for the return of Israel to the Jewish people 2,000 years after they were driven from their land. But perhaps there is some internal conflict here, as sovereignty for American indigenous people is not encouraged by U.S. policy even in this quincentennial year.

Hank Folson

Date: Mon Jul 06, 1992 7:10 am PST  
From: Richard Robertson  
MBX: urrobert@uxa.ecn.bgu.edu  
TO: \* Dag Forssell / MCI ID: 474-2580  
Subject: CONGRATS

[From Dick Robertson]

Hey Dag, I just read your latest brochure. It looks fabulous. Not only do you present a view of HPCT in an enlightening historical context, but your sales message seems to me to come through beautifully too (of course I'm no expert on that part). But I think a 1% return must be terrific. The csg group must be a heck of a lot less than 1% of all the scientists it has been offered to.

Best-, Dick

Date: Mon Jul 06, 1992 8:11 am PST  
Subject: Re: Reorganization

[Martin Taylor 920706 11:30]  
(Bill Powers 920704.0800)

Bill, as usual, caught my sloppy descriptions. In this case it was the error functions. I described a potential function, not a "force" function. The error should be the derivative of the functions I gave for the type 2 systems, so that the signs are different on the two sides of zero error.

I won't reply to the rest of Bill's postings until he comes back on the net.

Martin

Date: Mon Jul 06, 1992 8:44 am PST  
Subject: learning mail

Martin tells me that there have been a number of posts in response to my inquiry on learning. I would appreciate it very much if a copy of each message could be sent to my personal mailbox when a message is sent out. I can't seem to access the CSG news and my box is too small to be connected to the net. Thanks.

Mark

Date: Mon Jul 06, 1992 9:07 am PST  
Subject: Internet

>I can't seem to access the CSG news and my box is too small to be connected  
>to the net. Thanks.

Who told you this? Every PC and Mac in our lab is scheduled to be connected to Internet, and there is free software out there that gives you just about all the functionality of network access. Go bang on somebody's door!

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca



Date: Mon Jul 06, 1992 9:28 am PST  
Subject: Apologies to CSG

My apologies for send a personal reply (about Internet) to the entire CSG distribution. However, I noticed a problem in the headers that is at least partly responsible for my error. I always check the header for mail, and the original message looked like it had a legitimate header with a From: line giving the sender's address. Unfortunately the mailer looks for a From line without the colon. So my aplogies are tempered with a suggestion that perhaps the headers could be made a bit more informative? Other mailing lists for example have a line "Originally From:" or "Reply To:" which is much clearer.

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Mon Jul 06, 1992 10:47 am PST  
Subject: closed loop

from Ed Ford (920706:11:31)

To all - sent yesterday, but I don't think it got through....

Closed Loop should be ready for distribution in about 10 days. Even with the flooding problems Greg Williams has had in and around his home, he got it finished and he and Pat planned to give it a final editing today. Since there is nothing for the CSG newsletter and to save the \$1 per issue mailing costs, I plan to hand it out at the conference. The balance will be mailed Monday, Aug. 3rd.

Ed Ford

Date: Mon Jul 06, 1992 11:21 am PST  
Subject: Direct mail - update

[From Dag Forssell (920706-1)]

It appears the list server has been down over the weekend. That does not stop life, of course. I have massaged my letter almost constantly. Many small changes add up. Here is what looks like the final version 9. Production this pm.

\_word\_ means underline, right? Italicized titles not shown. The real thing looks better, with 11pt New Century Schoolbook typeface. Many paragraphs optimized (given the margins) to avoid orphans and unnecessary lines.

Copyright 1992 Dag Forssell. All rights reserved. Permission is granted for quoting within the mailing list CSG-L, academic (but not commercial) classroom use and for use in Closed Loop and other publications of CSG,

provided that credit is given. Hard copy will be sent upon request, as will the updated intro package (current revision June 20).  
(Letterhead) (Page 1)

Edward E. Ford, President  
Control Systems Group  
10209 North 56th Street  
Scottsdale, AZ 85253

July 6, 1992

Dear Mr. Ford:

You may be interested in the only fundamentally new perspective on people that has been proposed since 1637. Adopting it can mean improvements for your bottom line, productivity, quality and morale - particularly if you deal with knowledge workers and would like to lead them in an effective and mutually satisfying way.

Costly people problems exist at all levels in American industry. Dr. W. Edwards Deming, pioneer in Quality Management, writes in "Out of the Crisis," page 85:

"In my experience, people can face almost any problem except the problems of people. They can work long hours, face declining business, face loss of jobs, but not the problems of people. Faced with problems of people (management included), management, in my experience, go into a state of paralysis, taking refuge in formation of QC-Circles and groups for EI, EP, and QWL (Employee Involvement, Employee Participation, and Quality of Work Life)... There are of course pleasing exceptions, where the management understands... participates..."

At the core of the design of any social or business organization lies some assumptions about people. If you identify, question and then change these assumptions, the impact on the design and function of your business organization may be significant.

The basic perspective from 1637 which still dominates our science and culture is the cause-effect idea that events impinging on organisms cause them to behave as they do. The new perspective is called Perceptual Control Theory, or PCT. William T. Powers, who has developed the theory, writes in "Living Control Systems, Vol II":

"Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to Humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains why it is so hard for groups of people to work together even on something they all agree is important. It explains what a goal is, how goals relate to behavior, how behavior affects perceptions and how perceptions define the reality in which we live and move and have our being.

Over, please...

Edward E. Ford

July 6, 1992

Page 2

Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single testable concept of how living systems work."

Understanding people no longer has to be complex and confusing. PCT can be taught in simple form with a comprehensive management application in one day and in more detail with leadership applications in three.

An executive gains insight that allows him or her to inform, influence, align and lead people with mutual respect. S/he can teach people to be more effective and cooperative. Employees can be more satisfied, while the company as a whole responds better to the leader's direction and becomes more productive.

This perceptual control perspective will also make it much easier to understand and teach Total Quality Management programs. For instance, when you review the 14 points of the Deming Management Philosophy, you will see that each point describes aspects of a system of control. Constancy of purpose, pride in workmanship, lack of fear, quality, productivity.... all can be seen as manifestations of effective individual control.

Personally, I am convinced that PCT - once it is widely understood - will have the same kind of impact in the life sciences as Newton's theories did in the physical sciences. Besides a consuming interest in this new development, I have 25 years management experience in engineering, manufacturing, marketing and finance. My formal education includes an MBA from the University of Southern California and a Masters degree in Mechanical Engineering from Sweden.

I have developed the Purposeful Leadership<sup>TM</sup> programs to explain PCT and apply it to skillful use of diagnostic tools that give an executive the capability to work on productivity. That includes effective communication, teaching effectiveness, resolving conflict, supporting self-motivation in employees, team building, Total Quality Management, leadership insights, effective performance appraisals, effective selling concepts, and development of corporate and individual mission statements. The executive learns how to build confidence, build trust and develop caring relationships.

The basic principles can be taught in a day to any attentive person, who can also verify them. People trained in the "hard" sciences will appreciate the scientific approach and elegant simplicity of the theory, and everyone will be able to begin applying the principles as soon as they understand the underlying model and have had some instruction and practice.

I would like to describe this perspective so you get the point immediately, but this is an impossible Catch-22 challenge, because it is a different concept altogether from what predominates in our world today. Until you understand the principles, you cannot understand at all. I need a few hours in class to explain and illustrate the principles.

When you request it, I will send you a do-it-yourself concept demonstration /test. Until then, perhaps I can indicate how I believe this new perspective

fits into the scientific evolution of the life sciences with the following illustrative analogy:

In an era when "everyone knew" that the earth was flat, scientific explanations were developed for navigation and astronomy. Many problems with those explanations persisted, but people worked around them.

Continued....

Edward E. Ford

July 6, 1992

Page 3

I cannot say what "everyone knows" about human behavior, but experts on the subject employ the 17th century perspective of cause and effect to guide their research. Any book on experimental psychology tells you that the scientific method to learn about behavior is to set up an experiment, establish initial conditions and then vary the stimulus (independent variable) and measure the response (dependent variable).

(This would be a valid scientific method if in fact animals and people were cause-effect organisms. But our demonstration will show you in a few minutes that they are not).

With this scientific method our experts have done many experiments and reported explanations which are now part of our language, culture and management practices.

There have always been natural leaders, successful salesmen, wise parents and good communicators. But it is rare that they can explain what they do and why. Their insight and skill seems intuitive. Human behavior practitioners and many executives make an effort to master people skills. They depend on a variety of personal experiences, interpretations and training programs (based on the experience of others, not science), to develop effective personal approaches for dealing with people.

(Imagine how good they will be when they get good understanding that applies every time. With PCT, the executive can learn to function as well as those intuitively wise people. With practice even better, since s/he will have greater insight).

Many problems with the experts' scientific explanations persist despite all the research, but people work around them. Lack of success indicates that we lack a good model or "paradigm" to help us understand why people do what they do. In our ignorance, we often spend our energies in debilitating conflict instead of in productive cooperation.

When Copernicus and then Galileo introduced the fundamentally new insight that the earth is round\_ (it has always been round), the problems of navigation and astronomy were placed in a new light\_. Science started over from a new concept. The new insight did not invalidate the common sense observation that the earth appears flat locally.

Most experts on the old science could not comprehend the new paradigm, because they had already internalized the flat paradigm in all its details

as their personal reality. With time the experts died off, and new ones grew up, embracing the new paradigm on its merits because it solved many of those persistent problems.

Isaac Newton's "Principia Mathematica," published fifty years after Galileo, was resisted in the same way, just like all dramatically new approaches have been. It took fifty years for it to be fully accepted. The evolution of science is much more than a steady accumulation of knowledge! The process depends on creativity. The opportunity for a revolution arises when a current paradigm fails to solve problems and competing paradigms are offered to provide better explanations. A struggle of many decades typically takes place. Trained scientists continue the development of the existing paradigm as usual while outsiders and early converts champion a new one.

Over, please....

- 1 The phenomenon and process is described in Thomas Kuhn's seminal book: "The Structure of Scientific Revolutions," which introduced the term "paradigm."

Edward E. Ford

July 6, 1992

Page 4

The 20th century understanding of Perceptual Control Theory (people always control their perceptions) provides a fundamental new insight that puts the problems that result from human interactions in a new light\_. Science is starting over from a new concept.

The new perspective does not invalidate any wise common sense observation or practice. It does provide an enhanced understanding of seemingly intractable problems. It suggests new diagnostic tools and shows why cookbook rules for behavior (programs which tell you what to do under certain circumstances) do not always work.

Perceptual control is as incomprehensible at first glance to a person trained in cause-effect thinking (which we all are in our culture) as the idea that the earth is round was to a person trained in the details of a flat earth. The demonstration shows this clearly.

Of course, this is only because it has never been noticed or explained. When you understand the principles, you will be aware of perceptual control in operation in yourself and others. You will also notice that an understanding of PCT contains an explanation for the illusion of cause and effect in people, just like the understanding that the earth is round contains an explanation for the illusion of a flat earth.

Another illustrative analogy is to say that we live in a maze where only the walls and passages are visible to us. The perspective of Perceptual Control allows us to rise above the maze and see the structure. We can then set and reach our goal much easier.

Perceptual Control Theory is already well developed. But no doubt it will take time -well into the 21st century - before this breakthrough is known, understood and embraced by a majority of experts. You can take advantage of what "everyone will know" in the 21st century right now to improve your company's competitive position. But because it breaks with the past, you must be willing to think for yourself to do it. You will participate in a scientific revolution when you understand and adopt it.

Some people will think that the term "control theory" promises a new way to control other people. It is precisely the other way around. We show how people control themselves at all times. When you understand PCT, you can work with people rather than get into conflict despite the best of intentions.

Please request the free introductory 39 minute audio tape with script and illustrations. It demonstrates the basic concept and explains the benefits, applications, background and content of our programs. The demonstration/test allows you to find out if your associates can recognize simple control in action. (I bet they can't).

When you receive the introduction, I think you will find the demonstration enlightening and entertaining. Please feel free to share it with your technical, operations and sales managers at any level for their evaluation. This is a win/win program to increase the understanding and effectiveness of anyone who deals with people.

Sincerely,

Dag Forssell  
(Signed)

Date: Mon Jul 06, 1992 11:23 am PST  
From: Dag Forssell / MCI ID: 474-2580  
Subject: Groups - Canada

[From Dag Forssell (920706-2)]

Hank Folson (920705)

>The U.S. media/government appear to have a goal of breaking Canada up.

>What do you think the rest of Canada is controlling for?

Hank, I find your post most enjoyable. But you mix speculations about the controlled perceptions in individuals with rather large "groups" as shown above. May I recommend Bill's article in LCSII: CT psychology and social organizations.

As a Swede with some limited knowledge of what is printed "over there" as compared to "here at home," I have the notion that each individual reporter writes about something s/he thinks the readers will read. I doubt that the U.S. media is controlling for anything with respect to Canada. If "It" did, why would it be in collusion with "gouvernement?" When you have read Bill's piece, you will see "media" and "gouvernement" as well as "the rest of Canada" as a soup of individuals. Social control systems do not exist, except as a (pre-PCT) construct in your mind. There is only the soup of individuals!

Dag

Date: Mon Jul 06, 1992 11:31 am PST  
Subject: Re: plausibility of random reorganization

[Martin Taylor 920706 14:10]  
(Bill Powers 920704.1200)

I said I'd wait until Bill rejoined the net. But I guess he will get this anyway, so what the heck...

> So clearly, not each  
> individual input signal can come under independent control.

> What happens is that BUNCHES of input signals are perceived as a single  
> signal, and THAT SIGNAL is brought under control. This leaves an enormous  
> number of ways in which individuals within the bunch can vary, but that's  
> irrelevant to any given control system. In those dimensions, the bunch is  
> simply not controlled.

> This leaves, of course, many potential bunches that are never perceived,  
> and thus are never controlled. In the uncontrolled degrees of freedom,  
> those bunches simply change when they are disturbed.

Quite true. There are  $O(N^2)$  bunches of two inputs,  $O(N^k)$  possible bunches that involve  $k$  inputs, even if we limit the relationships to simple difference functions like  $x-y$ . If there are so many more sensory degrees of freedom than output degrees of freedom, how many more possible "bunches" (i.e. ways of choosing individual degrees of freedom) are there? Sagans!

> Most potential percepts never become actual percepts.

And

> Some people learn to perceive some functions of their inputs among all  
> those that are possible. This is why individuals are different -- they  
> perceive and control the same environment in slightly or greatly different  
> terms. They don't need to find THE way of perceiving and controlling; all  
> they need is A way that is sufficient to form a coherent hierarchy good  
> enough to sustain life.

All of which is true, and none of which is relevant to the degrees of freedom argument about the necessity for switching which percepts are being actively controlled at any moment. It is relevant to the question of the plausibility of the reorganization argument, but tangentially.

On modelling reorganization:

> I'd settle  
> for three environmental variables and three intrinsic variables variously  
> dependent on the environmental ones. I've wanted to get moving on this kind  
> of computer experiment for a long time but have always had some other  
> project that seemed more essential for teaching purposes.

I wouldn't settle for three and three. I think that would avoid the central issue. I might settle for thirty environmental variables that behaved in a redundant manner (unknown initially to the control system that had three intrinsic variables).

Let's try to design an experiment and see if we can get anyone (including ourselves) actually to try it.

-----

>>I'll leave for another day the complications that arise when the  
>>behaviour of the world is discontinuous.

>

>Not much of that in the physical world at the level where we experience it.

>

Depends what you mean by "much." I think a chair is discontinuously distinct from a table, even if you can sometimes use the same object for either function.

In that context, I would have said that most of what we perceive in the physical world is discontinuous. Even at lower levels, edges are one of the most important elements of perception (edge detectors seem to be very peripheral in most of our sensory systems, perhaps lying outboard of the control hierarchy). Edges signal discontinuities in things represented by percepts that must be controlled. You can move an object smoothly up to an edge, at which point it may fall off a table, be blocked by a wall, go up in flames, or in some other way demonstrate a discontinuity in the possibilities

for control. Even if we accept "much" in your statement, nevertheless I would think that the discontinuities are at least as important as the regions of smooth control.

Martin

Date: Mon Jul 06, 1992 12:26 pm PST  
From: Bruce E. Nevin  
MBX: bnevin@ccb.bbn.com  
Subject: letter

Dag,

A couple of quick comments based on a very fast scan before I race out:

\_Word\_ generally means italics. If you lack italics, you use underscore, hence the \_word\_ convention.

\_Principia Mathematica\_ is a book title and should therefore be in italics rather than in quotes.

Bruce

Date: Mon Jul 06, 1992 12:53 pm PST  
Subject: news/mail



Anyone at the U of I who has been using the news program to access Info.csg:

If you're using the program Gary gave you, is it working for you? It used to work at all CSO sites, but now it works nowhere. Any suggestions. I just click on News and then click on Info.csg. Nothing happens. A display of recent posts is what used to happen.

Mark

Date: Mon Jul 06, 1992 5:31 pm PST  
From: mmt  
MBX: mmt@ben.dciem.dnd.ca  
TO: \* Dag Forssell / MCI ID: 474-2580  
Subject: Re: Martin's diagram

What makes you think your posting "didn't take?" Didn't you get my response of July 3? It was distributed back to me. I suspect your feed is screwed up. Someone said that the CSG-L mailer was down over the weekend. I wouldn't know, because my link machine was also down from Saturday afternoon until this morning.

On July 3 and 4 I received 6 postings from CSG-L after yours: 4 originated by me and 2 by Bill Powers. I don't think the fact I originated them has any bearing on making me a preferred recipient, so if you didn't get all 6, something is wrong. Before your posting got here, there were 5 postings, 1 from Oded Maler, 1 other from you, 1 from another MCI source (Greg?), one from Dick Robertson, and one from me. So there were 12 postings in all that I received on July 3 and 4. (only one on July 4).

Martin

Date: Mon Jul 06, 1992 8:11 pm PST  
From: Dag Forssell / MCI ID: 474-2580  
Subject: List server

[From Dag Forssell (920706 2100)]

The following may be of general interest:

Martin Taylor (Jul 06, 1992 5:31 pm PST) Direct

>Subject: Re: Martin's diagram

>

>What makes you think your posting "didn't take?" Didn't you get my response  
>of July 3? It was distributed back to me. I suspect your feed is screwed  
>up.

>Someone said that the CSG-L mailer was down over the weekend. I wouldn't  
>know, because my link machine was also down from Saturday afternoon until  
>this morning.

>

>On July 3 and 4 I received 6 postings from CSG-L after yours: 4 originated  
>by me and 2 by Bill Powers. I don't think the fact I originated them has  
>any bearing on making me a preferred recipient, so if you didn't get all 6,  
>something is wrong. Before your posting got here, there were 5 postings,

>1 from Oded Maler, 1 other from you, 1 from another MCI source (Greg?), one  
>from Dick Robertson, and one from me. So there were 12 postings in all that  
>I received on July 3 and 4. (only one on July 4).

Thank you for this note. Ed Ford said something about "not taking" this morning. I know Hank Folsom got nothing all weekend. Like me, he is on MCI mail. The last post I got was Gary Cziko on Friday the 3rd at 12:42 pm. I did get Hank's of Sunday the 5th 22:26 next. Certainly some feed somewhere was screwed up.

A long time ago, Gary said that it is not doable to get your own post in return from CSG-L. I would like to get that kind of feedback. How do you get it?

I will wait for the first week to be over, then download 9207A. Thanks for telling me that I have something to look forward to. It certainly was too quiet.

Thanks again, Dag

Date: Tue Jul 07, 1992 7:24 am PST  
Subject: Ecological emergence and control

I'm on the thesis committee for a grad student who is studying the interannual variability of fish catches. Past studies have shown that the length-frequency distributions (curves showing the number of fish of each size) for all the fish in an area show much more regularity than the distributions for individual species.

Two questions arise:

1. Can this regularity be classified as an emergent property of the system, or, more precisely, what kind of tests would be required to show that the observed regularity is an emergent system property and not explainable as the result of combined single-species curves?
2. How would one be able to show that this regularity is the result of a control mechanism? Some of my colleagues feel that the existence of a constancy in behaviour (often referred to as a "conservative property", but that is confusing to people who studied physics) must imply control, but I disagree.

With regard to the second point, I've given him the example of finding that the temperature inside a house is relatively constant throughout the year. Does this constancy imply a control mechanism? It depends. If the house is in the tropics, probably not. If the house is in the arctic and it can be shown that the insulation is not perfect, there must be a control system. Even if the observer cannot actually identify how the heat is controlled (perhaps he/she is looking for a stove and icebox and doesn't know about central heating and air conditioning), the existence of a control system is clear. Is there a general method for identifying the existence of feedback and control if the precise mechanism cannot be identified?

All comments will be passed on to the student. I hope that we will be able to refine the objectives and methods of the thesis in this way.

Thanks in advance, Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Tue Jul 07, 1992 7:52 am PST  
From: mmt MBX: mmt@ben.dciem.dnd.ca  
TO: \* Dag Forssell / MCI ID: 474-2580  
Subject: Re: List server

>A long time ago, Gary said that it is not doable to get your own post in  
>return from CSG-L. I would like to get that kind of feedback. How do you get  
>it?

I think it is because the CSG-L mail for DCIEM is sent to a common point and redistributed to the participants, so the mailer doesn't know that it is sending my own things back. This is a GOOD THING, because I was always forgetting to save what I wrote.

Martin

Date: Tue Jul 07, 1992 9:01 am PST  
Subject: Re: Ecological emergence and control

[Martin Taylor 920707 11:45]  
(Bill Silvert Date? 920707?)

> Is there a general method  
>for identifying the existence of feedback and control if the precise  
>mechanism cannot be identified?

It has been called The Test. If you identify a variable that you suspect may be under control, you disturb it and see if it is returned to its former state. I'm not sure this is adequate, since if you disturb a marble sitting at the bottom of a bowl, it will return to the bottom. But that is the method that has been recommended by Bill Powers and Rick:

The following is an explicit description, from Bill Powers (920215 09:00).  
We were discussing whether the perception of "wetness" in rain was a controlled variable:

-----

>Correct. Here's how to go from there to identifying a control system  
>without a model.

>

>First, characterize observables in the environment. Note the degree of  
>wetness W. Note effect A of action on same variable. Note effect R of rain-  
>rate on same variable (measure both rain and action in terms of effect on  
>wetness). Note that  $W = R - A$ . Equation says that if  $R = 0$  and  $A = 0$ ,  $W =$   
>0. Also predict that if  $R > 0$  and  $A = 0$ ,  $W > 0$ . Prediction verified by

>observation: so far so good. If  $R = 0$  and  $A > 0$ ,  $W < 0$ . Prediction no good  
>-- can't have negative wetness. So equation applies only when  $W \geq 0$ .  
>  
>Now postulate that  $A = g(W - W^*)$ . Amount of action is based on degree of  
>wetness in relationship to some base wetness,  $W^*$ , to be determined from  
>data. This leads from  $W = R - A$  to  
>  
> $W = R - g(W - W^*)$ .  $g$  is the form of the unobserved organism function.  
>  
>Solve  $W = R - g(W - W^*)$  for the form of function  $g$  that will satisfy the  
>observed relationship between data points for  $R$  and  $W$ .  $W^*$  is the value of  $W$   
>at which action has zero effect on wetness. The constant  $W^*$  is the  
>reference level (not the reference signal, which would be a variable inside  
>the actor). We find a function  $g$  that will fit  $W = R - g(W - W^*)$  to the  
>recorded values of  $W$  and  $R$ , obtaining  $W^*$  from the data.  
>  
>The function  $g$  describes the effect of stimulus  $W$ , via the actor, on action  
> $A$ . Calculate the partial derivative  $g(W - W^*)$ , the slope of the stimulus  
>effect on behavior, and the partial of  $W$  with respect to  $A$ , the slope of  
>the effect of behavior on the stimulus (which is observable). If the  
>product of these two partials is a large number, and if that product is  
>negative, we have a control system. The ideal control system of that form,  
>in which this loop gain is negative infinity, predicts that  
>  
> $W = W^*$  (wetness equals reference level of wetness)  
> $A = -R$  (effect of action on wetness equal and opposite to effect of  
>rain on wetness)  
>  
>Notice: no input function, comparator, or output function. All variables  
>observable.  
>The only postulate is that the organism is repending to  $W$  by producing  $A$ ,  
>so that  $A$  is some organism-function of  $W$ . If this is false, the best form  
>of  $g$  found from the data will be  
>  
> $A = 0 * (W - W^*)$   
>  
> -- in other words, wetness depends on the organism's action and the rain-  
>rate, but the action does not depend on wetness. The appearance of stimulus  
>and response is due to statistical happenstance.  
>  
>If a systematic form of  $g$  is found, then we may or may not have a control  
>system. Only if the loop gain (product of the partials) is negative and  
>large is control of  $W$  relative to  $W^*$  verified.  
>  
>The above just shows the logical form of the procedure. The actual  
>measurements and mathematical relationships you use would be selected as  
>appropriate. For instanced, wetness might actually depend on rain rate as a  
>time integral minus an exponential decay (evaporation), and so on.  
-----

Back to your posting:

> Past studies have shown that  
>the length-frequency distributions (curves showing the number of fish of  
>each size) for all the fish in an area show much more regularity than

>the distributions for individual species.

>....

>

>2. How would one be able to show that this regularity is the result of a  
> control mechanism? Some of my colleagues feel that the existence of  
> a constancy in behaviour (often referred to as a "conservative  
> property", but that is confusing to people who studied physics) must  
> imply control, but I disagree.

Apply The Test. If you disturb the length-frequency distribution (e.g. by changing the distribution for one species), does it get restored? According to your specification, the answer is Yes. Try removing all the fish of one length range (a thought-experiment, I hope). Does the distribution return to its original state? I assume it would. According to The Test, there is control. The fact that no "organism" has been identified that could exert this control, and we cannot see the "actions" by which it is exerted seem to me to be irrelevant.

But this raises a key point about The Test. Interesting control, as opposed to control, occurs when the stable value of a percept can be altered by a change in the reference signal for that percept. Is there any way of doing that with the fish? Perhaps there is. What will happen to the length-frequency

distribution during the two-year moratorium on cod fishing? I understand that for cod, at least, the average length has reduced drastically over the last couple of decades. In two years, the living cod will have grown larger, and so, presumably will the other fish that might have been caught along with the cod. The overall distribution may change. Should this be called a disturbance that is not resisted, or a change in a reference, or a return to a pre-existing reference in a control system that had been in conflict with another control system (implemented by the fishermen) that had been pulling the distribution toward shorter average length?

Control in an ecological system is a very subtle issue, particularly if we couple it with Bill's (Powers) ideas on the role of intrinsic variables internal to an organism.

Martin

Date: Tue Jul 07, 1992 9:16 am PST

Subject: lists of interest

Note, for example, the list concerning autism.

Bruce  
bn@bbn.com

-----

\*\*\*\*\*

4266 0

Received: from BBN.COM by CCB.BBN.COM ; 7 Jul 92 12:33:12 EDT

Received: from vax3.sara.nl by BBN.COM id aa19385; 7 Jul 92 12:33 EDT

Received: from VAX1.SARA.NL by SARA.NL for bnevin@ccb.bbn.com;

7 Jul 92 18:37 MET

Received: from ALF.LET.UVA.NL by VAX1.SARA.NL with PMDF#10201; Tue, 7 Jul 1992 18:37 MET

Date: Tue, 7 Jul 92 18:30 MET

From: "Reply from Linguists name server." <LING\_REPLY@ALF.LET.UVA.NL>

Subject: Output from your request to Linguists@alf.let.uva.nl

To: bnevin@ccb.bbn.com

Message-id: <ACB49F7EF84000C5@VAX1.SARA.NL>

X-Envelope-to: bnevin@ccb.bbn.com

X-VMS-To: IN%"bnevin@ccb.bbn.com"

Comments: Sent using PMDF-822 V3.0, routing is done by SARA5

Welcome to Linguists@alf.let.uva.nl.

If you want information, use the HELP command (no arguments).

If you have other questions, contact nsmith@alf.let.uva.nl

This facility was developed by CCL, the Computer Department of the Faculty of Arts of the University of Amsterdam. It is managed by CCL and the Department of General Linguistics of the University of Amsterdam.

Note:

For an explanation of the various symbols that may appear in responses to a LIST request, consult our help-file, obtainable in response to a message consisting of HELP sent to LINGUISTS (N.B. do not reply to LING\_REPLY).

-----  
MET

list list\*

autism list: autism@sjuvm.bitnet

childes users' group - list: info-childes+@andrew.cmu.edu

colibri computerlinguistiek brief - list: colibri@kub.nl

electronic communal temporal lobe - list: dsleip@brocku.ca

fonetiks - list: llsfonet@cms.am.rdg.ac.uk

greek (new testament studies) - list: nt-greek@virginia.edu

kindertaal onderzoek nieuwsbrief kon - list: frank.wijnen@let.ruu.nl

natural language and knowledge representation - list: nl-kr@cs.rpi.edu

tolkien's languages discussion - list: tolklang@lfcs.ed.ac.uk

voynich ms decipherment list: voynich-request@rand.org

linguist list\_editors: linguist-editors@uniwa.uwa.oz.au

aztec studies discussion list - listserv: nahuat-l@fauvax.bitnet

aztec studies discussion list - listserv: nahuat-l@acc.fau.edu

communication and gender - listserv: gender@rpiecs.bitnet

computational linguistics list - listserv: ln@frmop11.bitnet

cross-cultural communication - listserv: intercul@rpiecs.bitnet

cross-cultural discussion group - listserv: xcul@albnyvm1.bitnet

deaf list - listserv: deaf-l@siucvmb.bitnet

english language discussion list - listserv: words-l@uga.bitnet

esperanto list - listserv: esper-l@trearn.bitnet

ethnomethodology/conversation analysis - listserv: ethno@rpiecs.bitnet

french language list - listserv: langues@uquebec.bitnet

gaelic language bulletin board - listserv: gaelic-l@irlearn.bitnet

greek tex discussion group - listserv: ellhnika@dhdurz1.bitnet

humanist - listserv: humanist@brownvm.bitnet

intercultural newsletter - listserv: xcult-1@psuvm.bitnet  
 interpreting and translation list - listserv: lantra-1@finhutc.bitnet  
 japanese language list - listserv: nihongo@finhutc.bitnet  
 language learning & technology international - listserv: llti@dartcms1.bitnet  
 language testing research & practice - listserv: ltest-1@uclacn1.bitnet  
 linguist - linguistic discussion group - listserv: listserv@tamvm1.tamu.edu  
 linguist - linguistic discussion group - listserv: listserv@tamvm1.bitnet  
 multilingualism list - listserv: multi-1@barilvm.bitnet  
 psycoloquy interdisciplinary refereed forum - listserv: psyc@pucc.bitnet  
 russian language list - listserv: russian@asuacad.bitnet  
 russian tex and cyrillic text processing - listserv: rustex-1@ubvm.bitnet  
 scholar - natural language processing on-line news - listserv:  
 scholar@cunyvmm.

cuny.edu (submissions to jqrqc@cunyvmm.cuny.edu)  
 second lang. acqu. research & teaching - listserv: slart-1@psuvm.bitnet  
 sign linguistics discussion network - listserv: asling-1@yalevm.bitnet  
 slavic and east european languages list - listserv: seelangs@cunyvmm.bitnet  
 speech disorders - listserv: commdis@rpiecs.bitnet  
 tact interactive text analysis software discussion list - listserv:  
 tact-1@vm.  
 utcs.utoronto.ca

Message ends here.

Date: Tue Jul 07, 1992 9:29 am PST  
 Subject: A Granddaughter

TO ALL (Y'ALL):

Sondra Walker Boyle, GrandDaughter of Sondra Walker and Charles Wright  
 Tucker was born on 3 July, at 2:24 PM and with a weight of 7 lbs. and 10 ozs.  
 a height of 19 3/4 inches (very poorly measured) in Baltimore, MD. If you  
 are interested in listening to people anthropomorphize about an organism you  
 should hear the video made 1 1/2 hour after the birth.

GrandDad is tired but doing fine.

Regards, Chuck

Date: Tue Jul 07, 1992 9:31 am PST  
 Subject: Ecological emergence and control

Jeff Borchers <borcherj@CCMAIL.ORST.EDU> writes:

> A couple of random thoughts on fishes and houses.  
 >  
 > 2. RE: The house in the arctic and the one in the tropics.  
 > This analogy must be getting at internal vs. external  
 > control systems. At the equator the planet provides  
 > external control, ja?  
 >  
 > Also: Is "insulation" just a passive control mechanism as

> opposed to an active thermoregulator that requires energy  
> input?

Does control involve feedback? I should think that systems that maintain constancy through the absence of perturbation can not be viewed constructively as control systems. Plus this use of the phrase "passive control" is very different from that used in control theory.

I consider the temperature constancy of a house in the tropics or the coffee in a thermos to have nothing to do with control. Similarly, I feel that Newton's First Law (that an object not subject to forces moves with constant velocity) is not a manifestation of control. Is this a common view, or am i in left field, and everything is control?

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Tue Jul 07, 1992 9:37 am PST  
Subject: Re: Ecological emergence and control

mmt@BEN.DCIEM.DND.CA [Martin Taylor 920707 11:45] writes:

>> Is there a general method  
>>for identifying the existence of feedback and control if the precise  
>>mechanism cannot be identified?  
>  
>It has been called The Test. If you identify a variable that you suspect may  
>be under control, you disturb it and see if it is returned to its former  
>state. I'm not sure this is adequate, since if you disturb a marble sitting  
>at the bottom of a bowl, it will return to the bottom. But that is the  
>method that has been recommended by Bill Powers and Rick:

This is clearly an example of control. Mathematically it is equivalent to virtually all negative linear feedback models. But...

>Back to your posting:

>

>Apply The Test. If you disturb the length-frequency distribution (e.g. by  
>changing the distribution for one species), does it get restored? According  
>to your specification, the answer is Yes. Try removing all the fish of one  
>length range (a thought-experiment, I hope). Does the distribution return  
>to its original state? I assume it would. According to The Test, there is  
>control. The fact that no "organism" has been identified that could exert  
>this control, and we cannot see the "actions" by which it is exerted seem  
>to me to be irrelevant.

I don't see how one can answer this by thought experiments, which imply models in which the question is already resolved. And your later suggestion, that we wait until the end of the 2-year cod moratorium to get one more data point, is probably impractical from the viewpoint of a student trying to finish a thesis (and from many other points as well).



In cases where one can apply "The Test" and carry out lots of experiments in a short period of time, these questions are easy to resolve. But this is a real world problem in identifying control systems in noisy data over which we have little experimental control.

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Tue Jul 07, 1992 10:43 am PST  
Subject: Re: Ecological emergence and control

mmt@ben.dciem.dnd.ca writes:

>I thought that The Test had already been applied, with a positive result:  
>If you disturb the length distribution by changing it in one species,  
>the overall distribution is returned to its original state. Wasn't that  
>the specification?

It isn't this clean and easy to interpret. Consider a simpler case. Several years ago Bill Sutcliffe and others looked at the total biomass of fish in the Gulf of St. Lawrence. If you assume that the biomass for each species is an independent normal variate (not a good assumption, but let's keep this simple), the total biomass should also be a normal variate whose variance can be calculated. If the observed variance is less than that predicted by the independent variate model, there is evidence for emergence. But how much less does it have to be? How do you take into account special events like El Ninos and changes in fishery management policy?

When the signal to noise ratio is very high, all of these questions are easy to answer. When the S/N ratio is low, as is generally the case in ecology, the problem of setting forth criteria for identifying control, feedback and emergence, become much more difficult.

Bill Silvert

Date: Tue Jul 07, 1992 10:56 am PST  
Subject: Canada/Collusion

[From Hank Folson (920707)]

Dag Forssell (920706-2) says:

>But you mix speculations about the controlled perceptions in individuals with rather large "groups"

I broke the population into politicians businessmen and the general public in awareness that the population is not homogeneous and different groups may have different goals. Some (sub)groups probably are controlling more

aggressively than others, and so may have more effect on the outcome. I was asking for an opinion, not a study. Martin's information, however statistically flawed because of his personal choices in friends and media, is still more and better information than is available in Los Angeles!

>I have the notion that each individual reporter writes about something  
>S/he thinks the readers will read. I doubt that the U.S. media is  
>controlling for anything with respect to Canada. If "It" did, why would it  
>be in collusion with "government?"

Living control systems not only try to reduce error signals, they control to avoid them in the first place, too. For this reason a liberal writer will probably choose not to join a conservative magazine. A liberal paper probably will not hire conservative writers. When I say: Wall Street Journal, Utne Reader, The Nation, New Republic, you have no trouble labeling them as liberal or conservative, because the system concepts of the owners and writers are consistent.

Self-selection also occurs when politicians solicit contributions from business. Oil companies do not contribute to politicians that campaign for a solar energy policy. I think, at the risk of over-generalizing, that the individual politicians, businessmen, and media owners will tend to have similar system concepts, and control for much the same things. Collusion is not needed because there is a natural self-selection. Isn't this a natural result of the nature our society and our nature as control systems?

>Social control systems do not exist, except as a  
>(pre-PCT) construct in your mind. There is only the soup of individuals!

True, but societies often have direction because there is some mechanism that averages out all the goals and controlling efforts of the individuals. If the majority are controlling for something, and there is no minority controlling more aggressively, the odds are the will of the majority will prevail, but not as effectively as a control system would prevail.

Hank Folson

Date: Tue Jul 07, 1992 11:29 am PST  
Subject: Re: Ecological emergence and control

[Martin Taylor 920707 14:45]  
(Bill Silvert Undated 970707)

>mmt@ben.dciem.dnd.ca writes:

>>I thought that The Test had already been applied, with a positive result:  
>>If you disturb the length distribution by changing it in one species,  
>>the overall distribution is returned to its original state. Wasn't that  
>>the specification?

>It isn't this clean and easy to interpret.

I was basing my answer earlier today on your statement that the length distribution had been found to be constant over all fish even though the

distribution within a species was not. Do I understand you now retract that statement. All I meant to imply was that if the claim were true, then The Test had been already applied, with a positive result. There is control. (Now you say that the data are not that clean, and you can't tell whether the Test has been applied. Right? Then the questions in your original posting are based on a hypothetical situation, and my answer applies to that hypothetical situation, not the real one as determined by inadequate measurements.)

Then I went on to suggest that maybe The Test in itself is inadequate to distinguish a clearly uncontrolled basin of attraction from what we would call control, the bringing of a percept to equality with a reference level.

Which could open up a whole new (and rather pointless, probably) thread of discussion.

Martin

Date: Tue Jul 07, 1992 11:33 am PST  
Subject: Re: Ecological emergence and control and Dirty Data

mmt@BEN.DCIEM.DND.CA [Martin Taylor 920707 14:45] calls me to task:

>I was basing my answer earlier today on your statement that the length  
>distribution had been found to be constant over all fish even though the  
>distribution within a species was not. Do I understand you now retract  
>that statement.

>(Now you say that the data are not that clean, and you can't tell  
>whether the Test has been applied. Right? Then the questions in your  
>original posting are based on a hypothetical situation, and my answer  
>applies to that hypothetical situation, not the real one as determined  
>by inadequate measurements.)

Did anyone else understand that I meant to imply that data on the total biomass of fish are exact numbers? To all who did so, my apologies.

Martin and perhaps others will perhaps be surprised to discover that we have yet to come up with exact methods for sampling fish in the ocean. I clearly erred in assuming that the readers of this group knew that.

Just for future reference, real data are rarely exact. Saying that this is due to "inadequate measurements" seems a bit unfair (anyone here ever hear of the uncertainty principle, which makes some physical measurements not exact?). But honestly, I thought that control theory had at least some relevance to the real world!

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Wed Jul 08, 1992 11:12 am PST

Subject: Re: Ecological emergence and control and Dirty Data

[Martin Taylor 920708 14:15]

(Bill Silvert, still not dating his postings)

me>>I was basing my answer earlier today on your statement that the length  
>>distribution had been found to be constant ove all fish even though the  
>>distribution within a species was not. Do I understand you now retract  
>>that statement.

>Did anyone else understand that I meant to imply that data on the total  
>biomass of fish are exact numbers? To all who did so, my apologies.

>Martin and perhaps others will perhaps be surprised to discover that we  
>have yet to come up with exact methods for sampling fish in the ocean.  
>I clearly erred in assuming that the readers of this group knew that.

Sorry, but I am confused. What does length distribution have to do with  
total biomass? You said, around 920707 (why don't you date your postings?):

> Past studies have shown that  
>the length-frequency distributions (curves showing the number of fish of  
>each size) for all the fish in an area show much more regularity than  
>the distributions for individual species.

That's all I have ever been responding to. Why do you keep bringing in  
other factors. If your statement is true, then my points about The Test  
hold. The Test has been performed, and if interpreted in the usual way,  
it shows control. If yours was a hypothetical statement, my points hold  
in the hypothetical case.

And in respect of exact numbers, may I ask what percept in the whole world  
is immune to sampling error? Only the magnitude of the error variance in  
relation to the perceptual magnitude is different from one percept to another.  
Even if we were talking about the perception of total biomass, which we never  
have been, I would have thought readers of this group knew that. I clearly  
erred.

It is nice, sometimes, to respond to arguments or points that are made, rather  
than to bring new ones out of thin air and use them to castigate.

Do you want to return to the thread about the informational limits on control?  
We let it drop last year, partly because it got played out without resolution.  
Ask Gary for the archives on it if you missed it the last time through.

I must have missed some postings on this issue, because I have received none  
(and certainly sent none) to which your last paragraph can refer:

>Just for future reference, real data are rarely exact. Saying that this  
>is due to "inadequate measurements" seems a bit unfair (anyone here ever  
>hear of the uncertainty principle, which makes some physical  
>measurements not exact?). But honestly, I thought that control theory  
had at least some relevance to the real world!

Could you copy back to us something that this connects with?

Martin

Date: Wed Jul 08, 1992 11:29 am PST  
Subject: Re: Ecological emergence and control and Dirty Data

mmt@BEN.DCIEM.DND.CA [Martin Taylor 920708 14:15] responds:  
>(Bill Silvert, still not dating his postings)

Actually my mailer puts the date on, but the CSG server removes it.  
This is being written on 92/7/8 at 15:40 ADT.

>Sorry, but I am confused. What does length distribution have to do with  
>total biomass? You said, around 920707 (why don't you date your postings?):  
>  
>> Past studies have shown that  
>>the length-frequency distributions (curves showing the number of fish of  
>>each size) for all the fish in an area show much more regularity than  
>>the distributions for individual species.

The mass of a fish is roughly proportional to the cube of its length  
(actually more like the 3.2 power, since fish change shape as they grow).  
When sampling on board ship it is faster to measure the length than to  
weigh the fish, especially in heavy seas.

>That's all I have ever been responding to. Why do you keep bringing in  
>other factors. If your statement is true, then my points about The Test  
>hold. The Test has been performed, and if interpreted in the usual way,  
>it shows control. If yours was a hypothetical statement, my points hold  
>in the hypothetical case.

Here is a counter-example. Suppose that there is only one kind of fish  
and its number is constant, simply because nothing ever happens to  
perturb it. I would not consider this kind of constancy a manifestation  
of control. Now suppose that teams of English, French and Dutch  
scientists go out to sample and report numbers of cod, morue, and  
kabeljau. They sample in different areas, and curiously enough the  
total of these three "species" is constant from year to year, even though  
the number of each "species" varies from year to year. Is this a  
manifestation of control? Or does it simply show that the samplers  
don't realize that cod, morue, and kabeljau are really the same fish?

This may sound absurd, but I spotted exactly this problem in the  
analysis of yellowtail flounder data in the 1970's (cf. Silvert,  
William, and Lloyd M. Dickie. 1982. Multispecies interactions between  
fish and fishermen. In "Multispecies Approaches to Fisheries  
Management Advice," M. C. Mercer, ed. Can. Spec. Publ. Fish. Aquat.  
Sci. 59:163-169). The actual abundance data for individual species were  
artifacts of the way the data were interpreted.

More generally, if the whole is not the sum of its parts, but rather the  
parts are simply components of the whole (i.e., the whole gets divided  
up in some randomly variable way), then I don't see evidence for control.

>And in respect of exact numbers, may I ask what percept in the whole world  
>is immune to sampling error? Only the magnitude of the error variance in  
>relation to the perceptual magnitude is different from one percept to another.  
>Even if we were talking about the perception of total biomass, which we never  
>have been, I would have thought readers of this group knew that. I clearly  
>erred.

I thought that's what I said.

>It is nice, sometimes, to respond to arguments or points that are made, rather  
>than to bring new ones out of thin air and use them to castigate.

OK, I'll agree with that. Do you?

>Do you want to return to the thread about the informational limits on control?  
>We let it drop last year, partly because it got played out without resolution.  
>Ask Gary for the archives on it if you missed it the last time through.

If everyone feels that these questions have been fully dealt with, I  
will be happy to drop it. I'm only responding to Martin's flames.

>I must have missed some postings on this issue, because I have received none  
>(and certainly sent none) to which your last paragraph can refer:

>  
>>Just for future reference, real data are rarely exact. Saying that this  
>>is due to "inadequate measurements" seems a bit unfair (anyone here ever  
>>hear of the uncertainty principle, which makes some physical  
>>measurements not exact?). But honestly, I thought that control theory  
>>had at least some relevance to the real world!

>  
>Could you copy back to us something that this connects with?

Sure. I sent out a posting with the date Tue, 7 Jul 92 16:05:24 ADT  
containing the following quote, although I didn't preserve your posting:

mmt@BEN.DCIEM.DND.CA [Martin Taylor 920707 14:45] calls me to task:

>I was basing my answer earlier today on your statement that the length  
>distribution had been found to be constant ove all fish even though the  
>distribution within a species was not. Do I understand you now retract  
>that statement.

>(Now you say that the data are not that clean, and you can't tell  
>whether the Test has been applied. Right? Then the questions in your  
>original posting are based on a hypothetical situation, and my answer  
>applies to that hypothetical situation, not the real one as determined  
>by inadequate measurements.)  
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^

Bill

Date: Wed Jul 08, 1992 12:01 pm PST  
Subject: Dating postings

[Martin Taylor 920708 15:15]  
(Bill Silvert 920708 15:40)

>mmt@BEN.DCIEM.DND.CA [Martin Taylor 920708 14:15] responds:  
>>(Bill Silvert, still not dating his postings)  
>  
>Actually my mailer puts the date on, but the CSG server removes it.  
>This is being written on 92/7/8 at 15:40 ADT.

I could have replied privately to Bill on this, but there are others in this group who make the same mistake, so it is as well to send it publicly. It's probably in Gary's new groupie's posting, too.

The convention in CSG-L is to start the text with your own name followed by the date in yymmdd format, followed by the time, so that whoever follows up the mailing can, in the second line of their text, repeat that group. The two postings can then be linked if back references are attempted automatically, or can be sought out by an editor if a reader is keeping an archive of postings on an issue. I happen to be keeping HyperCard stacks of all mailings, so I am particularly sensitive to the issue.

Martin

Date: Wed Jul 08, 1992 12:16 pm PST  
Subject: First attempt at PCT

[From Francisco Arocha; 920708;15:22]

Since I've been reading PCT literature for several months now I thought that probably the best way to learn something first hand is to do a PCT study (or at least do design one). The following is my first attempt at it, so I would appreciate if some more experienced PCTers help me "clear" my ideas. I acknowledge that these (my ideas) are "nebulous" now, so please, be patient.

The following then describes a very preliminary idea for a study of patient-management that I'm planning to conduct, both because I'm interested in the process and because I'd like to use this experiment to have hands-on experience in PCT. Before I start, I'd like first to explain what I'm doing and my research background. I've been studying medical expertise for the last 6 years from a cognitive science perspective. I have looked at the process of clinical diagnosis in terms of the generation of hypotheses and the use of knowledge in this process. The research I have conducted is mainly descriptive and is done with a few subjects whose verbal protocols are analysed in detail. Few, if any statistics given the small sample size.

Now, the patient/management process is more dynamic than diagnostics. And more suitable for applying a PCT approach, I think. Since my interest, and the main focus of the centre where I work, is the study of expertise, the study will be focused on a comparison between physicians at different levels of expertise, say entering residents, advanced residents, and expert physicians with many years of practice.

The physician treats and follows up a patient. The treatment, say a drug, has an effect on the patient; and this is supposed to "cure" him or alleviate his state. The patient state of health, however, is affected by numerous other factors. So the state of the patient is a result of the action of both the doctor's treatment/management and other, often unknown, factors. The problem I face is that the reference value may be different for different people at the same level of expertise, and may be difficult to determine what the different people are controlling for. Obviously, the physicians want to see the patient recover, so I'm assuming the the reference value is to keep the patient on a state of good health, so they must use whatever they have at their reach to keep the patient at that level.

One hypothesis may be that different physicians with different levels of experience control for different things, so the study will be set up to specify what are those things that differently experienced physicians control for. Methodologically, one can make things easier by presenting a description of the patient on a computer and using something like Hypercard, modify the patient's response to the physician's actions. An nice thing would be to have a quantitative display that can show the current patient condition, for instance some value of vital signs (blood pressure or some other).

So what I have thought so far goes like this: Present patient (may be give the diagnosis). The patient is in a state that needs continuous treatment and monitoring. The treatment is supposed to act on the patient in the expected way, but there are RrandomS disturbances to the patientUs condition. These disturbances act on the patient so that the patient's condition worsens.

r = reference signal determined by the preferred state of the patient in the subjectsU mind (if the study is done so as to keep a value, set this value to, say 1). This, hopefully, remains constant. This may coincide with the goal (at least in this study).

p = perceptual signal determined by the current state of the patient. This signal's value varies depending on the subject's action and a random disturbance.

d = random disturbance. This should randomly fluctuate between two values: let's say between .05 (close to death) and .75 (satisfactory condition).

To do this experiment I need a patient that is in chronic state and that needs constant attention. I'm not sure what to expect in terms of the controlled variables and the expertise of the subjects. Experts, nonetheless, should be able to keep the perceptual signal's value closer to the reference value than novices. This is because of the assumption (or the hypothesis) that expertise involves the learning of controlling perceptual inputs. Another possibility is that there may be a difference in terms of the quality of the reference signals maintained by experts and novices.

I would appreciate to have some feedback on this. All sorts of criticisms are welcome. Thank you.



Francisco

Date: Thu Jul 09, 1992 6:59 am PST  
Subject: Fishy Test; Medicine & PCT; 23-level diagram

[From Bill Powers (920709.0700)]

Martin Taylor and Bill Silvert (920708) --

The spillover from this little squabble seems to indicate the idea of looking for control processes in large natural phenomena, using The Test.

I remind one and all that there is more to the Test than looking for low or high correlations of various sorts. The idea is to track down an actual control system. Once correlations or other measures have suggested that a control system might be responsible for an unexpectedly low or unusually high correlation or for unnatural stability of some observable variable, it is incumbent on the investigator to go on to identify the system in question. An active system must be found producing outputs that specifically oppose disturbances, and it must be shown that the controlled variable is in fact sensed by the system, loss of the sensory link destroying control.

If you find that the parrot clings to the perch only because its feet are nailed to the wooden bar, I would advise not buying it as a live working control system. To call a system a control system should be, I think, a last resort after all other possibilities have been ruled out as being highly improbable, technically impracticable, or factually false. And, of course, after all the necessary aspects of control have been demonstrated to exist. There will always be cases in which the only defensible conclusion is "I don't know."

-----  
Francisco Arocha (920708;15:22) --

I think you have a nice idea there: what are doctors controlling for? This could turn out to be a rather larger project than you may now envision. It could also turn out to be a study of extraordinary importance.

The approach as you outlined it makes the assumption that doctors are controlling primarily for patients in a good state of health, and that they will pick treatments that do in fact help. If this is true, you should find that disturbances of the state of health that make it worse result in increased treatment activity or more drastic treatments, and disturbances that improve it result in decreased amounts of treatment or milder treatments. A simple study of this kind would make a nice demonstration package.

You should also find that the doctor's beliefs about the efficacy of treatments corresponds to the treatment used, as appropriate to the hypothetical error signal. A treatment that the doctor does not perceive as effective would tend not to be used, regardless of the statistical data concerning that treatment.

The perceptual side of the control process can be investigated in terms of diagnostic procedures used by the doctor and tests that the doctor orders. "How do you know that this patient is ill, the kind of illness, and the severity of the illness?"

Other hypotheses about other possible controlled variables can also be investigated. For example, many doctors are suspected of preferentially employing treatments or tests that they have special expertise in, or that use medical facilities in which the doctor has a financial investment, quite independently of what is wrong with the patient. The costs of various treatments (in terms of how much money the doctor can make by using them) should also come under scrutiny. Doctors may be controlling for perceptions of specific causes of illness, or for income, or for a chance to exercise special skills -- in addition, one presumes, to controlling for the restoration of health.

There is also the question of the efficacy of treatment. Large parts of modern medical practice consist of reading brochures put out by drug companies and trying out the latest magic bullet. There are many drugs whose positive effects are found in only a small minority of patients, yet because of statistical analyses they are considered "effective." Therefore when many drugs are used, the result actually to be expected is negative, most such drugs having adverse side-effects and only some of them benefiting a clear majority of patients.

An important question, therefore, is how doctors explain failure of a drug (or other) treatment to have the expected effect. If the doctor perceives the treatment as effective, and it fails to improve the patient's health, does the doctor try to correct the error by increasing the amount of the same treatment, or does he/she change the reference level for using the treatment? Is the failure blamed on invisible disturbances that were larger than anticipated, or on a treatment that is less effective than supposed? In short, does the doctor control for a continued perception of medical practice as being effective, or for improving the condition of the patients, one by one, or for not doing harm (a principle that at least some doctors profess to hold)?

Finally, the way in which data about many participants in this experiment are combined is of utmost importance. A complete control analysis for each participant has to be done BEFORE the data are combined. If you subtract the average perception of drug efficacy from the average reference signal for efficacy, you will get a average error signal that doesn't represent the error experienced by ANY individual. What must be done first is to look at the relationships among reference signal, perception, disturbance, and action for each individual; what should be found is that for all individuals, action opposes disturbance and perception is brought near to the reference signal. The possibility of bimodal or multimodal measures should be kept strongly in mind, because there will be large individual differences in reference settings and perceptions. The relationship between disturbance and action should be the same for all those in whom you have identified a true controlled variable. But the reference settings can be different in every individual.

There is a lot about medical practice that you can discover through PCT analysis. There is a lot that NEEDS to be discovered.

-----  
Dag Forssell --

I agree that you, I, and Bruce Nevin see the "23-level" diagram in the same way. But I also think that Martin Taylor does, too. Martin and I often go around and around in apparent disagreement, only to find that we are really talking about the same thing but in different words. As you say, establishing agreement is actually harder than establishing disagreement.

I think that often the parties to a disagreement, in explaining their views and arguing against others, experience a gradual shift in their perceptions and end up agreeing to a consensus position that does not exactly match what any of them began with. There may still be residual disagreement between the internally-held positions, but it no longer shows up in words. All we can do is keep iterating this process, in the hopes that all actual positions are moving in some meaningful direction.

Best, Bill (back on the air) P.

Date: Thu Jul 09, 1992 7:55 am PST  
Subject: Re: Fishy Test; Medicine & PCT; HyperCard

[Martin Taylor 920709 11:00]  
(Bill Powers 920709 07:00)

>I remind one and all that there is more to the Test than looking for low or  
>high correlations of various sorts. The idea is to track down an actual  
>control system. Once correlations or other measures have suggested that a  
>control system might be responsible for an unexpectedly low or unusually  
>high correlation or for unnatural stability of some observable variable, it  
>is incumbent on the investigator to go on to identify the system in  
>question.

You cut right to the place I was trying to get to in the discussion with Bill Silvert, except that I had thought that The Test consisted only of the initial disturbance-correction check, with the search for the actual control system to follow if the Test is successful. In a whole mess of private correspondence yesterday Bill and I got more and more tangled in what we had said and meant, and that blocked further progress. Thanks for pulling us out of the quagmire.

Which leads to:

>I think that often the parties to a disagreement, in explaining their views  
>and arguing against others, experience a gradual shift in their perceptions  
>and end up agreeing to a consensus position that does not exactly match  
>what any of them began with. There may still be residual disagreement  
>between the internally-held positions, but it no longer shows up in words.  
>All we can do is keep iterating this process, in the hopes that all actual  
>positions are moving in some meaningful direction.

Amen.

If the residual discrepancies affect some aspect of a problem still to be discussed, they can result in surprises when that problem comes up. We think

we have reached an agreement on something, but haven't at a deeper level. At present, I know Bill and I disagree on several issues, so that's good. We can continue to, as Bill puts it,

>go

>around and around in apparent disagreement, only to find that we are really >talking about the same thing but in different words.

With few words, false agreement is easy to reach. With many, confusion is easy to reach. There must be a happy medium. I wish there were a telepathic way to do this, but there isn't. And I think the issue of how PCT works is very important, so words and net bandwidth must be used.

-----

A sidenote, to avoid a separate posting. Last night I uploaded to Bill Silvert's site (biome.bio.dfo.ca) two HyperCard stacks that anyone with ftp and a Mac can download. Both were built for my private use, and I'd be interested in feedback as to whether anyone else is able to use them, and what improvements might make them more useful. One stack, called Mailsplitter, takes a text file on the Macintosh that is made from the catenation of a set of mail messages that have an Internet standard header (CSG-L postings do, at least when they arrive here); it separates the messages, one per card, and provides some simple ways of searching by author, by included text, by subject line, and so forth. The second stack is one made using Mailsplitter, for CSG-L messages in June 1992. Mailsplitter is about 104K, the June 92 stack (CSG.9206) is about 512K compressed self-extracting archive. Use BINARY mode when downloading.

Martin

Date: Thu Jul 09, 1992 8:51 am PST  
Subject: reorganization

[From: Bruce Nevin (Thu 92079 09:13:52)]

(Martin Taylor 920701 22:40) --  
(Bill Powers 920630.2000) --

Bill suggested that one way to get reorganization to target the locus of error

> is to define the  
> reorganizing system as a distributed system, a mode of operation of every  
> ECS, but one that is NOT concerned with the normal business of control.  
> This would solve the problem of specificity of the locus of reorganization,  
> in that this distributed system would sense error and act to correct it at  
> the place where it occurred. I have long held this concept in reserve -- I  
> think I even mentioned it in BCP -- but as I don't have any idea what this  
> special mode of operation might be (the Hebbian solution is not yet, to my  
> mind, worked out well enough to model) I have elected to go with a lumped  
> model that will work in essentially the same way. There are other possible

I am imagining such a distributed reorganization system. It seems to parallel the perceptual control hierarchy by virtue of pervading it. What I am imagining depends upon the recognition that an ECS is a

group of cells cooperating with each other by means of cell-level and molecule-level control mechanisms.

Suggestions of how this might work are in the discussion of the origins of life. Feedback loops between molecules and between cells does not go away with the advent of nervous systems implementing control systems as we usually discuss them. It seems extremely likely that they are ongoing in parallel with such higher "orders" (as distinct from "levels") of control. In particular, the kinds of feedback relations whereby cells take their "voluntary" places in colonies, or within fungi, or within various orders of plants, or within animals of increasing complexity, are probably the same kinds of feedback relations whereby they do so in the embryology and development of multicellular organisms, and persist in "vegetative" functions that preserve systemic integrity--what we call intrinsic error.

More specifically, it seems likely to me that the cooperating cells constituting each ECS can reorganize themselves so as to reduce or increase something in their environment, such as neuropeptides. That chronic error in an ECS might result in increased production and/or release of neuropeptides, with an influence on the cells of neighboring ECSs; might result in chemical or even neural signals that influence the number, location, or sensitivity of receptor sites associated with ECSs that are neurally connected but not physically close enough to be "neighboring" in the same sense. And so on.

The existence of non-neural inter-cellular communication and cooperation as a mechanism for reorganization does not preclude other mechanisms for reorganization operating in parallel. At a much higher level is reorganization within ECS pandaemonium, as I suggested earlier. For example, the input function I of an ECS may reject one candidate signal  $i_1$  (or reduce its value) because other signals are not present to complement  $i_1$  and so the input requirement of I is not met. As  $i_1$  is rejected, a signal  $i_2$  from another source becomes the leading candidate. Ex hypothesi,  $i_2$  was present all along, but had been hitherto ignored. One can think of various reasons a signal  $i_2$  might be held in abeyance by an input function I, for example, because it was weaker than  $i_1$ , or because I had complements of  $i_1$  satisfied, some of which were imagined and overridden later by real inputs, and so on.

There is another kind of resource to support the targeting of reorganization, and that is structured intervention from the organism's environment. Martin says:

> It is not by accident that the young of all species are either  
> incompetent or already organized with an effective sensory-motor control  
> hierarchy. If they do not have the control hierarchy, then the reorganizing  
> system must develop new ECSs, with the issues as raised above. If they do  
> have an existing hierarchy, then the reorganizing system cannot be randomly  
> linked to it. So, your argument seems to lead to the situation that you don't  
> want to allow: the reorganizing system DOES know about specific aspects of the  
> control hierarchy, whether the two hierarchies are separate or no, or even  
> whether the reorganizing system is a hierarchy at all.

Most organisms have other like organisms or symbionts or both as

cooperating neighbors. Among mammals at least, adults control for cooperation and control for their young learning how to cooperate successfully. This is not control strictu sensu--many bytes have spilled across the net concerning social control. But what I understand of reorganization is that it is probably not control either, but rather influence, exerting (strong) selective "pressure" as the cells of ECSs "in distress" try different changes in various aspects of their structure and function that they can change, such as:

Gain

Weights on various signals in input and output functions

Location, number, and activity of neuropeptide receptor sites

Input function "requesting" an imagined signal to complement

existing input signals--could lead to changed reference

signals higher up if error is reduced in imagination

Neural connections with other ECSs

Coming at this from a different direction: the analogy is often made between cooperation of cells in an organism and the social cooperation of animals, including people. If this analogy were valid, it might work like this:

1. Each cell in an ECS can control itself using intracellular control mechanisms such as ion exchange across membranes.
2. A cell cannot control another cell with these mechanisms.
3. Yet cooperating cells can together constitute a control mechanism of a higher order.
4. This higher order of control is invisible to and does not itself affect the constituent cells.
5. However, chronic error in the higher-order control system is perceived by the cells in the form of environmental factors that are distressing to the cells.
6. In consequence of such distress, the cells may alter their relations with one another, while still controlling for inter-cellular cooperation. This constitutes reorganization at the higher order of control.
7. One may substitute "person" for "cell" in propositions 1-6 (naming appropriate mechanisms in proposition 1).

Two questions:

Is (2) true of cells? Or can one cell truly control another (in the same intra-cellular terms in which a cell controls itself)?

Does the specialization of cells for cooperating functions have a parallel in human differences of temperament, talent, etc., as well as in educative specialization for social function?

Bruce

bn@bbn.com

Date: Thu Jul 09, 1992 9:19 am PST  
Subject: encodingism

[From: Bruce Nevin (Thu 92079 12:21:32)]

(Joel Judd 920630) --

> Martin (920729 & 0630) and Bill (920629 & 0701),

> >Knowledge can certainly be represented, but interactions do not require the  
> kind of representation that >implies regress.

> Right; \*transmission\* requires the kind of representation that involves  
> regress. Transmitting encodings requires that both receiver and sender  
> understand what the encodings represent.

I regret that I don't have the available bandwidth just now to show why  
"encodingism" and "representationalism" don't work. As Martin said, we  
covered this ground in some depth a year or so ago.

Briefly: distinctions between elements of language must be preset as  
learned social conventions in sender and receiver. They must have had  
experiences such that perceptions of words are associated with at least  
some nonverbal perceptions on which they (the language users) can  
successfully presume agreement--though they might not think of the  
"same" associations immediately, but through the negotiation toward  
agreement that constitutes much of communication. But sentences do not  
"represent" nonverbal experiences. They incorporate or, better,  
constitute linguistic information on the basis of which people can  
represent nonverbal experiences to themselves--if memory and imagination  
are taken to be "representations." From a PCT perspective, I believe  
that too is an inappropriate use of the word "represent."

Bruce bn@bbn.com

Date: Thu Jul 09, 1992 9:35 am PST  
Subject: re.: first attempt at PCT

To: Francisco Arocha  
From: David Goldstein  
Subject: re.: first PCT experiment  
Date: 07/09/92

I enjoyed reading about your proposed experiment. As a therapist,  
I am faced with a similar kind of problem.

Here are some thoughts on your proposal:

(1) I think it would be better to study the changes which take  
place in a person as the person becomes more expert than to compare  
different people who are supposedly at different levels of

expertise.

This means that the same sort of case would have to be presented over and over. The person would show changes in how they managed a case as the expertise grew.

Potential changes could occur in perceptions, goals and actions taken by the physician.

(2) I think that videotaping interactions between your physicians and patients would have a better chance of finding out what perceptions physicians are controlling than to study the interactions between physicians and a computer. I know that you are trying to simplify the situation. However, I think you are losing too much realism.

(3) At the residential treatment center where I work, I chair a committee which reviews the use of psychotropic medicines. Two nurses, two psychiatrists and I form the committee. One very interesting part of the medicine review is the question: Are there changes in the target symptoms? This seems to be the key question for the review and maybe for your study.

I have to run now. Good luck with your study.

David Goldstein  
internet: goldstein@saturn.glassboro.edu

Date: Thu Jul 09, 1992 9:53 am PST  
From: Dag Forssell / MCI ID: 474-2580

TO: csg (Ems)  
EMS: INTERNET / MCI ID: 376-5414  
MBX: CSG-L@VMD.CSO.UIUC.EDU

Subject: List server. Historic record incomplete  
Message-Id: 53920709175335/0004742580NA2EM

[From Dag Forssell (920709-1)]

I have downloaded "Get CSG-L LOG9207A." It was long @ 316,690 bytes. I have found a way with a simple dos wordprocessor to eliminate the ^z - end of record marker - that sometimes appears after Bill's posts, and make it impossible to read and edit the whole file. There are two of those in 9207A.

I'll read the many missing files with interest. Right now I'll share an impression about the list server.

The typical transition between one message and the next in this weekly file looks like this:

Reference) \_\_\_\_\_  
Martin

=====  
Date: Sat, 4 Jul 1992 11:39:39 -0600



Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>  
 Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>  
 From: "William T. Powers" <POWERS\_W%FLC@VAXF.COLORADO.EDU>  
 Subject: Reorganization

---

My suspicion that the file server had some hiccups this past weekend are fanned by the appearance of several faulty transitions as shown here:

1) \_\_\_\_\_  
 statements lead very naturally to simple experimental designs. What we're talking about is REAL falsifiability.

-----  
 Best

Bill P.all ECSs would be in conflict about almost every percept. I assume "apple controllers" come into play only with respect to things that are sufficiently like apples to make it reasonable to try to perceive them as apples.

I'll answer the plausibility post later. I think we are even closer to agreement there, but some questions do remain.

Martin

---

and:

2) \_\_\_\_\_  
 That was NOT saganesque. Sagan said "Billions and Billions." I am a thousand times less saganesque than Sagan.

-----  
 Best,

Bill P.  
 ems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>  
 From: "William T. Powers" <POWERS\_W%FLC@VAXF.COLORADO.EDU>  
 Subject: Off time

Starting at 6:30 on the 5th and continuing until some time on the 8th,

---

and finally:

3) \_\_\_\_\_  
 What this boils down to is that I'm impatiently waiting for you modeling guys--you, Rick, Martin Taylor and all--to get around to modeling the highest levels, so we can see how they have to work in an autonomous organism.

Best-, Dick Robertson  
 chart and fold it on

>>the mirror line, then interconnect the control systems across so we

>>control all the perceptions up and down. We are back to the diagram as  
>>we know it, but with an expanded understanding of it.

---

The file server holds past weeks for a limited time (two months or so).

Greg Williams as CSG archivist keeps a complete historical record. I know he appreciates that a record is kept by others as well, the more places the better, since disasters can strike any one location.

From the above, it is clear that the record Gary will download and mail to Greg upon his return from vacation is incomplete. In 1), it is clear that most of one post by Martin is lost. In 2), there may or may not be anything lost. In 3), the beginning of my post on Taylor's diagram is lost. It is of course conceivable that whole posts are lost in between, although from Martin's listing to me and the record, that appears unlikely to me.

Therefore, for the record, Martin may want to repost to Gary and Greg (and me, since I file everything without necessarily trying to keep up with it all as it unfolds). I'll send my file by snail mail to Greg.

Dag

Date: Thu Jul 09, 1992 7:24 pm PST  
Subject: flipping eggs

Something someone might want to check out (off a mailing list) ....

The breakfast-making aficionados out there might want to look at University of Rochester Computer Science Tech Report 416, "Contextually Dependent Control Strategies for Manipulation". This is by Polly Pook, one of Ballard's students. In the report she demonstrates a qualitative method for flipping a plastic egg in a frying pan using an MIT/Utah hand on a Puma. There is no configuration space planning, no stable grasp determination, and no CAD model of the either the pan or egg. Instead there are a bunch of phases which are either fixed action patterns or compliant moves. Transitions between phases are signalled by changes in tendon tensions. Since each phase has a specific goal and operates in a situation which has been highly constrained by previous phases, very little sensory interpretation is required. Pretty neat - check it out.

-- Jon (Avery.Andrews@anu.edu.au)

Date: Thu Jul 09, 1992 9:23 pm PST  
Subject: Modeling reorganization

[From Bill Powers (920709.2000)]

Bruce Nevin (920709.0913) --

I haven't responded to one of your (somewhat rare) posts recently. Mostly just because I agree with you, so with no error signal ... but don't take that for a blanket endorsement!

>I am imagining such a distributed reorganization system. It seems to >parallel the perceptual control hierarchy by virtue of pervading it.

This is a good starting point. As we take off from it, I'd like to restate a ground-rule that I try to adhere to in talking about reorganization.

It is that reorganization can't use any facility that develops out of reorganization. By this I mean that the processes behind reorganization themselves can't make use of knowledge about the outside world at any level, from kinesthetic to cognitive. This may turn out to be too restrictive a demand, but in my mind answers to the basic questions about reorganization must apply from the beginning of an individual's life, before the hierarchy has been significantly elaborated. And implicit in this view is the idea that the reorganizing system operates in the same way throughout an individual's life.

>Suggestions of how this might work are in the discussion of the origins >of life. Feedback loops between molecules and between cells does not >go away with the advent of nervous systems implementing control systems >as we usually discuss them. It seems extremely likely that they are >ongoing in parallel with such higher "orders" (as distinct from >"levels") of control.

Agree. With respect to the genetic-level control systems, my favorite example is the repair enzymes, a product of DNA that continually restores DNA to conform with some built-in reference pattern. This is just one of many mechanisms that renders the organism relatively immune to externally-induced mutations. In my proposals about the origin of life, the same organization in a much less elaborate form existed from the beginning -- and as you say, it is still here, working away at the most basic biochemical level. It's a complex control system, perhaps even a hierarchical one -- but it's not a reorganizing system.

> In particular, the kinds of feedback relations >whereby cells take their "voluntary" places in colonies, or within >fungi, or within various orders of plants, or within animals of >increasing complexity, are probably the same kinds of feedback >relations whereby they do so in the embryology and development of >multicellular organisms, and persist in "vegetative" functions that >preserve systemic integrity--what we call intrinsic error.

Excellent thought. PCT, it seems, suggests a whole new approach to phylogeny, cladistics, or what have you. Control supplies a theme that runs through evolution and in fact shows its direction (Stephen J. Gould would

be horrified). And the principle of reorganization may offer a parallel theme: the emergence of systematic control from nonsystematic control. Perhaps repeated at many levels (orders?).

>More specifically, it seems likely to me that the cooperating cells >constitutive each ECS can reorganize themselves so as to reduce or >increase something in their environment, such as neuropeptides.

Please elaborate on neuropeptides -- is this a generic term, or a specific substance? What do they do? What makes them? Are you speaking of neurotransmitters in general? Are these proteins with specific functions?

> That chronic error in an ECS might result in increased production >and/or release of neuropeptides, with an influence on the cells of >neighboring ECSS; might result in chemical or even neural signals that >influence the number, location, or sensitivity of receptor sites >associated with ECSS that are neurally connected but not physically >close enough to be "neighboring" in the same sense. And so on.

I've proposed previously that cell differentiation and the "turning off" of genes could be a simple control process in cells sharing a common environment. Imagine a set of cells all of which contain a gene that specifies a reference level for some substance in the shared environment. Imagine that there is a spread in the natural reference signals. At first, all cells experience a deficit of that substance, and all therefore begin creating it. As more and more cells appear through continuing stages of cell division, the controlled substance (no puns please) will eventually be brought to the reference level by the action of all the tiny control systems controlling for a specific concentration of that substance. No one control system can maintain the required concentration, but many working independently and in parallel can.

Eventually the concentration will reach the level specified by cells with the lowest reference settings. As the number of cells increases, that concentration will exceed those lowest reference settings, and those cells will cease to produce a contribution to the total concentration. Equilibrium will occur when there are just enough cells with the highest reference settings to produce enough of the substance to shut down all control systems with a lower reference level, and leave a steady-state population of cells with the highest reference levels just maintaining a steady concentration of the substance in question. These are one-way control systems; errors represent only deficits. Therefore there is no conflict.

Under the usual cause-effect interpretation, the genes responsible for producing the substance are "programmed" to "turn off" in some cells. The PCT version of this explanation is that the genes remain as "active" as ever, but feedback effects from the general concentration of the substance are higher than the reference level, making the error go negative. As a negative output is not possible, these control systems are effectively turned off. There probably isn't any serious difference between this interpretation and observations -- "repressor" enzymes are known, for example, which we would interpret simply as perceptual signals with a negative feedback connection.

My point is that there can appear to be coordinated actions and even apportionment of functions among systems of the same level without, in fact, any superordinate coordinating system existing. Your comment above is on the same track as my thinking.

>The existence of non-neural inter-cellular communication and >cooperation as a mechanism for reorganization does not preclude other >mechanisms for reorganization operating in parallel.

I agree ... however, I would not call what I just described above "reorganization." The reason is that it can be accounted for entirely in terms of the normal operation of control systems of the normal type. At a given level of organization, cells containing multiple copies of the same control systems will automatically divide the labor between them: those with the highest reference levels for a given substance will end up maintaining a specific concentration of that substance for ALL the cells. If you think there are parallels between this principle and the organization of social systems, so do I.

I think we have to use the concept of reorganization sparingly; it can too easily become a catch-all for unsolved problems of every kind. I would like to see as much of the growth of the organism as possible, and as much of the behavioral hierarchy as possible, accounted for by normal interactions among normal control systems. So I'm not in favor of ...

>For example, the input function I of an ECS may reject one candidate >signal i\_1 (or reduce its value) because other signals are not present >to complement i\_1 and so the input requirement of I is not met.

This is similar to a suggestion of Martin Taylor's that I also rejected, and for the same reason. You're proposing a very complex "E"CS, and I believe we should resist complexities until observations force us into accepting them because we can see no alternative. We haven't reached that point yet; we haven't proven that the normal operation of the hierarchy, and a SIMPLE principle of reorganization, won't solve the problem. And we may not be ready for such a proof for a very long time (I have no doubt that it will be forthcoming).

A serious problem is that the neural signals would somehow have to be interpreted in terms of the meaning they will be given in the new ECS, so "complementarity" could be detected despite the fact that all neural signals are basically alike -- just magnitudes.

In order to model what you propose, we would have to show functions in the model that could detect input signals and judge their complementarity WITHOUT combining them in the normal way. Then other functions would have to be provided that convert this judgment about the potential input signals into a process you call "rejection," which itself might be difficult to embody in a model.

Even if you could draw such an elaborated diagram of a CCS (complex control system), you might have problems with stating what a model based on such a diagram would actually do, using only the rules you have put into it. You can't just point to the function you want accomplished as proof, until you can show that the presented model will actually behave in the imagined way.

We understand quite well how an ECS works, given a black box to accomplish its input function. We can simulate such systems and discover what such a system will do, even if the input function is too complex to represent analytically or in any detail. but we know little about more complex systems.

Most of the basic rules of thumb of PCT and HPCT are based on known and demonstrated properties of simple control systems. If we start elaborating on the simplest organization, we must go very slowly and take small steps, because at every step we have to re-analyze the whole control system to find out how our changes and additions have changed its basic properties. Even the most innocuous change could alter the properties we are familiar with beyond recognition. The only way to handle this is to introduce small changes and re-do the analysis and simulations each time to find out whether we have actually made a qualitative change in the system -- whether we have created something with radically different rules of behavior. This isn't the sort of thing that can be done every day or every month -- perhaps not even every year.

>... what I understand of reorganization is that it is probably not >control either, but rather influence, exerting (strong) selective >"pressure" as the cells of ECSs "in distress" try different changes in >various aspects of their structure and function that they can change >...

Reorganization IS control: it uses, however, a unique kind of primitive output function, which acts at random but at variable intervals. The result is to bring a controlled variable to a reference level, the same result we get from any control system. The list of functions possibly subject to these random effects,

Gain

Weights on various signals in input and output functions

Location, number, and activity of neuropeptide receptor sites

Input function "requesting" an imagined signal to complement existing input signals--could lead to changed reference signals higher up if error is reduced in imagination

Neural connections with other ECSs

... is a good one, subject to preceding quibbles. What we have yet to demonstrate is that random variations in such parameters can actually result in organized semi-permanent systematic control systems.

In fact, none of your concepts amounts to a model yet, but all of them are good candidates for the primordial soup of concepts from which we will eventually evolve a more competent model. These are all things we must try out in simulation.

>... can one cell truly control another (in >the same intra-cellular terms in which a cell controls itself)?

Control systems control variables, not things. Your question thus really asks, can a control system in one cell control a variable inside another cell? And I think that pretty much answers itself: not if the same variable is already under control in the other cell. It doesn't seem likely that a chemical messenger representing a variable inside one cell membrane could

flow freely out of that cell and into a different one to provide a perception of the controlled variable, nor that an output signal from one cell could travel equally freely in the other direction.

>Does the specialization of cells for cooperating functions have a >parallel in human differences of temperament, talent, etc., as well as >in educative specialization for social function?

See comment a couple of pages ago.

-----  
RE: modeling reorganization

A preliminary report, especially to Martin Taylor who has begun making some noises about actually doing some joint simulation research on this subject.

I have been playing with reorganization as a way of solving a system of linear equations:

$$y[m] = \text{SUM}(a[m,n]*x[n]) \text{ where } m = n \text{ in all cases.}$$

I actually started with the inverse problem: given a set of inputs  $x_1..x_n$ , and a set of outputs  $y_1..y_m$ , with  $m = n$ , find the matrix of coefficients  $a[m,n]$  that will satisfy all  $m$  equations. So this is like perceptual learning.

The basis for reorganization is the sum of squared errors between  $r[m]$  and  $y[m]$ , where  $r[m]$  is the desired set of values of the functions, and  $y[m]$  is the actual set of values for any given set of coefficients  $a[m,n]$ . The vector  $x[n]$  is a fixed list of  $n$  numbers.

To reproduce the E. coli method, it's necessary to define a direction in  $m$ -dimensional space, using an auxiliary matrix  $\delta[m,n]$ . The entries in the  $\delta$  matrix are changed independently and at random between positive and negative limits, to create a "tumble." The matrix is normalized so the sum of its entries, squared, is 1. This makes the entries into direction cosines in  $m$ -dimensional space. A constant times  $\delta[m,n]$  is added to the coefficient matrix  $a[m,n]$  on each iteration. "Tumbles," however, occur at intervals determined by the error signal. Between tumbles, the hyperspace point  $a[m][n]$  moves at a constant velocity (actually, this works best if this velocity depends on the magnitude of error).

Also, instead of using the error (squared) itself, it's necessary to use the time rate of change of error as the controlled variable. The interval between reorganizations of the  $\delta$  matrix is proportional to the negative time rate of change of the squared error: the more rapidly the error is decreasing, the greater is the number of iterations before the next "tumble." When the error increases, there is a "tumble" on every iteration. So the loop gain is set quite high -- maybe too high.

With 10 equations in 10 variables, the required matrix emerges after somewhere between 2500 and 10000 iterations, and perhaps 1/10 that many reorganizations. The RMS error between the target vector  $r[m]$  and the actual-value vector  $y[m]$  is then about 0.001 of the initial error; the numbers  $r[m]$  and  $y[m]$  agree to one part in 1000 of the maximum or better. With  $m$  and  $n$  equal to 50, convergence occurs, but I haven't run it to

completion -- doing so would take days. This is NOT a parallel computer. Interestingly, the rate of convergence per iteration with 50 dimensions is not dramatically slower than that with 10 dimensions, given adjustment of parameters for best performance. It's just that each iteration takes a LOT longer with 50 than with 10 equations (with 50 equations, the delta[m,n] matrix involves 2500 random changes per iteration).

I've already learned some things about this method of reorganization. The best indicator for triggering tumbles is time rate of change of the error. The variables being randomly altered must be changed not directly, but by randomly choosing the rate at which they are altered on each iteration. The delta matrix effectively creates movement in hyperspace at a constant velocity or a velocity that decreases systematically with error, with only the direction being altered at random when there is a reorganization event. I think I can see now that directly varying the output values (in the a[m,n] matrix) at random would not lead to systematic approach to a solution, nor would simply using the magnitude of the error rather than its rate of change. I don't know that for certain, but it seems likely.

I've tried using the squared error, the RMS error, and the mean error as the basic error measure, and a constant velocity or a velocity that depends on linear or squared error. Everything tried works, although convergence rate is affected. Testing is so slow that I haven't really compared the different possibilities in any useful way, nor have I found any way of optimizing things like gain and step size. I'm sure that someone with a better grasp of n-dimensional mathematics and probability than I have could derive the optimum settings without all this experimentation.

I've also done one test in which complete control systems were used for each of 10 dimensions. The result converges. But I haven't tested yet with randomly varying reference signals. Neither have I set up any intrinsic variables (other than the error signal itself) which are affected variously by the controlled variables x[n], so that reorganization is based on an indirect effect of the controlled variables. It turns out that there is an enormous number of possibilities and variants to investigate; getting to them all will take some time.

I hope to make a little more progress on this before the meeting. I have to go to Boulder and Denver next week (a talk on What is Information, a panel at the meeting of the International Society for Systems Science, into which I was sweet-talked by Peter Corning -- I'll have a copy of my remarks for distribution at the meeting and will put it on the net, too, afterward, with permission from the ISSS). So I won't have a lot of time for this until later in the summer, after the meeting.

One thing's sure: there's still a lot to learn about this process of reorganization, even with simple linear systems. And it looks just as powerful as I thought it would be.

-----  
Best to all, Bill P.

Date: Fri Jul 10, 1992 7:49 am PST  
Subject: Flippin' eggs



[From Bill Powers (920710.0900)]

Avery Andrews (920709.2112) --

>In the report she demonstrates a qualitative method for flipping a  
>plastic egg in a frying pan using an MIT/Utah hand on a Puma. There is  
>no configuration space planning, no stable grasp determination, and no  
>CAD model of the either the pan or egg. Instead there are a bunch of  
>phases which are either fixed action patterns or compliant moves.  
>Transitions between phases are signalled by changes in tendon tensions.

What does it do if there isn't actually any egg in the frying pan?

Best, Bill P.

Date: Fri Jul 10, 1992 11:18 am PST  
Subject: Re: clarification

[Martin Taylor 920710 14:15]  
(Bruce Nevin 920710 13:42;31)

I don't want to obscure Bruce's main point, but he provides a wonderful opening for me to mention a reference I was planning to introduce to the discussion: "Evidence for a computational distinction between proximal and distal neuronal inhibition," E.T.Vu and F.B.Krasne, Science, March 27 1992, 255, 1710-1712.

Bruce:

>Consider an ECS with a sensory input function I, a reference input  
>function R, a comparator C, and an output function O. In this simple  
>and perhaps merely schematic example, I has coming into it a number of  
>neural fibers bearing sensory input signals, and it has leaving it one  
>neural fiber bearing a unified sensory input signal to the comparator.  
>Ditto for R, with respect to reference signals. Conversely, O has one  
>fiber entering and a number exiting. C has two entering from I and R,  
>and one exiting to O. Each function (I, R, O, and C) comprises a number  
>of cells, and each neural fiber is at least one cell. We want to say  
>that the input signal and the reference signal together determine the  
>error signal in the comparator, C. In this really rather complex chain  
>of inter-cellular relationships, does one cell in the chain control the  
>next, with respect to the transmitted neural current?

If I read Vu and Krasne correctly, all of these functions of an ECS can be executed in a single neuron, plus the control of gain from another neuron (interestingly, in their example, this gain control comes from the motor area of the cortex). They were studying "inhibitory control of the lateral giant command neurons for crayfish tail-flip escape behaviour," but I doubt that the specifics matter. They do claim more generality. Their claim is that inhibitory connections to the outer portions of dendrites have a subtractive relation to the excitatory connections, whereas the inhibitory connections to the roots of the dendrites have a multiplicative (I guess divisive would be a better term) effect, including total inhibition. I can see this in PCT terms, that the distal connections might, depending on where the input comes from, contribute either to I or to R, and the additive

(subtractive) relationship has the function of C. The normal firing function of the neuron represents G, and the proximal inhibitory connection is a control on G, including switching it off. Multiple neurons of this type working in parallel could be seen as generating the neural current in an abstract ECS.

This view is not exactly the same as the standard model, in that there is no clear distinction between I and R inputs to the ECS. But then, functionally there is no distinction other than that the I inputs are subject to a transformation providing a single value before being linked to R. I think that this conventional distinction has no formal effect. Even if I and R inputs converge on the same dendrite, nature might well maintain their connections in the way demanded by the standard model.

Let this not stand in the way of an answer to Bruce's original question, please.

Martin

Date: Fri Jul 10, 1992 11:36 am PST  
Subject: clarification

[From: Bruce Nevin (Fri 920710 13:42:31)]

(Bill Powers (920709.2000) ) --

Neuropeptides: rightly or wrongly, I am using this term in a generic sense. I posted some stuff about work of Candace Pert and others a while back, but this is certainly not my field.

Even so, I apparently included too much detail so that you missed the main point, which is that control systems of one order (cells), controlling for coordination, cooperation, or simply for a more stable and predictable environment (which happens to include numerous of their fellows), together can constitute what we recognize as control systems of another order (ECSs within a complex living control system).

Each constituent cell of an ECS is in itself blind to the functioning of the ECS. Nor does it in any direct sense control for helping to constitute an ECS. Nor does it have any means for perceiving the ECS of which it is a constituent.

Now let me restate the question about one cell controlling another, if I can. Consider two neural fiber cells connected by a synapse. Does the state of one control the state of the other, with respect to neural impulses (modulo the value of the synapse)?

Consider an ECS with a sensory input function I, a reference input function R, a comparator C, and an output function O. In this simple and perhaps merely schematic example, I has coming into it a number of neural fibers bearing sensory input signals, and it has leaving it one neural fiber bearing a unified sensory input signal to the comparator. Ditto for R, with respect to reference signals. Conversely, O has one fiber entering and a number exiting. C has two entering from I and R, and one exiting to O. Each function (I, R, O, and C) comprises a number of cells, and each neural fiber is at least one cell. We want to say

that the input signal and the reference signal together determine the error signal in the comparator, C. In this really rather complex chain of inter-cellular relationships, does one cell in the chain control the next, with respect to the transmitted neural current?

Perhaps with this context, what I said earlier turns out to say something a little different than it seemed to?

If it does, please consider again the analogy to the human situation.

Bruce bn@bbn.com

Date: Fri Jul 10, 1992 12:15 pm PST  
Subject: Martin's postings

BILL SILVERT 920710

Has anyone been able to transfer Martin's files csg.9206.sea and mailsplitter.bin to a Mac successfully? I tried to do this today but they came out as document files. If anyone knows how to do this, please send me an informative message that I can include with them.

The files have been moved to a subdirectory called mmt.

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.dfo.ca

Date: Fri Jul 10, 1992 2:13 pm PST  
Subject: Cellular to Social Control

[From Bill Powers (920710.1330)]

Bruce Nevin (920710.1342)

>In this really rather complex chain of inter-cellular relationships,  
>does one cell in the chain control the next, with respect to the  
>transmitted neural current?

If cell A controls the neural current in cell B, then according to the definition of control, an independent perturbation of the neural current emitted by cell B should result in a change in neural current from cell A that has an effect equal and opposite to that of the perturbation. As a result, the neural current emitted by cell B should prove resistant to such perturbations because of the action of cell A.

I don't think that this is how synapses work: nothing that happens to cell B affects cell A at all (i.e., backward through the synapse). If the neural current in cell B is disturbed, cell A will simply continue to send its own signal into cell B in the same way as before. So the signal from cell A INFLUENCES the neural current emitted by cell B, or where it is the sole

influence, DETERMINES B's neural current, but does not control it.

Now, dropping back a paragraph or two:

>Each constituent cell of an ECS is in itself blind to the functioning  
>of the ECS. Nor does it in any direct sense control for helping to  
>constitute an ECS. Nor does it have any means for perceiving the ECS  
>of which it is a constituent.

Somewhere in here is an important observation -- I can't carry it all the way through, but perhaps we're thinking along similar lines. You have some provocative ideas here.

The individual (neural) cells that constitute an ECS are themselves independent living entities. As you say, they know nothing of the larger system of which they are the components. The variables for which they control are only those that they can sense. The actions they employ for control are those that affect the same variables. Disturbances that alter the controlled variables are opposed by the actions of the system.

Guess: the controlled variables of a neural cell include the potential inside the cell at the axon hillock. Incoming signals disturb that potential. When the potential gets high enough, it is restored to the acceptable level by the cell's firing. In the continuous-firing situation, the cell's rate of firing has an effect on the average internal potential that is equal and opposite to the average effects of excitatory incoming neural signals. So neural cells react to disturbances by going through repetitive electrical convulsions that keep the mean internal potential near a reference signal specified within the cell, perhaps by its DNA. As a side-effect, the cell emits neurotransmitters from the end or ends of its axon. This may amount to getting rid of waste-products generated by its own error-correcting activities!

Neural cells clearly disturb the controlled variables in other nerve-cells, in the process of correcting for disturbances FROM other nerve-cells. One cell in the midst of a network thus acts on its environment, which in turn (through external feedback paths) acts on inputs to the same cell. These feedback paths have to be negative if the internal variable is to be controlled via the external loop rather than running away to one extreme or another.

That leaves me stumped for the moment. So what? It occurs to me that this approach is an attempt to deduce the existence of a higher order of control systems -- the neural heirarchy -- by referring only to the reference signals and control systems inside the cellular components of the larger system. Can we get there from here? I'm thinking of neurotaxis, which seems to be a phenomenon of a level higher than the cell. Can we express neurotaxis in terms of reference signals and controlled variables inside the cell itself? Could a cell that needs negative feedback from somewhere emit a chemical signal from the spot where it's needed? Could a cell with surplus neurotransmitter to unload grow itself toward a spot that wants and can accept that kind of transmitter?

I think there's a limit to how far we can go with this kind of emergence -- maybe. The behavioral hierarchy gets most of its negative feedback through

the external world, where physical phenomena foreign to the body get into the loop. And the effects of controlling for different external (that is, sensory) variables in different ways are important to the body in places remote from the controlling systems: in the stomach, the bloodstream, the gonads, and so on. Something has to link these remote effects back to the very organization of the behavioral systems (that's what my reorganizing system is supposed to do). Could these remote effects get into the loop in any meaningful way at the level of a single cell trying to maintain itself? Somehow I think not: the effects of a single neural signal on the external world, outside the body, would be lost in the general effects from all the nerve-cells that participate in behavior. We're talking, I think, about a much smaller-scale environment, including only a small volume around the nerve-cell.

I have a strong sense of something important here that's considerably beyond my reach.

-----  
The analogy to the human situation that you want to communicate is clear. There is a system whose components are individual human beings. The individuals know nothing of their role in the larger system; they see their actions as affecting themselves only, and don't realize that the side-effects are linking them to other people, loops and meshes of other people that end up affecting the same individual. As the individuals seek to gain control over their immediate environments, they adapt to the feedback effects that include all the other people with whom they interact. These adaptations, created by each individual simply for the purpose of controlling local variables, keeping the sign of the local feedback negative, give the system as a whole properties that aren't characteristic of any individual, but only of the whole network.

The "whole network" is probably a lot of small groups loosely linked to each other. Within the group with which an individual most closely interacts -- like CSG-L -- the feedback effects are strong and immediate, although in our case they're only verbal. The members of such small groups adjust their means of action, their mutual disturbances, so each person can control for what is important to that person, including how each person wants the group to appear. Close-knit groups take on a character that is recognizable, particularly to those in the group but also to others. Hard to describe, but familiar.

The groups interact with other groups. But there are fewer direct interactions than among individuals in a given small group. The world-as-a-whole group probably hardly merits the term.

Hearkening back to previous discussions, there doesn't seem to be any reason to think that at the group level and up the resulting organization is that of a hierarchy of control. But there's certainly an emergent organization, many of them. I think you're quite right in suggesting that whatever this organization is, individuals know nothing of their roles in it. Not that it's impossible to analyze -- but before it can be analyzed, one has to understand that it's there.

Once again, I feel that you've made an important point that is for now beyond my reach. So now I feel as if the bottom has dropped out and the ceiling has been removed and I don't know whether to fall or fly.

Thanks. Best, Bill P.

Date: Fri Jul 10, 1992 2:40 pm PST  
Subject: Re: Modeling reorganization

[Martin Taylor 920710 15:30]  
(Bill Powers to Bruce Nevin 920709.2000)

> You're proposing a very complex "E"CS, and I  
>believe we should resist complexities until observations force us into  
>accepting them because we can see no alternative. We haven't reached that  
>point yet; we haven't proven that the normal operation of the hierarchy,  
>and a SIMPLE principle of reorganization, won't solve the problem. And we  
>may not be ready for such a proof for a very long time (I have no doubt  
>that it will be forthcoming).

OK. You have a problem with some of my suggestions because they are not SIMPLE, which is a good reason. It is the same reason I have a problem with your principle of reorganization using a structure separate from the hierarchy, and why I tried to replace it with a mechanism inherent in the normal operation of the hierarchy. I think it is not simple. Simplicity is in the eye of the beholder, to some extent, but it nevertheless is the only legitimate way of choosing between theories that claim the same precision and range of description of observables.

Reading not so far between the lines of your other posting(s), I think you do not like the idea that intrinsic variables and their reference levels could be at the top level of the hierarchy, whereas I find that concept to be both evolutionarily and developmentally natural. If you reject that concept a priori, then you must see my concept of local reorganization as more complex than your concept of global reorganization. (I note that you reluctantly leave the door open for some degree of locality in reorganization, and you explicitly permit locality of level in reorganization; I allow no such exceptions in the structure I propose).

Let me reiterate what I am actually proposing, because you seem to have some slight misunderstanding of it. I start out by denying something you said, because it contradicts your basic principles, which I embrace:

(920704.0800)

>As I see the reorganizing system, it is concerned with controlling INTERNAL  
>variables only. To reply to a previous comment of yours, I see these  
>variables as including variables not available to the senses, even to  
>proprioception.

In my way of looking at things, ANY variable that can be controlled MUST be sensed. If its value is not determinable by the thing "controlling" it, then it is not being controlled by that thing. I think you must mean something different. Anyway, I start by assuming that the levels of intrinsic variables can be, and must be, sensed in some way (for example, one is taught in high school that an increase of around 5 degC usually doubles the rate of a chemical reaction. That reaction serves as a temperature sensor).

In the initial state, there are only intrinsic variables. Nothing else is sensed. If the entity is to survive and propagate, it acts in such a way as to keep the levels of these variables near some genetically set reference, which is to say that some control system is operating. Perhaps there is a wiggle-motor driven by the summed absolute deviation of the intrinsic variables from genetically set reference levels. Quick wiggling gets it away from places that drive its intrinsic variables to "bad" levels. Such a system sounds to me like the foundation of your reorganizing system. It also sounds to me like the foundation of the normal hierarchy.

Consider. Suppose some aspect of the intrinsic variable set tended to be correlated with some aspect of the environment. Say, for example, that one of the chemical chains started with the absorption of photons. It would be advantageous if the organism had some way of detecting when it was being bathed with photons (but not too many). Let's say that the photon-based reaction affected the concentration of CO<sub>2</sub> in the organism. Then if the wiggle-motor became more specifically sensitive to deviations in the CO<sub>2</sub> level, both above and below optimum, then the organism would begin to control for light level. It would, operationally, have a percept of light--part of the normal hierarchy. But, it still would have only one degree of freedom for control--wiggle rate--and therefore there might be conflict between setting optimum CO<sub>2</sub> and setting optimum values for the true intrinsic variables, for which CO<sub>2</sub> level was a surrogate discovered by chance. Generally (we have assumed) good CO<sub>2</sub> levels covary with good levels of the really important variables, but not always. Good light levels may be found in an acid bath lethal to the organism.

The organism would survive better if it discovered another degree of freedom for control--say, closing and opening membrane pores, assuming it to be bounded by a membrane. Every dissipative structure must have a way of ingesting material and energy, and of disposing of waste, so the organism has to have the equivalent of pores. Chemical evolution has, by now, discovered many specific degrees of freedom for controlling such channels, so it's not an unreasonable thing for our hypothetical wiggler-cell to discover. Given this other degree of freedom, a reorganization driven (by your mechanism or mine) would probably be most successful if it reduced the linkage between the intrinsic variables and wiggle-rate, increased the linkage between light-level(CO<sub>2</sub>) and wiggle-rate, and increased the linkage between the intrinsic variables and pore aperture. So now we have two control systems operating in parallel, one being largely what we might call sensory-motor (S-R terminology, forgive me), and the other being "intrinsic," dealing only with the chemical state.

In this scenario, the reorganization affects whatever input-output connections may be, implicit in the chemical reactions. There aren't any neural links yet. There aren't any neurons. But there are ECSs, firstly one and then two. The intrinsic variable levels are affected by the environmental surroundings of the wiggler cell, such as acidity, temperature, "food," and so forth, of which light level is one. With pore control, the wiggler cell can change its sensitivity to these things while seeking optimum light, but so long as the intrinsic variable errors are not totally decoupled from the wiggle-rate, it will still remove itself from well-lit lethal areas.

At this point, I suspect we already have a difference of opinion. If I read your reorganization concept aright, the intrinsic error would allow the coupling of CO<sub>2</sub> sensing to wiggle rate, but would require something else

to be a sensory surrogate that affected the output leading to pore aperture. And neither of these two control systems would have a reference level. I can't see how that could work, so I imagine I do not read you aright. But I can't see how else to interpret your decoupled reorganization and sensory-motor control structures. And you have said that the top-level reference are a mystery, or are always set to zero. Neither seems appropriate here.

You questioned >"How do you open up the connections from a higher to a lower >system to insert a complete control system with all its connections to and >from both the higher and the lower systems? This idea seems to me to entail >enormous difficulties, whereas building from the bottom up eliminates those >particular problems completely."

I have here proposed an example, as low-level and basic as I can imagine (at present). In general, I assume a complex hierarchy is normally built much as you describe, from the lowest levels of control upward, with one important difference: always new higher levels are inserted between the controllers for intrinsic variables and the levels already built. I assume that it is possible but rare that an established lower level gets seriously disturbed when new ECSs are built. In your system, as I understand you, this decoupling happens because of some (to me mysterious) organizing principle:

>I'm assuming now that reorganization is specific to each level.

I think it simpler to make no such assumption. But one assumption I do make is that each intrinsic variables has some genetically determined optimum value, deviation from which serves as an error signal that (after amplification) serves as a reference signal for some set of ECSs in the "normal" hierarchy.

The second assumption that I think you don't like is that each ECS, including those involving the perception of intrinsic variables, will "reorganize" at a rate depending on its error (or rate of change of error--the criterion is not at issue). Reorganization could mean any of the things you talk about as being subject to reorganization, so that's not an issue either. I am assuming that what you don't like is the complication of the ECS by providing a mechanism for it to reorganize. But is this more complex than having a separate reorganization system that can go into individual ECSs like a mad electrician and do the same rewiring job?

According to my scheme, ECSs that are able to control their percepts near their references will not reorganize much, if at all (we agree that there may be a threshold--I think there must be, for thermodynamic reasons). Low-level ECSs will normally maintain control, barring accidental damage. The most reorganization will tend to occur on higher levels that are newly being built. Reorganization will thus SEEM to be specific to each level, but that falls out naturally rather than being a design criterion for the system.

Occasionally, the reorganization of a high-level ECS will send to lower ones reference levels that cannot be satisfied, causing the lower ones to experience sustained error and to reorganize themselves. This can cascade, but is unlikely to do so for very long in a developed hierarchy.

I think my scheme is much simpler than yours, and should be expected to occur naturally right back to the very beginning of life/control. It is inherent in the structure of control, rather than being dependent on two interacting



structures, one of which exists only to disturb the other. And it fulfils your criterion that the reorganization should be random, knowing nothing about where error comes from or what it "means" to the hierarchy.

-----  
Reorganizing high-gain systems:

>If we think of reorganization as instituting small changes in the  
>parameters of control, there's no reason why a high-gain control system  
>can't go right on controlling. Even in category control (to pick up that  
>thread), slight changes in perceptual parameters would just move the  
>boundaries of the category a little. I think those boundaries are fuzzy  
>anyway, so we don't really have any discrete-variable problems here.

Aren't you shifting ground here? Hebbian perceptual learning is a very different kettle of fish from switching the sign of an output-reference link. The latter can't be smooth. Changing what an ECS is trying to perceive is a little different from changing its method of trying to perceive a given variable. Hebbian learning is important, I think, in developing control structures that control complex environmental variables that "really" do something coherent. If the perceptual functions are all of the simplest type, a weighted sum followed by a nonlinearity, then the perceptual side of the hierarchy is a multilayer perceptron, and if it has three or more layers, it is capable of subdividing the sensory input space in any way at all. It will divide it in a way that is consistent with coherences between the input and the desired output.

And I don't think category boundaries can ordinarily be fuzzy. I think they are usually catastrophic, so that although small changes usually have no effect, sometimes they have a dramatic effect.

>Look at it the other way around. If reorganization is to work as I propose,  
>it CAN'T work fast.

I agree, and I see that as one of its problems. My scheme can handle (will handle automatically) the control-reversal experiment you did with Rick. How do you explain the switch in the sign of control that happened within half a second of the reversal of the external connection? How did the change of control to a new, previously learned, system happen, and when and why did that learning to reverse take place? The world isn't full of reversing feedback connections.

I do think rapid reorganization can happen, at least within my local scheme. Any time a ECS has a large and rising absolute error, it should be expected to do some random sign flipping and/or reduce its gain. It is true, as you say, that each control system sees only one degree of freedom, so if it can flip and detect the results of the flip quickly, it can succeed through the e-coli procedure. But, and here's a big but, its changed requirements on lower-level ECSs may induce conflict with other same-level ECSs, causing more error in them and possibly failing to solve its own problem. They also may reorganize, and we are back in the same old dimensionality problem. The error threshold for reorganization comes into play here, with luck reducing the dimensionality to a feasible range.

-----

>However we end up talking about reorganization, we have to arrange for  
>somatic states of the organism to have a very powerful directing effect on  
>what control systems are acquired. Logic and principles often bow to  
>hunger. How does that work? I can see how my model would do it, but how  
>does yours?

Hunger is presumably a surrogate for some error in intrinsic variables (I believe it is largely caused by stomach contractions, so it isn't a sensation of intrinsic variable state). It should be capable of control by changing reference levels for principles or anything else. I don't know where the actual ECSs that accept the hunger sensation as part of their input might be in the hierarchy, but in my scheme, if hunger is not controlled over some long period, then intrinsic variables will depart from their reference levels, and they are the top-level ECSs, dominating all else. No problem.

>Perhaps what we need to do, as we're not likely to resolve this conflict  
>experimentally, is to find a way to talk about reorganization so it doesn't  
>matter which is the right model.

There's enough in common that we should be able to do that. Reorganization is blind and random in both cases. The only substantive difference is in the localization of the reorganization in my scheme as compared to its distribution in yours. That difference may often be unimportant.

Martin

Date: Fri Jul 10, 1992 2:42 pm PST  
Subject: Re: Martin's postings

[Martin Taylor 920710 17:30] (Bill Silvert 920710)

>Has anyone been able to transfer Martin's files csg.9206.sea and  
>mailsplitter.bin to a Mac successfully? I tried to do this today but  
>they came out as document files. If anyone knows how to do this, please  
>send me an informative message that I can include with them.

Did you transfer the files in MacBinary II format? If you use NCSA telnet, you set binary mode on the transfer, and MacBinary in the File menu. I use Fetch 2.1b4 from Dartmouth myself. I know one can do it with Zterm (shareware) because I've done it that way over the phone from home. I have never used Kermit, but MacTerminal 2 or 3 also will handle MacBinary files. (I recommend Zterm, available on sumex-aim.stanford.edu and other places if you are coming in over the phone, Fetch if you are ftp'ing from the Mac). If you are ftp'ing from a mainframe, use binary mode and then be sure you are downloading to the Mac in MacBinary mode.

Sorry if it's confusing. I just used Fetch directly from my Mac. It's lovely for ftp purposes. Does all these format checks and transfers in the proper mode automatically. Available on dartvax.dartmouth.edu.

Martin

PS. I just checked that the Mailsplitter (at least) was correctly uploaded by downloading it. It's fine. Runs normally. I used Fetch to download it.

Date: Fri Jul 10, 1992 7:03 pm PST  
Subject: Re: Flippin' eggs

>what does it do ... ? - Wouldn't have a clue...

Avery.Andrews@anu.edu.au

Date: Sat Jul 11, 1992 12:06 pm PST  
Subject: Reorganization

[From Bill Powers (920711.0800)]

Martin Taylor (920710.1530) --

What I like about your proposal is the idea that there's a random-output reorganizing process operating in the (potential) CNS hierarchy, gradually being replaced by (or paralleled by) systematic control. This reorganizing process could monitor certain built-in aspects of potential control systems in the CNS, such as the signals being emitted by comparators (which are so simple that we can assume them to be part of the initial complement of parts). Also, the brain is physically organized so that sensory computers tend to be lumped together, as do motor output computers. There may be aspects of these more or less localized networks that permit monitoring for invariance, negative feedback relations between motor and sensory signals, and so on -- without an implication that the reorganizing process knows in advance what perceptions or actions will be appropriate for the current environment.

One bit of confusion was caused by my using "sensing" in an ambiguous way. I agree, of course, that all control requires sensing. When I said that intrinsic variables were outside the sensory interface, I meant that they are outside the CNS sensory interface. I don't think that chemical sensors and comparators can be counted as part of the CNS hierarchy. If some biochemical states of the organism become parts of experience (as in emotions), they must affect neural sensors connected to the CNS. Under my concept of the reorganizing process, if biochemical states are to become a basis for reorganizing, they would be sensed in some other way, probably by chemical or autonomic sensors. To the CNS, sensed internal states are no different from sensed external states; they are simple data about the current state of affairs in the world outside the CNS. They have no value or any built-in reference states as far as the CNS is concerned.

>Let's say that the photon-based reaction affected the concentration  
>of CO<sub>2</sub> in the organism. Then if the wiggle-motor became more  
>specifically sensitive to deviations in the CO<sub>2</sub> level, both above and  
>below optimum, then the organism would begin to control for light  
>level.

You're making the photon reaction itself produce the CO<sub>2</sub>. What I'm

interested in is the case in which the photons arise from some object, the presence of which implies some OTHER effect on CO2 in the body. The photons might come from a warning light that indicates a leak in a CO2 container. Now the organism has to learn to get out of there when the warning light (whistle, vibration) is sensed, even though the light itself is not the cause of the elevated CO2 tension in the bloodstream, and position in space, not CO2 tension, is the variable that becomes controlled. So the organism must learn to control for a specific level of one perception as a way of controlling another variable inside itself that has no necessary connection to the controlled perception, save for happenstance properties of the external world.

The pigeon has to be able to learn to walk in a figure eight (itself a control process) as a means of making food appear. Walking has no effect on improving nutritional state, save for the mad scientist in the environment. Something has to make the learning of one control process depend on an indirect and arbitrary effect of that control process on the environment, and by that indirect route on the internal state of the organism. This is the kind of thing that my reorganizing system is meant to accomplish. I don't think that any localized reorganizing principle could do it.

-----  
 I think that before we get any further into complicated arguments and misunderstandings, we should do some work on simulating reorganization. I'm going step by step on this, and am not ready to show a whole indirect reorganizing process yet. At the moment, what I have is a method of solving simultaneous equations by reorganization. This is not meant to imitate any particular organismic process; that comes a few steps further on. What it does is illustrate some principles.

The basic setup is this:

There are 10 perceptual functions, each producing a linear function of 10 environmental variables. The form of the linear functions is generated by 10 weights for each of 10 perceptual functions, so that

$$p[i] = a[i,j]*v[j], \text{ where } 0 \leq i < 9, 0 \leq j < 9$$

The weights  $a[i,j]$  are generated at random in the range between -1 and 1 and are fixed. The initial values of  $v[i]$  are zero.

There are also 10 fixed reference signals  $r[i]$ , generated at random in the range -50 to 50. Thus we can compute 10 error signals  $e[i] = r[i] - p[i]$ .

The reorganizing system computes the square root of the sum of squares of all error signals, which is the distance in 10-dimensional hyperspace between the point  $p[i]$  and the point  $r[i]$ . On each iteration, the distance is compared with the distance on the previous iteration, to provide a measure of velocity toward or away from the point  $r[i]$ .

If the velocity is positive on any iteration, a reorganization takes

place. For all negative velocities, the direction of movement resulting from the last reorganization is applied over and over.

Reorganization is a two-step process.

First, a 10-element vector  $\delta[j]$  is filled with random numbers between -1 and 1. It is normalized by dividing each entry by the square root of the sum of the squares of all entries; this makes the entries into direction cosines in a 10-dimensional space.

Reorganization thus changes the direction of this vector randomly in 10-space. This normalization is not essential, but seems to make convergence faster.

Second, the 10 elements of  $\delta[j]$  are added to the 10 current values of the environmental variables  $v[j]$  after multiplication by "stepsize." This causes the complex environmental variable to move a distance "stepsize" in hyperspace.

The loop is now closed, because that movement in hyperspace results in a change in perceptions  $p[i]$  and therefore in the errors  $e[i]$ .

The variable stepsize can be a small constant, or (for faster convergence) can be proportional to the remaining hyperspace distance. I have used the latter method.

Here is part of a run arbitrarily cut off at 5000 iterations:

Iteration	Error left, fraction of max	# consecutive steps	# consecutive reorgs
n= 1	e/emax= 0.996	steps= 1	reorgs= 1
n= 14	e/emax= 0.978	steps= 12	reorgs= 1
n= 17	e/emax= 0.977	steps= 2	reorgs= 1
n= 24	e/emax= 0.975	steps= 5	reorgs= 2
n= 29	e/emax= 0.974	steps= 3	reorgs= 2
...			
n= 4968	e/emax= 0.011	steps= 8	reorgs= 1
n= 4974	e/emax= 0.011	steps= 3	reorgs= 3
n= 4984	e/emax= 0.011	steps= 8	reorgs= 2
n= 4990	e/emax= 0.011	steps= 1	reorgs= 5
n= 4994	e/emax= 0.011	steps= 3	reorgs= 1

-----  
DATA SUMMARY after run:

Environmental variables (target of reorganization):

49.84 13.35 -24.48 -3.29 25.79 14.76 -0.48 -9.68 67.15 -0.92

Perceptual coefficients (picked at random, normalized to 1.0):

-0.890 -0.460 -0.750 0.370 -0.850 0.840 0.610 0.630 -0.360 0.440  
-0.850 0.060 -0.950 -0.950 0.530 -0.910 -0.150 0.370 0.400 0.200  
-0.690 0.870 0.790 0.140 -0.570 0.100 -0.740 -0.390 -0.430 0.540  
-0.390 -0.760 -0.400 0.720 -0.450 -0.760 0.340 -0.930 -0.560 0.130

```

-0.460 0.700 0.170 0.190 -0.360 0.550 0.990 -0.090 0.390 -0.800
0.030 -0.810 -0.040 0.220 -0.440 -0.490 0.110 -0.280 0.560 -0.280
0.640 -0.650 0.310 0.670 -0.970 0.050 -0.220 0.310 0.730 0.930
0.650 -0.050 -0.240 -0.460 0.300 0.820 -0.750 -0.570 -0.430 -0.820
0.030 -0.140 0.840 -0.910 0.200 0.140 0.310 0.200 0.790 -0.910
0.750 0.870 -0.780 -0.360 0.470 -0.940 0.640 0.230 0.620 0.650

```

Perceptual signals:

```
-20.93 -30.02 930 10.08 33.01 -20.97 -0.87 15.00 36.93 -1.97
```

Reference signals (picked at random):

```
-21.00 -30.00 10.00 10.00 33.00 -21.00 -1.00 15.00 37.00 -2.00
```

Error signals:

```
-0.07 0.02 0.70 -0.08 -0.01 -0.03 -0.13 -0.00 0.07 -0.03
```

Count = 5001, RMS Err/max = 0.0106

-----  
What's interesting about this control system is that it is independently controlling 10 perceptual variables with respect to 10 arbitrarily-selected reference signals, by means of randomly altering 10 environmental variables on which all the perceptual variables depend in different ways. The control employs only the total error, not the error in each channel by itself.

As I mentioned in a previous post, I have also made this work by reorganizing the coefficient matrix  $a[i,j]$ , with the environmental variables fixed at random settings. So we have proof of principle for two major ways of reorganizing: reorganizing the input function, and reorganizing the feedback link (the present case). There's no reason why both modes of reorganization can't be going on at the same time. All that matters is whether those two points in hyperspace are getting closer together or farther apart.

The basic random process occurs between the total error signal and the set of environmental variables. This set of variables could represent whole control systems, with the random adjustments being made on the forms of their input functions, the signs of output connections, and the various factors influencing loop gain. This arrangement would look like the one you're suggesting, where the CNS is reorganizing itself.

Here's the C function that does the actual reorganizing. Most of the variables are globals. The `init()` routine and the routine that calculates error signals are also included:

```
/*below is called once per iteration */
```

```

void reorg(float *e,      /* pointer to list of error signals */
           int nerr,     /* number of error signals */
           float *delta, /* pointer to motion array */
           float *var,   /* pointer to controlled variable array */
           int nvar)     /* number of entries, delta & var array */
{
float temp;
static int numsteps = 0;

```

```

static int numreorg = 0;
static int newstate = 0;      /* state=0 means not reorganizing */
static int oldstate = 0;
  lastesq = esq;
  esq = 0.0;
  for(m=0;m<nerr;++m)
    esq += e[m]*e[m];
  distance = sqrt(esq);
  rate = distance - lastdistance;
  lastdistance = distance;
  stepsize = 0.7 * distance/maxdistance;
  if(rate > 0.0)
  {
      /* reorganize delta array */
      newstate = -1;
      ++numreorg;
      for(m=0;m<nvar;++m) /* pick all new deltas at random */
        delta[m] = (rand() - 0.5*RAND_MAX)/(0.5*RAND_MAX);
      temp = 0.0;
      for(m=0;m<nvar;++m) /* normalize to 1.0 */
        temp += delta[m]*delta[m];
      temp = sqrt(temp);
      for(m=0;m<nvar;++m)
        delta[m] /= temp;
      /* delta now is a set of direction cosines */
    }
  else {++numsteps; newstate = 0;}

  if(newstate != oldstate) /* print consecutive # steps, # reorgs */
  {
    if(newstate == 0)
    {
      printf("\x0d\x0an= %6d  e/emax= %7.3f  steps= %3d  reorgs= %2d",
        count,distance/maxdistance,numsteps,numreorg);
      numsteps = 0; numreorg = 0;
    }
    oldstate = newstate;
  }

  for(m=0;m<nvar;++m) /* move contr var in hyperspace */
    var[m] += stepsize*delta[m];
}
void calcerr() /* compute all error signals */
{
  for(m=0;m<numvars;++m)
  {
    p[m] = 0;
    for(n=0;n<numvars;++n)
      p[m] += v[m]*a[m][n]; /* compute perceptual variables */
    e[m] = r[m] - p[m]; /* compute error signals */
  }
}

void init()
{
  for(m=0;m<numvars;++m)

```

```

{
  r[m] = random(100) - 50;
  delta[m] = 0;
  for(n=0;n<numvars;++n)
    a[m][n] = 0.01*(random(200) - 100);
}
count = 0;
maxdistance = 0.0;
for(m=0;m<numvars;++m)
  maxdistance += r[m]*r[m];
maxdistance = sqrt(maxdistance);
}

```

Feel free to use, modify, etc.

-----  
 Best Bill P.

Date: Sun Jul 12, 1992 11:38 am PST  
 Subject: Back in the USSA

[From Rick Marken (920712.1200)]

Well, I'm back. Had a nice time in London; hardly though about control at all -- just did it, as best as we could given the peculiar reference levels over there.

I found out why Psych Review didn't even send my "Blind men" paper out for review. According to the editor it was because:

"It would need to speak more directly to current psychological issues and theorizing. One would need to see more clearly a connection between what you are talking about and the issues that dominate psychological theorizing today."

I guess the nature of the phenomenon they are theorizing about is not a current issue for psychologists. I think that the most direct connection between "current psychological issues and theorizing" and control theory is that from the latter perspective the former are complex rationalizations of non-existent phenomena. How do you tell psychologists that their theories which explain the effects of factors a,b or c on variables x, y or z are a waste of time because there are no such effects; just statistical noise?

People sometimes criticize control theory for not being based on a large enough data base. Most of PCT data comes from simple tracking experiments; we have looked at the control of many different types of variables -- but it seems like the amount of PCT data is small compared to the amount of data piled up in the psychology journals. I think that this is a misperception, however, because 90+% of the data in the journals is basically noise. The amount of real, worthwhile facts in the journals is probably far less than what is already part of the PCT literature. In fact, what real data exists in the psych literature has already been plucked up by PCT types. This includes some of the operant conditioning data



and existing tracking data. There is a lot of suggestive data in the psych literature -- but it will continue to be little more than suggestive until someone does the studies correctly -- so that the relationships between variables are consistently perfect -- less than 3% error variance always.

I think it is interesting that the only data in the psych literature that meets PCT standards of quality (in terms of error variance) is data obtained in situations where the subject is clearly controlling a variable -- as in the tracking and operant tasks (and some psychophysical tasks, especially those where the subject is controlling a relationship between variables). In fact, that may be one way of pressing the case for the value of the PCT perspective (assuming that you accept the idea that a science should be based on high quality data. I have found that this is NOT a universally accepted idea, especially in the social sciences; in fact, I have actually run into people who found the results of some of my studies of coordination to be suspect (or uninteresting) precisely because they were NOT statistical; the fact that the control model accounted for 99% of the variance in behavior also made the results seem "trivial" to these people. PCT has to deal with the fact that many social and behavioral scientists think that data are not interesting unless there is a sizable amount of error variance.) Only by viewing behavior as control and setting up situations where a person can control a variable can one get the kind of quality data one finds in "real" sciences like physics and chemistry. The fact that the IV-DV approach to research gives such crummy results suggests that it is based on the wrong model. Maybe it's time for that discussion of statistics now?

Best regards

Rick (Obviously not mellowed by Europe) Marken

Date: Mon Jul 13, 1992 6:29 am PST  
Subject: Trendy science

[From Bill Powers (920713.1730)]

Rick Marken (920712) --

Welcome back. I'm too flabbergasted by the reaction to your "Blind men" paper to speak of anything else.

>I found out why Psych Review didn't even send my "Blind men" paper  
>out for review. According to the editor it was because:

>"It would need to speak more directly to current psychological issues  
>and theorizing. One would need to see more clearly a connection  
>between what you are talking about and the issues that dominate  
>psychological theorizing today."

What could be a better illustration of the trendiness of psychology? Only 13 years ago, I published an article on control theory in that very

journal. Already it has passed over the horizon and is no longer an "issue that dominates psychological theorizing today."

Can you imagine what physics would be like with an event horizon of only 13 years? Physics would no longer be concerned with inverse square laws, optical refraction, the Hubble Constant, gas laws, electrical phenomena, lasers, transistors, or the Mossbauer Effect. Instead of building up a coherent and growing picture of nature, physicists would be worrying about whether they're working on things that are popular and current. There would be as many schools of physics as there are of psychology, sociology, or economics.

This is confirmation of my thesis that a science based on low-probability facts can't create a coherent picture of nature. When only specific effects under specific circumstances are studied, facts lie scattered around the landscape in disconnected confusion. No argument involving more than a small handful of facts can lead to deductions with a probability of truth greater than chance. Reasoning is limited to the fourth-grade level. This is why the great majority of observations that pour out into the literature are forgotten the day after they are published, not to mention 13 years later. There is no underlying body of understanding to which each new observation adds, or even potentially adds. Today's hot subject is tomorrow's phrenology, disappearing when the careers of those making a living off of it end. There is no science of psychology. There's only a Psychology Club, and its newsletter is called Psychology Today. I suppose that it might have a Nostalgia Column titled "Three years ago this month."

In disgust, Bill P.

Date: Mon Jul 13, 1992 7:08 am PST  
Subject: cellular/supracellular control

[From: Bruce Nevin (Mon 920713 08:03:56)]

(Bill Powers (920710.1330) ) --

The hypothesis: differences in the environment of an organism that make a difference within the organism (error, but especially intrinsic error and conflict resulting in chronic error) must also make differences that make a difference in the environments of the organism's constituent cells (error in the intra-cellular control system). The actions of cells to reduce intra-cellular error must (in part?) amount to the supra-cellular changes that we call reorganization, and perhaps also some forms of learning. As you put it:

>The individual (neural) cells that constitute an ECS are themselves  
>independent living entities. . . . they know nothing of the larger  
>system of which they are the components. The variables for which they  
>control are only those that they can sense. The actions they employ for  
>control are those that affect the same variables. Disturbances that alter  
>the controlled variables are opposed by the actions of the [cell's] system.

Nonetheless, as a byproduct of their autonomous self-control in an environment comprising other cells and their byproducts, the cells

together do in fact constitute higher-order control systems of which they can have no ken (because they lack the perceptual means).

>It occurs to me that this  
>approach is an attempt to deduce the existence of a higher order of control  
>systems -- the neural heirarchy -- by referring only to the reference  
>signals and control systems inside the cellular components of the larger  
>system. Can we get there from here? I'm thinking of neurotaxis, which seems  
>to be a phenomenon of a level higher than the cell. Can we express  
>neurotaxis in terms of reference signals and controlled variables inside  
>the cell itself? Could a cell that needs negative feedback from somewhere  
>emit a chemical signal from the spot where it's needed? Could a cell with  
>surplus neurotransmitter to unload grow itself toward a spot that wants and  
>can accept that kind of transmitter?

I am not trying to \*deduce\* the existence of a higher order of control systems in the case of cells and neural control systems. They are observational givens (within the theory). (Perhaps you are here looking ahead to deducing suprapersonal control systems, by analogy? We'll get to that below.) My question at this point is how supracellular control systems can come to be, using only the means that cells have at their disposal.

Taxis in general is construed as the movement of an organism toward a stimulus. We reject the explanatory framework presupposed in the word "stimulus," of course. There is some disturbance to a controlled variable within the cell, such that the observationally perceived behavioral output called taxis counteracts the internal effect of that disturbance.

>Could a cell that needs negative feedback from somewhere  
>emit a chemical signal from the spot where it's needed? Could a cell with  
>surplus neurotransmitter to unload grow itself toward a spot that wants and  
>can accept that kind of transmitter?

A "need for negative feedback" can be observed only from the point of view of the supracellular control system (actually, from a point "above" that even). How might such a need manifest as a disturbance to one or more cells in a control system? And which cell or cells?

We can see how complex interdependencies of cells can come to exist over evolutionary time. By establishing symbiotic relationships cells mutually experience the advantage of each providing a more stable and more easily controlled environment for the others. Observationally, "above" the cell's point of view, we see specialization in terms of the morphology and behavioral outputs of different kinds of cells.

As you suggest, a waste product of a cell's error-correcting behavioral outputs may serve as what we observationally perceive as a signal or as a neurotransmitter. Possible analogy: the particular smell of scat may signal to a predator that one of the deer in a herd is sick and ripe for culling, or the smell of testosterone-laced urine warns of territorial limits.

>I think there's a limit to how far we can go with this kind of emergence --

>maybe.

Absent an antecedently given teleology, this has to be the route for explanation of ontogeny. For phylogeny, symbiotic interdependencies (giving the benefit of a more stable and predictable environment for each cell) provide supracellular scaffolding that was not present for evolutionary forebears, but explanation still must take the point of view of the cell, not of the control system that it participates in constituting. (Social institutions, customs, traditions, etc. provide suprapersonal scaffolding that was not present for evolutionary forebears, cf. Bruner's Language Acquisition Support System. But I'm jumping ahead.)

>The behavioral hierarchy gets most of its negative feedback through  
>the external world, where physical phenomena foreign to the body get into  
>the loop. And the effects of controlling for different external (that is,  
>sensory) variables in different ways are important to the body in places  
>remote from the controlling systems: in the stomach, the bloodstream, the  
>gonads, and so on. Something has to link these remote effects back to the  
>very organization of the behavioral systems (that's what my reorganizing  
>system is supposed to do). Could these remote effects get into the loop in  
>any meaningful way at the level of a single cell trying to maintain itself?  
>Somehow I think not: the effects of a single neural signal on the external  
>world, outside the body, would be lost in the general effects from all the  
>nerve-cells that participate in behavior. We're talking, I think, about a  
>much smaller-scale environment, including only a small volume around the  
>nerve-cell.

You have identified the problem. I think you have also identified the solution to it, long since: the control hierarchy. Variables in the environment of the organism are far beyond the immediate environment of a cell in the neural control system. But the neural control system is so structured as to bring news of a difference in the organism's environment (a difference that might make a difference to the organism) from cell to cell, each making a difference that makes a difference to its neighbor (each time just in immediate cellular environment of its neighbor), until there is a difference in the environment of the particular cell that we have singled out for observation. Setting aside for the moment the question of how this marvellously articulated ramification of neural bucket brigades came into being (ontogeny and phylogeny), its existence explains how a difference in the environment of an organism can be transformed (through the control hierarchy) into a difference in the environment of any given cell in the hierarchy.

=====

You have understood and articulated the social analogy very clearly, and you have also stated a reason (I suspect) you have been reluctant to broach it in other discussions:

>I feel as if the bottom has dropped out and the  
>ceiling has been removed and I don't know whether to fall or fly.

There is no need to do either. Nothing has changed, this is just where you always were. More to the point, the same choice to limit your focus

to the proper purview of control theory is still available. The only difference is your acknowledgement that it is a choice and not a preconditional Reality.

Unsolicited homily #37: Limitation is the first step of any creative process. You define the scope of the work. You draw your magic circle, and you ignore everything outside it. But as any magician worth his salt knows, you don't just forget about what lies outside the ordered realm of the work, and you don't live there inside the circle. From time to time you dissolve the circle, take a lunch break, whatever. When you re-form the circle, you may well bring something in that wasn't there before, and throw out some baggage that turned out just to be in the way. All familiar process, by small trial-and-error steps. I'm not advocating that you expand your magic circle too far, beyond your (our) means to control. I'm just suggesting that you take your lunch breaks in interesting places. And that you not be so troubled by others drawing intersecting circles. In general you don't. There are just some strong commitments respecting social control that sometimes blindside you.

What this does first is to give license to say "I don't know" about whether there are suprapersonal control systems or not. And to be comfortable with that, since (on the proposed analogy) it is only by virtue of individual human "cells" controlling for just what matters most to each of them that such higher-order systems can be constituted. The intuitive grasp which I am seeking to articulate is that the pursuit and realization of one's (evolving) heart's desire turns out to be one's way of helping to constitute a healthy, well-controlling higher-order control system (though one inherently cannot control for constituting such a system per se--as with the cell constituting the ECS, it is beyond one's perceptual means), and that participation in one's particular capacity in such a control system turns out to be the most personally fulfilling thing one can find to do (all by trial and error, of course). This is related to Ruth Benedict's ideas on synergy (the post I sent to CSG-1 shortly after coming on board a little more than a year ago).

Of course it is possible for people to control for reference perceptions other than those that (by trial and error) they find most fulfilling. They can define personal achievement in terms of ability to deny fulfilment to others "under" them, for example. There's a lot of that going around. I understand that the precursors of symbiosis are competitive relations destructive to the rivals for the same niche and parasitic relations destructive to the host. How might such changes come about, from the cell's point of view? Our shared animosity for abuses and abusers of social relations and institutions must be separated from our advocacy of control theory as a science. Control theory (or an understanding and acceptance of it) does not preclude such abuse, alas. Indeed, a conviction that it does could, with appropriate missionary mind-set, support rationalization and outright ignoring of one's participation in such abuses, as human beings have demonstrated again and again, with remarkable creativity and imagination. (Al Capone, Dale Carnegie tells us, thought of himself as a benefactor to humanity. Dale Carnegie did too.)

The purview of control theory is limited to the same purview as one has for conflict resolution within the control hierarchy. What you can see is what you can "get," so to speak. Beyond that we can make suggestive analogies upward and downward. How does it all work from the point of view of a cell? Of a molecule? What does the process of cells evolving to a control system potentially tell us about our ongoing social evolution?

Unlike the cell, and to a much greater extent than other animals (so far as we know), we have limited capacity to extrapolate beyond our immediate perceptual means. From what we can determine from observation of cells and molecules (of their behavioral outputs constituting structures and systems), and of control systems on the scale of animals and humans, what analogies are there to groups, cultures, ideologies, etc., and are those analogies useful?

Between levels of control there is a control relationship. Between orders of control (e.g. cellular vs ECS) there is a constitutive relationship. The limits of our direct perception are related to the constitutive bounds of our perceptual control hierarchy. By using imagination and analogy to interpret "meter readings" of various kinds, we can extend our understanding lower in the constitutive hierarchy--and perhaps higher as well. (Science is imagination and analogy systematized.) It seems to me implausible that the constitutive hierarchy of control systems has its upper bound the order of control evolved by multicellular organisms such as humans. I am prepared to entertain the possibility of persons (or person-like organisms) whose "cellular" control participates in constituting a higher order of "supra-cellular" control. We can't test that notion, but it may be important context for devising and interpreting tests that are within our grasp.

Things to think about when control theorists are out to lunch.

Bruce bn@bbn.com

Date: Mon Jul 13, 1992 12:08 pm PST  
Subject: Individual vs. social control systems

[From Bill Powers (920713.1200)]

Leaving for Denver/Boulder tomorrow (Tuesday 14th) morning -- back Saturday.

Bruce Nevin (920713.0803) --

RE: individual and social hierarchies

In some forms of hierarchy theory, a hierarchy consists of a group of elements seen individually, in small groups, in groups of groups, and so on. The "levels" in such a hierarchy don't introduce any new individuals; they are more like the result of a single observer taking successively more abstract points of view toward the one set of individuals, in effect re-perceiving the original individuals in different ways.

In the PCT hierarchy as I have conceived it, new levels consist of new control systems; they aren't simply the same bottom-level systems looked at from farther and farther away. If a level of control is added, it is added EXPLICITLY, as input, comparison, and output functions physically distinct from the components of already-existing systems at lower levels.

This principle of explicitness distinguishes, I think, purely conceptual models from models intended to represent a physical system. In a conceptual model, you could say that a collection of intensity signals is conceptually equivalent to a sensation perception, that the set of conceptually-defined sensations is equivalent to a conceptually defined configuration level, and so on. After all, the sensation is implicit in the collection of intensities; the configuration is implicit in the collection of sensations, and so on for as many levels as you like. However many levels you add in this conceptual way, however, in the physical system there is still only the original set of elements; nothing has in fact been added.

In a model intended to be physical in nature, however, nothing that is merely implicit can have any effect. A collection of control systems controlling individual intensity signals will behave exactly the same way whether an observer attends to the individual systems or to the conceptually implicit control of sensations. In a physically-oriented model, there will be no control of sensations until some neural function receives a set of intensity signals and creates a new signal explicitly dependent on that set, according to some functional form. There will be no control until that explicit perceptual signal is compared with a reference signal, and the difference is routed to the reference inputs of some of the sensation-controlling systems. All those functions and signals must physically exist, distinct from the systems of the lowest level, before sensation control can become explicit, and not just represent a viewpoint of the observer.

Frank Rosenblatt elucidated this principle a long time ago, when he insisted that for any variable pertinent to behavior to have effects in a real system, it must be embodied as an explicit signal. This is my basis for saying that anything experienced has to exist as a neural signal, a perceptual signal. It's not enough that something COULD be perceived in a collection of elements; it must BE perceived, exist as an explicit signal, before the rest of the brain can do anything with it. That's why my system uses explicit signal paths even to represent imagined information. Without this principle of explicit representation, there would be no hope of committing any model to hardware.

So much for the preamble; now to your post.

>... as a byproduct of their autonomous self-control in an environment  
>comprising other cells and their byproducts, the cells together do in  
>fact constitute higher-order control systems of which they can have no  
>ken (because they lack the perceptual means).

If these cells constitute higher-order systems ONLY IMPLICITLY, by virtue of the way we conceptualize their interactions, then there is no physical higher-order system. The higher-order system exists only as a conceptual level in the mind of the observer, and is not actually part of the physical cells or cell assemblies. It is not part of the organization of the real

system.

In order for an actual higher-order system to exist, some cells must take on roles that the others do not; they must become concerned with an environment consisting of the other cells, and act on that same environment, explicitly sensing something about the other cells, explicitly acting on something that alters what is sensed about the other cells. In the brain, the configuration level of control consists of neurons in the midbrain physically distinct from those in the brainstem and spinal systems that carry out lower levels of control. The configuration system senses signals that are the perceptual signals of sensation systems; its outputs go not to the spinal systems, but to the sensation-control systems, as reference signals. This is an explicit new level of control that is physically distinct from and does something different from the systems of lower order.

So the mere existence of a collection of control systems of a given level, and the mere fact that they interact with each other in the course of their control actions, can't by itself result in a new level of control. There can be no higher levels of variables until some set of cells explicitly computes the higher variables; there can be no higher level of control action until some set of cells generates an explicit error signal that reaches specific members of the lower-order group of systems. New cells with new specializations must appear in the proper relationship to the old population of cells before control at a new level can explicitly exist.

>We can see how complex interdependencies of cells can come to exist >over evolutionary time. By establishing symbiotic relationships cells >mutually experience the advantage of each providing a more stable and >more easily controlled environment for the others. Observationally, >"above" the cell's point of view, we see specialization in terms of the >morphology and behavioral outputs of different kinds of cells.

But I claim that we cannot yet see how levels of dependency come into being. A special kind of specialization is needed. Valentin Turchin characterizes the required kind of development as a "metasystem transition," which is very different from a mere proliferation of systems at an existing level. This required specialization removes some cells from the population that previously existed, and puts them into a superordinate position, so now the old population is in their environment; they no longer share the same environment with the older systems of cells. A metasystem transition produces a physically new population of cells with different functions in the whole system: the functions we associate with higher levels of control.

We can see, of course, what is gained from such a metasystem transition, but we haven't a clue as to what is different about that kind of specialization.

So you can see that I must reject your extension of this analysis to social systems:

>...it is only by virtue of individual human "cells" controlling for >just what matters most to each of them that such higher-order systems >can be constituted.



The analogy to a metasystem transition within the cells of an organism would be the appearance of a new kind of organism which senses the condition of other organisms, compares that with the condition it wishes to perceive, and acts by setting the highest reference signals of other organisms. I know of no such superordinate organisms. Some people claim that they do, and call them gods, or God: creatures so superior to human individuals that they control for variables of kinds inconceivable to mortals, and who act by injecting reference signals into the highest levels of human consciousness, for purposes beyond human ken.

Perhaps such individuals exist. I wouldn't know, and neither would anyone else who isn't one of them. Whatever the case, it's certain that there are no social control systems constituted by the mere control behaviors and interactions of ORDINARY people -- and until dramatic evidence to the contrary appears, I will assume that all people now on earth are quite ordinary human beings. You can make up any stories to the contrary that you like, but they won't be germane to the point we have been disputing.

To sum up: your argument seems to depend on the emergence of higher levels from populations of systems of an existing level, without the addition of any new kinds of physical systems. I would claim that you are relying on IMPLICIT organization to create new levels. I, on the other hand, argue that the new levels must be EXPLICIT.

-----  
Best, Bill P.

Date: Mon Jul 13, 1992 12:12 pm PST  
Subject: Re: Trendy science

[Martin Taylor 920713 15:15]  
(Bill Powers welcome back to RICK Marken 920713.1730)

I'm still listening, but not contributing unless tweaked, until probably this weekend or next week.

>>"It would need to speak more directly to current psychological issues  
>>and theorizing. One would need to see more clearly a connection  
>>between what you are talking about and the issues that dominate  
>>psychological theorizing today."

>

>What could be a better illustration of the trendiness of psychology? Only  
>13 years ago, I published an article on control theory in that very  
>journal. Already it has passed over the horizon and is no longer an "issue  
>that dominates psychological theorizing today."

I have a certain sympathy with both sides of this. If the editor really meant what he said, then Bill's comment is quite justified, along with the rest of it. But I prefer to read the comment in a more social sense (and I don't mean Psychology Club and Nostalgia column). As I see it, the requirement is to get people to read. To do that, one has to give them something they can perceive as relevant to their interests. I have no doubt that they SHOULD find PCT relevant to their interests, no matter which facet of psychology they work on. That doesn't mean that they know that they should. So the

editor's comment makes a great deal of sense. As far as he/she is concerned, to use up the valuable print space on something that the readers will simply pass over is a waste, no matter how valuable the article might be found to be ten years from now.

Seen in this light, I don't think the "trendy science" comment is justified. All science is trendy. Science is a sociological function, based on the belief structures held by scientists, and that includes beliefs about where important new things will happen. You can't fault them for not agreeing that your own position is such a place, just on your own say-so. You have to show them in their terms, not yours.

Also, editors are people, and control for the stability of their own views. I have had papers refused by editors who insisted on the use of significance statistics, which I abhor (doesn't that sound funny, given my insistence that statistics are/is very important). I refuse to publish a significance level, and if an editor won't accept that, I go elsewhere or keep the paper in a drawer. If I see a significance level in a published paper, my first thought is that the author doesn't know what the data say.

>This is confirmation of my thesis that a science based on low-probability  
>facts can't create a coherent picture of nature. ...  
> No argument involving more than a  
>small handful of facts can lead to deductions with a probability of truth  
>greater than chance.

An amusing self-contradiction! Also a take-off point for an argument. Low probability facts do not contribute to a logical argument, but they can pool to generate high-probability facts. Most perception works that way, I strongly believe. Perception is not deduction.

Martin

Date: Mon Jul 13, 1992 12:29 pm PST  
Subject: Trendy science

[From Rick Marken (920713.1300)]

Bill P. -- Thanks for the sympathetic comments on the "Blind men" paper.

I'm tempted to forward your "trendy science" remarks to Kintsch (the editor of Psych Review). But I'm pretty tired of arguing with these people. Unfortunately, I'm also running out of places to send my papers; places that might be willing to publish "non-trendy" science. Maybe I'll just start saving these papers up for "Mind readings II".

By the way, does anyone out there have any idea what Kintsch might mean by "current psychological issues and theorizing"? Maybe if I knew for sure what these were I could add a sentence or two to the "Blind men" paper relating control theory to these current issues and theorizings.

Hasta luego Ricky

Date: Mon Jul 13, 1992 2:32 pm PST  
Subject: Re: Trendy science

[From Rick Marken (920713.1500)]

Martin Taylor (920713 15:15) says

> As I see it, the requirement is  
>to get people to read. To do that, one has to give them something they can  
>perceive as relevant to their interests. I have no doubt that they SHOULD  
>find PCT relevant to their interests, no matter which facet of psychology  
>they work on.

Unfortunately, the relevance of PCT to their interests is quite negative. PCT shows that most psychologists are interested in an illusion of one kind or another -- the illusion of control by reinforcement, the illusion of external causation of behavior, the illusion of internal programming of behavior.

> So the  
>editor's comment makes a great deal of sense.

From a PCT perspective, I suppose so.

> As far as he/she is concerned,  
>to use up the valuable print space on something that the readers will simply  
>pass over is a waste, no matter how valuable the article might be found to be  
>ten years from now.

I agree that this appears to be the basis for the rejection -- but I find it appalling; I can't believe that the role of an editor of a scientific journal is to attract readers; National Enquirer, yes; Psychological Review, no! I would hope that the editor of a scientific journal would have the integrity to "waste" precious print space on non-trendy science if s/he thought the article actually made a valuable contribution. I get the impression that this is indeed the way some editors go about their business -- Estes for one, bless his fair (but unquestionably conventional s-r) heart.

>Seen in this light, I don't think the "trendy science" comment is justified.

In this light it seems just as trendy -- the editor is evaluating on the basis of what will "sell" in the current market, not on scientific merits.

>All science is trendy. Science is a sociological function, based on the  
>belief structures held by scientists, and that includes beliefs about where  
>important new things will happen. You can't fault them for not agreeing  
>that your own position is such a place, just on your own say-so. You have  
>to show them in their terms, not yours.

There is no place in the "blind men" paper where I ask the reader to accept, on the basis of my say-so, that mine is a new, important position. Instead, I present an analysis of a closed loop negative feedback system -- an analysis that the reader is free to question and test -- and show that aspects of the behavior of this system look like s-r, reinforcement or cognitive behavior (the latter being "their" terms for these types of behavior). So I think I have tried to show "them", in their own terms, that what they consid-

er behavioral phenomena are (possibly) different perspectives on closed loop perceptual control. I suggest that, if this is the case, then they can only understand behavior if they start trying to figure out what perceptual variables the systems is controlling. That's what the paper is about. It was short, sweet and to the point. It seemed to be "relevant to what psychologists care about" -- understanding behavior. So I don't understand what Kintsch could have meant by his justification for not having the paper reviewed. My guess is that he felt that the issue raised by my paper was not interesting to him (and possibly to most other psychologists who already KNOW what behavior is). That is a very poor basis for deciding what get's disseminated to the scientific community.

I don't really know what I could have done to make the "Blind men" paper more publishable, given Kintsch's criteria (and your interpretation of them). Do you think there is a way to re-write the paper so that it could meet these criteria? Any suggestions would be greatly appreciated.

Regards Rick

Date: Tue Jul 14, 1992 10:21 am PST  
Subject: orders of vs levels in HPC

[From: Bruce Nevin (Tue 920714 13:14:09)]

>To sum up: your argument seems to depend on the emergence of higher levels  
>from populations of systems of an existing level, without the addition of  
>any new kinds of physical systems. I would claim that you are relying on  
>IMPLICIT organization to create new levels. I, on the other hand, argue  
>that the new levels must be EXPLICIT.

I take it you mean orders of control systems (not to be confused with levels of control within one hierarchical control system).

Each new order is explicit from a higher-order perspective (the perspective of the new order or of a yet higher order), but not from the perspective of any orders out of which the new order is constituted.

For example, a reference signal is explicitly present for a cell as an electrical potential, ion concentration, whatever, but it is not a reference signal for the cell. The ion concentration is explicit for the cell. The reference signal is not explicit, and cannot be, because reference signals as such do not exist in the cell's universe. As the cells (by whatever evolutionary process) come to constitute control systems of a higher order, an ion concentration within a cell can take on a new identity as a reference signal (or error signal, etc.), in addition to its value as a variable within the cell. All of that is invisible to the cell, which is only controlling its own variables in its own terms.

I am not claiming to demonstrate the existence of social hierarchies. I am arguing for agnosticism regarding them.

I can flesh this out with more words. Do I need to? You say.

Bruce bn@bbn.com

Date: Tue Jul 14, 1992 10:33 am PST  
From: Dag Forssell / MCI ID: 474-2580

TO: list (Ems)  
EMS: INTERNET / MCI ID: 376-5414  
MBX: LISTSERV@VMD.CSO.UIUC.EDU  
Message-Id: 84920714183348/0004742580NA3EM

```
get csg-l log 9207a
set csg-l ack
review
info genintro
index csg-l
info database
```

Date: Tue Jul 14, 1992 10:37 am PST  
Subject: Re: Trendy science

[from Joel Judd 920714.1255]

Rick laments:

```
>Unfortunately, I'm also running out of places to send my papers;
>places that might be willing to publish "non-trendy" science.
```

I for one would enjoy publishing ANYPLACE. I am awaiting what will no doubt be interesting replies to an attempt to break into the SLA literature. Which brings me to an alternative: you could try to cross fields and co-publish with someone else on CSG-L, for example. Do you want to try out language acquisition? There are still some editors that will entertain pretty far-out ideas (for which label PCT seems to qualify). I'd be happy to entertain any possibilities that interest you.

[I would have sent this direct but I'm not using my own e-mail disk. You can reply just to me if you wish]

Date: Tue Jul 14, 1992 11:06 am PST  
From: Revised List Processor  
EMS: INTERNET / MCI ID: 376-5414  
MBX: LISTSERV@vmd.cso.uiuc.edu

TO: \* Dag Forssell / MCI ID: 474-2580  
Subject: Output of your job "0004742580"

```
> get csg-l log 9207a
File "CSG-L LOG" from filelist 9207A is unknown to LISTSERV.
```

```
> set csg-l ack
Your distribution options for list CSG-L have been successfully updated.
```

> review

Missing argument - specify at least 1.

All subsequent commands have been flushed.

Summary of resource utilization

```
-----  
CPU time:          0.282 sec          Device I/O:      34  
Overhead CPU:     0.036 sec          Paging I/O:      6  
CPU model:        3081              DASD model:     3380
```

Date: Tue Jul 14, 1992 12:36 pm PST

Subject: Re: Trendy science

[From Rick Marken (920714.1300)]

Joel Judd (920714.1255) says:

>I for one would enjoy publishing ANYPLACE. I am awaiting what will no doubt  
>be interesting replies to an attempt to break into the SLA literature.  
>Which brings me to an alternative: you could try to cross fields and  
>co-publish with someone else on CSG-L, for example.

You are right Joel -- I should be glad that my PCT stuff has been published at all. I would be happy to try to cross fields and/or co-publish (I've done that with Bill P.; I'd like to work with others on CSG-L too). I'm just kvetching because I'm tired of trying to get published by wriggling through hoops in order to make PCT palatable to editors and reviewers. When editors or reviewers catch factual errors then I'm happy to change stuff. But when they want changes that change the meaning of the paper (like the reviewer who suggested that I make it clear how the target guides behavior) it becomes tiresome. I think PCT should have reached the point by now where we can just put our work before the public and assume that that public either has been or can get educated about PCT. I don't want to have to write a "what PCT is about and why you should care" section every time I submit a PCT paper for publication. That's why I won't resubmit the "Behavior of perception" paper to journals that say "we'd accept it if you could just explain how this fits into conventional psychology". I don't think people who try to publish papers based on other, "trendy" theories of behavior are made to jump through such hoops. I'm just tired of having PCT treated differently than the "trendy" stuff; and I'm tired of jumping through the hoops.

Although I'm willing to cross fields in order to publish -- I also think that it is important to hit the right audience. I submitted to Psych Review, for example, because I thought the "blind men" paper was most relevant to an audience of theoretical psychologists. Obviously, my thoughts were not consistent with those of the editor.

Hasta luego Rick

Date: Tue Jul 14, 1992 2:58 pm PST

Subject: Re: Trendy science

You people seem to have lost sight of the fact that you can still be thought of as working within the broader systems science/cybernetics fields. I can point you to a number of journals and conferences that would probably be happy to have you, and can personally refer good papers to some editors.

Many of you are aware of Bill's and others' experiences with the ASC. But there are others as well.

It is true that the Systems literature is pretty much a ghetto, with a low signal/noise ratio and a relatively high crackpot ratio. But at least they're open to otherwise far-out views (not that I think CT is REALLY far-out). But if you're REALLY desperate, it IS available.

```
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
| ^^^^^^      !NOTE NEW EMAIL!      ^^^^^^
V All the world is biscuit shaped. . .
```

Date: Tue Jul 14, 1992 5:52 pm PST  
Subject: Hierarchies, explicit and implicit

[Allan Randall (920714.1900)]

I'm one of the contractors who does work for Martin Taylor. I'm still just becoming familiar with control theory, and the following is part of my attempt to mesh what I've learned so far on this group with my background in AI and neural nets.

Bill Powers (920713.1200) writes --

```
< In the PCT hierarchy as I have conceived it, new levels consist of new
< control systems; they aren't simply the same bottom-level systems looked at
< from farther and farther away. If a level of control is added, it is added
< EXPLICITLY, as input, comparison, and output functions physically distinct
< from the components of already-existing systems at lower levels.
```

I think I see with your basic point here, but that big capitalised "EXPLICIT" still bothers me somewhat, and I'd like to use it as a starting point to express some of my concerns with the hierarchical control notion. Could your argument be summarised as follows? There are two ways one can talk about different "levels":

- (1) Conceptual: perceived levels.
- (2) Physical (architectural): perceiving levels.

In (1) the "levels" are not actually in the control system under discussion, but are in the type (2) perceiving levels in the mind of the scientist building the model. The scientist is, hopefully, controlling for these perceptions to square with reality (or to get him grant money, whatever). Type (1) perceived levels are IMPLICIT in the control system under study, while type (2) are EXPLICIT.

Now I have two (possibly related) concerns with this breakdown. (Forgive

me if most of this is in AI/connectionist terms, rather than PCT, but I'm still in the process of relating the two. Feel free to translate/refute any of this stuff in terms of PCT).

(1) EMERGENT PROPERTIES AND DISTRIBUTION:

How do emergent properties (a la connectionism) fit into this scheme? 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. A "higher level" does not necessarily exist explicitly in the network. E.g.: the generalised concept of PERSON could be an \*implicit\* emergent property at the same \*explicit\* (i.e. architectural) level as the less general concepts of JOHN and MARY. The fact that it is at a higher level can only be determined by studying its dynamical/informational properties. Also, to get interesting distribution, as opposed to localised classification, you need a recurrent network.

Now if I extend this principle straightforwardly to networks of control, then the "level" at which an ECS is controlling is a dynamical/informational property of the net, and would not be explicit in the architecture. A "high-level" ECS could provide its perceptual output to a lower level in the "hierarchy," so the network would not be a strict hierarchy in its architecture. A node that might be called "low-level" in one context, might be controlling at a higher level in another context.

Two questions: (1) Is this coherent within the framework of control theory (PCT)? (2) Can such a system be said to be hierarchical, within the paradigm of HPCT?

(2) INFORMATION THEORY:

More generally, how valid is it to make this distinction between implicit and explicit hierarchy? It seems to me that the describability of a system in terms of an implicit hierarchy does not necessarily mean that the hierarchy is only in our heads and not a real property of the system. A random system, for instance, cannot be any more compactly described with a hierarchy than without. A more organised system, however, may be describable with fewer bits of information using a language with a notion of hierarchy. So isn't the hierarchy a real informational property of the system? Just where is the dividing line between an explicit and implicit hierarchy? Can't one always claim the hierarchy is not really a property of the system,

but rather of the language used to describe the system? Even then, can't I claim the hierarchy isn't a property of the language, but of the language used to describe the language? Etc, etc, etc...

I guess talk of explicit hierarchies just strikes me as wrong. I've always thought of "lower" levels as also controlling for things in the "higher" levels. At least this has been my notion of "control" before running across PCT. I do not know for sure how consistent this is with the HPCT paradigm. I have a feeling you will say that such a system would not be hierarchical?

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



NTT Systems, Inc.  
Toronto, ON

Date: Wed Jul 15, 1992 6:55 am PST  
From: Dag Forssell / MCI ID: 474-2580  
TO: list (Ems)  
MBX: LISTSERV@VMD.CSO.UIUC.EDU

get csg-1 log9207a  
review csg-1  
index csg-1  
info genintro

Date: Wed Jul 15, 1992 7:17 am PST  
From: Revised List Processor  
MBX: LISTSERV@vmd.cso.uiuc.edu  
Subject: Output of your job "0004742580"

> get csg-1 log9207a  
> review csg-1  
> index csg-1  
> info genintro

Summary of resource utilization

-----  
CPU time: 10.276 sec Device I/O: 500  
Overhead CPU: 3.124 sec Paging I/O: 10  
CPU model: 3081 DASD model: 3380

Date: Wed Jul 15, 1992 7:42 am PST  
From: Dag Forssell / MCI ID: 474-2580  
Subject: Hierarchies

[From Dag Forssell (920715-1)]

Allan Randall (920714.1900)]

>I'm one of the contractors who does work for Martin Taylor. I'm still  
>just becoming familiar with control theory,...

Welcome to the net. What does familiar mean? Why don't you provide a  
little more of a personal introduction, including some more specifics on  
what you have studied of PCT.

Specifically, have you read Behavior: the Control of Perceptions? In  
the chapter on A hierarchy of Control Systems you find on page 71 a  
representation which I believe answers many of your questions.

>I guess talk of explicit hierarchies just strikes me as wrong. I've  
>always thought of "lower" levels as also controlling for things in  
>the "higher" levels. At least this has been my notion of "control"  
>before running across PCT. I do not know for sure how consistent this  
>is with the HPCT paradigm. I have a feeling you will say that such a  
>system would not be hierarchical?

PCT is a model of how human beings control themselves. This is a specific suggestion, subject to a great many real world constraints. It lives up to the requirements of a hard science. It is not a general fuzzy "concept" with any imaginable connection allowed but a suggestion of a physical arrangement of analog circuits that can work in a real body.

It is quite unnecessary to attempt to explain HPCT in terms of analogies with hierarchy in language or implicit / explicit hierarchies. (All of which are questionable in themselves). You can study HPCT directly in the book.

I have no connection or much interest in AI, so I am reluctant to interpret many of the fancy words used in your post. I just have a vague feeling that much of the AI terminology relates to phenomena that can be created (computed) in the artificial world of the computer but have no bearing whatsoever on the structure of a human as represented by HPCT. Therefore, it is hard to relate to AI terms and to translate them into PCT terms. When this is attempted, as it has been on this net, it leads to interminable exchanges with iteration after iteration of misunderstanding. Please read Behavior: the Control of Perception to get a feel for the nature of the physical model of HPCT.

Just the same, much of conventional talk about human phenomena concerns stuff that has been defined by appearances without any understanding. People ask us all the time to define for them with PCT all kinds of non-existent phenomena which they have constructed out of almost nothing (Rick Marken identifies it as noise) in their imagination and cannot explain themselves, even to themselves.

Welcome aboard Dag

Date: Wed Jul 15, 1992 8:26 am PST

Subject: orders of vs levels in HPC

[From Rick Marken (920715.0900)]

Bruce Nevin (Tue 920714 13:14:09) says:

>As the  
>cells (by whatever evolutionary process) come to constitute control  
>systems of a higher order, an ion concentration within a cell can take  
>on a new identity as a reference signal (or error signal, etc.), in  
>addition to its value as a variable within the cell. All of that is  
>invisible to the cell, which is only controlling its own variables in  
>its own terms.

I think that what is required for the cell to participate as a component of a control system is what I will call "functional specificity". To avoid high-falutin' language, let me just say why I think neurons work as components of control systems and why people don't (even though neurons, like people, are probably control systems in and of themselves). A neuron, as a cell, is probably busy controlling many variables -- such as concentrations of K+ and Na- ions, etc. The systems controlling these variables are made

out of cell components (like RNA and DNA molecules) that may be control systems themselves. But one thing a neuron cell does is generate axon potentials (spikes) at a rate proportional to the integrated (over time and number of dendritic inputs) electrical charge at the cell body. This functional property of the neural cell (rate of spiking proportional to charge on cell body) is not "controlled" by the neuron itself; it is a cause-effect property of the cell's activity. That is, the variables involved in this functional relationship (electrical charge on the cell body, spike rate) are not (as far as I know) perceived and controlled by the cell. For example, the cell does not have a preferred (reference) spike rate that it tries to maintain; it just fires at a rate dependent on the charge at the cell body (up to the limit of saturation -- the cell just physically cannot produce spikes faster than a certain rate regardless of the input charge). It is this input-output characteristic of the cell's electrical behavior that makes it a useful component of a control system. The cell responds to input cell body charge with a certain rate of firing; this is a "dedicated" cause-effect characteristic of the cell; the cell cannot change the way it responds (firing rate) to input (cell body charge) -- there is no control involved in this functional relationship; that is what I mean by functional specificity. In terms of its electrical response to electrical stimulation of the dendrites the cell functions like a wire in a circuit (with firing rate the analog of current and cell body charge the analog of voltage). A control system must be built out of such "functionally specific" components.

People could also be components of control systems; but the only aspects of human behavior that could be a component of such a system are those that are "functionally specific" in the same way that the cell's electrical behavior is functionally specific. In other words, only cause-effect aspects of human behavior could be a reliable component of a control system. Controlled aspects of human behavior could do the job -- but this would not be reliable because the reference for the controlled variable could be changed in a way that bombs the function of the control system of which the person is a component. For example, suppose that a control system depended on people (as the components) producing a response that was proportional to some perceptual input. Consider this "social control system"; one person (the sensor) responds "yes" when a lion appears and "no" otherwise. A second person (the comparator) says "help" when he hears "lion" and nothing otherwise. A third person (the output) shoots a gun into the air when he hears "help", and does nothing otherwise. People can set their references so that they act this way -- but they are free to change those references at any time (free in the sense that nothing outside of the person controls the setting of the reference -- directly, anyway). Most obviously, any of the people in this "social control system" could decide to leave (heeding nature's call) thus changing, substantially their input/output function in the control system.

>I am not claiming to demonstrate the existence of social hierarchies.  
>I am arguing for agnosticism regarding them.

Agnosticism is what science is all about. I'm certainly willing to believe that social control exists if it is demonstrated to me -- otherwise, I go with the null hypothesis, which is based on my understanding of how control works and how people work as control systems. This understanding leads me to believe that the only way that people can be components of

a control system is in terms of cause-effect aspects of their behavior. And there is not much that you can just cause people to do without the use of extreme physical force -- and force is the typical way that people are made to function as components of a social control system. And even force doesn't succeed for long.

People can temporarily arrange themselves so that they function as a control system -- this is true and I've seen it happen. So in this sense, social control systems can exist; but these systems are quite transient. I don't think you would want your nervous system to work the way a social control system works.

Best regards Rick

Date: Wed Jul 15, 1992 10:54 am PST  
Subject: re: orders of vs levels in HPC

[From: Bruce Nevin (Wed 920715 13:05:23)]

(Rick Marken (920715.0900) ) --

You are very close to a point that I was trying to communicate.

You claim that the relationship between adjacent orders of control systems (such that CSs of order n are constituted of CSs of order n-1) evidently must include the following:

Variables used to implement control for order n must not be controlled variables for order n-1.

The level n variable is not controlled on level n-1, and it is a cause-effect property of the cell's behavioral outputs on level n-1. As you say, a neuron cell

> generate[s] axon potentials (spikes) at a rate proportional to the  
> integrated (over time and number of dendritic inputs) electrical charge  
> at the cell body. This functional property of the neural cell ( rate of  
> spiking proportional to charge on cell body) is not "controlled" by the  
> neuron itself; it is a cause-effect property of the cell's activity.

The relationship of ECS function to intracellular control, you claim, is an incidental byproduct. It is only the physics of the cell's body (level n-1) that the ECS (level n) uses, cellular metabolism being only used to maintain the viability of a segment of "wire" in place.

But it is a byproduct of behavioral outputs ("the cell's activity") which presumably are variable means for achieving uniform results that matter to the cell. A cell in a given state controls a disturbance of a given sort reliably with behavioral outputs of a corresponding given sort. These behavioral outputs (including internal changes) can have cause-and-effect consequences remote from the controlled variable and the disturbance, which are not themselves controlled. Thus, it is possible that these changes in the cell's observed behavioral outputs are cause-effect byproducts of controlling other internal variables

(such as Na concentrations) against disturbance. Exploring this possibility might lead to some explanations of how learning and reorganization work.

I am proposing (920709 09:13:52) that reorganization is carried out in populations of entities of order n-1. If control of order n results in chronic error in CSs of order n-1, then the CSs of order n-1 control to correct the error, with behavioral outputs which for a nerve cell might include growing new axons, detaching or moving or withdrawing (atrophying?) existing ones, changing receptor sites around, etc. Local error in a few order n-1 CSs results in learning. Error in many CSs results in reorganization.

> A neuron, as a cell, is probably busy controlling many variables -- such  
> as concentrations of K+ and Na- ions, etc. The systems controlling these  
> variables are made out of cell components (like RNA and DNA molecules)  
> that may be control systems themselves. But one thing a neuron cell does  
> That is, the variables involved in this functional relationship . . .  
> are not (as far as I know) perceived and controlled by the cell. For  
> example, the cell does not have a preferred (reference) spike rate that  
> it tries to maintain; it just fires at a rate dependent on the charge at  
> the cell body . . . . It is this input-output characteristic of the  
> cell's electrical behavior that makes it a useful component of a control  
> system. The cell responds to input cell body charge with a certain rate  
> of firing; this is a "dedicated" cause-effect characteristic of the  
> cell; the cell cannot change the way it responds (firing rate) to input  
> (cell body charge) -- there is no control involved in this functional  
> relationship; that is what I mean by functional specificity. In terms of  
> it's electrical response to electrical stimulation of the dendrites the  
> cell functions like a wire in a circuit ( with firing rate the analog of  
> current and cell body charge the analog of voltage). A control system  
> must be built out of such "functionally specific" components.

To this I would add that the cell does not \*want\* to change the way it responds. The firing rate per se does not matter to the cell. Indeed, it probably does not even "know" that it is changing its electrical potentials, that these changes constitute "spikes," and that they are occurring at a variable rate. All of that is invisible to the cell, \*and\* \*must\* \*remain\* \*so\* for the higher-order function to maintain its integrity, as you also point out.

Substitute humans for cells:

A person is busy controlling many variables. The systems controlling these variables are made out of cells, neural structures, and organs of perception and execution (probably there's a better word, but I'm in a rush). But one thing a human does is change the color of its aura. That is, the variables involved in this functional relationship . . . are not (as far as I know) perceived and controlled by the human. It is this input-output characteristic of the human's auric behavior that makes it a useful component of a control system. The human responds to changes in color of a neighboring human's aura by changes in its own aura. This is a "dedicated" cause-effect characteristic of the human; the human cannot change the way it responds (aura color) to input (aura color) -- there is no control involved in this functional relationship;

that is what I mean by functional specificity. In terms of it's auric response to auric stimulation of the etheric body the human functions like a wire in a circuit ( with auric change rate the analog of current and color the analog of voltage). A control system must be built out of such "functionally specific" components.

Exploring this possibility might also suggest ways of explaining input functions and output functions. Must these be separate multi-cellular structures, or might the metabolism of a single ramified nerve cell be such that it is not a simple cause-effect "wire" passing neural current through, but is actually doing the weighting (and the changes of weighting) of signals? In the case of an input function, the weighting and changes in weighting of input signals that get combined in the unified output signal; in the case of an output function, the weighting and changes in weighting that get applied to the different copies of the input signal in the process of making them into specific output signals. In this case, the cell is controlling variables that matter to it, and as a byproduct differentially weighting electrical potential, which does not matter to it, in its several branches. The electrical potential in the dendrite of another cell, on the other side of a synapse, does matter to it and may disturb variables that it controls; the rate of peaks "firing" in itself does not matter to it, and is a byproduct of that control.

>People can temporarily arrange themselves so that they function as a  
>control system -- this is true and I've seen it happen. So in this sense,  
>social control systems can exist; but these systems are quite transient.

Such things are not control systems of an order above that of humans, for the simple reason that humans (level n-1) control for creating and maintaining them (level n). If there are such supra-human organisms, the variables that matter for them are incidental for us. It may be that we have a craving to belong to groups and get off on group participation that is working well (and are frustrated when it is not) because some aspects of a transhuman organism's function require human intercommunication as a vehicle. But human intercommunication itself can be of no import to it, nor can its functioning have any import to us. Except when it is experiencing error and conflict (chronic error). That's when we reorganize. Maybe that's when we try to create or change social institutions.

Bruce bn@bbn.com

Date: Wed Jul 15, 1992 11:07 am PST  
Subject: Hierarchies

[From Rick Marken (920715.1030)]

Allan Randall (920714.1900) says:

> I'm still  
>just becoming familiar with control theory,

I agree with Dag's suggestion that you read Powers' "Behavior: The control of

perception" for starters. I would also suggest that you become familiar with the phenomenon of purposive behavior (control) while you learn about the theory. I recommend Powers' Demo1/2 program and (of course) my "Mind readings" book.

>and the following is part  
>of my attempt to mesh what I've learned so far on this group with  
>my background in AI and neural nets.

I'm sure that there is much of technical value in AI and neural nets. But I'd put it aside for now and not try to mesh it with PCT until you understand what control theory is about -- or you'll get a mish mash. AI and neural nets are just a part (and a small one at that) of what PCT is about.

> How do emergent properties (a la connectionism) fit into this scheme?

Purposive behavior is an emergent property of the organization of the control model.

>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. I suggest that you take a look at my spreadsheet hierarchy (Martin has it) to get a feel for the "distributedness" of the control hierarchy.

> Now if I extend this principle straightforwardly to networks of control,  
>then the "level" at which an ECS is controlling is a dynamical/informational  
>property of the net, and would not be explicit in the architecture.

Beg pardon??

> Two questions: (1) Is this coherent within the framework of control  
>theory (PCT)?

No.

>(2) Can such a system be said to be hierarchical, within the  
>paradigm of HPCT?

Yes, you can say whatever you like; in PCT it's how the model WORKS that counts, not what you say about it.

> More generally, how valid is it to make this distinction between  
>implicit and explicit hierarchy? It seems to me that the describability

>of a system in terms of an implicit hierarchy does not necessarily mean that  
>the hierarchy is only in our heads and not a real property of the system.

A model is built out of real (explicit) components in the hope that, when it is fired up, it will exhibit the same behavior (in the same environment) as the system being modelled. A control model contains (explicitly) a sensor, comparator and output component. When these are hooked up properly they act to keep the output of the sensor at a reference level -- this purposive behavior is implicit in the connections in the control system (or it's an emergent property of the system or whatever you want to call it). The hierarchical arrangement of control systems is built as it is because 1) it produces behavior like that of real people (see my JEP article) and 2) it works. Other arrangements might be better but such changes should be motivated by discrepancies between the performance of the model and that of the subjects (at least in my opinion) otherwise we are doing scholasticism rather than science.

> I guess talk of explicit hierarchies just strikes me as wrong.

That's where the science part comes in. It should strike you as wrong if the behavior of the hierarchical model is a poor match to the behavior of living control systems -- not just because it sounds wrong. It strikes me as wrong that c (the speed of light) will be measured as the same value in any inertial frame of the observer -- but it works that way, apparently. I highly recommend that you develop experimental tests of the hierarchical nature of the behavior of real control systems -- then build your models to match this behavior. It may be that you will find that an explicit hierarchical arrangement is not necessary to produce the behavior that you observe. That would be a great discovery and a worthwhile development in modeling control. And it would cut a lot of useless gab.

Best regards Rick

Date: Wed Jul 15, 1992 11:53 am PST  
Subject: re: orders of vs levels in HPC

[Martin Taylor 920715 15:00]  
(Rick Marken 920715.0900, Bruce Nevin 920715 13:05:23)

I posted 920710 14:15 a pointer to an article that to me suggests that the whole function of an ECS (possibly more than one) can be performed in a single neuron, plus the control of gain that Bill Powers thinks of as an unnecessary complication. The reference, once again, is:

"Evidence for a computational distinction between proximal and distal neuronal inhibition," E.T.Vu and F.B.Krasne, Science, March 27 1992, 255, 1710-1712.

Maybe you don't see in this article as much as I do, but as I said in my earlier posting, I think all of the elements are there. If this is true, then Bruce's analogies to human-based control systems take on more force.

Rick, I don't think humans can control other humans by force. All they can do is to alter the range of control available to the other humans. But humans CAN fairly reliably get other humans to perform actions that affect the world



in such a way that their (the masters') percepts come closer to their references. Usually, all one has to do is to ask, but there are more subtle ways. In a normal PCT-world mirror diagram, one would say that an ECS in the master is controlling a perception that the slave is doing something, and that ECS in the master is being provided a reference by another ECS that is controlling a perception of something affected by the slave. If for master and slave one reads partner A and partner B, one still comes up with control.

The master does not control the slave, no; I do not control my keyboard. But both slave and keyboard are aspects of the world that provide me with percepts I can control. That my actions differ in achieving the same result on different occasions is fully in accord with normal PCT, and so is the fact that they sometimes fail because of external world disturbances or barriers that go beyond my range of control. It is clear that no other human can be part of MY hierarchy, but the perception of what another human is doing can be the signal in an ECS that IS in my hierarchy.

In other words, I think that the whole discussion of social control is and has been following a kipper dragged across the trail.

Martin

Date: Wed Jul 15, 1992 12:03 pm PST  
Subject: re: orders of vs levels in HPC

[From Rick Marken (920715.1300)]

Bruce Nevin (Wed 920715 13:05:23) says:

>Substitute humans for cells:

>A person is busy controlling many variables. The systems controlling  
>these variables are made out of cells, neural structures, and organs of  
>perception and execution (probably there's a better word, but I'm in a  
>rush). But one thing a human does is change the color of its aura.  
>That is, the variables involved in this functional relationship . . .  
>are not (as far as I know) perceived and controlled by the human. It is  
>this input-output characteristic of the human's auric behavior that  
>makes it a useful component of a control system. The human responds to  
>changes in color of a neighboring human's aura by changes in its own  
>aura. This is a "dedicated" cause-effect characteristic of the human;  
>the human cannot change the way it responds (aura color) to input (aura  
>color) -- there is no control involved in this functional relationship;  
>that is what I mean by functional specificity. In terms of it's auric  
>response to auric stimulation of the etheric body the human functions  
>like a wire in a circuit ( with auric change rate the analog of current  
>and color the analog of voltage). A control system must be built out of  
>such "functionally specific" components.

I agree with this analogy. If, while controlling, people give off an aura output that is lawfully related to an aura input (where aura is some unperceived variable and where the lawful function relating auric input to output in some way depends on the structural and functional characteristics of the human control system) then this auric function could be used as part

of a meta control system, with "people" as its components. But this control system would be pretty much invisible to us.

Actually, if the neuron analogy is carried through, then aura is equivalent to cell potential -- and this should be perceivable by the system itself; people should be able to perceive (possibly with instruments) the aura just as the neuron (if it had any brains) could perceive its cell potential. The neuron could even measure the cell body charge to spike rate transfer function; and I think people would be able to measure an auric transfer function if it existed (indeed, this auric function might be a relationship between any ol' physical variables, as long as it was a cause effect relationship that depended on the presence of a person (as the cell charge/spike relationship depends on the presence of the neuron. The neuron might even be able to figure out that it was a part of a large control system (the nervous system) just as people might be able to figure out (or, at least fantasize) that they are components (or pawns) in a higher order control structure. But I don't see any way that the neuron could figure out what the control system of which it is a part was controlling; the variables we control would be "cognitively impenetrable" to the individual cells that make up the nervous system. A "perceptual" neuron would never be able to tell that its firing rate was proportional to the degree of "squareness", say, of an image on the retina ("what the hell is a retina?", says the neuron). So if people are part of a higher order control system (which uses uncontrolled "auric" relationships are components of control, we MIGHT be able to figure out that this is the case (A BIG MAYBE) but even if we do I don't think we could figure out what the hell that control system was about; what would it matter, really, anyway. Why would a neuron care that it's part of a control system that controls what its owner (us) calls "the real world". All it cares about is whatever variables a neuron cares about.

So I think the kind of "meta control" that Bruce is talking about here is a bit deeper (and more impenetrable) than "social control". Meta control would involve control of variables that we (humans) will never know about-- placing them outside the realm of science. Social control, to me, implies a system that controls social variables -- like relationships between people, programs of action involving several people, etc --and these variables are easy to perceive. It is also easy to test to see whether these variables are controlled by anything other than the people involved (in which case there would be evidence of some kind of superordinate "social" controller). I think the evidence suggests that there is no such social controller; consideration of the existence of a meta controller I leave to the Pope (when he feels better).

Date: Wed Jul 15, 1992 12:58 pm PST  
Subject: Re: Hierarchies

[Martin Taylor 920715 15:45] Taking a breather from writing, to write.  
(Rick Marken to Allan Randall 920715.1030)

I can't speak for Allan, but knowing a little bit of the background, I suspect Rick has a little misapprehension of what Allan is getting at. Or of what I would have been getting at had I used Allan's words.

Allan and I had been working with recurrent neural networks before we got interested in PCT, and had found that quite simple networks could exhibit varieties of behaviour dependent on their history, without changing any of their structure or weight patterns. The self-same network might be stable, might show a simple short-period oscillation or more than one such, and might go into long-period or chaotic behaviour. Now it is a tenet of HPCT that there exists the possibility that within-level connections among ECSs might exist, but that we won't worry about them yet because things get complex. I think that Allan, knowing this, nevertheless wanted to consider how one would look at a net with recurrent connections.

>> I'm still  
 >>just becoming familiar with control theory,  
 >  
 >I agree with Dag's suggestion that you read Powers' "Behavior: The control of  
 >perception" for starters. I would also suggest that you become familiar with  
 >the phenomenon of purposive behavior (control) while you learn about the  
 >theory. I recommend Powers' Demol/2 program and (of course) my "Mind readings"  
 >book.

I know he has read BCP, or at least much of it, and has played with the Powers Demo programs. We have also had many discussions of PCT, so I think that his comment about "just becoming familiar" is one we could all justly accept. I know Bill Powers often makes such comments about his own understanding. Allan is not the kind of novice who we often see introducing themselves to CSG-L.

>> How do emergent properties (a la connectionism) fit into this scheme?

>Purposive behavior is an emergent property of the organization of the control  
 >model.

I have to disagree here. Purposive behaviour is built into the elementary control system that is the basis of the structure. There are different possibilities for emergence (a la connectionism). I think coordination may be one (possible forthcoming thread, but not now, please). Purpose is not.

>>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.

>> Now if I extend this principle straightforwardly to networks of control, >>then the "level" at which an ECS is controlling is a dynamical/informational >>property of the net, and would not be explicit in the architecture.

>

>Beg pardon??

>

>> Two questions: (1) Is this coherent within the framework of control >>theory (PCT)?

>

>No.

Well, the wording may not be coherent, given that Rick said "Beg pardon," but Rick's "No" refers to that rather than to the question Allan asked (as I interpret it). I think that the answer is "that's a difficult question, and we can't answer it in our present state of understanding. Some day, we will have to deal with it.

>Yes, you can say whatever you like; in PCT it's how the model WORKS that >counts, not what you say about it.

Right. But at the level we are talking about, there aren't any working models. Any form of HPCT that goes beyond three or so levels is pretty much a talk show. That doesn't stop us talking about (and in some cases using the results of talking) psychotherapeutic uses of HPCT.

>> I guess talk of explicit hierarchies just strikes me as wrong.

>That's where the science part comes in. It should strike you as wrong >if the behavior of the hierarchical model is a poor match to the behavior >of living control systems -- not just because it sounds wrong.

Again, a comment correct in the abstract. If, and only if, it is found that there are no recurrent perceptual or reference connections in a real-life control system, Allan's intuition would have failed. We are a long way from being able to assert that it has failed or that it has not.

> It may be that you will find that an explicit >hierachical arrangement is not necessary to produce the behavior that you >observe.

Again putting myself into Allan's words, I would answer Rick by saying that it is accepted that a hierarchic arrangement is necessary. But is it sufficient? We have had this discussion about configurations of configurations. Bruce Nevin likes them. Bill Powers does not. I think that experiments on reading processes lead to the impression that configurations of configurations are used in parallel with the grand configurations that Bill likes: a chair is a chair at the same time that the chair legs are chair legs, etc., according

to Bill. With configurations of configurations, a chair is a chair at the same time that it is composed of chair legs that are chair legs etc. The difference is subtle, but if one accepts that configurations of configurations are necessary, then Allan's intuition holds true.

Maybe I have misread where Allan is coming from. But I think Rick and Dag have, too.

Martin

Date: Thu Jul 16, 1992 6:07 am PST  
Subject: between orders of control

[From: Bruce Nevin (Thu 920716 08:19:31)]

(Rick Marken (920715.1300) ) --

Hooray! At last one of my analogies didn't sound too far fetched!

>Actually, if the neuron analogy is carried through, then aura is equivalent  
>to cell potential -- and this should be perceivable by the system itself;  
>people should be able to perceive (possibly with instruments) the aura just  
>as the neuron (if it had any brains) could perceive its cell potential. The

We don't have to stick with auras, of course, there's lots that passes us right by. Odors, pheromones produced and perceived by other species and those produced and perceived by humans, modulation of electromagnetic fields, demonstrated and postulated particles, etc.

But In fact, many people do perceive auras (sight is not the only modality). And my wife, for example, teaches people how who initially don't. There is good intersubjective agreement among them as to what they perceive. And there has been some instrumental research (elsewhere, not with Sarah), though if you think funding for PCT is precarious you should look at the history of parapsychology (interesting summary in Margins of Reality, a recent book concerning issues of consciousness).

But most of the time most people don't perceive auras consciously, though my informal and subjective impression is that people integrate subliminal glimmerings into their perceptions of mood, "vibes," etc. Highly subjective stuff. This would fit, because there could be selective "pressure" for subliminal perception (early warning of error, for example--hunches, intuitions about social situations) even in the presence of much stronger selective "pressure" against any perception. (If the neuron started perceiving spikes or rather spike rates as such and there was a correlation between intracellular error and spike rate, it might reorganize to control those perceptions, and that would create error at the higher order, which would adversely affect the would-be controller through its environment--what goes around comes around? Analogies to cancer, the immune system, AIDS, epidemiology in general, all invite speculation. It could get pretty woolly fast, but there might be some tufts of solid ground in the swamp.)

A different point: if order  $n$  are neurological CSs, then order  $n+1$  organisms need not be constrained to humans, or even to being constituted of living control systems all of the same species.

As Martin says, most of our talk about social control systems has to reorganize itself to accommodate the \*necessary\* disparities between control systems of one order (e.g. cells) and control systems of another (ECSS and neurological control systems constituted of them). A final point (reiterated): perhaps the postulated supra-animal CSs can provide some explanation for why people perennially try so hard to (re)create and make to work various social institutions that (often crudely) mimic living control systems. And it might perhaps be fruitful to reconsider histories of social change from the point of view of cells carrying out higher-order reorganization. (Kurt Vonnegut's wry imaginings of unexpected interdependencies spring to mind, too, of course--\_Sirens of Titan\_ perhaps.)

>So I think the kind of "meta control" that Bruce is talking about here is >a bit deeper (and more impenetrable) than "social control". Meta control >would involve control of variables that we (humans) will never know about-- >placing them outside the realm of science. Social control, to me, implies >a system that controls social variables -- like relationships between >people, programs of action involving several people, etc --and these >variables are easy to perceive.

First, I have proposed that there is more connection than that between two orders, by way of learning and reorganization (and maturation and evolution).

Secondly, I think we humans have more resources to bring to bear for a scientific investigation of a higher order of control than cells have. (I don't know that, of course, but from a human perspective it sure seems that way . . . the impenetrability works both ways, we don't really have a solid grasp of even all the external variables (those of which a cell is controlling its perceptions), and much less understanding of the internal variables in the cellular CS and of the structure of the cellular CS. Is it hierarchical? Hell, can we even describe the non-neurological CS of a plant? That would be a useful exercise--such "vegetative" perceptual control mechanisms probably exist in animals alongside the neurological mechanisms.)

Such an investigation would have to proceed with care and tact. It has been argued that much of our current troubles arise from increased human control of variables whose "transmission" through the ecosphere (partly) by way of humans was hitherto unconscious. Bateson on purposive action, for instance. But it could well be that the processes of learning, or maturation, or evolution (whichever analogy turns out to be appropriate) for a higher-order organism are carried out in part by organisms of the animal order taking on control of variables that previously had been available to be functional elements in the higher-order control system itself. Note that such functional elements are normally not diddled with by either the higher or lower order. (We don't normally intervene and change the spike rate of a neuron instrumentally, though we have the equipment now to do so I think, crudely.)

>A "perceptual" neuron would never  
>be able to tell that its firing rate was proportional to the degree of  
>"squareness", say, of an image on the retina ("what the hell is a retina?",  
>says the neuron).

The observer's perception of squareness of the observer's perception of an image on the observers perception of a retina, where "of" and "on" are the observer's perceptions of relationships. Somehow I think these things matter only to a psychologist, probably a PCT psychologist at that. They don't matter to the a normal human any more than they do to the neuron. (Just twitting you about problems of perspective--don't take it seriously.)

Bruce  
bn@bbn.com

PS--Sarah is off on a lecture tour Friday, her first, for two weeks, going to Chicago, Kansas City, Boulder, someplace in Texas I think, San Francisco, Seattle, with former clients or students organizing things in each place. I hope she breaks even, otherwise the mortgage company will be enquiring again. Emmanuel has for some time been ready whenever she is, and says not to worry. By all conventional ways of reckoning we should have gone into personal bankruptcy a dozen times in the past five years, and he has always said not to worry. You'd think I'd eventually catch on to the habit of not worrying.

Date: Thu Jul 16, 1992 6:41 am PST  
Subject: Gopher, anyone?

In the interest of making the CSG files (and others) more easily available, I've installed a Gopher server on biome. If you have a gopher client running on your system, you only need to gopher to biome.bio.ns.ca, port 70, to access it. If you don't have a gopher client on your system, ignore this message.

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.ns.ca

Date: Thu Jul 16, 1992 8:02 am PST  
Subject: Emergence, parapsychology

[From Rick Marken (920716.0830)]

I was going to let this one go; but my curiosity has gotten the best of me.

Allen Randall asked:

> How do emergent properties (a la connectionism) fit into this scheme?

And I replied:

>Purposive behavior is an emergent property of the organization of the control  
>model.

And Martin Taylor (920715 15:45) said:

>I have to disagree here. Purposive behaviour is built into the elementary  
>control system that is the basis of the structure. There are different  
>possibilities for emergence (a la connectionism). I think coordination may  
>be one (possible forthcoming thread, but not now, please). Purpose is not.

I guess my understanding of "emergence" differs from Martin's. I thought "emergence" referred to a behavior or other characteristic of a system that is qualitatively different than the behavior or characteristics of its constituents. One example of "emergence" that I learned was the salt crystal that is qualitatively different than its constituents, Na and Cl. I think that this same kind of "emergence" happens with a control loop. The constituents of this loop are s-r (cause effect) components. This includes the neurons, their synaptic connections and the physical entities that are in the environment. The equations that describe the two basic components of the control loop are s-r equations:

- (1)  $o = f(p^*-p)$  and
- (2)  $p = g(o+d)$

(1) describes the behavior of the neural component that produces variations in output as a function of input. (2) describes the behavior of the environmental component that produces variations in input as a joint function of outputs and independent environmental events.

There is no purpose in either of these components; they behave according to good old fashioned cause effect laws. Purpose emerges (to the considerable surprise and satisfaction of those who first built these systems) when these components are hooked up properly -- ie. so that there is negative feedback and dynamic stability.

So I have to disagree with Martin's disagreement; I think the negative feedback control system is a perfect example of emergence (as I understand the term) -- where a phenomenon (purpose) emerges from a system of components that, themselves, do not exhibit the phenomenon.

Bruce Nevin (Thu 920716 08:19:31) says:

>though if you think funding for PCT is  
>precarious you should look at the history of parapsychology (interesting  
>summary in Margins of Reality, a recent book concerning issues of  
>consciousness).

I think this comparison would be more interesting if parapsychology had data that were anywhere near the quality of the data we obtain regularly in studies of control. To my knowledge, the parapsychological results are even noisier than the results achieved in standard psychological experiments. Since I consider the latter to be of no value, it is difficult for me to characterize value of the former (negative, perhaps). The only thing that is surprising to me is that funding agencies are willing to pay for standard psychological research but not for parapsychological



research, which is of only imperceptibly poorer quality. I believe it must be a religious thing -- the religion of standard psychology being the currently dominant form. When psychology becomes a science (ie. when it stops turning the statistical crank and just goes for the quality data -- something that can only happen when it starts to focus its research efforts on trying to understand how organisms control) such religious territorial squabbles should diminish considerably.

Best regards Rick

Date: Thu Jul 16, 1992 10:38 am PST  
Subject: Gopher

It turns out that you don't need gopher software to access the CSG files on biome. If you have telnet access, then just telnet here and login as gopher. In other words.

telnet biome.bio.ns.ca  
Login: gopher

The rest should be self-explanatory.

Bill

--

Bill Silvert at the Bedford Institute of Oceanography  
P. O. Box 1006, Dartmouth, Nova Scotia, CANADA B2Y 4A2  
InterNet Address: bill@biome.bio.ns.ca

Date: Thu Jul 16, 1992 11:33 am PST  
Subject: Gopher -- PS

>It turns out that you don't need gopher software to access the CSG files  
>on biome. If you have telnet access, then just telnet here and login as  
>gopher. In other words.

>  
>telnet biome.bio.ns.ca  
>Login: gopher

>  
>The rest should be self-explanatory.

It turns out that on some systems (like VAX/VMS) the telnet program doesn't automatically pass on the terminal type. You may need to use a command like:

telnet -t vt100 biome.bio.ns.ca

in this case, since gopher is a full-screen program and needs to know your terminal type.

Bill

Date: Thu Jul 16, 1992 2:36 pm PST

Subject: Re: Emergence, parapsychology

[Martin Taylor 920716 18:00]  
(Rick Marken 920716.0830)

We have a pure misunderstanding of words here. Easily fixed, I think.

>>Purposive behavior is an emergent property of the organization of the control  
>>model.

>

>And Martin Taylor (920715 15:45) said:

>

>>I have to disagree here. Purposive behaviour is built into the elementary  
>>control system that is the basis of the structure. There are different  
>>possibilities for emergence (a la connectionism). I think coordination may  
>>be one (possible forthcoming thread, but not now, please). Purpose is not.

>

>I guess my understanding of "emergence" differs from Martin's. I thought  
>"emergence" referred to a behavior or other characteristic of a system that  
>is qualitatively different than the behavior or characteristics of its  
>constituents. ... Purpose emerges (to the considerable surprise and  
>satisfaction of those who first built these systems) when these components  
>are hooked up properly -- ie. so that there is negative feedback and dynamic  
>stability.

We agree about what "emergence" means. We had a misunderstanding about what was the unit under discussion. I had taken Allan's question about organization to be asking what emerged from the organization that is not inherent in negative feedback control systems. I think my comment that you quoted made that assumption explicit. You took the question to refer to the constituent components of control systems, and the emergent to be control. I agree that if we are talking about the s-r components of a feedback loop, then of course purpose is an emergent from the organized structure.

I assumed that the unit was a complete ECS, and the question was what emerged from their hierarchic organization. I proposed coordination, which seems to me to be a collective function different from collaboration. I don't know how strongly I would support that proposal if pressed.

Martin

Date: Thu Jul 16, 1992 3:00 pm PST  
Subject: Re: Emergence, coordination

[From Rick Marken (920716.1600)]

Martin Taylor (920716 18:00) says:

>We have a pure misunderstanding of words here. Easily fixed, I think.

Yes, indeed.

>I assumed that the unit was a complete ECS, and the question was what emerged  
>from their hierarchic organization. I proposed coordination, which seems to

>me to be a collective function different from collaboration. I don't know  
>how strongly I would support that proposal if pressed.

Thank you. Now I understand. I think coordination can, indeed, be an emergent property of a hierarchy of Elementary Control Systems (ECSs) and if you don't want to, I will (and have) strongly support that proposal if pressed (and so will Tom Bourbon if he'd get his butt back on the net). In other words, I vigorously agree with you -- maybe even more than you agree with you.

Best regards Rick

Date: Fri Jul 17, 1992 11:41 am PST  
Subject: The Relevance of Blind Men

[from Gary Cziko 920717.1000]

I'm back after a wonderful two weeks in France and Switzerland. I ran into some very strange reference levels in France, though. The truckers were controlling for stopping traffic on the highways to protest new driving regulations. It took us over 12 hours to drive from Paris to Geneva (usually a 5 1/2 trip), and we were lucky to have gotten through at all.

Rick Marken (920712.1200)

>I found out why Psych Review didn't even send my "Blind men"  
>paper out for review. According to the editor it was because:  
>  
>"It would need to speak more directly to current psychological  
>issues and theorizing. One would need to see more clearly a connection  
>between what you are talking about and the issues that dominate  
>psychological theorizing today."

Would it really be all that hard to show more explicitly how your paper is relevant to current psychological theorizing? Why not take some "trendy" perspectives in cognitive psychology today and discuss how it fits one part of the elephant. And there are lots of stuff still being published in the behavioral journals to show the other parts of the pachyderm. You could also refer to the growing body of experimental results which support PCT. I think you have the framework for a dynamite paper. PCTers can see how it is relevant without this added stuff, but the mainstream psychological community cannot. You have to give them a bit more help. I say, give Kintsch another shot before giving up on Psychological Review.--Gary

Date: Fri Jul 17, 1992 12:24 pm PST  
Subject: Interest in PCT

from Ed Ford (920717:12:50)

Rick and all -

My sympathies concerning your Blind Men paper. I guess they perceive that what you have to offer won't improve the quality of what they produce. I've heard

this frustration from many other CSGers over the years. Recently, I've been more fortunate.

I've been dealing with the former superintendent of Johnson City Schools in N.Y. who created and still teaches The Outcomes-Driven Developmental Model (known as ODDM) which is presently in nearly 200 school districts in the U.S. It really is a dynamite program. Six months ago, he asked me to develop (and I did) a program of instruction for teachers and/or parents who in turn would teach interested school district parents the necessary techniques for raising children as a tie-in to his model. I called it the Outcome Based Parenting Program. I've named my program Teaching Responsibility To Youth.

Lately, I have been helping him in several areas where he saw needed improvement within certain aspects of his model. With my using PCT as the theoretical basis and as we planned and thought things through in the development, delivery, and the testing of new ideas for his model, some really great new vehicles came to life. He has watched me thinking and suggesting, using PCT and thinking PCT, and obviously, his interest in PCT has grown immensely. In short, he saw the wisdom of using PCT as the basis for his model.

For example, he has always been looking for a way to implement what Deming called statistical process control into his model. I showed him how he could set up individual charts that everyone at their own level could use, from the school board down to and including the student. I showed him how the reference signal and the controlled variable both must be measurable and specific for any human control system to work effectively and achieve what it wants. The more immediate the feedback, the better. The more delayed the feedback, the less efficient the system will operate. The lower goals (students, teachers, principals, superintendents) had to reflect and tie-in to the highest goal (school board) for the system to work in a coordinated way.

To all -

Closed Loop arrived yesterday from Greg and I took it to the printers today. As I said earlier, I plan to hand it out at the conference. The balance will be mailed Monday, Aug. 3rd.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU  
10209 N. 56th St., Scottsdale, Arizona 85253

Ph.602 991-4860

Date: Fri Jul 17, 1992 2:12 pm PST  
Subject: The Relevance of Blind Men, Model-based research

[From Rick Marken (920717.1300)]

Gary Cziko (920717.1000)

Welcome back!! And thanks for the excellent advice, viz.;

> Why not take some "trendy"  
>perspectives in cognitive psychology today and discuss how it fits one part  
>of the elephant. And there are lots of stuff still being published in the  
>behavioral journals to show the other parts of the pachyderm. You could  
>also refer to the growing body of experimental results which support PCT.

> I say, give  
 > Kintsch another shot before giving up on Psychological Review.--Gary

Shall do. I guess it's off to the library this weekend. Luckily, I will only need to look at what's been published in the last couple of months. In fact, that will be my rule of citation for the new version of the "Bind men" paper -- I will only cite papers published since June 1992. I hope to have a new version ready for the meeting in a couple weeks. Thanks for the encouragement.

-----  
 Here is another thought I had on the topic of psychological research. I have been arguing that most of the results of psychological research based on the current paradigm is of little (or no) value because the data is of such low quality; it is basically noise. I think that this low quality data is acceptable to many psychologists because it is actually consistent with their model of behavior. The s-r model that is the basis of most psychological research is not just a cause - effect model; it is also a stochastic model. Effects are expected to be only probabilistically related to causes. The model that this research is based on is the "general linear model" which says

$$y = a.1x.1 + a.2x.2 + \dots + a.nx.n + e$$

(ignoring the interactions and other higher order terms). The important term is e which is the error or noise term. So y (the subject's response) is EXPECTED to be a function of the input variables (x.1 .. x.n) as well as a noise component. The e component can be thought of as unexplained variance that could be accounted for if the right variables were discovered to account for it). But, in fact, e is rarely accounted for -- and usually remains as a hefty proportion of the variance of y (I'd say it's usually about 50% of the variance since the correlations in most experiments are rarely higher than .7).

The control model, simple as it is, makes precise predictions of what a subject will do in an experiment -- and there is no error term in these models. When we do research based on PCT we KNOW EXACTLY what the results should be -- in advance (because we can run the model in the same conditions as the subject BEFORE the subject is run). We usually find that the subject behaves EXACTLY as the model does. If not, our first guess is that we have done the experiment incorrectly. In PCT research, behavior that does not match the model exactly is a BIG surprise and requires (or would require -- it rarely happens) a very serious look at the way the experiment was done. In other words, in PCT we consider the model to be more important than the data; this is the same as in "real" sciences like physics. When you do an experiment in physics you know EXACTLY what should happen. If it doesn't happen, you first check the apparatus, the procedure, everything -- BEFORE CHANGING THE MODEL. A dramatic example of this was the Michelson-Morley experiment where they found no difference in the measured speed of light regardless of the orientation of the beam with respect to the "ether". This was a VERY surprising result that was inconsistent with the current model of light. Rather than immediately moving on to a new theory, Michelson and Morley did the experiment over and over again, checking everything they could think of, in an attempt to get the results that the KNEW (based on the model) that they should have gotten. Even when they convinced themselves that the results were real they didn't advocate any substantial change in the current model; luckily, Einstein had been inventing a new model (for other reasons) that happened to handle this result.

In conventional psychology, any experimental result that happens to be statistically significant will drive many people from one "model" to a new one; thus, we get "trendy science". In PCT, the model comes first; we do research to test the model -- but we expect exact matches of our research results to model predictions. If we don't get those matches, we KEEP TRYING until we are driven, by the sheer stubbornness of nature, to revise the model.

So in PCT, the model determines what one expects the data to be; in conventional psychology, the data (which is always changing) drives what one thinks the model should be. I think psychologists have a hard time with PCT because they are so used to their models being secondary to the current "discovery" that drives the invention of those models. In PCT, the control model IS the discovery and every experiment that tests the model should produce results that are EXACTLY consistent with the model. If the experiment does not produce such results, the first thing to become suspect is the EXPERIMENT, not the MODEL -- and steps should be taken to get the results predicted by the model. Only if you CAN'T get the expected results, after you've tried like hell to, do you say "we need to change the model" -- and if you do change the model, it should still produce behavior that exactly matches all the data that was matched by the previous model AS WELL AS the new data. It's an ideal -- but if it can't be done then the study of behavior can't be a science; just a floating crap game (floating to new models with each roll of the dice).

Best regards Rick

Subject: Control = Manipulation

[From Dag Forssell (920717)]

I have run into objections to my program. It suggests manipulation, I am told. I think it matters little if I talk about cybernetic control, perceptual control, individual control or whatever. A reader knows that control means manipulation. It matters not what I say or write.

I do not wish to call it something else. Therefore, I have written this appendix to be enclosed with my promotional materials.

Copyright Dag Forssell 1992. All rights reserved. Hard copy sent on request. Suggestions for improvement will be appreciated.

Headline: CONTROL: WHAT IS IT?

Subhead: A WORD WITH MANY MEANINGS:

The first meaning that comes to mind is the negative connotation of authoritarian dominance, command and manipulation; management by threats and coercion. If this were the only meaning, then a management program based on the concept of control would be incompatible with total respect for your fellow man.

But control also has a positive connotation of self-sufficiency, expertise, consistency and quality. Control is a phenomenon of life. The person happily "doing his/her own thing," is happily controlling what happens to him/her. The highest tribute of respect I can imagine is to support your fellow man in his/her exercise of effective control. An

understanding of control explains the "laws of nature" for life and shows you how to support and lead people in productive teamwork - without coercive manipulation.

Subhead: AN ENGINEERING CONCEPT:

Control is also a 20th century engineering discovery. A certain arrangement of inter-dependent elements creates a control system. The system influences the perception of a variable until it agrees with a reference (want, purpose, plan), by acting as best it can. Control in this sense is a characteristic of all living organisms.

Diagram of a control system

A complete diagram shown here:  
Perception, Want  
Compare, Instruction  
Action, Variable  
Disturbance  
Other effects

A control system is a building block which helps explain some of the smallest interactions in the nervous system. A single diagram can serve as a summary of the interactions of millions of control systems which together enable us to walk and talk.

Subhead: PERCEPTUAL CONTROL THEORY: A BREAKTHROUGH IN UNDERSTANDING

Control systems control what they perceive. A comprehensive explanation of how people function has been created, based on this new insight. It is called Perceptual Control Theory, or PCT. William T. Powers, who has developed the theory, writes:

"Perceptual Control Theory explains how organisms control what happens to them. This means all organisms from the amoeba to Humankind. It explains why one organism can't control another without physical violence. It explains why people deprived of any major part of their ability to control soon become dysfunctional, lose interest in life, pine away and die. It explains why it is so hard for groups of people to work together even on something they all agree is important. It explains what a goal is, how goals relate to behavior, how behavior affects perceptions and how perceptions define the reality in which we live and move and have our being.

Page 2

Perceptual Control Theory is the first scientific theory that can handle all these phenomena within a single testable concept of how living systems work."

Subhead: TRADITIONAL INFLUENCE:

Most psychological theories are based on the 17th century perspective of cause and effect. (The only perspective known in any science until earlier this century). Any book on experimental psychology tells you that the scientific method to learn about behavior

Focus of traditional science  
  
outline of control diagram, but control section gutted and

is to set up an experiment, establish initial conditions and then vary the independent variable (cause, stimulus) and measure the dependent variable (effect, stimulus). Our language and culture reflect this heritage.

only Effect (action) and Cause (disturb) shown along with Other effects

This would have been a fruitful scientific method if in fact animals and people were cause-effect organisms.

But if you accept for the moment (without a detailed explanation) that organisms are control systems, then you realize that by experimenting with cause and effect, we have learned little about the organism itself and what is important to it.

Subhead: AN EFFECTIVE PERSPECTIVE:

I believe most people intuitively realize that people are autonomous living control systems, full of desires and insisting on controlling themselves. Good managers and counselors respect and support control in their fellow humans. We talk about the elements of control, but lack a framework.

Most important control elements

Complete control chart with graphic emphasis on Wants, Perception and Variable.

When you realize that a person can be thought of as a control system, your own focus shifts. What is important about that person is what she wants and why, how she perceives and what might be interfering with her ability to achieve what she wants. Action is almost incidental. By openly helping her control, you can achieve your goals without manipulation.

The Purposeful Leadership™ programs explain PCT and applies it to management and leadership. The all-important wants and perceptions are explored in depth.

I am personally convinced that once the phenomenon of control becomes widely understood, this new insight will be a great force for the good of man. PCT shows the way to cooperation, productivity and high quality with satisfaction and respect for all.

End Dag

Date: Sat Jul 18, 1992 8:53 am PST  
Subject: Conditional stability

[Martin Taylor 920718 12:00]

I am hesitant about starting a new and possibly long thread before dealing with outstanding issues, but in writing up my extended abstract for the Paris talk, I came across a phenomenon I think could be quite important



in various areas. I think specifically of the mass murderer who is always the kindest, quietest kid on the block. Since it hasn't been talked about while I have been on this group, I'm wondering if it is a new discovery or one that has been long dismissed. (By the way, the Paris paper is the one for which I asked assistance some months ago, on HCI, where I have to give a keynote presentation that I intend to base around PCT. It is the reason I have been (relatively) quiet for the last couple of weeks.)

Consider two ECSs at the same level, labelled Y and Z. Their perceptual input functions are, say,  $y(a-b+c)$  and  $z(b-c+d)$  where  $a, b, c, d$  are sensations coming from A, B, C, D which may be lower level ECS perceptions or objects in the real world. Notice that Y and Z both control for functions that contain  $b$  and  $c$ , but also have another variable.

The outputs of Y and Z affect A, B, C, and D. Y affects A, B, C, and Z affects B, C, D. The signs of the output connections determine whether Y and Z have proper negative feedback control (getting them that way is what reorganization is all about). Let us assume that they do, so that all of them move in such a way as to oppose the error in both Y and Z.

Y and Z have their reference signal levels set from above. If their reference levels change oppositely, then Y tries to (say) increase  $a$  and  $c$ , but decrease  $b$ , whereas Z tries to decrease  $b$  and  $d$  while increasing  $c$ . B and C are affected the same way by both, and probably  $b$  and  $c$  change dramatically while  $a$  and  $d$  change only a little.

If the reference levels of Y and Z change in the same direction, then their outputs affect B and C oppositely. If they have roughly equal gains and equal shifts of reference level,  $b$  and  $c$  will not change at all, and Y and Z will reduce their errors by causing large changes in  $a$  and  $d$  respectively. Everything is still under control, but the outside observer will not see that B and C are in play. It will seem as though Y is concerned only with A and Z is concerned only with D.

Now comes the interesting part. If you haven't drawn the situation on paper yet, I suggest you do so, to make the argument easier to follow.

Let us change the sign of one of the outputs of Z, specifically the one affecting D. Now if Z acts alone, it will cause  $b$  and  $c$  to change so as to correct any error, but  $d$  will change so as to increase the error. Assuming they all have equal responsiveness, the changes in  $b$  and  $c$  will dominate the change in  $d$ , so that Z maintains control and all looks normal. Neither Z nor an outside observer can detect that the loop through D has positive feedback.

With this changed Z, reconsider what happens when Y comes into play. If Y collaborates with Z in its effect on  $b$  and  $c$ , the situation is normal. All is well controlled. But what happens when Y opposes Z in respect of  $b$  and  $c$ ? The levels of  $b$  and  $c$  do not change. A moves in such a way that the error in Y is reduced, but D moves to increase the error in Z. There is positive feedback. Z goes out of control. And supposing that Z contributes  $z$  to the perceptual input function of higher level ECSs that contribute to its reference level, there is the possibility that they, too, will go out of control.

There are mechanisms that would reduce the probability of propagation of control loss up the hierarchy. If the connections from above, say from ECSs labelled P and Q were such that Z contributed only a small amount to the perceptual input function of either, they might be able to retain control through Y and other ECSs at the level of Y and Z. Also, saturating non-linearities in the perceptual input and/or the gain function of Z would reduce the likelihood of a runaway loss of control, since the contribution of D to the propagated error would be limited to a fixed maximum that could be countered by control exercised through other variables.

But sometimes, the error induced by a positive feedback loop would propagate. The hierarchy could work flawlessly, with near zero error, for a long time. Then a particular pattern of reference signals might occur at some level of the hierarchy, preventing some ECS from exercising control through loops having negative feedback and making it act through a loop that had unsuspected positive feedback. And then, the hierarchy might go crazy.

I haven't gone too much further with this analysis to see whether the propagation of errors induced by hidden positive feedback loops will normally proceed, be suppressed, or whether it is a matter of chance. But it does seem to show that a hierarchy that has reorganized and has near zero errors in all ECSs (and in its intrinsic variables), still has the potential for violent and maladaptive (non-controlling) action. It can lose its temper.

A system in this state would be hard to detect, either from outside or from inside. It exhibits no errors most of the time, but goes mad under some ill-defined kind of stress. I cannot see how any reorganizing principle we have yet discussed (including my own) can avoid producing such a brittle system. But it is possible that it could be avoided by an annealing process in which the developing system is exposed to many different kinds of minor stress, in conjunction with either Bill's or my reorganizing principle. The minor stresses might sample the space of reference level changes, and could trigger these conditional instabilities, inducing reorganization that should remove the dangerous loops. But that kind of sampling might not always work, leaving around a potential mass murderer who does not know that he is one. And he would be the quiet, well-behaved type who has always been in control.

I don't know if the above is correct, but it seems so to me at the moment. I am still trying hard to write this paper, which had a deadline of last Wednesday, so don't expect too much follow-up from me for a while.

Martin

Date: Sat Jul 18, 1992 4:42 pm PST  
Subject: Catching up

[From Bill Powers (920718.1600)]

Back again after 5 days. Gave a talk (part of a panel) to the International Society for Systems Science, called "Information: a matter of perception." It was well received; a couple of people asked for more info on the CSG. I'll post the text in a day or so. Mary and I camped one night on the way there and another on the way back, and saw a lot of the interior of

Colorado including some preprocessing dirt roads. Fun.

-----  
 Naturally, there was 100K of mail when I got back. I'm not going to answer it all directly, particularly as the discussions that went on were pretty fruitful. I'll just focus on what for me are the biggest error signals.  
 -----

As I had figured, there are some dissenters from my concept of the hierarchy, and my discussion of implicit versus explicit functions turned up some more.

Bruce Nevin:

>I take it you mean orders of control systems (not to be confused with  
 >levels of control within one hierarchical control system).

I meant levels of control within one hierarchy, whether it be the CNS hierarchy or the biochemical one, or whether (as Martin Taylor proposes) one considers the whole thing just one big hierarchy (I don't).

Suppose you have a set of control systems of the same level within a hierarchy. These systems will be controlling for specific levels of specific perceptions. If there are no higher levels, the reference signals can only be random or fixed; there is no systematic means coordinating the independent controlled variables.

Let's suppose that these controlled variables are vector forces being generated by a set of brainstem or cerebellar control systems that control approximately along the axes of the mechanical degrees of freedom that are available. Martin's comments suggest that these systems would implicitly produce a vector in this space. I agree, they do. In fact, ANY combination of reference signals for these systems will result in a vector force. All possible resultant vectors are IMPLICIT in this set of n component vectors.

However, neither the individual systems nor the set of all systems at the same level control EXPLICITLY for any specific vector in this space. In order to produce control of a specific vector (such as "twist-and-push"), some higher-level system must perceive  $Av_1 + Bv_2 + \dots + Zv_n$ , where the v's are the individual vectors. The reference signal for that system then specifies that the component of the total vector in that direction be maintained at a given level. Now the lower-level world is represented, in this one higher-level system, as a component of force in a specific direction. There is a neural signal that represents the magnitude of this force; the direction is set by the perceptual weightings. No other component of the force is controlled. The actual multidimensional force can vary in many ways, but this control system will see to it that the magnitude of the component in one direction matches the reference signal that this new system receives (whether from a higher system, from a random process, or from a fixed property of a "floating" comparator).

When one thinks of a specific higher-level variable, it's easy to see that this variable is implicit in a set of lower-level variables. But it's also easy to miss the fact that SO ARE ALL OTHER POSSIBLE HIGHER-LEVEL VARIABLES OF THE SAME HIGHER TYPE. Without a specific higher-level control system to define a projected direction and to control for the amount of a perception in that direction, there is no reason to think that the lower-level

variables will spontaneously take on just the magnitudes needed to produce the state of the higher variable you have in mind.

-----  
When I said in BCP that higher levels in the CNS hierarchy are physically distinct from lower levels, I said this not for any abstract reasons or principles, but because it seems that the CNS is physically connected in exactly this way. The lowest level of behavioral control consists of the stretch and tendon "reflexes," with reference signals arriving via the alpha and gamma efferent signals. This level of spinal-cord control systems forms a package, fully functional without the remainder of the CNS. What it does is make sensed effort depend reliably on certain reference signals, with the internal connections serving also to create dynamic stability. It doesn't matter what supplies the reference signals. The first level of control is a physical entity.

The second level of control receives many sensory signals including copies of at least the tendon signals and possibly the stretch signals as well (the bifurcation of the dorsal roots). It perceives, through functions embodied in the sensory nuclei of the brain stem, variables that are functions of many of the signals arriving from the first level, both those under control and those not under control at the first level. Its outputs, which come from the motor nuclei of the brain stem, are identically the alpha and gamma reference signals reaching the first level of systems. Thus the second level is physically distinct from the first level, and acts exclusively through setting reference signals for the first level.

The third level is found in the thalamic regions, the midbrain. The sensory nuclei here receive signals coming from the brainstem sensory nuclei; the motor nuclei around the thalamus send their outputs to the motor nuclei of the brainstem (where comparison takes place -- at all levels, there are "collaterals" that carry sensory information into the motor nuclei and they synapse with a sign opposite to the sign of signals entering those nuclei from higher systems. Comparators of the second level are physically located in the motor nuclei).

All of these first three levels of systems perceive and control variables that seem consistent with my definitions of the first three levels in the HPCT model. In addition, the perception and control of higher levels, as I have defined them, seems consistent with the kinds of functions that have been found at higher and higher layers in the physical organization of the brain. In every case, moving to a higher level of control in the hierarchy seems to go with moving to new collections of neurons distinct from those concerned with lower levels of control.

This basic progression isn't quite exact, because the second and third levels appear to be repeated in the motor cortex and in the cerebellum. If you count synapses, however, the systems arrange themselves physically into levels even though a given level may have components in the brainstem, the motor cortex, and the cerebellum. A signal that passes through two intermediate nuclei, wherever it ends up, is involved in the perception and control of configurations, and so on.

While I can't prove this general theorem, therefore, I think there is excellent reason, based on neuroanatomy, to say that different levels in the human hierarchy correspond to physically distinct neural functions. Of

course without the constraints under which I worked, it's possible to define functions in such a way that this wouldn't appear to be true. One can invent all sorts of abstract hierarchies with different internal connectivities, with levels representing abstract concepts. Most such inventions, if not constrained to correspond to the organization of real brains and the organization of real behavior, would show no relationship to neuroanatomy and would suggest no intuitively-pleasing ways of parsing experience. If you simply look at hierarchies as mathematical constructs, anything becomes possible -- but if mathematics is the only constraint, the chances of describing the actual human hierarchy of organization are, to my mind, negligible.

-----  
Allan Randall comments:

>There are two ways one can talk about different "levels":

>(1) Conceptual: perceived levels.

>(2) Physical (architectural): perceiving levels.

>

>In (1) the "levels" are not actually in the control system under >discussion, but are in the type (2) perceiving levels in the mind of >the scientist building the model. The scientist is, hopefully, >controlling for these perceptions to square with reality (or to get him >grant money, whatever). Type (1) perceived levels are IMPLICIT in the >control system under study, while type (2) are EXPLICIT.

What I have attempted to do with my proposed hierarchy is to make (1) and (2) the same thing. In the course of developing the levels, I tried to catch myself using perceptual capacities that were not yet in the model. It took me about 35 years of observation to arrive at 11 levels. I think that you will find in this model all the kinds of perceptual functions that a scientist uses in conceptualizing levels (even if the conceptions aren't the levels I define, or even if the scientist doesn't believe there are any levels at all). If my project has succeeded, you will find the same things in the model that you find in the observer of the model.

Observers and theoreticians who do not use my levels as a description of brain organization nevertheless take the same perceptual elements that I propose for granted in their arguments. They all speak of objects undergoing transitions and making patterns we call events. They all explore relationships among these things. They all categorize. They all consider sequence or ordering significant. They all use some form of rule-driven logic in developing their symbolic arguments. They all derive general principles that guide their specific programs of reasoning. They all have coherent system concepts that give form to the collections of principles, programs, and so on that they treat in their investigations.

My claim is that these types of perceptions, and the systems that control in these terms, represent the real basic organization of the human brain. The content of brain activity at these levels -- for example, specific kinds of taxonomies, specific mathematical analyses or verbal arguments, specific principles -- do not represent basic organization. They are simply examples of what this kind of organization can produce. Their main significance is in their existence, not in what they appear to say. They are themselves evidence about how the brain is organized -- any brain, including the brain of a theoretician.

To me, the task of understanding how human beings work depends on putting direct experience, anatomical and functional knowledge, and observations of behavior together into a single self-consistent model that looks the same from any of these viewpoints. It even depends on producing mathematical and functional analyses -- but not just any old analyses. Whatever is analyzed, it must be consistent with the model in all respects, not just internally consistent. For example, many people are analyzing perception as if it consisted of chaotic oscillators. I don't object to exploring chaotic oscillators for the sake of their own fascination as phenomena, but how is this concept consistent with the way the world looks to direct experience, with simple facts of neural function, with the architecture of the brain, with the kinds of control processes in which we see people engaged?

I have tried to develop a model that says the same things about all phenomena no matter what point of view you take toward them. There's a lot left out of this model, but as far as it goes, I believe that it adheres to these principles.

You say, "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."

But what good does it do the brain to have the theoretician know this? I've heard this view before, and have always wondered how it connects to the fact that we perceive specific things separately -- what is doing that perceiving? I think this approach is the ultimate in implicit functions. As long as one doesn't require the brain actually to do something, such as reach out and press a specific button of a specific color, it seems to hang together. But I predict that this concept of perception will prove completely useless when it comes to trying to get such a system to act, to control specific variables relative to specific reference states. When a distributed perceptron has to produce an actual output, something has to recognize the distributed state of the system, and conclude that it is one state rather than a different one. All that's accomplished by this idea is to postpone the day when a specific perceiver has to be designed and built, saying "this is an 'A'".

> I guess talk of explicit hierarchies just strikes me as wrong.

It's wrong in relation to what some people believe about distributed functions. It's not wrong in relation to neuroanatomy.

Best to all, Bill (unbeliever) P.

Date: Sun Jul 19, 1992 10:51 am PST  
Subject: More catching up

[From Bill Powers (920719.0800)]

A somewhat more detailed response to the backlog of posts:

Bruce Nevin (920714.1314) --

>For example, a reference signal is explicitly present for a cell as an  
>electrical potential, ion concentration, whatever, but it is not a  
>reference signal for the cell.

By this I trust you mean "is not perceived by the cell as being a reference signal." Control systems don't perceive any of their components or signals; without consciousness and higher-order perceptions in a system, a given control system simply operates as it operates.

Without external inputs that create a internal bias in the perception-comparison-action process, the effective or default reference signal is determined by properties of the neurons or chemical reactions in the control system. At some level of sensory input, and in the absence of disturbances acting independently on the input, the feedback effect of the action on the input will be zero, tending neither to increase nor to decrease the input. That locates the effective reference level of the input quantity.

>The ion concentration is explicit for  
>the cell. The reference signal is not explicit, and cannot be, because  
>reference signals as such do not exist in the cell's universe.

The cell's knowledge of the reference signal is not what I mean by "explicit" and "implicit." The reference signal, being a bias on the input-output function of the control system, is always explicit and always has some value, although its value may not be variable by agencies outside the cell.

The terms explicit and implicit are intended to refer to coordinations or organizations made of independent systems. In systems of the same level that operate independently, interactions among the systems can give the appearance of coordination even though each system is unaware of any other systems, treating all others simply as environmental disturbances. Reference signals for these systems are also uncoordinated, by definition (if there are no higher systems to coordinate them). An observer can perhaps see that the systems are interacting through disturbing each other, but that knowledge is not available to the collection of systems and can't influence its behavior.

>As the cells (by whatever evolutionary process) come to constitute  
>control systems of a higher order, an ion concentration within a cell >can  
take on a new identity as a reference signal (or error signal, >etc.), in  
addition to its value as a variable within the cell.

The operation of a control system is not altered by an observer's analysis of its parts. If anything in a cell "takes on a new identity," that is only because of the observer's seeing it in relation to higher-order systems. Reference signals, error signals, etc. remain variables within the cell, related in the way required to create negative feedback and control, even if the cell becomes an element in a larger control system.

>All of that is invisible to the cell, which is only controlling its own  
>variables in its own terms.

Agreed.

>I am not claiming to demonstrate the existence of social hierarchies. >I am arguing for agnosticism regarding them.

I am arguing against the concept because I do not think there can be hierarchies of systems whose internal organization is essentially identical over the systems at all levels. Adding a superordinate system requires a new TYPE of controlled variable, a new dimension of control. Adding such a system in a control hierarchy requires the ability to set directly the highest reference signals in the subordinate systems through direct access to the highest-level comparators (i.e., not via the sensory inputs). This ability exists between levels inside single organisms; it does not exist between organisms.

-----  
RELAY OF A POST DIRECTLY FROM MARTIN TAYLOR:

From: VAXF::IN%"mmt@ben.dciem.dnd.ca" 14-JUL-1992 13:35:11.94  
Subj: RE: Reorganization

Bill, conflicting references sometimes give weird results. My conflict leads to the state of "I can't resist any longer." So here's a partial answer to yours of July 12. Only partial, because my other reference is still to write my Paris abstract.

-----  
>Still missing from your concept is the idea that errors in systemic >variables that are not part of the CNS can come to modify the organization of the CNS hierarchy.

Yes, but I conceive the CONTROL hierarchy as incorporating the CNS hierarchy, and I make no commitment about whether the CNS even constitutes a complete hierarchy on its own. At present, I treat all controlled aspects of the body as part of one hierarchy. CNS events affect the output of pheromones, the levels of hormones in the blood, and all sorts of things more or less directly. I don't see why ECSs for such things should not be treated in the same hierarchy as ECSs that use neural currents as information carriers. There's nothing sacred in my mind about one specific transmission system.

>This modification brings perceptual variables under >control that in themselves have no importance to the behaving systems >in the CNS; no inherent value or meaning.

But I think this is true of ALL controlled variables. None of them have any meaning to the behaving systems. There are only levels of perceptual signals, reference signals, error signals, and so forth. No ECS knows where these come from, or where they go. All it can know is their values.

> The value of controlling such a >variable is that controlling it has a side effect, through external >paths, on the systemic variables that (by departing from their >reference levels) instituted the process of reorganization.

Yes, in your concept of a separated reorganization system. Irrelevant in mine, since if the intrinsic variable departs from its reference, the relevant ECS will produce output that effects control in some way. If the



error is maintained, that ECS will probably reorganize--change what and how it links to other ECSs.

> So we can learn to control ANY  
>perceptual variable as a means of maintaining life, without the >slightest understanding of WHY THAT variable has to be controlled to >avoid thirst, hunger, pain, illness, emotional upset, and so on. >Neither do we have to be able to sense (CNS-wise or awaredly) the >intrinsic variable that is restored to its reference state by learning >to control a specific perceptual variable.

True in both our approaches to reorganization.

>All we know consciously is that controlling that  
>variable is more pleasing, feels better, than not controlling it. That  
>variable attains a value for us: high if the required reference setting  
>is high, and low if the required setting is low. It's a "good" >perception or a "bad" perception.

No problem with that.

Where there IS a potential problem with my approach is in the transduction between the chemical realm in which most of the intrinsic variables roam and the electrical realm of the CNS. I suspect that the intrinsic variables do not normally provide an explicit reference signal that is physically subtracted from a perceptual signal. More immediately, I suspect that they affect the excitation of some neurons. This can serve directly as a transducer under some conditions. For example, suppose that the pulse frequency of a free oscillator were under hormonal control. That frequency would be a neural current that could be used as a reference in a standard ECS of the CNS.

So there is at least a plausible way the intrinsic control system could be coupled into the CNS control hierarchy as a reference provider. But it is also true that changes in intrinsic variables can have global effects on the CNS hierarchy without any reorganization, by affecting the gain in moderately large parts of the network. I know of no evidence for this, but it seems a possibility to me. I know you don't like plausibility discussions, for good reason. All I am trying to do here is to show that there is no need to consider the control of intrinsic variables as different and separate from the control of other percepts in the CNS.

-----

I think it would be profitable to echo this interchange to CSG-L, starting back two or three messages. If you agree, could you do it? I'd like to get back to CSG-L, and that will be more quickly done if I finish this Paris paper. You'll get a copy when it is done.

Martin

-----  
This exchange gives the flavor of the last two rounds.  
-----

I agree that it's possible to see the entire organism as a collection of control systems of all kinds, ignoring the differences between chemical and

neural systems and treating reorganization as a process that occurs everywhere. In fact, I do see it that way. But for me, that degree of generalization, while allowing some true statements to be made, amounts to backing away from the problem of understanding human nature instead of getting into it more deeply. I do not think that true understanding arises from finding the most general possible way to view a system. I think it comes from noticing differences that make a difference, to quote Bateson. It was through looking at details that I arrived at the hierarchical model in the first place, differentiating the general process of control into levels and the levels into specific functions. The concept of reorganization arose from seeing that there is a kind of control that is different from systematic control of perceptual variables.

To lump the CNS hierarchy with the biochemical one is to ignore time-scales and relationships to the immediate world of the senses. It is the CNS hierarchy that produces the overt behavior of organisms and our own ability to observe and make sense of that behavior. Most of human experience is concerned with relationships to an outside world unperceived by the biochemical systems, and even our conscious relationships to processes inside our bodies are limited to what few aspects of the milieu interieur are represented as sensory neural signals. For the most part, the biochemical systems provide a controlled environment in which the CNS can live, acting as a life support system that largely takes care of itself and supports the life of the CNS, and which is adjustable in a few regards by the CNS where needed to support different modes of action. While the principles of control apply in the biochemical systems just as well as in the CNS, this is not sufficient reason to overlook the fact that the biochemical systems are NOT the CNS.

When you say that no controlled variables have meaning to the behaving systems, you commit a peculiar reductionistic contortion, using knowledge and meaning to deny knowledge and meaning. Isn't meaning one of the phenomena we're trying to explain in terms of the organization of a complex system? Don't higher systems perceive, and even think about, information contained in the signals arising from lower-level systems, thus giving meaning to these lower-level processes?

I agree that "if the intrinsic variable departs from its reference, the relevant ECS will produce output that effects control in some way." What I am concerned with is HOW this happens. How is it that an intrinsic variable, which represents the internal state of the organism, can lead to effects that alter the details of the way the organism controls variables external to the organism? This, not control of local conditions by local action inside the organism, is the heart of the problem of reorganization. I am trying to explain how conditioning works, how it is that an animal can come to modify control of arbitrary variables in the world of exteroception as a way of assuaging hunger, thirst, illness, and so on -- states of variables inside the organism. I want to explain how this kind of reorganization can work on any level in the hierarchy, regardless of the cause of the error, so that a child can learn to add 2 and 2 as a way of alleviating chronic physical discomfort, or can learn to exercise and build up muscles as a way of doing the same thing.

I resist the pressure, which comes not just from you but from large segments of modern science, to search for global generalizations that will

wrap up the grand principles of behavior as Einstein expressed the principles of General Relativity in four simple (-looking) equations. Generalizations are dangerous, because they are usually evaluated with respect to specific situations one is trying to explain, but when actually applied outside those situations almost always prove to be false. They aren't really general; their truth almost always turns out to depend on establishing the particular conditions under which they were derived, which of course means that they aren't generalizations at all even though the language in which they are expressed suggests a vast scope of applicability. It is very difficult to arrive at a generalization that is true in all situations, all of the time, with no counterexamples or exceptions. True generalizations, I believe, arise from considering details, not generalities. And even the best of them eventually fail and must be modified.

-----  
Cliff Joslyn (920714) --

>You people seem to have lost sight of the fact that you can still be  
>thought of as working within the broader systems science/cybernetics  
>fields. I can point you to a number of journals and conferences that  
>would probably be happy to have you, and can personally refer good  
>papers to some editors.

I don't want to be thought of in that way. I don't consider system science/cybernetics to be a broader field, but a very narrow one based on many unexamined assumptions and containing almost no experimental tests of hypotheses. I think that control theory introduces principles of organization that are mostly unknown in that field (you are an outstanding exception).

>It is true that the Systems literature is pretty much a ghetto, with a  
>low signal/noise ratio and a relatively high crackpot ratio.

Yes, it is true. I don't really want to be associated with it. But then, I am under no pressure to publish.

-----  
Allan Randall (920714.1900)

More on the concept of distributed functions.

>A "higher level" does not necessarily exist explicitly in the network.  
>E.g.: the generalised concept of PERSON could be an \*implicit\* emergent  
>property at the same \*explicit\* (i.e. architectural) level as the less  
>general concepts of JOHN and MARY.

This is an example of what I mean by using the products of the brain as if they were explanations. Why is it that PERSON seems more general to us than JOHN and MARY? Is XRPT more general than ZLTF or DRAQ? We are so used to hearing words and immediately converting them to their meanings that we think the terms themselves ARE the perceptions to which they point. If we look at the perceptions behind PERSON and JOHN, we see immediately what the difference is: JOHN refers to a specific configuration, while PERSON refers to a class. The level of perception involved (or perhaps I should say the lowest level) is simply different. PERSON is more general because it refers to a class that can be exemplified by any number of people

recognizeable at the configuration level; JOHN and MARY refer to elements of the class, which is to say, individuals distinguishable by their appearance at lower levels. So for the vague term "concept" we substitute a more specific idea, level of perception, and thus redefine the problem in terms of HPCT.

In order to perceive PERSON, it is not sufficient that perceptual functions exist for reporting the presence of JOHN, MARY, JOE, KWAME, and MIDORO. It is perfectly possible to perceive and distinguish individuals on the basis that they look different without assigning them to any class at all. In fact, if one has only the ability to perceive configurations, classes do not exist. Only a person who has developed perceptions at the classification level can see that in this collection of configurations there are potential classes. There can be no particular classes, such as PERSON or CIVILIAN, until there is a computing function that receives information about the elements that are present and detects the degree to which a particular class appears present. If one perceives in terms of FEMALE, then that class is present in the group named above (two individuals). If three of the people wear a certain kind of uniform, then the class MILITARY is also present. And so on. This sort of analysis demystifies "concepts," and encourages us to look behind the words for their perceptual meanings (if any).

>The fact that it is at a higher level can only be determined by >studying its dynamical/informational properties. Also, to get >interesting distribution, as opposed to localised classification, you >need a recurrent network.

Possibly this is so -- but if so, the recurrence would be entirely within the category level of perception. I rather doubt, however, that recurrence actually happens in the brain. This concept, I suspect, arises partly from computing techniques or mathematical concepts (recursion), and partly from words lifted without understanding from neurology ("recurrent collaterals," most of which prove to be feedback paths which are not recursive at all in the sense of either computing or mathematics).

>Now if I extend this principle straightforwardly to networks of >control, then the "level" at which an ECS is controlling is a >dynamical/informational property of the net, and would not be explicit >in the architecture.

The problem here is that you are extending the abstract principle according to the rules of a particular intellectual game, and are ignoring the meaning of the word control, the observed architecture of the brain, and control phenomena demonstrable in behavior. The AI literature has not been concerned with negative feedback control systems, alone or in networks. What you mean by "networks of control" is mostly likely not at all what I mean. I would rather suspect that such networks of control would be difficult to relate to an action like standing up, or aiming a gun at a target, or adding more salt to the soup. Has anyone, to your knowledge, simulated such a system that can control specific variables in an external environment, as animals and people can? By "control" do you mean the ability to stabilize external variables at specific reference levels despite the application of unpredictable disturbances?

>A node that might be called "low-level" in one context, might be  
>controlling at a higher level in another context.

This would be hard to apply to a system like the stretch/tendon reflex, which controls the force and length sensed in a muscle system. Such a simple system can't even control a sensation, which is defined as a weighted sum of receptor signals. How could this system "control at a higher level" under ANY circumstances?

>It seems to me that the describability of a system in terms of an  
>implicit hierarchy does not necessarily mean that the hierarchy is only  
>in our heads and not a real property of the system.

It does seem to me to mean exactly that. There are infinitely many ways to describe any complex system, depending only on your initial assumptions and the principles you choose to apply and not apply. In order to establish that any one description is NECESSARY, you must show that all the assumed components of the description do in fact exist, that they are connected as they must be, and that if so connected they behave as the real system would behave. The whole problem in modeling is to rule out alternative descriptions, which can't be done by logic alone. The question must be put to nature at some point, or all you have is a point of view, a possibility, a plausible conjecture.

>Can't one always claim the hierarchy is not really a property of the  
>system, but rather of the language used to describe the system? Even  
>then, can't I claim the hierarchy isn't a property of the language, but  
>of the language used to describe the language? Etc, etc, etc...

That depends on what you mean by "really" and how you establish the credentials of a "claim." One can, of course, claim anything. But making the claim acceptable involves testing its implications against the behavior of a real system, and when possible against the internal construction of that system. Behavioral tests will rule out most claims simply by demonstrating that their implications are contrary to observation. Examination of physical structure will distinguish between claims that pass the behavioral test, eliminating some claims by showing that the components of the real system don't have the required properties or are not connected in the implied way. Mere internal consistency of an abstract analysis is never enough to establish the acceptability of a claim. The only real credentials are to be found in showing that the premises of the argument are supported by observation.

>I've always thought of "lower" levels as also controlling for things in  
>the "higher" levels. At least this has been my notion of "control" >before  
>running across PCT.

Then your conception of the dimension in which lower and higher are measured is different from mine. "Lower" means to me "more peripheral in the nervous system" and "higher" means "derived from the lower." I wonder whether your usage of the term "control" here really takes into account the closed-loop kind of organization. Sensations can be altered IN ORDER TO CONTROL configurations (a higher level of variable), but this does not mean that sensations control configurations. It means that configurations depend on sensations, so that to alter a configuration it is necessary for the

configuration control system to tell the sensation-control systems to control for different sensations.

-----  
Bruce Nevin (920717.1305) --

>... it is possible that these changes in the cell's observed behavioral  
>outputs are cause-effect byproducts of controlling other internal  
>>variables (such as Na concentrations) against disturbance. Exploring  
>this possibility might lead to some explanations of how learning and  
>reorganization work.

In general I agree with this, but the scope of such reorganizations must be small. And anyway, if some aspect of a set of cells is to come under control, whatever the process of "coming under control" may amount to, something that receives information about ALL the cells must exist first -- some set of cells must appear that has a new kind of function. The old cells can't become superordinate to themselves, nor in fact is this how things work in the body. New functions always entail the construction of new physical devices distinct from the previous ones: a pituitary gland appears, which is part of a control system that uses existing organ systems to control variables that are more general in type than those controlled by the individual organs. These new systems are, of course, cells of the same old kind, but specialized to deal in the new type of variable. Explicitly.

>I am proposing (920709 09:13:52) that reorganization is carried out in  
>populations of entities of order n-1.

Propose away. But I ask that you specify all the functions and variables needed for this reorganization to happen, in addition to those needed for the control processes already going on. You've indicated an overall result that you want; it remains to demonstrate that such a result is possible, by some means that we can grasp. What does each part of this new reorganizing process have to accomplish, sense, do ... ? And what is your reason for believing that what such a system will do will actually solve the problem?

Things are getting awfully loose around here. What we need is a little more DISCIPLINE. So if you'll hand me that leather thing with the knobs on it ...

>But one thing a human does is change the color of its aura.

No, it doesn't. The observer does that.


It's not fair to bring up things your wife does as examples of real behavior, because anyone (like me) who doubts that these phenomena exist outside the imagination of believers is put in the position of criticizing your wife -- and you are put in the position of defending an idea in which you have a personal emotional investment and a loved one to defend. I think that the human brain is capable of supplying itself with any experiences it has reason to want, even with a vividness and to an extent that normally would be dangerous because of substituting too much imagined information for real-time sensory information. My opinion of psychic research is that it is sloppy, credulous, selective, and often dishonest. If you want to get into this subject that's up to you. But the result will not be an enhancement of our sense of intellectual companionship.

>... suggest ways of explaining input functions and output functions.  
 >Must these be separate multi-cellular structures, or might the metabolism  
 >of a single ramified nerve cell be such that it is not a simple cause-  
 >effect "wire" passing neural current through, but is actually doing the  
 >weighting (and the changes of weighting) of signals?

See [Martin Taylor 920715 15:00]

A single ramified nerve cell is quite capable of being an entire control system, with input function, comparator, and output function. A sensory nerve whose output goes directly to a muscle can be a control system, provided only that there is a negative feedback connection from the muscle, through the environment, back to the sensory ending (and that the loop gain is significant). I draw control systems in a conventionalized form in which the various functions are shown separately so we can consider their roles separately. In real systems, these functions can be embodied in different ways, and can be physically separate or combined into one computing unit like a neuron. I think that it's pedagogically useful to begin with a standard diagram, but there's always the risk that the overall relationships will come to be associated with that diagram instead of being understood in themselves.

For example, this is a biochemical control system:



A is the substrate; X1 ..X3 are reactions that result in a product concentration, X4. X4 influences the concentration of y1, which converts an allosteric enzyme, at a certain rate, from the active form Ea to an inactive form Ei. The enzyme catalyzes the reaction X2--> X3, determining the overall rate of production of X4 from the substrate A. The concentration of y2, which depends on another substance B, converts the enzyme at a certain rate from the inactive form Ei to the active form Ea. The ratio of inactive to active forms of the enzymes changes very rapidly when y1 is nearly equal to y2.

As a result, X4, the concentration of the final product, depends very precisely on the concentration of B, a substance which sets the reference level for X4 via y2. The concentration of X4 is nearly independent of drains (shown as an arrow to the right) and on the concentrations of precursor substances indicated by A. (Hayashi and Sakamoto, Dynamic analysis of enzyme systems , reference details lost, but it's a fairly recent book).

This diagram doesn't look like my standard diagram, but it's a control system. If you want a perceptual function, it would be somewhere in the link from X4 through y1 to the effect of converting the enzyme to the active form. The reference signal comes in via y2. The comparator is the

conversion from active to inactive form of the enzyme; the output function is the catalytic effect of the enzyme on the conversion of X2 to X3. The controlled variable, the output quantity, is X4. The disturbance consists of all the extraneous processes that can alter X4 independently of the conversion X3 --> X4. Notice that this control system doesn't consist of cells at all (although cells, or processes inside them, are responsible for making the enzymes, and probably for bringing some of the substances into position for reactions to take place).

>If there are such supra-human organisms, the variables that matter for >them are incidental for us.

Worse than that: we wouldn't even know that such variables exist, as we are incapable of constructing them as perceptions. We would know no more of such manipulations than a bacterium would know it has been sucked up into a pipette and deposited 1 cm away from its original position. Spatial "position" does not exist in the bacterium's world. A variable that is a function of a set of system concepts does not exist in our experiences, if my guess is right, so we would have no way to know that such a function is being controlled by some other agency.

-----  
Martin Taylor (910715.1545, 920718 ...) --

>Allan and I had been working with recurrent neural networks before we >got interested in PCT, and had found that quite simple networks could >exhibit varieties of behaviour dependent on their history, without >changing any of their structure or weight patterns. The self-same >network might be stable, might show a simple short-period oscillation >or more than one such, and might go into long-period or chaotic >behaviour.

I sense a cart pulling a horse down the road. Recurrent networks are certainly designable, and they will certainly have behavioral properties (every network will do something). But if one is simply investigating what networks with different designs will do, what is there to say that any particular one has something to do with behavior or perception? I think I'm more interested in the problem you were trying to solve by designing these networks than in the networks themselves (as an old designer of electronic systems, I'm used to the approach that says you first define what the system is to do, and then design it). Why, for example, did you consider oscillatory circuits? Is there something about perceptions that suggests oscillations? Chaotic oscillations? If so, what is it?

Your point about the control systems that control variables that are potentially in conflict is an excellent one. I presume that when such a situation arises, reorganization would eventually be generated because of the large errors. There's another class of problem similar to the one you mentioned, in which the environment takes on properties that are outside the design limits of a control system. Mary and I followed such a control system for a while yesterday, driving home. This control system, I assume, could drive a pickup truck with perfect competence. But yesterday it was driving a pickup truck with a largish four-wheeled hay trailer attached to it. When a disturbance made the hay trailer veer a little to one side, it dragged the rear end of the pickup truck the same way, causing the truck to move opposite to the deviation and pulling the hay trailer in the opposite direction. The dynamics were evidently almost beyond the control system's



capacity to stabilize the whole truck-trailer system; the driver's path was a continuous sine-wave oscillation showing no signs of ever stopping. Mary managed to get us past this incipient disaster safely, but I wonder where that hay trailer is today.

>Purposive behaviour is built into the elementary control system that is  
>the basis of the structure.

You're correct, of course, but Rick's remark is also correct, as you later noted. Emergence is the appearance of a property or function that isn't an aspect of any one component of a system. The failure of many people to realize this can be seen in lots of diagrams of control systems in which one of the components is called the "controller."

>I suspect Allan was getting more at the question of whether a group of  
>ECSS at a level can act as a population vector ...

Not in a way that specifically controls the vector rather than one variable contributing to it (see my post of yesterday).

>... "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.

>If, and only if, it is found that there are no recurrent perceptual or >reference connections in a real-life control system, Allan's intuition >would have failed. We are a long way from being able to assert that it >has failed or that it has not.

The term "recurrent" is pretty general. Recurrent connections in a perceptual function could belong to a system using negative feedback in a perceptual analogue computer; they could also cause oscillations, or bistable or multistable states. I don't think there's anything general you can say about ALL recurrent connections; their behavior depends fundamentally on the values of parameters such as the sign of the feedback and the time-delays involved, and on whether the underlying variables are discrete or continuous.

I can't visualize "recurrent reference connections." What does that mean?

Sorry about mixing together comments on various posts.

-----  
Dag Forssell (920717) --

>I have run into objections to my program. It suggests manipulation, I >am told. I think it matters little if I talk about cybernetic control, >perceptual control, individual control or whatever. A reader knows that >control means manipulation. It matters not what I say or write.

Well, I think what you have written will matter. That's a nice orderly exposition, and I don't see how anyone could get "manipulation" out of it. It is difficult to get people to start using an old word in a new way, but unless you want to talk in a language full of jargon and fake-sounding terms, I don't see any alternative. One advantage of succeeding in getting an old term redefined is that the old definition which conveys wrong meanings is automatically junked.

At the ISSS meeting where I gave a talk, I listened to Howard Odum going on about "transformity" and "emergy" [sic] which are made-up terms for something pretty simple: conversion of energy in one form to a result or outcome or energy in another form. He was actually talking about some pretty useful ideas, such as how much energy it costs to do things in different ways. In his case, the jargon wasn't even necessary, because the underlying concepts are quite conventional.

-----  
Well, I'll never be any more caught up than that.

Best to all, Bill P.

Date: Sun Jul 19, 1992 11:52 am PST  
Subject: A Bomb in the Hierarchy

[From Rick Marken (920719.1300)]

Martin Taylor (920718 12:00) suggested an arrangement of control systems where one might exist as a "hidden" positive feedback system just waiting to go off.

The proposal was to have two systems, Y and Z, controlling different linear combinations of lower level perceptual variables A, B C and D. The idea was that it would be possible to hook the outputs of Y and Z to A,B,C and D (or to the references of the systems controlling these variables) so that a positive feedback connection from Z would be concealed through "collaorative" behavior of Y (see Martin's post for details). Not being a linear algebraicist, I set up Martin's suggestion as a spreadsheet model. The result was that there is no "collaborative" way for systems Y and Z to control their perceptions when system Z has an inappropriate connection to one of the variables it is controlling. The only collaborative relationship is the one where system Y is just not controlling its perception (the one that partially overlaps the one controlled by system Z). It is true that system Z can control it's perception even with the wrong output connection to one component of that perception -- at least over the range of references I investigated. Actually, now that I think of it, there would be no way for Z to control for negative reference values. I'll try that in a second (I don't have multifinder at home dammit).

So modelling shows that Martin's problem is no problem -- at least for the specific case he proposes. An inappropriate output connection would lead, quickly, to lack of control (especially if two systems are controlling two variables that are not completely independent) and this is something that would lead to reorganization. So the hierarchy would be strongly biased against maintaining such systems.

Regards Rick

Date: Sun Jul 19, 1992 12:00 pm PST  
Subject: Bomb in the Hierarchy II

[From Rick Marken 920719.1305]

Well I tried the Z system alone with one incorrect connection (of the three possible) and it works fine. I think its because the negative connection still exists for the C system. So you have still have enough DF available for control of the Z perception.

Hasta Luego Rick

Date: Mon Jul 20, 1992 7:06 am PST  
Subject: ISSS talk on information

Information: a matter of perception  
William T. Powers  
The Control Systems Group

(Abstract)  
Information and perception

Information is a term that requires definition, rather than having one. The intuitive concept that information is carried by an energy flow from source

to destination is incorrect, for information and energy can flow in opposite directions. Even the transmission of information is not a straightforward process. A simple demonstration shows that even transmitting the answer to a question cannot be associated with any particular motor action by the transmitting person. The subject of information is closely connected with the subject of perception, and perception is associated with behavior most realistically through the concept of control: the behavioral control of perceptual variables.

#### Introduction

I had toyed with the idea of titling this paper "Noitamrofni," as a way of emphasizing that "What is information?" not only asks a question, but begs one. When we hear a word like information, intelligence, behavior, or perception, there is a sense of familiarity that we can easily mistake for meaning. An empty word like noitamrofni lacks that sense of familiarity, so we don't expect it to mean anything. It's unfortunate that we can't begin that way with all words that point toward abstractions, and decide what they are to mean rather than assuming that they already must mean something because we use them in conversation.

The term information is so utterly familiar that it has become reified -- we speak of information flowing from a transmitting person to a receiving person almost as if it is a physical substance. This leads, in turn, to making mistakes which any thinking person could discover and avoid if this unfortunate image didn't get in the way.

#### Information and energy

I think we all use the term information when we mean that a person gains knowledge from another person or from the physical world. Information is meant to indicate what passes from source to destination that results in this knowledge. Shannon and Weaver examined a physical signal flowing from source to destination, and analyzed it into bits as a way of comparing the flow of information to channel capacity. From this start it was a small step to to link information flow to energy flow and the measure called entropy: information decreases the entropy of the receiver. One can then identify information with a decrease in uncertainty at the receiver. And so on and so on -- each convincing analogy leads to the next analogy. Unfortunately, this is a house of cards and one of the bottom cards is imaginary. Information does not necessarily travel in the same direction as physical energy.

If you're reading these words in black print on a piece of white paper, you can see immediately what I mean. The energy that flows into your eye (a process to which we refer as "looking" and treat as if it goes the other way) comes from the reflection of light off the white paper. This energy flows from the page to your retina everywhere except where the letters are printed. If the letters are dark enough, a net energy flow may actually travel from diffusely illuminated points on your retina, through the lens, and to the light-absorptive ink on the paper. That is how you see the letters and gain whatever information they are carrying into your brain.

Or consider an old-fashioned telegraph. When a station agent out in Dodge City sends the sheriff's message to a Federal Marshall in Chicago, he does

so by closing and opening a key that connects by wire to Chicago. When the key is open, no information is traveling. When the key closes in long and short patterns, a message in Morse Code is sent to Chicago. The dots and dashes represent momentary drains of energy on the battery in the central telegraph office in Chicago; the average direction of energy flow is from the Chicago office to the short circuit in Dodge City, heating the wires along the way.

Of course if the battery is in Dodge City, the same message is sent by a flow of energy in the other direction. My point is that the direction of energy flow is unrelated to the direction of flow of information, so that concepts like entropy (positive or negative) have no actual physical relationship to whatever it is that we mean by information. That whole line of argument, which sounds so airtight when you play the game according to the background assumptions of Shannon and Weaver, is really just a case of mistaking an analogy for an analysis.

I hope that these examples have created a little doubt that reception of information is as straightforward a process as we commonly assume it to be. I will leave the examples to ferment and pass on to some even more doubt-provoking thoughts, this time about the transmission of information by one human being communicating a message to another.

#### Information and behavior

It's commonly assumed that transmitting a communication is done by commands originating in the brain. If you ask a person "How much is 2 and 2?" the person's brain computes a response that is produced by the muscles of the mouth and diaphragm, creating the sound "Four." So to communicate the answer, 4, the person emits a response that means 4. A straightforward behavioristic interpretation of this situation is that one has learned to respond to the question by emitting the answer: for every question successfully answered, there is a learned pattern of motor actions that is triggered off to produce the answer.

I have brought with me a demonstration device to show you that this is not how answering questions works. I hope you will be able to generalize from this simple demonstration and see that the transmission of information, whether in answer to question or not, is not just a matter of initiating a simple chain of causes and effects travelling from source to destination.

The apparatus consists of a board with large numbers on it arranged in a semicircle and a pointer that can swing to indicate any of the numbers. Attached to the pointer is an elastic band; pulling on this long rubber band can deflect the pointer from one end of the number scale at +5 to the other end at -5. I have arranged for a volunteer, who will answer my questions by pulling on the rubber band. That is a motor behavior just as surely as speaking the answer is.

With very little practice, the volunteer can learn to emit the correct answers to simple questions by using this unusual but simple kind of response. For example, I now ask the volunteer "Counting only you and me, how many people are actually putting on this demonstration?" And the pointer swings to 2. The actual response was to pull back on the rubber band to bring the pointer from its resting position at +5 down to the

number +2.

Now I ask a series of questions. What is 2 plus 2? What is 3 minus 5? What is 7 minus 4? In each case, the volunteer moves the hand holding the rubber band to a new position and the pointer comes to the right answer. With enough experimentation, we could make a table showing how the person's hand will move in response to all questions of this kind that have answers from -5 to plus 5.

I would like to point out as an aside that there's an energy-flow problem here, too. Whenever the answer changes in the positive direction, the person is moving the hand in the same direction that the rubber band is pulling. The rubber band is therefore doing work on the hand, so the energy flow is from the apparatus to the volunteer. The other transitions involve energy flow in the other direction, so the net energy transmission over a long series of answers approaches zero.

So far this looks just like the behavioristic concept of answering questions. Each question elicits a hand movement that is appropriate to the question. But now, while I ask a few more questions, I will reach behind the apparatus and do something invisible to the volunteer. How much is 4 minus 2? How much is 8 minus 6? How much is -1 plus 3?

You can see that each question elicited the right answer, 2, but that the motor response of the volunteer was different each time. What I was doing, behind the board, was pulling by different amounts on another rubber band that also influences the position of the pointer. I did this just as I finished each question. Maybe the volunteer could detect some small motion of my arm, but certainly not accurately enough to align a pointer within a few degrees of the right answer by balancing out my added pull.

So transmitting the answer to a question is clearly not a simple matter of emitting a motor response that corresponds to the question, at least not in this demonstration. Yet the questions are clearly and correctly answered with no hesitation. How does this work?

#### Information and control

As a step to the answer, I'll ask the volunteer now to close his/her eyes and, in the same way, answer one more question. The eyes are now shut; volunteer, please keep them shut. How much is five minus 8?

Just to make sure this is impossible to do, I am pulling on the hidden rubber band. The answer indicated is clearly wrong. Now the volunteer can look, and try again. Immediately, the right answer. I think that concludes our need for the volunteer, with thanks.

What you have just witnessed was not a series of "responses" but a process called control. The final part of the demonstration showed that in order for the position of the pointer and its relationship to the numbers to be controlled, it was essential that the volunteer be able to see the pointer and the numbers. What I was doing behind the board was disturbing the pointer through another rubber band. In order for the right answer to be indicated, the volunteer had to change the hand position just enough to counteract my invisible disturbance, and just enough more to bring the

pointer to the place where the volunteer could see it indicating the answer. The only way to do this was to see where the pointer was in relation to the numbers, and change the hand position in whatever way was required to bring about the right relationship.

The "right" relationship, of course, was internally known to the volunteer as soon as I asked the question. But this knowledge couldn't be translated into any set amount of action, because the amount of disturbance I was injecting was unknown. What had to happen was a smooth adjustment of hand position that brought the pointer toward the number known to be the right answer, until the pointer was perceived by the volunteer as being aligned with that number.

The volunteer, therefore, transmitted the right answer by controlling his/her own perception of the pointer's position relative to the number indicating the right answer. When the pointer came to rest, of course, we, too, could see the indicated number, and we, too, obtained the answer to the question. If I had asked "How many living children do you have?" we might have obtained some new information, some knowledge we didn't have before. But that gets us into different subjects, such as the significance of an answer of -3, which is certainly among the possibilities. I will leave questions about that kind of "information" for other speakers.

#### Information and perception

A question that is quite normally omitted from discussions of information is "How do you know what information you're transmitting?" This seems like a trivial question -- we simply transmit what we intend to transmit, and hope the message gets through. But how do you know that your transmission matches what you intended to transmit?

From what we have seen demonstrated, we can construct an answer. We perceive the immediate consequences of our own acts, and compare those consequences with an intended consequence. Before we initiate those acts, we know what the intended message is but no one else does. After we emit those acts, even while we are emitting them, we may or may not perceive that the message sent is the one intended. We may see that we have typed an "a" where we meant a "w" to be sent. This calls for corrective action of some sort -- rubbing out the "a" and retransmitting a "w", or simply retyping the whole word or packet. The point is that we can perceive what we're sending, and compare it with what we intend to be sending, and make sure the two things match.

Significantly, perception seems to be involved at both ends of transmitting information. The sender perceives what is being sent, and varies the actions involved to correct any differences between the actual transmission and the intended transmission, even while the transmission is in progress. The receiver also perceives the same result, give or take any channel noise. So the transmission of information from one person to another would seem to involve both people perceiving something, while one of them alters what's being perceived to keep it in a match with an inner intention.

I'll conclude by opening another subject which you can pursue on your own. When we transmit information, we hope we are transmitting more than words; we hope we are transmitting meanings. But meanings are a private matter

which each of us must learn alone. What we're hoping is that the meanings we perceive as we hear the words leave our mouths will be the meanings that the listener will construct when our words get to their destination -- the brain of the other person. Meanings aren't carried by words; they're evoked by words from the personal experiences of the transmitting person and the receiving person. Those experiences are not identical.

Even when the experiences are reasonably similar, words are slippery objects and can often be construed in different ways. I like an example from linguistics: "The shooting of the hunters awakened me." I may intend to transmit a complaint about inconsiderate hunters; you may receive information about a massacre. Did I transmit the information you received? As far as I was concerned, the meaning I transmitted was the one I heard as I sent the message. But you got a different message.

Information exists at many levels in the same message, because information is perception, and perception in both sender and receiver occurs in a hierarchical system with many levels. I tell you that I have a check made out to you in my hands; you hear that I am making another excuse for not having mailed it. I am using a message to control for a certain response from you; you are responding to control for a certain action from me. The physical transmissions that pass between us are like any other physical transactions in which we engage: we use them to control our own perceptions.

Date: Mon Jul 20, 1992 8:13 am PST  
Subject: No bomb in hierarchy

[From Bill Powers (920720.0930)]

Martin Taylor (920618) --

RE: Bomb in the hierarchy.

OK, Rick Marken is right. The two control systems, one with a wrong connection to one environmental variable, will be stable and will control their respective perceptual signals independently. Rick proved it by simulation: here is the algebraic proof.

LET:

a,b,c,d = environmental variables

Rx = reference signal, x system

Px = perceptual signal, x system

Ex = error signal, x system

Gx = output gain, x system

Ry,Py,Ey,Gy = ditto for y system

Perceptual signals:

$P_x = a - b + c$

$P_y = b - c + d$



Error signals:

$$\begin{aligned} Ex &= Rx - Px \\ Ey &= Ry - Py \end{aligned}$$

Output effects on environmental variables:

$$\begin{aligned} a &= Gy*Ey \\ b &= Gz*Ez - Gy*Ey \\ c &= Gy*Ey - Gz*Ez \\ d &= Gz*Ez \quad \text{(right connection to d)} \\ d &= -Gz*Ez \quad \text{(wrong connection to d)} \end{aligned}$$

Calculation of perceptual signals:

$$\begin{aligned} Py &= 3*Gy*Ey - 2*Gz*Ez \\ Pz &= 3*Gz*Ez - 2*Gy*Ey \quad \text{(right connection to d)} \\ Pz &= Gz*Ez - 2*Gy*Ey \quad \text{(wrong connection to d)} \end{aligned}$$

Substitute  $Ey = Ry - Py$  and  $Ez = Rz - Pz$ , solve for  $Py, Pz$ :

$$Py = \frac{3*Gy*Ry - 2*Gz*(Rz - Pz)}{1 + 3*Gy}$$

$$Pz = \frac{3*Gz*Rz - 2*Gy*(Ry - Py)}{1 + 3*Gz} \quad \text{(right connection to d)}$$

$$Pz = \frac{Gz*Rz - 2*Gy*(Ry - Py)}{1 + Gz} \quad \text{(wrong connection to d)}$$

Nifty shortcut to solution:

Assume solution:

$$\begin{aligned} Py &= Ry \text{ and} \\ Pz &= Rz \end{aligned}$$

Then

$$Py = \frac{3*Gy*Ry}{1 + 3*Gy}$$

$$Pz = \frac{3*Gz*Rz}{1 + 3*Gz} \quad \text{(right connection to d)}$$

$$Pz = \frac{Gz*Rz}{1 + Gz} \quad \text{(wrong connection to d)}$$

Let gain be very large:

$$\begin{aligned} \lim 3*Gy/(1+3*Gy) &= 1 \text{ as } Gy \rightarrow \text{infinity} \\ \lim 3*Gz/(1+3*Gz) &= 1 \text{ as } Gz \rightarrow \text{infinity} \\ \lim Gz/(1+Gz) &= 1 \text{ as } Gz \rightarrow \text{infinity} \end{aligned}$$

Then

$$\begin{aligned} Py &= Ry \text{ and} \\ Pz &= Rz \text{ as assumed.} \end{aligned}$$

Assuming infinite gain is the same as using a time-integrator in the two output functions (as I believe Rick does in his simulation).

We should have seen immediately that if a pair of linear systems work stably under ANY conditions, they work stably under ALL conditions.

If anyone feels like doing it, it could be shown that the same solution will result if we add Da, Db, Dc, and Dd to the four environmental variables (arbitrary disturbances).

Best, Bill P.

Date: Mon Jul 20, 1992 11:00 am PST  
Subject: No bomb in hierarchy

[From Rick Marken (920720.1130)]

Bill Powers (920720.0930)

>OK, Rick Marken is right. The two control systems, one with a wrong  
>connection to one environmental variable, will be stable and will control  
>their respective perceptual signals independently. Rick proved it by  
>simulation: here is the algebraic proof.

Actually, what I found in the simulation is that two control systems, one with the wrong connection to one environmental variable, DO blow up. I took this to mean that there was no way to "hide" a "positive feedback system" in the hierarchy, as Martin suggested. When two systems are connected as Martin suggested (and as you set up in the algebraic derivation) then there is no stability (contrary to the results of your algebraic proof). This means that either the simulation is wrong or the algebra is wrong. I have just re-run the simulation at work and it still blows up -- when the output connection to the d variable is made negative instead of positive. It only blows up when BOTH the Y & Z systems are operating simultaneously. When just the Z system is operating alone with the wrong d connection it works just fine.

I don't have time to check the algebra; but I'm pretty sure that the simulation is correct. What the simulation suggests to me is that you can have a system, like Z, that works just fine with a "wrong" connection. But when another system is added at the same level it may lead to instability of the existing system (and itself) because of the "bad" connection in the existing system. This suggests

that learning something new (the Y system) can screw up what you already know how to do (the Z system). If you need to control the Y variable, then you will have to learn a new way to control the Z variable (change its connection from negative to positive, for example). The Y system could never be "organized" into the hierarchy until the appropriate "relearning" in Z had occurred too. That's why I suggested that there is no Bomb in the hierarchy -- there is no way to add a new system so that the "bug" in the existing system remains hidden. Instead there will be sudden loss of control in the existing system. This would STOP the reorganizing system (which is adding the new system) thus preventing further loss of control in the existing (Z) system. I think it's an interesting case -- and would be great if it could be set up as a reorganization experiment. But I'm puzzled by the discrepancy between my simulation results and your algebra. If you're right again I'll, I'll...tear up my spreadsheet.

Best regards Rick

Date: Mon Jul 20, 1992 11:00 am PST  
Subject: so long, farewell, auf weidersein, etc.

[from Joel Judd 920720.1345]

Don't know when my account will officially "end," so I thought I would say goodbye before it does. A permanent surface mail address for me is the following:

Joel Judd  
1103 Alderbrook Ln.  
San Jose, CA 95129  
(408) 252-8740

We'll be leaving Illinois (Motto: The only good histamine is an anti-histamine) on August 7th and going to the above address for a couple of weeks. Hopefully I will surface in the southwest sometime after that. If not (or if my employment has not moved into the e-mail age yet, or if I end up self-employed), I can always be reached at the above address.

I'm gonna miss the daily ritual of reading the list, and will no doubt fall behind during the interim. So please--no major paradigm shifts for a couple of months, OK?

Hasta la vista--Joel

Date: Mon Jul 20, 1992 1:10 pm PST  
Subject: PCT Seminar

Hi! Professor Robertson gave me a copy of your version 9 letter on PCT and I found it interesting. I will graduate in December 1992 under the BOG Program and with a major in Psych. Control theory is basic to my interest in the field and I will be needing something to do soon. I have many years of business management including the conduct of seminars in sensitive areas (equal employment opportunity). I understand you may be looking for someone to deliver your program in the Chicago area. Please contact me if you have an interest in discussing this.

Date: Mon Jul 20, 1992 1:13 pm PST  
Subject: Bombs back again

[From Bill Powers (920720.1500)]

Rick Marken (920720.1130) --

Well, I got to feel clever for a couple of hours, anyway. I think I trust the simulation more than the algebra. I started to solve the equations the long way -- they get pretty horrendous -- but then saw that I could use a shortcut, by assuming the solution and then showing that the rest was consistent with it. I haven't redone it all, but I think I know where the problem is: you can't get there from here.

The algebra doesn't know anything about the path by which a solution is reached from a set of initial conditions. I should have known this. If you write the algebraic equations for a system with positive feedback, you get a perfectly good-looking set of equations that can be solved for the variables. They indicate a state, however, that can't actually be reached by a real system: there's a singularity in the way. The only way to discover this is either to solve the differential equations, or to simulate the system as you did. As you say, the system as a whole blows up -- heads for an infinity or a pole that it can't get past. Maybe there's a solution somewhere on the other side, as there is for the Lorenz equations with velocities greater than the velocity of light. But not in THIS universe.

So Martin's intuition was better than mine, and your spreadsheet was better than his intuition. I'll have to remember not to use that slick way of reaching a solution again.

Humbled, Bill P.

Date: Mon Jul 20, 1992 4:04 pm PST  
Subject: Re: List server

[from Gary Cziko 920720.0900]

Still catching up from my two-week absence.

Dag Forssell (920706 2100) said:

>A long time ago, Gary said that it is not doable to get your own post in  
>return from CSG-L. I would like to get that kind of feedback. How do you get  
>it?

The listserver does not send a posting back to the source from which it came. But it is possible to get an acknowledged that a post was distributed. Just send the following command as the first line of a text in a message to LISTSEV@VMD.CSO.UIUC.EDU

set csg-l ack

This should give you an acknowledgement that your post was distributed--very useful for all you feedback-seekers out there.--Gary

Date: Tue Jul 21, 1992 9:41 am PST  
Subject: data analysis

[From Pat Alfano]

July 21, 1992

I am not very happy with the data analysis I am doing on some recently collected data. I was wondering if PCT might offer me a new way of looking at the data.

Briefly, I used two neuropsychological tests of spatial orientation and had subjects perform them under different conditions. Subjects were timed and errors counted. After the test subjects were given a list of strategies and asked to check off any of the strategies they may have used to perform the test, and to note which strategies worked and which did not work for them. I thought that strategies used may have had an effect on time and errors. I did a multiple regression on the Stick Test strategies and found that 5 of the 11 strategies accounted for most of the variance in time to completion.

The Stick Test uses 10 patterns of 5 inch long by 1/2 inch wide, flat sticks with one white tip. The examiner arranges 2, 3, or 4 sticks in a pattern and the subject who is sitting across the table must arrange her sticks to look to her as the examiners sticks look to the examiner. In other words, the subject has to up-down and left-right reverse the pattern. Five of the patterns use only sticks placed in a vertical/horizontal orientation and 5 patterns use sticks placed on a diagonal orientation only.

When I did a pilot study of the Stick Test and plotted each subjects time on the two orinetation dimensions subjects times seemed to fall into 1 of 2 patterns, but, on further testing, this did not hold up.

If anybody has any ideas about a better way to analyze the data I would appreciate hearing them.

Thanks in advance.

Pat Alfano

Date: Tue Jul 21, 1992 10:03 am PST  
Subject: Portable Demos

[from Gary Cziko 920721.1215]

For the Durango meeting next week, I am compiling a list of portable demonstrations of perceptual control and will demonstrate a few of my own.

I think I have already included most of the demos mentioned in the net plus

the classic ones, but I fear I may have missed a few, particularly ones involving cooperation-conflict in social situations.

If you have a portable demo that you don't think I know about, please clue me in. I know that Clark McPhail uses some demos involving BIG rubber bands in his sociology class, but I don't know exactly how this work.

I will eventually compile all of these and make them available on the fileserver. --Gary

Date: Tue Jul 21, 1992 11:37 am PST

Subject: Stick patterns

[From Bill Powers (920721.1200)]

Pat Alfano (920721.1130) --

Hi, Pat.

>The examiner arranges 2, 3, or 4 sticks in a pattern and the subject  
>who is sitting across the table must arrange her sticks to look to her  
>as the examiners sticks look to the examiner.

Did you verify somehow that the subjects actually understood the instructions? The phrase "arrange your sticks to look to you as the examiner's sticks look to the examiner" can't be understood unless the person grasps the idea of transforming from one perceptual point of view to another. It seems to me that the very ability you are testing for is required in order to give meaning to the words of the instruction.

Assuming that all subjects succeeded, the next question is what time has to do with this ability. If it takes some subjects longer than others to accomplish the task, what is the difference between them?

I would start with a much simpler task, using just one stick. The examiner places a stick at some random angle, and the subject places another stick according to the instructions. This tests for the ability to rotate the frame of reference by 180 degrees in the simplest possible way.

A rotation of 180 degrees is an ambiguous task, because it can go in either direction. It's possible that time differences are due to the indecision as to which way to commence the mental turning (in control-system terms, the left and right errors are balanced, like a stick balanced on end, so only some small chance deviation can start the correction process in one direction or the other). If so, even the single-stick task would prove variable in time to completion.

To test that hypothesis, I would seat the subject at right angles to the examiner, to left and right, instead of directly opposite. Now the rotation required is only 90 degrees, and there is no ambiguity in the direction of mental movement required. Rotation times should now be shorter, and less variable. I would repeat this experiment for many angles between 0 and 180 degrees plus and minus.

As baseline measurements, you should determine for each subject how long it takes to reorient one or more sticks by 180 (or 90) degrees from the original position. This time should be subtracted from the total time during the real task, or the

total time should be measured in units of this baseline time. The time should be reduced to manipulation time per stick before the data are combined.

Another dimension to test for would be the ability to remember the target pattern (shown briefly) long enough to reproduce it as it is, or in the opposite orientation. This would test for the accuracy of the reference signal and the decay of accuracy with time delays.

Did you check to see whether the different strategies inherently required different amounts of time to carry out? How about some examples of those strategies?

Notice that what you are interpreting as two dimensions of reversal (discrete variables), I am interpreting as an angular rotation (a continuous variable).

Given that subjects persist in taking different lengths of time for this task, I would begin to look more closely at how they accomplish the task, recording each discriminable movement -- how it begins, runs its course, and ends, and how much delay there is before the beginning of the next movement. I would note how often the subjects look at the experimenter's array (you could have them press a button to illuminate the experimenter's array, using the same hand they use to move the sticks). If the subjects take different times to do the task, it is for a reason. The reason may be different for every subject, or there may be common kinds of reasons. It is highly unlikely that variations in a task on this macroscopic scale are due to random noise in the system; you are not using threshold stimuli, nor is there any uncertainty to speak of in perceiving orientation or executing the gross movements required. Randomness in the data is most likely to mean that there's nothing to measure, as you've defined the problem.

I'd videotape the proceedings and then write up a detailed description of each person's behavior during each repetition of the task. This would probably show you why there are variations in the time to completion.

Is this a phenomenon you really want to understand? Or is it just that you'd like to get some significant statistics without doing more experiments? If the former, I'll be glad to come up with whatever other suggestions I can think of. If the latter, forget it. Why waste your time on phenomena that you have to use statistics even to see? Once you get the right slant on the problem, the phenomena will be big and obvious.

By the way, I have a start on the motion-illusion thing, but won't do more with it until after the meeting.

Best, Bill P.

Date: Tue Jul 21, 1992 1:45 pm PST  
Subject: Re: A Bomb in the Hierarchy

[Martin Taylor 920720 10:30]  
(Rick Marken 920719.1300)

Rick says that his spreadsheet modelling indicates that one can't have a masked positive feedback loop within the hierarchy. Modelling is certainly superior to wordsmithing, but I remain unconvinced. Since I came across

the idea, I seem to see it in a lot of apparently inappropriate human behaviour including suicide, which seems to correspond exactly to this situation.

Rick, I guess I will have to do a little more analysis. But are you absolutely sure about your spreadsheet analysis? There are issues of computational order and the gain regime within which masking will happen, aren't there? In the situation I envisage, if Z acts alone, it will display negative feedback and be able stably to control its percept, because if B, C, and D have the same impedance, the effects of Z on b and c will swamp those on D. You say that your model agrees with this. If the effective impedance of B and C is reduced by Y, this should make Z more stable, but if it is increased by Y opposing Z's action on them, then the effect of Z on D should be relatively more important, and at some point the overall Z loop should go positive. If this doesn't show up when modelled, I'd like to understand why not. How is the effective gain of Z affected by the action of Y?

I don't see how you can reconcile:

```
> It is true that system Z can control it's perception
>even with the wrong output connection to one component of that
>perception -- at least over the range of references I investigated.
```

with

```
> The result was that there
>is no "collaborative" way for systems Y and Z to control
>their perceptions when system Z has an inappropriate connection
>to one of the variables it is controlling.
```

You mean Z can control if it acts alone, but never if Y is also acting, even if Y is affecting B and C in the same sense that Z is?

Martin

Date: Tue Jul 21, 1992 1:32 pm PST  
Subject: Conditional intuition

[Martin Taylor 920720 16:15] (Rick Marken 920719.1300)

Earlier today I replied to Rick Marken's claim that his spreadsheet model refuted my intuitive claim that there could be hidden positive feedback loops in the hierarchy. I now think I understand why he got his results, and have a better feel for what is going on (Bill is right to warn against thinking one can intuit what goes on in a hierarchic control system, and to argue for modelling instead).

The example situation can be reduced from my earlier description. Here's a simpler one that is easier to understand.

There is an ECS called Y whose output affects something in the outer world called B that provides a sensory input b, which Y controls as the percept  $y(b)$ . There is another ECS called Z whose loop goes through B and separately through D. Z controls for the percept  $z(b+d)$ . Take B and D to represent gain functions such that  $\delta(b) = B \delta(O_y + O_z)$  where  $O_y$  and  $O_z$  are the outputs of the two ECSs, and  $\delta(d) = D \delta(O_z)$ .

Even if the sign of D would lead to positive feedback, Z can be stable if the



sign of B is correct and  $B > D$ . Z can move b more with a given output than it moves d. But if Y is acting, REGARDLESS OF THE DIRECTION IN WHICH Y TRIES TO MOVE B, B becomes less susceptible to the actions of Z. The algebraic expressions get horribly complex, but we went through this a while back with one of Rick's studies that showed it to happen. So if the gain of Y is large enough with respect to the gain of Z, then Z can go unstable.

It is the effective gain of Y as applied to B that is important in this, not the direction in which Y moves B, cooperatively or in conflict with Z, just as Rick says. There can still be hidden positive feedback links in the hierarchy, but they can become manifest only when some gain changes, not just when some reference level shifts.

My Degree of Freedom argument (920701 0240) suggests that gains must be changeable, even if it is only in the binary sense that some ECS turns off to allow another to take control (please don't complain about the sloppy wording here; good wording takes pages). So the potential problem of masked positive feedback still exists in the hierarchy.

Models Rule, OK! Martin

Date: Tue Jul 21, 1992 2:47 pm PST  
Subject: Re: More catching up

[Martin Taylor 920721 17:30] (Bill Powers 920719.0800)

We were disconnected from the Internet for a couple of days, and have just been reconnected. Nothing I posted since (including) Saturday 18 July has shown up here, and I don't know whether it got out. I suspect it did, from some comments in what Bill and Rick wrote. If so, it probably bounced back to the Mail Server. I hope it returns some day, for my HyperCard stacks, if nothing else.

=====

>To lump the CNS hierarchy with the biochemical one is to ignore time-scales  
>and relationships to the immediate world of the senses. It is the CNS  
>hierarchy that produces the overt behavior of organisms and our own ability  
>to observe and make sense of that behavior.

Well, this gives me a perfect intro to a paper I was going to mention anyway: Behavioral Hypothermia and Survival of Hypoxic Protozoans *Paramecium caudatum*, by Gary M. Malvin and Stephen C. Wood, *Science*, 13 March 1992, 255, pp1423-1425.

Malvin and Wood start from the proposition that hypothermia is believed to be a protection against reduced oxygen conditions (hypoxia), and that many species there is either a behavioural response or a physiological response that cools the organism if it gets hypoxic. "In mammals, hypoxia induces hypothermia by decreasing heat production and increasing heat loss. In ectotherms [I think that means lizards, insects, and the like. MMT], hypoxia induces behavioral selection of a lower ambient temperature."  
(Right there is a suggestion that control can transfer between CNS and non-CNS modes over an evolutionary time scale.)

M&W wanted to see whether an organism without a nervous system would control its temperature differently as a function of depressed oxygen levels. They chose a single-celled protozoan. "Two specific hypotheses were tested (i) Hypoxia causes paramecia to select a lower temperature in a thermal gradient, (ii) survival of hypoxic paramecia is increased at lower temperatures." Both hypotheses proved true.

I do not see where your reference level comes from for segregating effects that act solely through the CNS from effects that occur by other media. CNS events can affect physiological states directly, and vice-versa. Why should the systems have to be considered as separate. Applying your own criteria of faithfulness to what is observed in life, I would have thought that separating CNS and biochemistry into two systems would have been difficult for you. The Paramecium exerts behavioural control of an intrinsic variable without a CNS. Pre-mammals presumably used behavioural regulation of temperature, as lizards now do. Mammals developed physiological means to do the same thing. The controlled percept is the same. Isn't it easier to imagine that we added a physiological ability to a behavioral ability that we continue to use, within a single hierarchy rather than to imagine that we developed a new ability in a separate hierarchy to duplicate an ability we had anyway?

>When you say that no controlled variables have meaning to the behaving  
>systems, you commit a peculiar reductionistic contortion, using knowledge  
>and meaning to deny knowledge and meaning. Isn't meaning one of the  
>phenomena we're trying to explain in terms of the organization of a complex  
>system? Don't higher systems perceive, and even think about, information  
>contained in the signals arising from lower-level systems, thus giving  
>meaning to these lower-level processes?

I suspect we are talking at cross-purposes here. I think of an ECS as dealing only with neural currents or similar measurable variables. Within the ECS they don't have meaning... "An ECS has to do what an ECS has gotta do." Whether the perceptual input functions of high-level ECSs themselves "perceive, and even think about, information contained in the signals arising from lower-level systems" is, I think irrelevant. Maybe they do, maybe they don't. What would it mean to the ECS if they did? It wouldn't change the result that a perceptual value must be compared to a reference value and the resulting error transformed into an output that is distributed the ECS knows not where.

> How is it that an intrinsic  
>variable, which represents the internal state of the organism, can lead to  
>effects that alter the details of the way the organism controls variables  
>external to the organism? This, not control of local conditions by local  
>action inside the organism, is the heart of the problem of reorganization.

I agree, except that I am not convinced that one has to go beyond control of local conditions to obtain the systemic result.

> I want to explain how this kind of  
>reorganization can work on any level in the hierarchy, regardless of the  
>cause of the error, so that a child can learn to add 2 and 2 as a way of  
>alleviating chronic physical discomfort, or can learn to exercise and build  
>up muscles as a way of doing the same thing.

Yep, that's what I am getting at. I can see how my system would do that, though I haven't modelled it, and my experience with the time-bomb in the hierarchy leads me not to trust intuition too far. But I can't see how yours would do it for the reasons already discussed. Any way you get out of the degrees of freedom problem (at least so far proposed) strikes me as Ptolemaic.

>I resist the pressure, which comes not just from you but from large  
>segments of modern science, to search for global generalizations that will  
>wrap up the grand principles of behavior as Einstein expressed the  
>principles of General Relativity in four simple (-looking) equations.

Maybe so, but I think you failed in this resistance some 40 years ago. I think that you have indeed produced a "grand principle of behaviour," perhaps the principal principle. All else is symmetry-breaking, and the specification of the specific symmetries that have been broken.

> True generalizations, I believe, arise from considering  
>details, not generalities. And even the best of them eventually fail and  
>must be modified.

Probably true. All scientific laws are provisional descriptions. But they are helpful descriptions unless they guide us away from even more helpful ones. Failure is relative.

=====

On my commentary with respect to Allan Randall's posting...

>>I suspect Allan was getting more at the question of whether a group of  
>>ECSSs at a level can act as a population vector ...  
>  
>Not in a way that specifically controls the vector rather than one variable  
>contributing to it (see my post of yesterday).  
>  
>>... "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.

Coarse coding is very robust against loss of a detector (even in the colour system, a person missing one of the systems still can work reasonably well with the other two), as well as giving fine possibilities for control. It is at the heart of my proposal that coordination be considered an emergent of the hierarchy--that no single ECS actually controls a specific percept at any level, though each acts as if it does.

This is different from the notion of a vector representing the multiple features of a percept, since in coarse coding the detectors all represent different central values of the same features. I have tried to distinguish three ways ECSs can relate within a level, based on their connections at lower levels: conflict, cooperation, and coordination. (I'm not sure I have done it on this mailing list, but it is in the current draft of the Paris paper, which is still not finished.)

>I can't visualize "recurrent reference connections." What does that mean?

I was thinking within a level, if the output of ECS A contributed to the reference level of ECS B (even indirectly) and vice-versa. Such connections do not exist within your scheme, but there are all these little niggly bits of data that suggest it might happen among percepts. If it did, then it should be expected among references, too. All I wanted to say was that we can't rule it out, even if we don't want to think about it yet.

Best for the Durango meeting. Durango was featured on the front page of the Travel section of the Toronto Globe and Mail on Saturday. Sounds like a nice place, if somewhat touristy. Wish I could come.

Martin

Date: Tue Jul 21, 1992 2:56 pm PST  
Subject: Re: No bomb in hierarchy

[Martin Taylor 920721 18:30] (Rick Marken 920720.1130)

>Actually, what I found in the simulation is that two control systems,  
>one with the wrong connection to one environmental variable, DO blow  
>up. I took this to mean that there was no way to "hide" a "positive  
>feedback system" in the hierarchy, as Martin suggested.

Actually, I think they needn't blow up. I don't know whether my message in response to your simulation result ever got distributed, but your result is "intuitively" correct. If Z is the ECS with the bad connection, and Y affects some variable also affected by Z, then the fact of Y controlling will tend to stabilize or "stiffen" that variable. Z will be acting more through any overlapping variables, increasing the contribution by the one with the positive feedback, and reducing Z's effective negative gain. If Y stiffens enough the overlapping variables, then Z will go into positive feedback. Whether this happens depends both on the gain of Y and the degree of overlap between Y and Z. Of course, if Y overlaps on the variable with the positive feedback, it will help Z to stabilize.

I think that there are parts of this situation that require a dynamic analysis rather than the DC analysis Bill produced. And yes, Bill, the expressions get horrendous. After getting Rick's result, I went through the algebra to get the result discussed above. Martin

Date: Wed Jul 22, 1992 4:47 am PST  
Subject: stick patterns

To: Pat Alfano From: David Goldstein  
Subject: re.: data analysis Date: 06/21/92

Your question raises, in a concrete form, the issues involved in how HPCT research differs from standard psychological research which Rick Marken and others have been discussing in general terms.

Bill Powers gave his answer. I will not try to improve upon it but just observe from his answer:

(1) The experience which a subject was controlling was identified in its most simple form. Bill identified it as the angle of a single stick. Pat, do you agree with this alteration of the task?

(2) The seating arrangement of the subject with respect to the experimenter was identified as a disturbance. Another one was the angle of the experimenter's stick. Others would call these independent variables. I note that the two disturbances identified have an obvious impact on the experience. In that the focus is on the control of experience, it is important to select disturbances which will, without a doubt, impact on the experience being studied. This is different than the standard approach in which the effectiveness of the independent variable is uncertain.

(3) Lots of data is collected from a single person. A person is intensely studied. The expectation is that the phenomenon will be apparent without statistics. I am not exactly sure what the

phenomenon is in your case. I guess it is that the ability of people to control in this task varies with the angle of the lines making up the pattern; diagonals present special problems to people. Is this the phenomenon?

I was surprised that Bill chose time to be the main measurement. I thought he would have chosen error, the difference between the subject's stick angle and the experimenter's stick angle. Maybe he can explain this choice.

(4) The task strategies which you identified are higher level experiences which the person is controlling in the course of performing the task. See Dick Robertson's comments about this in the textbook he and Bill edited.

Just as an aside, I did a reasearch project related to yours when I was a graduate student. There are interesting developmental changes in the ability of people to perform such tasks. This was in my pre-HPCT days.

Goldstein, D.M. & Wicklund, D.A. (1973). The acquisition of the diagonal concept. *Child Development*, 44, 210-213.

Best of luck with your research project.

Date: Wed Jul 22, 1992 7:34 am PST  
Subject: Simple bomb in one ECS

[Martin Taylor 920722 11:00]

(Too many references to quote, by myself, Bill and Rick)

It is interesting how understanding evolves. I originally proposed that a potential bomb existed in a system with 2 ECSs controlling their percepts through 4 CEVs (complex external variables), of which each controlled one independently and two in conjunction with the other ECS. One of the ECSs had a positive feedback loop through one of the CEVs, but this positive feedback was masked by the negative feedback through the other two CEVs, until the second ECS opposed the action of the first. Rick showed that this analysis was faulty, which led me to a deeper understanding that it is the gain of the second ECS that hides or reveals the bomb, not the direction in which it controls. The bomb would work with 2 ECSs working through 2 CEVs with one overlap.

I now see that the bomb can be demonstrated in a system with one ECS acting on two CEVs, provided that at least one of the CEVs has a non-linear impedance. I think this is the simplest bomb condition apart from one in which a single CEV moves into a positive feedback mode. That's a situation that can hardly be called masked positive feedback, which is what I am getting at.

Consider a CEV with output gain  $G$  controlling the percept  $x+y$ , where  $x$  is based on the CEV "X" and  $y$  on "Y". Disregarding disturbances on "X" and "Y" the value of  $x$  and  $y$  depend on the output  $O$  of the ECS.  $x=XO$  and  $y=yO$ . The percept  $p=O(X+Y)$ . Let us make the sign of the output such that positive

signs mean negative feedback (i.e. choose the comparator sign appropriately in the ECS). As described, everything is fine.

Now change the sign of the output relation to Y, so that  $y = -Y_0$  and  $p = O(X - Y)$ . Still everything is OK so long as  $X > Y$ . Now comes the bomb. Suppose that for some values of x,  $dx/dO$  (i.e. X) is large, whereas for other values it is small--the ECS has, for example, pushed an object off a slippery surface onto a sticky one. Then the ECS will control fine so long as x stays in the high compliance (large X) region, but will go into a positive feedback condition when "X" becomes stiffer (the object goes onto the sticky surface).

Nonlinearity in the external world can have the same effect as overlapping control, by reducing the compliance (increasing the impedance) of a CEV contributing to negative feedback, thus reducing the loop gain. If there is a CEV contributing positive feedback to the controlled percept, the loop may as a whole go into positive feedback.

In all the above, a CEV could equally well be a lower-level ECS, and the relation between the reference sent to it and the percept it returns is the X and Y in the above.

So, the bomb is there. It can be masked, and my original problem remains unsolved: What kind of developmental methods can avoid the construction of masked positive feedback loops? None of the proposed methods of reorganization seem to accomodate this sort of situation, since an ECS that maintains control is not going to contribute to the triggering of a reorganization episode under either Bill's scheme or mine. It is true that an ECS with masked positive feedback will have a lower gain than it would if the masked loop were reversed, and perhaps this can be used in some way to detect the existence of such problems. But unless the positive loops are unmasked, their effects will be very subtle, affecting mainly the precision and speed of control, not its success.

When the bomb goes off in one ECS, the situation changes for all ECSs to which it contributes a percept. For them, the situation is not as if an object had been pushed from a slippery surface to a sticky one, but more as if the object had acquired a jet engine to propel it the way it was being pushed. Their overall loop gains will be reduced and perhaps go positive, and we have a potential avalanche in which the front of stability moves up the hierarchy, just as the front of an avalanche moves up the snowfield or sand dune. Reorganization should then fix the problem, if the organism survives.

Martin

Date: Wed Jul 22, 1992 7:42 am PST  
Subject: Fancy nonlinearities; gain control

[From Bill Powers (920722.0800)]

Martin Taylor (920720.1615) --

>There is an ECS called Y whose output affects something in the outer  
>world called B that provides a sensory input b, which Y controls as the

>percept  $y(b)$ . There is another ECS called Z whose loop goes through B >and separately through D. Z controls for the percept  $z(b+d)$ . Take B >and D to represent gain functions such that  $\delta(b) = B \delta(Oy + \delta(Oz))$  where  $Oy$  and  $Oz$  are the outputs of the two ECSs, and  $\delta(d) = D \delta(Oz)$ .

Can't picture it. Can you set up the equations for the two ECS in an environment, or draw a diagram so I can do it? After my experience with clever algebraic solutions I think I want to see a simulation of this.

>My Degree of Freedom argument (920701 0240) suggests that gains must be >changeable, even if it is only in the binary sense that some ECS turns >off to allow another to take control

When you introduce gain changes, there are two ways to do it. One is to have a control system specifically concerned with sensing and controlling the gain of another system. The other is just to make the system nonlinear. The first method is tricky, in that you have to decide what perception the separate gain-controlling system is monitoring (how do you make gain a controlled variable?). But Ve Haf Vays.

There is a way of turning control systems off in a nervous system that is very simple. It starts with the realization that neural functions are always one-way -- that is, neural signals can't go negative. In the standard diagram, we have  $\text{error} = \text{reference} - \text{perception}$ . This means that the perceptual signal is inhibitory, the reference signal excitatory. If the reference signal is simply set to zero, there is no excitation of the comparison neuron, and no amount of inhibitory perceptual signal will ever make it fire. So this comparator will produce zero error signal if the reference signal is zero, regardless of the amount of the perceptual signal. The control system is turned off.

To get two-way control about zero, a pair of control systems is always required in the nervous system. The pair of systems treats opposite directions of change of the perception as positive, and the error signals in the paired systems have opposite effects on the controlled variable. The simplest example is a pair of opposed muscles and their associated spinal control systems for controlling force. If the arm exerts a leftward force, the left-controlling control systems sense and control a positive force to the left. If the force is to the right, the right-controlling control systems sense and control a positive force to the right. This much you'll find in BCP.

To think of this pair of control systems as a single control system that can exert a continuum of forces passing through zero, we must think of the reference signals as a balanced pair. If the rightward reference signal is nonzero, there is a force to the right. As this reference signal declines toward zero, the rightward force declines toward zero. Then, just as the rightward reference signal reaches zero, the leftward one begins to rise, and the force begins to rise toward the left.

If both the rightward and leftward reference signals are zero, this pair of systems is turned off. A disturbance may cause an inhibitory feedback signal to arise, but because there is no excitatory reference signal reaching either the leftward or rightward comparators, there will be no error signal to drive either of the pair of outputs. The system will not



resist disturbances in either direction.

In order to achieve control of an arm in the state of zero net force, it's necessary to add a common-mode signal to the pair of reference signals. Now the "resting" state is that in which both control systems contain error signals, causing the muscles to pull against each other. The net left-right force is zero, but any force disturbance will cause one error signal to decrease and the other to increase, so there is control in the vicinity of zero NET force. Both control systems in the pair are receiving nonzero reference signals now, with only the difference in magnitudes showing up as a net left or right force.

The common-mode force is, of course, called muscle tone. A control system that controls for muscle tone controls to keep the SUM of the two positive force signals at a constant level. A second control system can then control to keep the DIFFERENCE between the same two force-signals at another reference level, which sets the net sideward force to left or right. The difference-controlling system has to emit a pair of output signals that vary in a complementary way; in fact it must also have a balanced pair of comparators in order to handle positive and negative errors. The higher-level muscle-tone control system can be a one-way system, because the sum of the muscle tensions can never be less than zero.

In our work on the arm model, Greg Williams found a reference that provided force-length curves for various muscles. These curves can be fitted quite closely with a second-power function over most of the force range (below the saturation level of tension). Muscle tension is produced by stretching the passive component of the spring, so muscle tension goes very nearly as the square of the driving signal and the amount of contraction.

When you oppose two such muscles, the net force as a function of length is represented as the difference between two offset square functions. Let  $c$  = common-mode contraction, and  $d$  = differential contraction. Then

$$F1 = (c + d/2)^2 \text{ and}$$

$$F2 = (c - d/2)^2$$

As a result, we have

$$F1 - F2 = (c^2 + 2cd/2 + d^2/4) - (c^2 - 2cd/2 + d^2/4) \text{ or}$$

$$F1 - F2 = 2cd.$$

This says that the differential force produced by a differential contraction is proportional to the common-mode contraction: that is, the output sensitivity of this force-generator is so determined. If the rest of the system is linear, the loop gain of this force-control system is linearly proportional to muscle tone, and the differential force at constant muscle tone is a LINEAR function of the differential contraction in the two muscles (until one muscle or the other totally relaxes).

This is why there is no control when you totally relax all your arm muscles, which means setting muscle tone to zero. An external disturbance will not produce any reaction; the arm will just give way and swing like a

dead fish.

In order to get control, you must raise the muscle tone from zero, so there is some mutually-opposing force. Up to a point, the greater the muscle tone, the higher the loop gain gets. That's why you tense your muscles when you have to do something delicate. But if you tense them too much, the loop gain will get too high, and you'll lose fine control again as system noise gets amplified and also as dynamic instability approaches.

I think this principle of gain control may apply generally in neural control systems. In one-way control systems, gain is zero when reference signals are near zero, becoming high only when signals are increased into a more linear part of the input-output functions. To get tight control of neural signals near zero, the reference signals and control systems must be present as balanced pairs, with a common-mode signal determining gain and a difference-signal doing the controlling. Whether single-ended or double-ended, a system is turned off when ALL reference signals associated with it are set to zero.

In a balanced system, "zero reference signal" really means EQUAL reference signals in a balanced pair of systems. In an unbalanced (one-way) system, zero reference signal is zero reference signal.

-----  
David Goldstein (920721) --

>I was surprised that Bill chose time to be the main measurement.

I didn't; Pat did.

Best to all, Bill P.

Date: Wed Jul 22, 1992 8:04 am PST  
Subject: Thanks, Gary

from Ed Ford (920722:0847)

Gary,

I think the idea of listing all the portable demonstrations of perceptual control theory is really great. Those things really do help and I for one could use them.

Also, the system I'm on through ASU here in Phoenix has always sent to me through E-mail what I've sent out. There are several advantages. First, it tells me that what I sent out was received (although my system tells me that as well) and secondly, it shows me how quickly what I sent got out.

Finally, just a note of appreciation for you continued work in keeping the CSGnet functioning. Thanks. Ed.

Date: Wed Jul 22, 1992 8:21 am PST  
Subject: The new IEEE Transactions on Control Systems Technology

[ The following is a cross-post from the E-LETTER on Systems, Control,

and Signal Processing ISSUE No. 52, PART 1, 15 July 1992. You are encouraged to subscribe by sending mail to Bradley W. Dickinson at bradley@princeton.edu or bradley@pucc.bitnet - Moderator ]

Manuscripts are solicited for the IEEE Transactions on Control Systems Technology, a new journal of the IEEE Control Systems Society to be published quarterly beginning March 1993.

#### Scope

The IEEE Transactions on Control Systems Technology will publish papers on new developments in all areas of control systems technology, including, but not limited to, new sensor and actuator technologies, software and hardware for real-time computing and signal processing in control systems, tools for computer-aided design of control systems, new approaches to control system design and implementation, experimental results, distributed architectures, intelligent control, and novel applications of control engineering methods. Survey and tutorial articles on timely topics of general interest will also be considered for publication. The Letters section provides for the rapid publication of brief reports on new research results and technical developments, and comments on previous papers.

#### ASSOCIATE EDITORS

X. Cao, Digital Equipment Corporation, USA  
D. Clarke, Oxford University, ENGLAND  
J. Cloutier, USAF Armament Lab, USA  
E. Collins, Harris Corporation, USA  
P. Dersin, GEC Alsthom, FRANCE  
B. Egardt, Goteborg University, SWEDEN  
S. Engell, Dortmund University, GERMANY (Co-Editor)  
H. Geering, ETH Zurich, SWITZERLAND  
E. King, ALCOA, USA  
S. Kumagai, Osaka University, JAPAN  
K. Lorell, Lockheed, USA  
J. Maciejowski, Cambridge University, ENGLAND  
R. Ravi, General Electric Co., USA  
D. Repperger, Wright Patterson AFB, USA  
M. Spong, University of Illinois-Urbana, USA  
M. Steinbuch, Philips Labs, THE NETHERLANDS  
G. Suski, Lawrence Livermore Labs, USA  
S. Williams, Cambridge Control Ltd., ENGLAND  
J. Winkelman, Ford Motor Co., USA  
N. Yoshitani, Naoharu, Nippon Steel Co. Ltd., JAPAN

#### ADVISORY BOARD

K. Astrom, Lund Institute of Technology, SWEDEN  
S. Kahne, MITRE Corporation, USA  
P. Kokotovic, University of California-Santa Barbara, USA  
M. Masten, Texas Instruments, USA  
G. Schmidt, Technical University of Munich, GERMANY  
L. Sweet, Asea Brown Boveri, USA  
T. Ueyama, Nippon Steel Co. JAPAN

Information for Authors

Three types of contributions will be considered: Papers, Brief Papers, and Letters. Papers and Brief Papers go through the same review process. Letters go through a shorter review process to facilitate rapid publication.

Manuscripts must be double-spaced. The first page should include the title, the names and affiliations of all authors, an indication of the corresponding author, and a one-paragraph abstract which briefly describes the contribution of the paper. Papers should be no longer than 32 double-spaced pages, including figures. Brief Papers should be no longer than 16 double-spaced pages, including figures. Letters should be no longer than 8 double-spaced pages of text, plus figures. In general, manuscripts should follow the standards of the IEEE as described in Information for IEEE Transactions and Journal Authors, available on request from the IEEE Publications Department, 345 East 47th Street, New York, NY 10017-2394, USA.

Seven copies of the complete manuscript with a cover letter stating the type of contribution (Paper, Brief Paper, or Letter) and the name and address of the corresponding author should be sent to one of the following editorial offices:

Bruce H. Krogh, Editor  
IEEE Transactions on Control Systems Technology  
Department of Electrical and Computer Engineering  
Carnegie Mellon University Pittsburgh, PA, 15213-3890  
USA

OR

Sebastian Engell, Co-Editor  
IEEE Transactions on Control Systems Technology  
FB Chemietechnik  
University of Dortmund  
Postfach 50 05 00  
D-4600 Dortmund 50  
GERMANY

Submissions should be sent to the editorial office which is most convenient to minimize time and cost of postage. Questions should be directed to Bruce H. Krogh at the above address, or computer mail: krogh@galley.ece.cmu.edu; phone: 412 268 2472; or fax: 412 268 3890.

Date: Wed Jul 22, 1992 8:41 am PST  
Subject: Simulating a bomb

[From Bill Powers (920722.0930)]

Martin Taylor (920722.1100) --

>Consider a CEV with output gain G controlling the percept  $x+y$ , where  $x$   
>is based on the CEV "X" and  $y$  on "Y". Disregarding disturbances on "X"  
>and "Y" the value of  $x$  and  $y$  depend on the output  $O$  of the ECS.  $x=XO$   
>and  $y=yO$ . The percept  $p=O(X+Y)$ . Let us make the sign of the output  
>such that positive signs mean negative feedback (i.e. choose the  
>comparator sign appropriately in the ECS). As described, everything is fine.

I interpret this to mean:

```
Initialize: X = 0; Y = 0; O = 0; G = 0.01;
           r = 100;
do
{
x = X; y = Y;
p = x + y;
e = r - p;
O = O + G*e;      /* integrating output function */
X = 0; Y = 0;
printf(variables)
}
while (!kbhit());
```

>Now change the sign of the output relation to Y, so that  $y=-Y_0$  and  $p=O(X-Y)$ . ...

>Suppose that for some values of x,  $dx/dO$  (i.e. X) is large, whereas for other values it is small

I interpret these changes to mean:

```
Initialize: X = 0; Y = 0; O = 0; G = 0.01;
           r = 100;

for(k=0.25; k <= 2.0; k += 0.25)
do
{
x = X; Y = y;
p = x - y;
e = r - p;
O = O + G*e;      /* integrating output function */
X = k*O; Y = -O;
printf(variables)
}
while (!kbhit());
```

>Still everything is OK so long as  $X>Y$ . Now comes the bomb.

>Suppose that for some values of x,  $dx/dO$  (i.e. X) is large, whereas for other values it is small--the ECS has, for example, pushed an object off a slippery surface onto a sticky one. Then the ECS will control fine so long as x stays in the high compliance (large X) region, but will go into a positive feedback condition when "X" becomes stiffer (the object goes onto the sticky surface).

Is this what actually happens? I'm not going to tell you. You tell me. If my interpretations above are OK, make the programs runnable and run them and let me know what happens. Otherwise, substitute your own program, run it, and let me know what the program is and what happens.

You don't know what will happen until you simulate the situation. Are you right or wrong? Find out.

Best, Bill P.

Date: Wed Jul 22, 1992 9:07 am PST  
Subject: Bernhard on fuzzy control

Apologies to those on the cybsys list who have already seen this.

Bruce bn@bbn.com

[ The following is a cross-post from the E-LETTER on Systems, Control, and Signal Processing ISSUE No. 52, PART 1, 15 July 1992. You are encouraged to subscribe by sending mail to Bradley W. Dickinson at bradley@princeton.edu or bradley@pucc.bitnet - Moderator ]

Pierre Bernhard, INRIA Sophia Antipolis, France, May 1992

--editor's note:

This is reprinted, with permission and slightly updated, from the European Control Newsletter. We thought it would be of great interest to our readership. Short replies can be sent to the Eletter editors and will be posted in the next issue.

#### FOREWORD

This is a slightly updated version of an older memo in French, which was never intended to be published in a French journal, let aside in a European one. The idea was rather to settle my mind, and have an answer ready to the very many requests I recieved about fuzzy control, mainly due to the abundant advertisement it enjoyed in the non technical press. A few things I wrote about where fuzzy control is being applied are not compltely true anymore. But I believe that globally the idea remains correct.

The original version bared a foreword acknowledging the help of Jean-Marie Nicolas and Michel Grabisch, both of Thomson-Sintra, France.

#### FRAMEWORK AND LIMITS

The general theory of "fuzzy" logic currently enjoys a rapid developpement with many applications, specially in Japan. What I write here is narrowly confined to fuzzy control . This is only one of the many applications, although often advertised as the most prominent one. It is in no way the only one. I know, and say, nothing about applications to such things as knowledge representation (which was the original motive behind fuzzy set theory), expert systems and the like.

#### 1) FUZZY CONTROL IN JAPAN

The basis of fuzzy control is to express a control law in terms of expert rules. The rules define the control value, or its rate of change, for some (range of) values of the measured variables or their rate of change. The specific techniques of fuzzy set theory can be seen as a systematic way of interpolating the data points.

The language used is one of sequential decisions, and as such is always applied to control problems which are fundamentally conditional sequencing problems, and where the continuous control part is completely elementary. It

is symptomatic that the yardstick used to judge the efficiency of this control is always the PID. Take the often quoted example of a bathtub hot/cold water mixer. It takes into account the fact that the water that first flows when one opens the hot tap is cold, and therefore reaches the desired temperature faster than a fixed gain PID. A "success" of fuzzy control.

In its original form at least, fuzzy control shares the ideology of expert systems to automatise what an expert knows how to do, not to do things no human expert can do. The motive of research in fuzzy control is therefore not to push back the limits of what automatic control can perform, even less to prove things about the performance of a control mechanism, such as stability, optimality, sensitivity. As in expert systems, experimentation is the means of validation.

The single stick balancing problem is also often quoted as test case. I consider it unfair to fuzzy control. As a matter of fact, it is a simple problem, with no sequencing involved. As a consequence, for a single boom, adjusting the coefficients of a PID that would do the job is much faster than using fuzzy control, and for the double boom with no measurement of the upper boom's angle with the lower one, an human expert cannot do it, nor fuzzy control either.

I think fuzzy control is a good tool where it applies, and I shall come back to that point in the next section. However it has been oversold on unjustified grounds, which obliges us to review some of the claims made.

-1) Gentleness. "Because it is fuzzy, fuzzy control is more gentle to the user than classical control which, for lack of fuzziness is by its essence bang bang". Do not laugh, this has often been said. It impresses the ignorants and the newsmen. The people who said that may have been themselves more ignorant of what control is than outright dishonest.

-2) Ease of implementation. This requires a more careful examination. The proponents of fuzzy control acknowledge that there are very many parameters to choose to setup such a control law. If the comparison item is PID, then the later is clearly easier to implement. If the comparison item is a problem that the PID would not solve (or a PID with, say, cubic terms added to it), then one has to look at the boundary of the possibilities of fuzzy control. And the simplicity is gone. (It requires something like 49 rules to balance a single stick while maintaining control of its translation). As a matter of fact, the very idea of what is simple depends very much on one's educational background. What is true is that fuzzy control lets one solve control problems with no mathematical education whatsoever. Where a more fundamental simplicity comes in is when the overall problem contains both conditional sequencing and simple continuous control. Again we shall come back to that.

-3) Robustness. I have seen no publication that scientifically substantiates the claim of greater robustness of fuzzy control as compared to modern control, nor any that disproves it for that matter.

-4) Lower computational requirements. This I consider as a false claim. The method of interpolation used is computer intensive (all rules are continuously evaluated and their conclusions weighted according to their degree of truth in a sophisticated way). What is true is that this is of no real importance, because thanks to specialized chips, it is cheaply done.

A definite weakness of this approach is that the inherent complexity of the interpolation process induced makes it essentially impossible to prove anything about the control laws generated. Anyhow, this proof would not be in the spirit of the method: the human controller does not "prove" his know-how either.

Let us quote the three reasons Dr Sugueno (scientific director of Laboratory for International Fuzzy Engineering) gives for the success of fuzzy control in Japan:

- i) The careful choice of the applications
- ii) The quality and the efficiency of Japanese engineers
- iii) The good fit with Japanese way of thinking

We leave it to the reader to interpret these explanations. The last one should not be underestimated, coupled with an "invented here" syndrome, in a more nationalistic society than ours.

One could deduce from the above that there is little more than a regression from mathematical analysis to empirical imitation of the human operator, and disregard the whole story. I believe that this would miss the point.

## 2) THE EUROPEAN RESPONSE

The challenge is less scientific than industrial. It is threefold.

The first striking fact is the wide range of elementary applications that have been widely quoted as success stories for fuzzy control. The good idea there is not to have included a fuzzy digital controller, it is to have included a digital controller. Japanese industry has been the first to understand that digital devices are from now on cheap and reliable, and to draw the practical consequences, that they can be put to use in cheap home appliances and other apparatus.

The response of Europe here should be to encourage our industry to use digital devices more extensively to improve consumer products.

A second remark is that qualifying simple control problems as "research" (since fuzzy control was new) has given the Japanese university scientists an opportunity to discover the pragmatic questions that standard industry had to face. What they discovered were problems where the practical difficulty to use commercially available tools was to make coexist simple continuous time controls with complicated sequencing tasks. What fuzzy control brought them was a single language to describe both, in terms of expert rules.

A European response might build upon the clear European lead in synchronous programming. But then such tools as the new real time languages (ESTEREL, SIGNAL, LUSTRE, to quote the three that cooperate in France) should be carefully hidden to the user, deeply buried in a system providing an elementary interface, devised to let the user solve elementary control problems of that type, with little control knowledge.



The genial feature of the Japanese fuzzy control culture has been to bring a tool well suited to their engineers (often with less control engineering education than their European counterpart) to solve simple problems. (And fuzzy control has been a good excuse, because it is unable to solve advanced, multivariable, control problems).

There is a niche for fuzzy control, or any tool sharing the peculiarities we described, (and better ones might be devised : fuzziness is not unavoidable in that respect. The real important feature is rather rule based control ) that we would be foolish to ignore, mainly since larger economic dividends may be at stake with simple problems than with advanced ones.

### 3) CONCLUSION: INDUSTRIAL ISSUES

The formidable advertisement that fuzzy control has enjoyed in the (mainly non technical) literature is of course not devoid of commercial aims. This is not the place to analyze them in details. Let us just recall that since consumer products are concerned, the non technical press was indeed the place where this commercial drive had to be carried out. Later will come the market for the specialized chips.

Finally, my friends in industry drew my attention to a last point which is probably not the least important one. This very article serves the purpose of entrenching the idea that there is a completely new theory behind fuzzy control, since it is being debated in scientific circles and universities, in Japan first and now in the US and Europe. If this is a completely new theory, nothing that is constructed referring to it can fall under old patents. Therefore, Japanese industry (or, for that matter, any industry clever enough to seize that opportunity) is instantly freed from all previous patents. It is straightforward to program (approximately) a PID controller with saturation using fuzzy control. Because it will be a fuzzy controller, it cannot be challenged by an old patent. And of course this is true of many other devices.

This is a matter for industry to address, not academia.

Date: Wed Jul 22, 1992 3:29 pm PST  
Subject: S

Hello, Martin --

I have an ulterior motive in throwing that problem back into your lap. Right now the CSG has only three people actively developing simulations and doing HPCT experiments: Rick Marken, Tom Bourbon, and me. That isn't enough. Rick and Tom both began, when they started working with me, with only a rudimentary understanding of BASIC and no real-time simulation experience at all. Both of them have since become quite expert in this field and no longer need me to show the way.

You are probably farther ahead than either of them were, by a long distance, in all regards. But I'm going to be frank with you: you have a tendency to sit on your butt and try to figure things out in general and in the abstract without carrying the ideas to the stage of actual testing. I would love to get you out from behind that desk and into some programming of real experiments and simulations. You would obviously be extremely good

at it. You are already very proficient at analysis; there would be less cut and try for you than there is for me. I expect that you already do, or at least supervise, a lot of this kind of experimentation and stimulation -- but not in HPCT.

I understand what you mean about simulations vs. analysis. But an analysis contains far more in it than you will ever notice (even when it's done right). The analysis stage is extremely useful for looking at sensitivities, singularities, extremums, and so on. I use it a lot that way, within my capacities. But a simulation of something you think you already understand through analysis, I guarantee you, will produce both surprises and shocks -- and sometimes delight. Analysis can't give you a grasp of control theory anything like what you can get out of writing and running even a SIMPLE simulation of a SIMPLE control process. There is nothing in analytical forms to suggest solving for different variables, re-expressing the forms in ways that reveal new relationships, and so on. No matter how much you know already, a simulation will teach you more.

I hope you'll brush up your programming techniques by trying some simple control system simulations before you get into the complexities; do some five-finger exercises first. I'm here to help.

I've had more than one excited phone call from Rick and Tom over the years, saying "You won't believe this, but this stuff ACTUALLY WORKS!" Join the party.

Best, Bill

PS: I wish you were coming to Durango, too.

Date: Thu Jul 23, 1992 7:23 am PST  
Subject: reorganizing

[From: Bruce Nevin (Wed 920722 13:24:12)]

The above time stamp is wildly bogus with respect to the prior and subsequent times of actually writing this. I'm stealing snatches of time whenever my workstation loads files, saves files, prints, etc. I hope coherence doesn't suffer too much.

(Bill Powers (920718.1600) ) --

>As I had figured, there are some dissenters from my concept of the  
>hierarchy, and my discussion of implicit versus explicit functions turned  
>up some more.

I'm not dissenting from your concept of the hierarchy, I didn't understand your distinction between implicit and explicit functions and how it applied to what I was exploring. Now I do, and it doesn't.

I was not talking about

(Bill Powers (920713.1200) )  
>the emergence of higher levels

>from populations of systems of an existing level, without the addition of  
>any new kinds of physical systems.

I was talking about the emergence of a new order of control hierarchy from populations of control systems of an existing order. Whether this involves the addition of new capacities or functions within control systems of order n (e.g. cells) as they evolve in their function as constituents of control elements (e.g. neurological ECSs) within control systems of order n+1 (e.g. neurological control systems, such as those of humans and cockroaches) is an interesting question, and perhaps I was wrong not to have raised it, but I didn't.

Bill Powers (920719.0800) --

> It's not fair to bring up things your wife does as examples of real  
> behavior, because anyone (like me) who doubts that these phenomena exist  
> outside the imagination of believers is put in the position of criticizing  
> your wife -- and you are put in the position of defending an idea in which  
> you have a personal emotional investment and a loved one to defend. I think  
> that the human brain is capable of supplying itself with any experiences it  
> has reason to want, even with a vividness and to an extent that normally  
> would be dangerous because of substituting too much imagined information  
> for real-time sensory information. My opinion of psychic research is that  
> it is sloppy, credulous, selective, and often dishonest. If you want to get  
> into this subject that's up to you. But the result will not be an  
> enhancement of our sense of intellectual companionship.

I introduced the topic of "auras" only make the notion of an analog to neural currents somewhat less abstract. Insofar as some humans can see them, the analogy of "auric changes" to neural current is a poor one. My first recasting of Rick's paragraph to show the analogy (which I did not post) used the "fu" of "fubar," somewhat like this:

> A person is busy controlling many variables. The systems controlling  
> these variables are made out of cells, neural structures, and organs of  
> perception and execution (probably there's a better word, but I'm in a  
> rush). But one thing a human does is change its fu. That is, the  
> variables involved in this functional relationship . . . are not (as  
> far as I know) perceived and controlled by the human. It is this  
> input-output characteristic of the human's fu behavior that makes it a  
> useful component of a control system. The human responds to changes in a  
> neighboring human's fu by changes in its own fu. This is a "dedicated"  
> cause-effect characteristic of the human; the human cannot change the  
> way it responds (fu) to input (fu) -- there is no control involved in  
> this functional relationship; that is what I mean by functional  
> specificity. In terms of it's fu response to fu stimulation of the body  
> the human functions like a wire in a circuit ( with fu change rate the  
> analog of current and intensity of fu the analog of voltage). A control  
> system must be built out of such "functionally specific" components.

This seemed needlessly abstract. Aside from being more specific and less mysterious, putting it in terms of auras had a nice fit precisely because most people have no awareness of them, as cells necessarily must have no perception of neural currents as such if they are to function as neural "wires." Even the very strong prejudice in our culture against

taking auras seriously or even discussing them led, I thought in an interesting way, to the question, does Darwinian selection suffice to keep cells in the dark about the neural currents they are transmitting, or might some additional factors be involved?

I have nothing vested in the choice of aura over fu as a term in the larger analogy. By saying it's not "fair" of me to bring such things up, you indicate (I think) that you feel a need to forebear from expressing your doubts, so as to avoid putting me in a position of defending my wife etc. Why don't we set all that baggage aside. It is that sort of fencing off of forbidden topics, to my mind, that interferes with "intellectual companionship." If you doubt the reality of a perception of mine, or reported by another person (my wife or anyone else), please feel free to say so. Given the indeterminately extensive role of imagination in ordinary perception, and the unsettled controversy over epistemological implications of control theory, that would put you on rather unsure footing, but you are welcome to occupy it. I most probably would not be too much provoked by your doubt. My intellectual M.O. is integrative, and I value CT in part because it supports the integration of things that so far as I can tell had been rejected only because they did not fit official canons. You saw stones fall from the sky? Nonsense! Ptolemy has shown that to be impossible. But I saw them fall--something more than Ptolemy must be going on. With this constitutive relationship between hierarchies, CT may come to provide an explanation for some aspects of "psi" perceptions. Perhaps psi is all imagination (a great deal of it surely is). But perhaps something more is going on.

Incidentally, psychic research may be so dreadful in part because it tries to apply conventional methodology that is inappropriate even for perceptions that people can control well to perceptions that people control less well if at all. (Just as the cell must not control neural currents as a cellular perception if it is to transmit them in an ECS.)

But let's get away from areas that provoke allergies.

What is important in all this, to my mind, is some possible insight into the character of the transition from one order of control to another, and into the character of reorganization (including learning, development, and evolution). You argue for the existence of distinct but in some way contiguous hierarchies, and that seems right to me. You have held in reserve a "distributed" model of reorganization that appears to accord with what I have been suggesting. Therefore, I was surprised at your response:

> >I am proposing (920709 09:13:52) that reorganization is carried out in  
> >populations of entities of order n-1.

> Propose away. But I ask that you specify all the functions and variables  
> needed for this reorganization to happen, in addition to those needed for  
> the control processes already going on. You've indicated an overall result  
> that you want; it remains to demonstrate that such a result is possible, by  
> some means that we can grasp. What does each part of this new reorganizing  
> process have to accomplish, sense, do ... ? And what is your reason for  
> believing that what such a system will do will actually solve the problem?

I proposed that the mechanism resides in the responses that cells make to consequences, for each cell in its intracellular environment, of chronic error in the higher-order control system that the cells constitute. The error itself is as invisible to them as the conflict that engenders it, and is as invisible to them as

the elements of the higher-order control system (in conflict) that they constitute; only the consequences of error in their intracellular environment are in their perceptual universe. Roughly speaking (perhaps you see me as always and only speaking roughly, that is, without adequate discipline), I suggest that higher-order error has as a byproduct lower-order environmental "toxicity" of some kind, something that individual cells control for.

Individual cells, like all control systems, do whatever it takes to reduce the error (the "toxicity"). They change their shape and proximity to other cells (tropism), etc. As it happens, however, only when they do something that reduces the higher-order error for the higher-order control system undergoing reorganization does their environmental "toxicity" go down.

In cellular terms, they seek only to reduce unwanted cellular perceptions. Byproducts of this seeking may include the appearance (to the observer) of cellular behavior that exploits and enhances environmental stability, that promotes stable inter-cellular relationships and structures (with greater environmental stability as a payoff). But whatever reduces error in the higher-order system reduces toxicity in the intracellular environment. The toxicity may be relatively local or distributed, depending on the higher-order elements (ECSSs) involved in conflict. Therefore (here, I contradict you), the appearance of reorganization from the higher-order perspective may be more than local, though each cell's motivation and action is of course local.

Distasteful though they may be, there are obvious analogies to human social unrest and distress.

Another positive product of the analogy is that it may help put to rest certain kinds of recurrent discourse about social control. The argument is that there can be no hierarchical social control in which humans directly receive reference signals from superordinate control systems. Any supra-human control system must be of a different order, a separate hierarchy from that of neurologically based perceptual control systems, an order of control systems which can "communicate" with humans and groups of humans only in ways that are perceivable by control systems of the neurological order as environmental conditions, including especially their perceptions of their relations to their peers.

It may be that humans (unlike cells, so far as we know) can learn to participate more consciously in such a suprahuman organism (if it exists), or that other complex living control systems of human scale in the universe have so learned. That would be an interesting approach indeed to questions of ethics and social order. However, it seems beyond us here.

Finally, this conception of different orders or hierarchies, and of reorganization, is capable of being modelled. Only that would properly ground and "discipline" the discussion.

Gotta run. Bruce bn@bbn.com

Date: Thu Jul 23, 1992 9:26 am PST  
Subject: Closed Loop

from Ed Ford (920723:0935)

To All -

Closed Loop Summer 1992, Volume 2, Number 3, has been printed. It is entitled Statistics vs. Generative Models. As I said a few days ago, these will be distributed (thus reducing our postage costs) at the conference and Mary Powers will be given additional copies for non-members wanting to buy this edition or prior copies. The cost per copy is \$6.00 which includes shipping costs. The other editions available from this year are: Winter 1992, Volume 2, Number 1, Social Control; and Spring 1992, Volume 2, Number 2, Epistemology. Closed Loop can be obtained for \$6 per copy by sending your check payable to: Control Systems Group and mailing it to our business office (Mary Powers) at: 73 Ridge Road, CR 510, Durango, CO 81301.

Ed Ford

Date: Thu Jul 23, 1992 7:16 pm PST

Subject: What is information?

[Allan Randall (920723.2100)]

From Bill's ISSS talk on information:

< Information: a matter of perception  
< William T. Powers  
<...</p></div>

<Shannon and Weaver examined a physical signal flowing from  
<source to destination...information decreases the entropy of the  
<receiver...Unfortunately, this is a house of cards and one of  
<the bottom cards is imaginary.</p></div>

Ouch! Well I wasn't going to post anything else here until I found the time to respond to the implicit/explicit hierarchy thread, but this I can't let pass. The entropic formulation of information theory, along with the related algorithmic formulation, has a pretty firm mathematical basis, and has shown itself to be quite useful, so you will be going some to convince me it is a "house of cards."

<Information does not necessarily travel in the same direction as  
<physical energy.</p></div>

I think you are confusing the concept of energy and that of entropy. They are related, but not the same. It is the latter, not the former, that information theorists associate with information content. There is no need in traditional information theory for the kind of "net flow" of energy in the direction from source to destination that you talk about in your examples, both of which I think are pretty easy to refute.

<If you're reading these words in black print on a piece of white  
<paper, you can see immediately what I mean. The energy that flows  
<into your eye...comes from the reflection of light off the white  
<paper. This energy flows from the page to your retina everywhere  
<except where the letters are printed. If the letters are dark  
<enough, a net energy flow may actually travel from diffusely  
<illuminated points on your retina, through the lens, and to the  
<light-absorptive ink on the paper.</p></div>

Okay. Correct me if I'm wrong. Are you suggesting there is a net

flow from my retina to the words on paper, rather than vice-versa? The "words" on the page are more than just the ink, they include the white background. The energy from the white background is what transmits the information to my brain. Just because it is the ground rather than the figure that I receive as energy is irrelevant. The energy from the white page is structured in such a way as to correspond to the ink markings (after all, that's why I'm able to "reconstruct" the black ink markings from the light I receive from the white background). This is all perfectly in accord with traditional information theory. There is no flow of energy in the "wrong" direction.

<Or consider an old-fashioned telegraph...When the key is open, no <information is traveling. When the key closes in long and <short patterns, a message in Morse Code is sent to Chicago. The <dots and dashes represent momentary drains of energy on the <battery in the central telegraph office in Chicago; the average <direction of energy flow is from the Chicago office to the short <circuit in Dodge City, heating the wires along the way.

Although I disagree with this as well, it is a better example than the first one, since here there really does seem at first glance to be a net flow of energy in the "wrong" direction. The mistake you are making here is using the energy output of the battery as the transmitting energy flow. This is incorrect. You are treating the battery as the information transmitter. But the battery is NOT the originator of the message. It matters not a wit whether we consider the battery to be part of the sender or the receiver. The battery is thus more justifiably considered as part of the medium of transmission. The message actually comes from the human being who is putting out the dots and dashes. This \*is\* a flow of energy from the human, and \*does\* decrease the entropy of the receiver and increase the entropy of the source (and the universe).

Compare what happens to the case of a transmitter that outputs dots and dashes due to chaotic or random forces in the world around it. These messages are less ordered, and thus higher entropy, than the messages put out by the human. The destination would end up with higher entropy and the source would have lower entropy than in the case of the human operator.

<...the direction of energy flow is unrelated to the direction of <flow of information, so that concepts like entropy (positive or <negative) have no actual physical relationship to whatever it is <that we mean by information.

Perhaps not what \*you\* mean by information, but the entropy concept is very rigorously defined mathematically, and \*is\* what most information theorists, computer scientists and physicists mean when we speak of information. Shannon himself was very clear in his original article that he was not giving a "semantic" definition of information, or "meaning". He quite frankly admitted that this was not covered by the theory he laid out. It is not entirely clear to what degree a semantic or a perceptual theory of information

would be built on the concepts of traditional information theory a la Shannon/Weaver/Chaitin, which does not pretend to speak to such issues.

<When we transmit information, we hope we are transmitting more  
<than words; we hope we are transmitting meanings.

I have no problem with your tying in perception/control with this notion of information, and if this is what you want to call "information," I'm not going to argue over definitions. But what Shannon and Weaver showed was that there is a key aspect of communication, which is now usually called information, that is independant of this "meaning" or semantic content and has nothing to do with perception. I think they succeeded. I don't think you have, as yet, given me any reason to doubt this. Your rubber band experiment merely shows one example of a case where control is necessary for information to be transmitted. It says nothing about whether such control is necessary for information transmission in general. I don't think it is. Give me reason to believe otherwise.

Allan Randall  
NTT Systems, Inc.  
Toronto, ON

PS: I'm actually a lot more favourable to PCT than I appear in my posts. I find it to be elegant with a lot of explanatory power. Its just that I'm, as Bill would say, controlling for the higher error signals.

Date: Thu Jul 23, 1992 8:31 pm PST  
From: Dag Forssell / MCI ID: 474-2580  
Subject: Seminar

[From Dag Forssell (920723)]

>Date: Mon Jul 20, 1992 1:10 pm PST  
>Subject: PCT Seminar

>

>Hi! Professor Robertson gave me a copy of your version 9 letter on PCT  
>and I found it interesting. I will graduate in December 1992 under the  
>BOG Program and with a major in Psych. Control theory is basic to my  
>interest in the field and I will be needing something to do soon. I  
>have many years of business management including the conduct of seminars  
>in sensitive areas (equal employment opportunity). I understand you may  
>be looking for someone to deliver your program in the Chicago area.  
>Please contact me if you have an interest in discussing this.

How do I contact you? I have no name, no address.

Glad you found my letter interesting. Thanks!

What is BOG program? Is your major in PCT?



I have yet to get my first customer. If I get one in Chicago, I will teach myself. I am sure my program will be modified as it meets reality beyond trial runs with groups of acquaintances. I am not looking for associates. In fact, I am sure that for the first year at least I will not have any, if ever. I know other consultant /trainer people who prize their independence and control. It seems to me that a program like I am designing by its nature is quite personalized. My convictions are rooted in my engineering background and extensive reading and living with PCT. I don't know how anyone else could just deliver my program and be prepared to answer questions that are bound to come up about why it is the way it is, or what I mean by a certain illustration.

As you may conclude, I have no great interest in discussing this for the purpose of an association, at least not now. However, the world can use thousands of consultants who teach PCT and do it justice. I am quite willing to discuss that in general.

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

Date: Fri Jul 24, 1992 5:00 am PST  
Subject: Energy, Entropy, Info

[From Bill Powers (920723.2130)]

Allan Randall (920723.2100) --

><Information does not necessarily travel in the same direction as  
><physical energy.

>I think you are confusing the concept of energy and that of entropy.  
>They are related, but not the same. It is the latter, not the former,  
>that information theorists associate with information content. There is  
>no need in traditional information theory for the kind of "net flow" of  
>energy in the direction from source to destination that you talk about  
>in your examples, both of which I think are pretty easy to refute.

It's been a long time since I studied anything having to do with entropy,  
so remarks like yours create instant insecurity. I went back to my old  
books, and found in the Handbook of Physics:

The increase in the entropy of a body during an infinitesimal state  
of a reversible process is equal to the infinitesimal amount of  
heat absorbed divided by the absolute temperature of the body. Thus  
for a reversible process

$$dS = Q/T$$

... where Q is the infinitesimal quantity of heat.

I have seen a somewhat different and more general-seeming definition,

$$dS = k(dQ/Q),$$

where  $dQ$  is simply a signed amount of energy absorbed and  $Q$  is the amount already present (of the same form). Schroedinger uses this form of definition in "Order, disorder, and entropy" in What is life?

It seems clear that the change in entropy is a signed quantity, and that the sign (for the receiver of the energy) is the same as the sign of the direction of energy transfer  $dQ$ . (But see correction below -- I have this backward).

You say,

>The entropic formulation of information theory, along with the related  
>algorithmic formulation, has a pretty firm mathematical basis ...

A firm mathematical basis does not mean the same thing as a firm physical basis, and a firm physical basis is not the same thing as a firm experiential or semantic basis. What do we gain by calling  $Q$  "order" and  $1/Q$  "disorder?" Schroedinger proudly declares that only a physicist can understand his definition of negative entropy and its relation to order and disorder. If that is so, then physicists have created a systematic delusion which can be shared only with people who have been painstakingly trained to believe in it. We should not assume that everything we ordinarily call order and disorder is what a physicist would mean by such a term -- for example, the difference between THEDOG and TDHOEG. What a physicist means by order is not what other people mean by it, or even what the physicist means by it when, attempting to make physics explain life, he or she substitutes the ordinary meaning for the special meaning as if there were no difference. There is a certain arrogance in this proprietary attitude toward understanding that has always put me off physics -- even when I was a physics student.

Browsing through my old Buckley, I find Raymond's article on "Communication, entropy, and life." Here he defines "The rate of increase of thermodynamic entropy during communication" as

$$dS/dt = W/T,$$

where "W is the average power expended in the communication device...", a neat way of avoiding saying which way this power (energy per unit time) is traveling. The assumption, of course, is that it is traveling into the receiver. The idea that you can affect a receiver by draining energy from it never occurred to him, or as far as I can tell, any other information theorist.

RE: the telegraph example.

>The mistake you are making here is using the energy output of the  
>battery as the transmitting energy flow. This is incorrect. You are  
>treating the battery as the information transmitter.

No, I am only assuming that the battery is the ENERGY source. Information is transmitted by draining the battery, which is located at the destination end of the circuit in the first form of my example. So if we include the battery as part of the black-box receiver in Chicago, it's clear that during transmission of a message, the wires at the Dodge City end are warmed when the key is closed ( $dQ$ ), which increases their entropy by an amount depending on their initial temperature-energy,  $Q$ . The

entropy has to "flow" in the same direction as the energy flow. However, the message "flows" in the opposite direction.

[Here I discovered my error]

Actually, now that you pin me down, I realize that I've made a mistake, but not the one you mention. If a constant current flowing in a wire heats it, the temperature ( $Q$ ) rises, and as it does so, with  $dQ$  constant,  $dQ/Q$  must be falling. So in fact, entropy flows OPPOSITE to the direction of flow of energy. I told you it's been a while. All this does is change my examples so that entropy flows in the reverse direction -- it still doesn't necessarily flow in the same direction as information.

>But the battery is NOT the originator of the message. It matters not a  
>wit whether we consider the battery to be part of the sender or the  
>receiver. The battery is thus more justifiably considered as part of  
>the medium of transmission. The message actually comes from the human  
>being who is putting out the dots and dashes. This \*is\* a flow of  
>energy from the human, and \*does\* decrease the entropy of the receiver  
>and increase the entropy of the source (and the universe).

This "deduction" depends on insisting that energy DOES flow in the direction of the message -- you're begging the question. If energy or negative entropy flow is NOT the same thing as information flow, your argument is false. You can't (legally) assume your conclusion and then use it to prove that your conclusion is true. The battery is NOT, as you say, the originator of the message. But it IS the originator of the energy flow, and entropy flows in the opposite direction to energy.

(By the way, if the telegraph operator is using a bug, there is no longer a single movement for each dot or dash, because the operator can simply hold the bug paddle sideways until the correct number of consecutive dots or dashes is perceived. And there's no energetic, or entropic, difference for the operator between making a dot and making a dash)

So in this case the entropy and the information are travelling in the same direction. By moving the battery to the sending end, you can make the entropy and information flows go the opposite way (as normally assumed in transmitting messages by wire, sound, light, or radio waves by sending energy through a medium from a transmitter to a receiver).

>Compare what happens to the case of a transmitter that outputs dots and  
>dashes due to chaotic or random forces in the world around it. These  
>messages are less ordered, and thus higher entropy, than the messages  
>put out by the human.

Why are they less ordered, when they are telling us in detail about some very complex processes that present an endlessly new pattern? Does a message contain less information when it is about a more complex process? This concept confuses the atomic type of random-seeming disorder with macroscopic disorder, a completely different proposition. I consider a phase plot of a chaotic system to contain information.

The relation of order to entropy at any level but the atomic is an analogy, not an equivalence. "Order" is an experiential term based on our capacity to perceive pattern and sequence; physicists have attempted to appropriate it to mean only the reciprocal of statistical disorder, and then have turned around to say that this

restricted meaning is the ONLY meaning, thus invalidating the ability to perceive pattern and sequence. Physicists, like behaviorists and other psychologists, thus have blamed our ignorance on nature. The moment they did that, physics ceased to progress and started to disintegrate (expensively) into particles.

For a clearer example, just think of transmitting a dot-dash message by touching an ice-cube to someone's skin. The body loses heat to the ice-cube, decreasing the ice-cube's entropy and increasing that of the body and its "cold receptors" -- I hope I still have my signs right. Information being defined as the negative of the entropy change, the formal definitions of information theory would say that we are taking information out of the body and putting it into the ice-cube. If instead we use a warm soldering-iron, at a temperature well above the skin temperature, then the entropy of the body is decreased by each brief touch and that of the soldering iron is increased. So in that case formally defined information is flowing from the soldering iron into the body. In both cases, information (semantic) is being transmitted into the body, for sensory nerves respond in either case.

If you want yet another example, consider sending a message from ground level to someone two stories up by opening and closing a valve that lets water out of a hose. There's no way that energy can be transmitted up the hose, or entropy down it, using the valve.

I don't think that the originators of information theory were thinking very much in terms of nervous systems. I don't think that they were looking for counterexamples, either. Physicists pay little attention to the properties of human perception. Especially at the higher levels, they simply project them into an objective universe. When HPCT gets into physics, physics, too, will undergo a(nother) revolution.

>But what Shannon and Weaver showed was that there is a key aspect of  
>communication, which is now usually called information, that is  
>independent of this "meaning" or semantic content and has nothing to do  
>with perception. I think they succeeded.

They succeeded in analyzing the physical situation under the assumption that the source of energy would always be at the source of the transmission, and that the energy would then travel to and have an effect on the receiver. They made a blunder in assuming that you can only affect the receiver by putting energy INTO it, but that doesn't make much difference under the circumstances they were trying to analyze. They didn't even have to worry about PNP vs NPN transistors -- just vacuum tubes.

They didn't have to use the word "information" at all, except that they hoped to draw a parallel between the physical interactions and the psychological or semantic world. They never considered any of the details of sensory perception or neural transmission, so it never occurred to them that energy entering the nervous system didn't simply proceed into neural channels and make its way to higher centers, like electron flow in a wire.

>Your rubber band experiment merely shows one example of a case where  
>control is necessary for information to be transmitted. It says nothing  
>about whether such control is necessary for information transmission in  
>general. I don't think it is. Give me reason to believe otherwise.

If you consider information transmission to consist only of objective signals traveling through a physical channel independently of human knowledge, you're talking about physical "information" -- simple lineal cause and effect. But that kind of

information transmission (whichever way the energy and entropy go) does not explain communication among human beings, which is a closed-loop process. All it does is set the limits of accuracy in transmitting the level-zero message, as in Martin Taylor's Layered Protocol scheme. As I said in my talk, the meaning of a communication must be supplied by the receiver, and it is not likely to be identical to the meaning intended by the transmitter. The difference is not due to channel noise, but to the different experiences of the human sender and the human receiver.

In fact, symbolic communication is an iffy way of getting meaning from source to destination. Experience is always far more detailed than our communications about it. When a mover struggles into the living room carrying a chair, the owner may say "Just put it down anywhere." But that is impossible: the behaving system has to put it down EXACTLY SOMEWHERE, to the limit of perceptual resolution. Our actual control processes are quantitative to the limit set by system noise; our symbolic communications are vague and fuzzy in comparison, admitting of many variations in the fine details of meaning that would still fit the message. So in interpreting communications, we always add enormously more detail by way of meaning than the message can possibly carry. This is why we misunderstand each other so easily despite all the acks and naks and multiple-bit error-detection and correction that goes on between keyboard at one end and screen at the other. Even despite the dictionaries we keep at our elbows. We do not mis- receive or misread the letters; we translate them into the wrong meanings.

That is why control is required: we must not just emit our messages blindly and assume that the intended meaning shows up at the other end. We must not just assume that what we read into messages we receive was intended to be launched. We must get information back -- first from our own fingers as they blunder about over the keys, then from our own screen that shows what code was actually produced by our own flakey keyboard (displayed in a form we easily recognize), and then from the recipient of the message, to see, if we can, what meaning the recipient assigned to the strings of symbols we stuffed into our end of the wire. Many rounds of this closed loop must be traversed before a wise transmitter will admit that the intended meaning may just possibly have been noticed at the receiving end. Isn't that what's going on here?

>PS: I'm actually a lot more favourable to PCT than I appear in my posts.

I knew that. Once you understand PCT, you can't un-understand it again. It's a trapdoor.

Best, Bill P.

Date: Fri Jul 24, 1992 9:38 am PST  
Subject: Re: Energy, Entropy, Info

[Martin Taylor 920724 10:30]  
(Bill Powers 920723.2130 and Allan Randall 920723.2100)

I'm not going to comment directly on these two postings on information, energy, and entropy. And I had intended to let Bill's original report pass by, as well, though I sympathize with Allan's complaints about it.

Information theory is much more subtle than either Bill's or Allan's postings would suggest. I think it would help both of them to read at least the first few pages of Nicolis' "Dynamics of Hierarchical Systems" (Springer Verlag, 1986), where he derives

explicitly the relation between thermodynamic and informational entropy, and shows that the thermodynamic requirements on entropy per bit are so small that they can be ignored in any current practical application.

Information flow can never be determined uniquely, because it is represented by a change in probability structure as seen by the observer identifying the information. That observer may be a party to the transaction, but need not be. Information is a perceptual construct, not a physical one. We had a long go-around on this last year, but if you go back to the archives and re-read that discussion, you might find it clarified by referring to my recent "mirror diagram" of PCT.

Later, perhaps, if I find time, I'll try to write something that can be deposited in Bill Silvert's ftp system, with an abstract to the net. But please don't take either Bill's or Allan's assertions about what Shannon said too seriously. They may both be on the side of public received wisdom, but we all know how likely that is to be correct, don't we?

I'm not trying to diminish discussion of information. Just be aware that the direction of both Bill's and Allan's postings is very like that of S-R psychology, and is just as valid.

Martin

Date: Fri Jul 24, 1992 10:15 am PST  
Subject: Re: Energy, Entropy, Info [Martin Taylor 920724 13:40]

Correction. Nicolis does say that the fastest computers are approaching the thermodynamic limit in terms of energy per bit of information. But how close that is, I am not sure. I thought there were still several orders of magnitude to go, but I guess I got that from other sources and imputed it to Nicolis.

Sorry. Martin

Date: Sat Jul 25, 1992 3:04 pm PST  
From: g cziko  
EMS: INTERNET / MCI ID: 376-5414  
MBX: g-cziko@uiuc.edu

TO: Hortideas Publishing / MCI ID: 497-2767  
CC: \* Dag Forssell / MCI ID: 474-2580  
CC: marken  
EMS: INTERNET / MCI ID: 376-5414  
MBX: marken@aero.org  
Subject: How it was possible, and longer is

Greg (direct)

>In going through some of the log files, I came across a weirdity. You posted, >on the 20th of this month in reply to Dag, that boomeranging of one's own >posts isn't possible on the net (but that getting an ACK IS possible). Well, >throughout the time I was connected, I ALWAYS got BOTH an ACK AND my post >WHENEVER I posted. How come?

It took my brain a few hours to conjure this one out. It remains

invariably true and nonetheless quite factual that the LISTSERVER will not send out posts to the same address from which they came. Therefore, in normal cases, there is no boomerang. But your case is different case. That is because I am not fond of leading zeros. Following zeros is a whole different wall of sacks (especially on checks and stuff like that), but leading zeros from my own personal point of view is just null and void kinda stuff. So I entered your address on the magic list as 4972767@mcimail.com. Since the return address stamped on your posts is 0004972767@mcimail.com, the LISTSERVER thinks you are two very different people (when in fact, as you probably realize, you are only ONE very different person) and sends you your message back to you, thinking you are not you but a non-you somebody else.

You may like this (But why do you want your own messages sent back to you? Do you have a bad memory? Do you pick up your phone and talk to yourself, too? Can you hear yourself speak over the dial tone?) for some odd reason, but it is not good and to be avoided at all reasonable costs since this way you have no control over your status on CSGnet (and is it always very important for one to be able to control his or her status, high or low, as the case may or may not be). You cannot signoff or ask for nomail status or do other neat things which you never do anyway since the address on the list won't match your return address and you have to ask ME to do all this stuff.

Now, the trick to the listmanager keeping a least partial part of his or her remaining sanity is to get SUBSCRIBERS to do this stuff themselves. Therefore, I have (or soon will) put you back on the WITH with the leading zeros (voila!). Das heisst (I like French better, but a little German makes it sound more official), that you will no longer get boomeranged, but you will still get ACKed, which is really all that any partially sane person really needs, really.

>My curiosity is aroused. Why not reconnect me to the net ASAP and I'll see if  
>I still get boomerangs. (IF I DON'T, YOU HAD BETTER FIX IT SO I CAN -- we  
>know that it is possible, now, don't we?)

Nein, leider ist es NICHT (mehr) moeglich. Because from now on I will be gentler and kinder to leading zeros, even they I still consider them IN MOST CASES null and void.

See you in Kolorado.--Gary

P.S. Sorry to take so long to respond to this. Since my kids are currently very far away, I took advantage of the fact and took my wife on a romantic get-away second-honeymoon-type trip to Terre Haute, Indiana! (Our original honeymoon was in France and Terre Haute is the closest French city to Urbana, Illinois). Did you know that the great basketball player Larry Bird (of the Boston Celtics) is from Terre Haute. I may know now why he left. Nobody even UNDERSTANDS French there, notwithstanding the fact that they cannot even say Terre Haute properly! My goodness, what's this country coming to? That's why I want to vote for Perot. At least he doesn't say the final "t" of "Perot" the way that the people in Terre Haute say the initial "H" of "Haute." (This is why it takes French people much longer to write than talk since they don't say half the letters they put down, but these letters are still important since they make the words LOOK

French, even if nobody actually says them.) So what if he doesn't want the job anymore. Force him to be president, I say. That'll teach him to start things he can't finish.

P.P.S. If you're still reading, I hope you realize by now why it is important for the me to try to retain all partial remaining sanity. I couldn't even THINK of writing a note like this before I started managing CSGnet. Now it's hard for me to think at all. That is why you cannot get boomeranged anymore so you can control your own status. But I fear I am starting to repeat myself (or somebody else much like me) so I need to go home and feed my home computer (I am in my university office now, or at least one that looks much like mine or one that I knew).

P.P.P.S. I will send you the current csg-1 log9207d when I put you back on the net so you will be caught up, and then will send the log again when it is full and just history.

Gary A. Cziko

Date: Sun Jul 26, 1992 12:36 pm PST  
Subject: Got them lonely, low-down PCT blues

[From Rick Marken (920726)]

Gary -- Still working on revision of Blindmen. By the way, why did you send me that direct note about Greg's net feedback problems?

Martin -- Did you try to simulate the single ECS "hidden bomb" yet? I think I know what happens -- no bomb again I'm afraid; but try it and see. I am working on developing an experiment based on my observation (stimulated by one of your earlier "bomb" posts) that the addition of a new control system can create problems if control of that variable requires the inconsistent use of outputs that are already being used to control another variable -- but that this depends on how the person originally learned to control the first variable. I think this should be fairly easy to set up -- and I think it could be pedagogically interesting also.

The subject line of my post refers to the fact that I was having a bit of the PCT blues yesterday. I was feeling blue because I was getting tired of PCT being perceived as such a "fringe" approach in the life sciences. While I dearly love and enjoy working with the dozen or so people I know who really understand PCT, it gets a bit lonely out here without them (though the net helps). I guess I just don't really like the fact (though I understand why it happens) that people (often colleagues) have such an allergic reaction to PCT. I want to cry "What's wrong with PCT? What don't you like? Why don't you want to just give it a chance? Why don't you want to at least TRY to understand it?" I know that PCT contradicts much of the basic dogma of the life sciences. But people seem so eager to overthrow dogma - to embrace ANY "brave new approach" to understanding life. Why don't they spend some time trying to understand PCT?

Of course, I know the answers to these questions (at least, from a PCT perspective) -- but it's still depressing sometimes. My current depression was set off on Friday when I had a meeting with a fellow human factors engineer from another company. He was a very nice, charming person. He was also a person who had done research on control models of people (from the engineering perspective -- trying to discover how



the input -- our disturbance -- was related to output). He was interested in my work and asked for a reprint. The article looked familiar to him and it turned out he had been a reviewer on it; it was kind of embarrassing because he had given it a relatively negative review (not real bad, luckily, and I didn't look up how I replied to him -- the article was published, after all -- but that should be a lesson to me to be a lot nicer in my replies to reviewers; they are just people and I might even know 'em). The depressing part of this encounter came from the realization that I was doing research from what I'm sure this fellow saw as such a "fringe" perspective. Despite what it might seem like, I DON'T LIKE BEING PART OF A REVOLUTIONARY MOVEMENT. It is NOT fun being part of a psychological movement that is viewed as "fringe", "radical" or whatever by 99.9% of my colleagues. The only reason I am a member of this weird group is because I value intellectual integrity even more than I value being part of the majority (the group I REALLY want to belong to). So please -- save me from the clutches of this radical cadre. TELL ME WHAT IS WRONG WITH PCT -- PLEASE!! Then I can go off and be in a big, popular group like the neural net group or the artificial life group (they have more famous people too; and a glossy covered newsletter).

One helpful therapy for these blues of mine would be if someone could explain (and show, through modelling and experimentation) why a particular theory is BETTER than PCT. Randy Beer tried to help me on this some time ago but failed rather miserably. Maybe it was my fault -- being too dumb to understand him. So, for my sake, please keep arguments against PCT simple and clear (and, hopefully, written in BASIC or PASCAL). Please, NetNiks, help me figure out what is wrong with PCT -- let me know what every other psychologist seems to know -- so I can rejoin that happy (and moral) majority. If you don't know what is wrong with PCT (possibly because you are already part of CSG) then ask a friend who knows enough about PCT to know that it's wrong.

Then have the friend explain it to me (why PCT is wrong, that is).

Thanks. Rick

Date: Mon Jul 27, 1992 1:08 am PST

Subject: Durango

[Wayne Hershberger 920726]

Rick Marken:

Will you please bring the IBM version of your spreadsheet program to Durango this week? I would like to get a copy of it and some coaching from you regarding its potential use in a teaching lab.

Regarding your lonely lament earlier today, are there things the CSG Group is not currently doing that it could be doing that might hasten the general acceptance of the "truth" we are championing. As Ed Ford might say, "Is it working; are we getting what we want?" Perhaps the issue should be discussed seriously this week in Durango.

Warm regards, Wayne

Wayne A. Hershberger  
Professor of Psychology  
Department of Psychology

Work: (815) 753-7097  
Home: (815) 758-3747

Northern Illinois University  
DeKalb IL 60115 Bitnet: tj0wahl@niu

Date: Mon Jul 27, 1992 5:09 am PST  
Subject: Justifying the Blues

From Greg Williams (920727)

Rick asked what is wrong with PCT. From the point of view of non-PCT psychologists, there are several possibilities, including the following:

Its name sounds illiberal: "Control! Ugh!"

Its origins are suspect: "Sprung virtually full-blown from the musings of a physicist who never paid his dues in psychology."

Its origins are dated: "This is supposed to be revolutionary, using WWII technological ideas?"

It is too bold: "Where is the empirical data for the details of the mechanisms of hierarchical control proposed in HPCT?"  
(and/or)

It is not bold enough: "PCT (no H) consists of truisms already well elucidated by nonPCT theories."

It has few supporters, even after many years: "So how could it be important?"

Several PCTers are Rank Outsiders, even Rebels: "Rabble rousers!!"

No Significant Psychologists have prominently expressed their support for PCT: "So how could PCT be Significant?"

Several of its supporters have associated with "fringe scientists," particularly "cyberneticians": "The implication is that PCT probably is as sterile as most of cybernetics with regard to important issues in psychology."

Its supporters are often arrogant, unaccommodating, disrespectful, and impatient with nonPCTers: "Screw 'em!!!"

In critiques of conventional psychological ideas, PCTers have focused on stimulus-response theories, which are dead anyway: "When will they catch up to the state-of-the-art?"

Some PCTers seem ignorant of much of the conventional psychological literature: "Why don't they read more and say less?"

PCTers appear bent on emphasizing the differences and minimizing the similarities between their ideas and those of others: "Foul! We're supposed to all be in this together."

Its early publications have not led to a sustained research program evidenced by a series of publications in leading psychology journals: "If PCTers can't get their stuff past peer reviewers, why should anybody else be

interested in it?"

It isn't useful for clinical work: "We don't want to tell our clients that THEY are in control!"

Or for applied psychology: "We don't want to tell ourselves that the SUBJECTS are in control!"

Or for government-sponsored studies: "PCT talks about individuals, not masses, and the grant overseers don't give a damn about individuals."

Finally, and perhaps most significantly, it provides a basis for folk-psychological beliefs in individual autonomy and lack of environmental determinism: "PCT is wishful thinking -- dream on, PCTers!"

Balanced against all of the above is one "right":

PCT (but not necessarily the details of HPCT) has greater explanatory ability than any other contemporary theory in psychology (at least that I've seen, and I've looked pretty far and wide): "A minor point."

I'm sure other PCTers can add to the list. Rick's blues appear amply justified -- but nobody promised a rose garden for psychological revolutionaries. My prescription: ignore the "wrongs" and concentrate on publicizing the "right." In particular, I suggest showing how PCT can explain numerous empirical findings already in the literature. That should keep PCTers too busy to be depressed.

Greg

Date: Mon Jul 27, 1992 5:45 am PST  
Subject: Re: Justifying the Blues

From: Oded Maler (920727)

[I promised myself to refrain from participating in this kind of things, at least till I get a position, but it's stronger than me].

I don't know if it will comfort you but in the Aix summer school, Tom Bourbon's talk, and PCT ideas were some of the \*least\* controversial. Many people accepted them (though some without understanding them [not to mention \*really\* understanding them..]) {I'm not sure whether this would have been the case if Rick was presenting them :-}

While explaining what I know to CS/robotics people, the usual misunderstandings came by (the name "control theory" confused with the math/engineering discipline, B is the C of P which is really a truism for roboticists (also not for cognitivists), etc.)

Concerning Bill P.'s criticism of Brooks, the latter didn't agree with Bill's insistence that higher-level CS don't have direct access to low-level sensors (he indicated also biological evidence for direct pathways). I think the recent CSG discussion (which I didn't read)

about functional vs. anatomical connections is relevant. He was not convinced either that it is useful to see everything in terms of serving and references, especially in the higher-levels.

And a last advice for some dedicated PCTers: you shouldn't adopt a too PCT-centric point of view - there's much work going on in biology and psychology which is neither pro nor contra-PCT, but rather orthogonal, by asking different questions, by investigating cognitive functions where HPCT is not more than a reasonable guess, etc. One's frustration of being rejected from a certain esoteric community called Psychology should not be taken too seriously.

Best regards --Oded

Date: Mon Jul 27, 1992 9:33 am PST  
Subject: Misc

[From Rick Marken (920727.1000)]

Wayne Hershberger (920726)

> Will you please bring the IBM version of your  
>spreadsheet program to Durango this week?

Will do; though can't we download it from Silvert's list when we are there?

> Regarding your lonely lament earlier today, are there  
>things the CSG Group is not currently doing that it could be  
>doing that might hasten the general acceptance of the  
>"truth" we are championing.

Definitely worth talking about. But I think a few beers and good company will fix me up in no time.

Greg Williams (920727)

Great list of "what is wrong with PCT" ideas. Some we can do nothing about, of course, but I have some quick answers to a couple:

> It is too bold: "Where is the empirical data for the details of the  
>mechanisms of hierarchical control proposed in HPCT?"

Answer: Where is the empirical data supporting the details of other models of behavior??

> Its supporters are often arrogant, unaccommodating, disrespectful, and  
>impatient with nonPCTers: "Screw 'em!!!"

Oops. Sounds like me when I'm not singing the blues.

> Some PCTers seem ignorant of much of the conventional psychological  
>literature: "Why don't they read more and say less?"

I've read it (and referred to it). Where's the beef? Actually, this has been a bit of a sore spot for me; when I do try to refer to the existing literature reviewers very often say that I have "got it wrong" (the theory, that is). This is annoying (and one reason why I have resorted to word for word quotes when referring to the theories of others) but it is an expected result of having a non-working implementation of a theory; it's just words, and we all know how ambiguous these can be.

> PCTers appear bent on emphasizing the differences and minimizing the >similarities between their ideas and those of others: "Foul! We're supposed >to all be in this together."

Good point. But it's hard to compromise in science, as Galileo discovered ("and yet, it turns").

> My prescription: ignore the "wrongs" and concentrate on >publicizing the "right." In particular, I suggest showing how PCT can explain >numerous empirical findings already in the literature. That should keep PCTers >too busy to be depressed.

I would be happy to do this. But what I have discovered (recently) is that most (say 90%+) of these empirical findings are of such LOW QUALITY (being statistical results -- the highest degree of relationship I found while perusing published studies was .90; not bad, but still not good enough for doing modeling) that they are really NOT RESULTS. Trying to use PCT to explain the empirical findings in the psychological literature would be like using PCT to explain Rhine's mental telepathy data or the winner's of the last five races at Hollywood Park. An unfortunate implication of PCT is that there are almost NO empirical findings of any use to the PCT modeller in the current, standard psychological literature. PCTers may not have collected huge amounts of research data in big research projects but I've looked through a lot of journals lately and unless one considers noisy statistical results to be data, conventional psych ain't got much data either.

Ah. That felt better.

Best regards Rick

Date: Mon Jul 27, 1992 10:06 am PST  
Subject: There is hope, Rick

from Ed Ford (920727:1100) Rick Marken, Greg Williams, et al

Not all is lost concerning the lack of acceptance of PCT. I think PCT will (and in my case is) find acceptance from those who are struggling to produce some positive results with the people with whom they work. I am working with several organizations and lots of people who at first haven't the faintest idea what I'm talking about. As they begin to find success in turning their own lives around or in getting their organization to function cooperatively and effectively, they begin to respect the theoretical basis (PCT) upon which I'm basing my ideas for helping them. What I'm trying to say is that people have to first succeed at what ever they are trying to do using your help or your

ideas as the basis for their success. You have to go from the practical success or practical examples to respect for the basis upon which those practical ideas come from. That's how I get another's attention. I'm working with a Catholic priest who runs a largely volunteer organization called Andre House, which feeds in excess of 800 homeless every day besides providing clothing, shelter (for some), and help in gaining skills and finding work. I had him read Freedom From Stress, then we worked on how people can work together while respecting each other's internal worlds. I talk about measurable goals and controlling variable perceptions and providing feedback and ways of working with the staff and they with the homeless that would help every one involved take more control of their internal worlds and the areas for which they are responsible. The staff hadn't the faintest idea what I was talking about at first, but my practical suggestions seemed to have made sense and they tried what I suggested cause it made sense. As they began to succeed, respect for my ideas spills over to respect for PCT, and for some, a desire to learn what the heck I'm talking about. I guess you have to go from practical success to respect for the theoretical basis and then, for some, a curiosity as to what the heck is PCT.

Ed Ford ATEDF@ASUVM.INRE.ASU.EDU  
10209 N. 56th St., Scottsdale, Arizona 85253 Ph.602 991-4860

Date: Mon Jul 27, 1992 10:39 am PST  
Subject: anyone get back to van Loon?

Really-From: John\_Van\_Loon.XRCC%xerox.com@BINGVMB.cc.binghamton.edu  
Date: Wed, 22 Jul 1992 08:30:08 PDT

Hi,

My name is John van Loon and I am currently in my last year of electrical engineering at McMaster University in Ontario Canada. I am most interested in robotics and control, that is, I love working on all aspects of robots from Neuronet integrated intelligence to motor control and installation. The problem is that my experience is severely limited to what I can afford to fund myself. As a thesis I will be responsible for the design and construction of a mobile base and manipulator system for a self-controlled robot. In addition to this I will also be working to help develop the supervisory controls of this robot. The more that I see and learn, the more interested and fascinated I become with the whole art of cybernetics.

If there are any discussion groups established for walking robots, sensory development, and hand/manipulator control please send me a short note about your group and an address where I can E-Mail you.

Thank you,  
John van Loon.

-----  
John van Loon: John\_Van\_Loon.XRCC%xerox.com@uunet.ca

Date: Mon Jul 27, 1992 12:06 pm PST  
Subject: Blues, 2 (%)

From Greg Williams (920727-2)

>Rick Marken (920727.1000)

GW>>My prescription: ignore the "wrongs" and concentrate on  
GW>>publicizing the "right." In particular, I suggest showing how PCT can  
GW>>explain numerous empirical findings already in the literature. That should  
GW>>keep PCTers too busy to be depressed.

>I would be happy to do this. But what I have discovered (recently) is that  
>most (say 90%+) of these empirical findings are of such LOW QUALITY  
>(being statistical results -- the highest degree of relationship I found  
>while perusing published studies was .90; not bad, but still not good  
>enough for doing modeling) that they are really NOT RESULTS.

I would be surprised if high quality results amount to as much as 2% of the total.

>An unfortunate implication of PCT  
>is that there are almost NO empirical findings of any use to the PCT  
>modeller in the current, standard psychological literature. PCTers may  
>not have collected hugh amounts of research data in big research projects  
>but I've looked through a lot of journals lately and unless one considers  
>noisy statistical results to be data, conventional psych ain't got much  
>data either.

Still, the good 2% (or maybe more) would keep PCTers busy for a long time in ways which could be perceived as relevant by non-PCT psychologists. In fact, there's more than enough to keep PCTers busy for a long time just in the sub-sub-sub-field of limb trajectories!

Greg

Date: Mon Jul 27, 1992 12:23 pm PST  
Subject: stick patterns ctd.

[From Pat Alfano]

Monday July 27, 1992

To: Bill Powers and David Goldstein

Thanks for the input on the analysis for my stick patterns.

I guess I should have supplied more information with my request for help on analyzing my stick pattern data.

My main goal in analyzing the data is to show that the stick test is not necessarily testing what neuropsychologists say it is testing, and to show, at least in part why this is so. If I could also tell them what it is testing, that would be a bonus.

Basically the stick test is thought to be a test of spatial ability and to require mental rotation. It is thought to test ones' ability to find

one's way around in his or her environment, to read maps, etc. If a brain injured person has trouble doing this it is thought that the part of the brain that performs mental rotation has been injured, and that part of the brain is most likely on the right side, probably in the temporoparietal area.

My argument is that there is no evidence that this task requires mental rotation (whatever that is) in order to be successfully completed, and that there is no evidence that it is a task of spatial ability (whatever that is).

I know that to simply ask subjects what strategies they used to complete the task does not get at what they actually did, but, I was hoping to show those who use the test that they need to discover exactly what it is people are doing when performing the task; and to point out to them that they are relying on assumptions when interpreting the results. I also wanted to point out that normal people have trouble doing the stick test and therefore no conclusions can be drawn about a brain injured person who cannot do the test.

By the way, I did count errors. I also videotaped subjects doing the stick test during a pilot study.

I would like to make suggestions about how neuropsychologists should look for what it is exactly people are doing when performing the stick test or similar tasks. You both have given me some ideas on what I might suggest.

At the present I am not planning on studying the phenomenon myself, but may be interested in doing so in the future. I would like to know what goes on in someone's head when performing a task such as this. Have I set my sights too high?

Pat

Date: Thu Jul 30, 1992 3:28 am PST  
Subject: possible venue for CT books

[From: Bruce Nevin (Thu 920730 07:14:50)]

Possible venue.

From:  
IRLIST Digest July 29, 1992 Volume IX, Number 26 Issue 122

Fr: EINA@ccvax.unicamp.br  
Re: Studies in Artificial and Natural Intelligence

STUDIES IN ARTIFICIAL AND NATURAL INTELLIGENCE

PURPOSE: The relevance of the study of the human brain to the study of artificial intelligence has long been an issue of debate both in the AI community and in the Neuroscience field. On the one hand, it has been claimed that the study of the brain is too



complex and mysterious to yield useful guides for the construction of intelligent machines and that these machines may be developed under different guidelines and hardware. On the other hand, it is claimed that the behavior of the brain is the very inspiration for the study of artificial intelligence.

Of course, when neurons agglutinate into brains, new (emergent) properties arise, not possessed by each of these neurons themselves, but which derive from their association. This could be the argument to support the claim that Neuroscience is irrelevant to AI. But besides these emergent properties, the high order system (brain) also inherits properties from the unitary elements (neurons) composing it. This could be the justification for the attempt to reduce reasoning to the physiology of the neuron.

As a matter of fact, this discussion is an old issue in the history of the western philosophy, since it refers to the mind and brain dualism. The modern connectionism starts to provide the ways to approach experimentally this mutual correlation under the optics of science, that is, putting it as a workable hypothesis which may be falsified by empirical data. The development of new techniques to knowledge acquisition and analysis in the AI field, was very much encouraged by the success of the Expert System technology. They were initially developed with the purpose of obtaining the contents of the knowledge base of the expert system. However, these methodologies may be applied to investigate the human thinking. In this way, AI may provide Neurosciences with a very strong tool to empirically test their hypotheses about the human reasoning.

The purpose of the present series of books is to be a forum for this kind of scientific debate. Any work contributing to the comprehension of the correlation between Artificial and Natural Intelligence; the experimental approach of Intelligence; the simulation of Intelligent Activities, and the multidisciplinary approach of Natural and/or Artificial Intelligence, is welcome to integrate the Studies In Artificial and Natural Intelligence forum. The series may also be the adequate vehicle for publishing selected papers from congresses and conferences on the above topics.

AUDIENCE: This series intends to be of interest for people working on distinct fields of science but with interest on both (or either) Natural and (or) Artificial Intelligence, such as: Neurophysiology, Neurochemistry, Neurogenetics, Psychology, Mathematical Biology, Intelligent Control, Expert Systems, Logic, Machine Learning, Connectionism, Hybrid Intelligent Systems, Robotics, Pattern Recognition, Philosophy, etc.

DISTRIBUTION: The books are printed in Poland by Omnitech and are distributed around the world by Physica Verlag.

CONTACTS:  
Editor-in-chief:

Prof. Armando F. Rocha  
RANI - Research on Artificial and Natural Intelligence  
Rua Tenente Ary Aps, 172  
13200 - Jundiai - Brasil  
e-mail: eina@bruc.bitnet

Publisher:  
Omnitech Press  
ul. Chmielna 16-2  
00-020 Warsaw - Poland  
Phone/FAX: (48) (22) 27-34-94  
-----

Armando Freitas da Rocha  
Dep. Physiology and Dep. Computer Engineering and Automation  
UNICAMP BRAZIL

Date: Thu Jul 30, 1992 5:21 am PST  
Subject: thanks for Closed Loop

[From: Bruce Nevin (Thu 920730 08:54:12)]

Thank you, Greg, and all others involved in putting together the current number of Closed Loop. This may both prompt and help me to remedy my educational deficit re statistics.

Bruce