A skeptical engineer joins CSGnet. Albus.

Unedited posts from archives of CSG-L (see INTROCSG.NET):

A SKEPTICAL ENGINEER CHECKS OUT PCT AND CSGnet.

This is a rather complete thread of Paul George discussing Albus with CSGnetters and appreciating _Behavior: The Control of Perception_ towards the end. As always, every discussion provokes thoughtful posts. Note Tom Bourbon's post: Albus articles and Bill Powers' post: PCT models and simulations.

I hope that this thread will prove helpful as an introduction to PCT and CSGnet.

March 21 1995, Dag

------------------------------------------------------------

Date: Thu Jun 23, 1994 4:39 pm PST
Subject: PCT models

I have been lurking for a while listening to the PCT debates, and think you may be using too simplistic a view of a control system, which is biasing your discussion (in fairness I haven't had access to the books/papers on the subject, just posts on BPR_L by Dag).

Sophisticated control systems don't use reference variables (e.g. setpoints or alarm thresholds) they use reference _models_ (reflected in control logic). In a sense we have a continuously running simulation of the 'real' world to which we compare our perceptions (sensor inputs). As PCT stresses we also filter our perceptions based upon our current model, trying to separate signal from noise.

The basic components of this view of a control system are Sensory Processing (filter/transducer), Value Judgement (comparator), World Model/database, and Behavior Generator. Part of the model is some set of goals or desirable states we wish to maintain or approach. We modify the model to better reflect 'outside reality' (as perceived) so that we can better predict the results of our actions (actuator outputs).

For a good exposition of this model I suggest:

"Outline for a Theory of Intelligence" James S. Albus IEEE Transactions on Systems, Man, and Cybernetics, vol 21 #3, May/June 1991, p 473-509. IEEE log # 9042583

Date: Fri Jun 24, 1994 7:28 am PST
Subject: Re: PCT Models

[Paul George 940624 0900] >[Martin Taylor 940623 20:30]

> Sometimes these more complex control systems do get talked about, but one of the questions of interest is what can be done with the simplest versions. For example, if we assert that the simple elementary control system controls a scalar variable, it is possible to imagine that variable being the output of a neuron. Is there any evidence to suggest a need for more complex controlled variables in real systems? Is it not always possible to emulate the effect of a complex control system with a hierarchy of scalar control systems?

{snip}

Basic general systems theory and gestalt psychology indicate that the simple reductionist approach doesn't work well for complex systems, particularly biological ones. The whole is more than the sum of the parts. There is a qualitative difference between a Cray and a flip/flop (bistable multivibrator for the picky).

I may be covering old ground, but simple scalar variables are not sufficient for control of dynamic processes, except in a simplistic case. We are here dealing with human behavior. Even for simple astrodynamics or object tracking you need vector systems such as Kalman Filters. A model need reflect pertinent aspects of the real world."Keep it is simple as possible, but no simpler".

Actually Albus posits a hierarchy of simple control systems where each control component interacts with like nodes both up and down the hierarchy. Effectively nodes form a network, incidently mapping to the structure of nervous systems.

> The basic attitude of a scientist when confronted with an explanation more complex than the one in current use is "show me what it explains beyond what I already explain, or more precisely than I already explain it."

Albus's model adds several things that allow separation of concerns and an acknowledgement of internal complexity. It takes us beyond what might be unfairly called a black box model.He alleges that it has been validated in terms of actual biological and behavioral structures (I haven't backtracked his references, but matches my understanding of such mechanisms).

1) a specific locus for a atonomic perceptual filter & sampler. The simplistic diagram seems to take a 'then a miracle occurs' approach. This helps deal with the fact that we don't mentally control all perception. Some things are wired in. Second we do have some control on sensitivity levels and which sensory inputs we pay attention to at a particular time. This has been experimentally verified.

2) An explicit concept of both reference variables and some kind of a state machine (or other rule/history mechanism) to allow prediction and complex behavior. "World model" overstates the concept somewhat, especially at the lowest level of the hierarchy. It is a history/prediction mechanism adequate to the task at hand. It also helps provide a locus for simple fight/flight response mechanisms and 'habitual' responses. The Model node can interact directly with sensory perception and behavior generation nodes with out the mediation of Value Judgement.

3) A specific locus for goal seeking behavior and evaluation internal model vs external inputs and planned outputs. This is where either simple threshold evaluation or complex rule processing takes place. It is, if you will the seat of 'consciousness' or 'will'.

4) An explicit locus for patterns of action or outputs. This cover both instinctual or otherwise hard wired actions (pulling a hand back from a flame), learned patterns (riding a bike), and consciously defined action plans. It also provides a mechanism for Lilly's behavioral metaprograms.

{World model component description}

> I accept that such modification is required, and have several times introduced it explicitly into diagrams. But I do not know that it cannot be equivalently represented in the interconnections of a hierarchy of control systems without explicit models. To introduce an explicit model is to introduce the question of mechanism of data storage and data retrieval.

I am not sure you need to specify a mechanism, though it is an interesting question, just recognize that it appears to exist. We acknowledge that we possess memory and can use language, even if we don't understand how it works at the neural level. Distributing the memory through the control system hierarchy just means a distributed model which incorporates the control logic. This is not wrong, but appears to hide the issue. I aesthetically prefer the concept of a hierarchy of models and value judgement systems. Whether you use this model or a black box +ghost in the machine' approach depends on the issue you are discussing. From the standpoint of observing external behavior, it doesn't matter. But the internal structuring may help our understanding of what is going on inside, from a conceptual standpoint.

Paul C. George

Date: Wed Jun 29, 1994 5:17 pm PST
Subject: Albus and Powers

From Tom Bourbon [940629.1731]

I'm just now working my way through the thread on "PCT models" initiated by Paul George (23 June 1994). Paul cited a 1991 article by Albus, and I notice that Martin Taylor has copied the article to read. I think a good place to compare Albus and Powers is in their independently written series that ran in _BYTE_ magazine, in June, July, August and September, 1979. Bill Powers wrote about the CST model (as PCT was known back then) and told how to actually write programs to test the model. Bill's descriptions of the model and the programs were crystal clear. His programs, hence his model, actually _worked_ -- a feat that some of you know is of great significance to the modelers on this net.

In contrast, Albus began the fourth article in his series (which was titled, "A model of the brain for robot control") by saying, "The essence of a hierarchy is that control is top-down. The ultimate choices are made at the top, and the goals selected at this level are decomposed into action as they filter down through the various levels of the hierarchy." That's about as far as you can get from the idea about a control hierarchy in PCT; top-down hierarchies of that kind don't work as models for living systems. The remainder of the Albus article contained no working model for a robot, or for a brain. It did contain lots of speculation and guesses about the future of robots.

I recommend the two parallel, but widely diverse, series of articles to anyone who wants to compare the ideas of Powers and Albus.

Later, Tom

Date: Wed Jun 29, 1994 5:24 pm PST
Subject: Re: PCT models

From Tom Bourbon [940629.1757]

I've finished reading the new thread on PCT models. Here's a first small set of comments.

In Message Thu, 23 Jun 1994 12:04:28 EDT, Paul George writes:

> I have been lurking for a while listening to the PCT debates, and think you may be using too simplistic a view of a control system, which is biasing your discussion (in fairness I haven't had access to the books/papers on the subject, just posts on BPR_L by Dag).

OK. You haven't read any of our work but you think it is too simplistic and it doesn't work. Since you are new to this net, we don't know about your work, either. Can you tell us something about how you decide whether a model works, or not? It would be helpful to me if you were to say some things about that question in general, then tell us specifically how you apply those criteria to quantitative predictions made with the Albus model and why you think the results lend credence to it.

> Sophisticated control systems don't use reference variables (e.g. setpoints or alarm thresholds) they use reference _models_ (reflected in control logic). In a sense we have a continuously running simulation of the 'real' world to which we compare our perceptions (sensor inputs). As PCT stresses we also filter our perceptions based upon our current model, trying to separate signal from noise.

Don't worry. We make no pretense of sophistication, so that field is wide open to anyone who wants to claim it!

Later, Tom Bourbon

Date: Thu Jun 30, 1994 7:05 am PST
Subject: Re: PCT Models

From Tom Bourbon [940630.0823]

>[Paul George 940624 0900] >>[Martin Taylor 940623 20:30]

Martin:

>> Sometimes these more complex control systems do get talked about, but one of the questions of interest is what can be done with the simplest versions. For example, if we assert that the simple elementary control system controls a scalar variable, it is possible to imagine that variable being the output of a neuron. Is there any evidence to suggest a need for more complex controlled variables in real systems? Is it not always possible to emulate the effect of a complex control system with a hierarchy of scalar control systems?

Paul:

> Basic general systems theory and gestalt psychology indicate that the simple reductionist approach doesn't work well for complex systems, particularly biological ones.

Could you tell us something about _how_ systems theory and gestalt psychology "indicate" that a "reductionistic approach doesn't work for complex systems" and _why_ you believe those ideas would apply to our work -- which you have not read? I am not familiar with the process of "indication" and do not know how to use it when I compare the behavior of various generative models. I'm in the dark about what you mean and will remain there until you say a few things about how you test the accuracy of quantitative predictions from the Albus model.

> The whole is more than the sum of the parts.

I believe the original statement was closer to, "The whole is _different_ from the sum of the parts." Whichever way we say it, so what? How can that truism guide us in evaluating the performance of generative models of behavior? Precisely how does it help us compare, to pick an example out of the air, the two models described independently by Albus and Powers, in 1979?

> There is a qualitative difference between a Cray and a flip/flop (bistable multivibrator for the picky).

Yes. So?

> I may be covering old ground, but simple scalar variables are not sufficient for control of dynamic processes, except in a simplistic case. We are here dealing with human behavior.

Yes, we are indeed dealing with human behavior here. When we use our simple (and insufficient) PCT scalar models to predict the performance of humans in a range of tasks, the predictions account for over 99% of the variance in the data from the humans. How well does the (sufficient) Albus model do in such cases? It must be pretty near perfect, if your remarks about it are close to the facts.

> Even for simple astrodynamics or object tracking you need vector systems such as Kalman Filters. A model need reflect pertinent aspects of the real world."Keep it is simple as possible, but no simpler".

_Who_ needs vector systems? We (PCT modelers) don't, at least not just yet. How much more of the variance does the Albus model "explain" when it is equipped with them than without them?

"A model need reflect pertinent aspects of the real world." Why? What is the evidence for this statement, from testing models? What if there are models that perform very well but do not "reflect pertinent aspects of the real world." Which do we question or reject, the models, or the dictum?

> Actually Albus posits a hierarchy of simple control systems where each control component interacts with like nodes both up and down the hierarchy. Effectively nodes form a network, incidently mapping to the structure of nervous systems.

Albus posits a top-down hierarchy. Have you tested a generative version of a top-down model? In our experience, they don't work as models of living systems. We would like to see evidence to the contrary if it is available.

Paul:

> Albus's model adds several things that allow separation of concerns and an acknowledgement of internal complexity. It takes us beyond what might be unfairly called a black box model.

Martin Taylor asked you some good questions about those remarks.

> He alleges that it has been validated in terms of actual biological and behavioral structures (I haven't backtracked his references, but matches my understanding of such mechanisms).

I noticed that, at several points in your posts on this thread, you said (or implied) you hadn't actually checked out Albus's allegations and claims. Have you seen any simulations or quantitative predictions from him? Have you run any of your own, using his model? I'm asking seriously -- you are not the first person to tell us that Albus says something important about the phenomena we study in PCT. We are still looking for good evidence to support those assertions.

Your enumeration of the following items from the Albus model makes it clear that the parallel articles in _Byte_, in 1979, are still one of the best sources for a comparison of Albus and Powers. You are talking about a different breed of cat from the PCT model.

> 1) a specific locus for a atonomic perceptual filter & sampler. .

> 2) An explicit concept of both reference variables and some kind of a state machine (or other rule/history mechanism) to allow prediction and complex behavior. . . .

> 3) A specific locus for goal seeking behavior and evaluation internal model vs external inputs and planned outputs. . . .

> 4) An explicit locus for patterns of action or outputs. . .

No thanks. I'll stick with functional, generative PCT models. Call me backward and inadequate; I don't care. ;-) (I _will_ rethink that position if I see examples of generative models that adhere strictly to the Albus architecture _and_ produce fits to quantitative human data that are as good as or better than the fits from simple PCT models.)

Later, Tom

Date: Thu Jun 30, 1994 8:51 am PST
Subject: Re: PCT models

I didn't say that PCT models wouldn't work. I just suggested a more complex internal organization might increase understanding, particularly in terms of where communications errors may occur. I do not judge PCT, as I don't yet know a lot about it. On the surface it seem very attractive. I laud the communities work in attempting to validate the theory via modeling.

A model is valid if it is useful for understanding and mimics the behavior or a portion of the real system. To simple a model can blind you to complexities of the real system and subtleties of its behavior. Perceptual filtering tends to make you ignore things you weren't expecting or were outside the parameters of your mental model.

Modeling is a useful technique but subject to limits, particularly in selecting the 'important' characteristics of the real system. Observing living systems and man-made control systems and generalizing to produce theory is also valid. Validation by inspection is a valid technique. Albus didn't make 'quantitative predictions', he proposed structure and semantics.

On the whole you appear very defensive, overreacting to perceived criticism. I would hope this group is not limited to TRUE BELIEVERS.

There are two kinds of a fool:
One says +This is old and therefor good+
The other says +This is new and therefor better+

Paul C. George

Date: Thu Jun 30, 1994 11:43 am PST
Subject: Re: Albus and Powers

[Paul George 940630 10:30]

>From Tom Bourbon [940629.1731]

> In contrast, Albus began the fourth article in his series (which was titled, "A model of the brain for robot control") by saying, "The essence of a hierarchy is that control is top-down. The ultimate choices are made at the top, and the goals selected at this level are decomposed into action as they filter down through the various levels of the hierarchy." That's about as far as you can get from the idea about a control hierarchy in PCT; top-down hierarchies of that kind don't work as models for living systems.

Note that the article I cited is the only Albus work I have read. It was circulated as a theoretical basis for a hierarchical distributed object control system architecture the company is working on.

Aldus's thinking appears to have evolved since 1979 (one would hope). The article is not oriented around top down control as you appear to think of it. Indeed, it is far more bottom up in nature, and focused on coordination. Control policies or goals _may_ flow down the hierarchy, but summarized data and model results (error signals) pass up. Both also move laterally in the network. It is not gospel, just a nice description of the workings of control system networks (once you get past the definitions and philosophy). I will not defend it as TRUTH. As Martin commented, his definition of language seems way off base to me, and there were a number of other things that caused raised eyebrows. I would encourage you to read the article and comment on the wheat and chaff.

BTW, I do not see how PCT could refer to a control system _Hierarchy_ is there is no kind of command relationship, or hierarchical filtering and summarization of data. That is part of the definition of 'hierarchy'. The quote you cite is fundamentally accurate. However, control systems need not be _strictly_ hierarchical. And equally clearly living systems do involve hierarchical structures.

If you deny the existence of hierarchical control, perhaps you should use 'tree network' instead. I think it is more likely that you are having a NIH reaction, focusing on minor distinctions as the most important, for the purpose of differentiation. PCT aficionados do not have a monopoly on understanding, even if they may have produced working models.

Paul C. George

Date: Thu Jun 30, 1994 11:48 am PST
Subject: Re: Albus and Powers

Tom Bourbon [940630.1201]

>[Paul George 940630 10:30]

>>From Tom Bourbon [940629.1731]

>> In contrast, Albus began the fourth article in his series (which was titled, "A model of the brain for robot control") by saying, "The essence of a hierarchy is that control is top-down. The ultimate choices are made at the top, and the goals selected at this level are decomposed into action as they filter down through the various levels of the hierarchy." That's about as far as you can get from the idea about a control hierarchy in PCT; top-down hierarchies of that kind don't work as models for living systems.

> Note that the article I cited is the only Albus work I have read. It was circulated as a theoretical basis for a hierarchical distributed object control system architecture the company is working on.

Understood. But recall that you appeared on this net with an announcement that, ". . . you may be using too simplistic a view of a control system, which is biasing your discussion (in fairness I haven't had access to the books/papers on the subject . . .", followed by, "Sophisticated control systems [TB: by implication, PCT is unsophisticated] don't use reference variables (e.g. setpoints or alarm thresholds) they use reference _models_ (reflected in control logic). [TB: PCT uses all of these things you say are not used by sophisticated models.] You concluded with:

=================
For a good exposition of this model I suggest:

"Outline for a Theory of Intelligence" James S. Albus IEEE Transactions on Systems, Man, and Cybernetics, vol 21 #3, May/June 1991, p 473-509. IEEE log # 9042583

==============

which looked like a pretty strong endorsement of Albus as an example of a sophisticated modern model of control.

I hope you will look back over your initial posts and see if their appearances might have something to do with what you see as defensiveness on our part. We _do_ have a history of coming under direct and heavy "attack" by people taking positions close to those you stated in your initial posts. But as you can see, we have been eager to ask questions and learn more about your ideas, while at the same time admittedly putting up a semblance of a defense against what looked at least a little bit like another attack.

I hope you will find the interest and the time to reply to some of the specific questions I addressed to you.

> Aldus's thinking appears to have evolved since 1979 (one would hope). The article is not oriented around top down control as you appear to think of it.

In fact, in his series in 1979, Albus included some very vague and general ideas about other kinds of connections in his "hierarchy," but he never developed a model -- just a set of conjectures on how organisms (brains) produce specified outputs. His model was, and from what you say of the newer article it remains, a model for the production of predetermined outputs. I'll need to look at the 1991 article before I say much more about that, though.

> Indeed, it is far more bottom up in nature, and focused on coordination. Control policies or goals _may_ flow down the hierarchy, but summarized data and model results (error signals) pass up.

The error signals go up? From what? To what? What role(s) do they play? That's the opposite direction from the flow of error signals in the elemental PCT model. (A question and an observation; nothing defensive intended.)

> Both also move laterally in the network.

From where to where, with what effects? Are there quantitative implementations of this model? Does it actually behave, or is it only required to sound plausible? (I hope you don't find my questions "defensive;" I really do want to know.)

> It is not gospel, just a nice description of the workings of control system networks (once you get past the definitions and philosophy).

This comment seems to suggest that the Albus model really _hasn't_ changed all that much since 1979, at least not with regard to its nature (a descriptive "model") and its role (sounding plausible -- nice). Again, I hope you don't find my comments defensive; I'm merely stating my impressions drawn from the 1979 articles and your description of how you use the ideas from the 1991 article.

> BTW, I do not see how PCT could refer to a control system _Hierarchy_ is there is no kind of command relationship, or hierarchical filtering and summarization of data. That is part of the definition of 'hierarchy'.

Ah, that's why you might want to read some of the PCT literature before you say much more about the PCT model. You see, you've identified one of the amazing things about a hierarchy of simple control loops, each of which controls only its own perceptual signals -- such a hierarchy accomplishes feats thought impossible by those who are familiar only with traditional notions about hierarchies. Look at some of the literature, then let us know what you think.

By the way, if you do look at the PCT model, you'll find there are _no_ "commands," but that there is indeed a hierarchical "summarization" of perceptual signals. Living systems have their own "definition" of hierarchy and it doesn't much resemble the more common definitions in the behavioral and life sciences.

> If you deny the existence of hierarchical control, perhaps you should use 'tree network' instead.

But we don't deny the existence of hierarchical control, only of "command" hierarchies that control outputs. Living systems hierarchically control their own _inputs._ "Tree network" doesn't do, in that case; it's a full blown hierarchy.

> I think it is more likely that you are having a NIH reaction, focusing on minor distinctions as the most important, for the purpose of differentiation.

Not really. Read the literature; test the models; then let us know what you think about this idea.

> PCT aficionados do not have a monopoly on understanding, even if they may have produced working models.

That's for sure! But we are certainly in a position to reply to anyone who tells us our model is too simplistic and unsophisticated to work, or that it cannot really be of the kind we say, or that it is equivalent to word-models that have never been required to prove their worth by actually behaving in however simple a setting. And we will certainly reply if the someone also tells us he or she hasn't read any of our work or tested any of our models. That's hardly the NIH reaction.

Happy reading. ;-))

I'm serious. You have gone to the trouble to emerge out of the pool of lurkers and engage in some discussion. I'm sure all of us who have been on the net for a while would enjoy seeing your thoughts about the PCT model after you have learned a bit more about it.

Later, Tom

Date: Thu Jun 30, 1994 3:56 pm PST
Subject: Back in service; a few comments

[From Bill Powers (940630.1330 MDT)]

Paul George:

Welcome to CSG-L! I echo Tom Bourbon's comments with, perhaps, a little more diplomacy: I, too, think it would be a good idea if you were to become familiar with the PCT model before drawing any important conclusions about it.

It's important to understand the control-of-input idea, and why control of output doesn't work for real control systems in real (i.e., variable) environments. Human beings NEVER compute outputs: doing so would be futile. To see why, just imagine driving your car by planning the motor forces you would apply to the steering wheel while driving to work. When we speak of planning, what we're always speaking of are planned perceptual consequences, not the detailed actions that will bring them about. When you plan a route for driving to the store, the route isn't an output, but a series of perceptions you will achieve by whatever means is required when you get to the appropriate part of the plan. There's no way to predict exactly what output will be required to achieve a given perceptual result. Fortunately, control systems base their actions on error, not planned outputs, and so can produce the right action over a wide range of external conditions.

> BTW, I do not see how PCT could refer to a control system _Hierarchy_ if there is no kind of command relationship, or hierarchical filtering and summarization of data. That is part of the definition of 'hierarchy'.

The "command" relationship in PCT is not one of commanding actions, but of specifying the states of perceptual signals that are under control by lower systems. This specification is done by sending a reference signal to a lower system which is like a sample of the lower perceptual signal in the desired state: make it look (feel, sound, taste) like _this_. This is quite different from the way other theorists have treated the relationship between levels. You really need to read BCP to see how this works, and particularly how it can work at many levels simultaneously.

Best to all, Bill P.

Date: Fri Jul 01, 1994 8:24 am PST
Subject: Re: PCT models

[Paul George 940701 10:00}

A last comment on this thread. Thanks to all for the welcomes. From my point of view, what you are discussing in PCT is not particularly strange. The patterns are familiar, if uncommon in the psychological and medical arenas. There seems to be mostly a communication and terminology problem, at least with Tom. I am mildly amused when I am told I don't understand PCT, immediately followed in another post by an example of what I mentioned.

>[Bill Powers (940630.1330 MDT)]

> The "command" relationship in PCT is not one of commanding actions, but of specifying the states of perceptual signals that are under control by lower systems. This specification is done by sending a reference signal to a lower system which is like a sample of the lower perceptual signal in the desired state: make it look (feel, sound, taste) like _this_. This is quite different from the way other theorists have treated the relationship between levels. You really need to read BCP to see how this works, and particularly how it can work at many levels simultaneously.

<Bob Clark (940630.1450 EDT)>

> [For convenience, here is a quick summary of HPCT concepts: 1) "control of perception" and 2) a "hierarchy of levels of control" 2a) "hierarchy" is defined as a relation between "levels" in which "higher levels" achieve their objectives by selecting and activating goals, "reference levels," for lower level systems.]

Hey guys, this is exactly how distributed control systems and software systems are architected, _and_ is the architecture described by Albus in the article I cited. It's not particularly revolutionary in the engineering field. In a sense the model has been validated not by simulation, but rather by construction. Of course that just means the architecture works, not that is of necessity the mechanism used by nature. However, your low level simulations do appear (from discussion and description) to be well on the way to 'proving' the PCT approach at the neuro/muscular and sensory perception level.

>[From Bill Powers (940630.1330 MDT)]

> It's important to understand the control-of-input idea, and why control of output doesn't work for real control systems in real (i.e., variable) environments. Human beings NEVER compute outputs: doing so would be futile.

> From Tom Bourbon [940630.1645] {offline communication} The PCT model does not control its outputs. Neither do living things. Again, your comments are uninformed, as they must be, given that you have read none of our work.

Perhaps there is confusion about the term 'output'. Control software (or firmware) produce outputs (signals) to actuator mechanisms which produce the 'environmental' output. I know of no one in engineering who suggests that the control system directly influences the external environment. If I try to catch a ball, I do not control eye tracking or arm/finger muscles. When I attempt to, I miss.

If I am constructing control system software, I have sensors and data collectors that sample and translate (A to D) the sensor output. Even the sensor does not directly provide the environment's condition, just some kind of transducer reading. Similarly if I wish to move a robot arm, the control system might output a command to a particular joint to move to a particular coordinate. Some mechanism translates that to providing a current or fluid flow for a period of time at certain points (this may be deemed another control node if you wish). This translation is based upon some model of what _should_ happen to the arm when this occurs, in terms of position or velocity. However, if there is no power or hydraulic fluid, or if the arm is blocked, nothing actually happens, while the control system 'thinks' it has. This is why we have other sensors to detect actual motion, though they too can fail or be spoofed.

Further, the control software actually has objects that represent the sensors and arm components being controlled. It is these 'simulations' or 'models' of the real world entities with which the control logic interacts. This, it appears, is PCT's controlled perceptions. For process control systems there are models of the physical process being controlled, and of the control system itself. One of the hard engineering problems is making sure this model stays synched with reality. As the old saw goes "the map is not the territory".

Enough for now. You guys aren't as alone as you might think, though perhaps a bit insular. Other domains deal with the same kind of issues, and have come up with solutions and approaches which may apply.

Paul C. George

Date: Fri Jul 01, 1994 12:08 pm PST
Subject: Re: PCT models

[From Bill Powers (940701.1100 MDT)] Paul George (940701.1000)

It's fascinating how one can read into expositions what one expects to find there. You seem convinced that PCT is just the same thing that control engineers are already doing, yet in explaining back to us what they do you are describing something very different. Not your fault; this is a common property of human beings. In fact, a few control engineers have told us that the PCT approach in control engineering is new to them and suggests a very different way to design control systems -- even though the equations might be exactly the same, as they sometimes are!

Maybe I'm misunderstanding you, but here is what I based my comments on:

> Perhaps there is confusion about the term 'output'. Control software (or firmware) produce outputs (signals) to actuator mechanisms which produce the 'environmental' output. I know of no one in engineering who suggests that the control system directly influences the external environment. If I try to catch a ball, I do not control eye tracking or arm/finger muscles. When I attempt to, I miss.

Good guess, but not the point we are making. The actual point comes up in the following:

> If I am constructing control system software, I have sensors and data collectors that sample and translate (A to D) the sensor output. Even the sensor does not directly provide the environment's condition, just some kind of transducer reading. Similarly if I wish to move a robot arm, the control system might output a command to a particular joint to move to a particular coordinate. Some mechanism translates that to providing a current or fluid flow for a period of time at certain points (this may be deemed another control node if you wish). This translation is based upon some model of what _should_ happen to the arm when this occurs, in terms of position or velocity.

This is exactly the conventional control-of-output conception of control that we are deviating from. In the conventional engineering view, the problem is to cause some objective effect to appear in the environment (like moving a gripper to a preselected objective position). One way to do that is to compute the outputs that must be generated in order to have that effect, including inverse kinematics and dynamics and any effects from disturbances that might be anticipated, and then to send the output signals to the transducers to produce the effect. Feedback then is used largely to trim the dynamic characteristics of the system for stability and to provide resistance to disturbances.

The PCT approach starts at a different point in the loop: not the objective effect, but the perceptual representation of it. This means that feedback is more than just a way to trim the system response and take care of disturbances: it is the only way a living organism can know ANYTHING about the effects its outputs are having in the external world. When a living control system controls, all it can control is the perceptual representation; it has no auxiliary channels through which it can know the objective effects of what its outputs are doing to the world.

This means that when, as engineers, we design a system from the PCT standpoint, the first step is to create a perceptual signal that reflects the state of the external variable to be controlled. As designers we have the advantage over the organism in that we can preselect what the external variable is, but in the spirit of PCT we follow the same strategy to which the organism is limited: we set up a perception that is the real controlled variable. In fact, what is perceived is what is controlled; if we want to make sure that some objective variable is controlled, we have to make sure that the perception accurately reflects its state. This is not quite the same problem that the organism has, for from the organism's point of view, what is perceived is always what is controlled, and the organism usually doesn't know what external situation corresponds to control of that perception.

Given that the perceptual signal is what we want to control, we now have a simple problem on the input side. The reference signal specifies the desired state of the perceptual signal, and the comparator detects the error. If the perceptual signal exactly matches the reference signal, the external counterpart of the perceptual signal -- which the engineer probably calls the controlled variable -- will be in a particular state. All of the design problem that remains is then that of devising an output function that will (a) create the necessary environmental effects to complete the control loop, and (b) introduce the dynamic functions needed to achieve stability in the presence of disturbances and changes in the reference signal.

In a PCT control model, two features of the standard engineering designs are missing. First, there is no model of the environment contained in the control system (although of course there must be one in the simulation as a whole). Second, there are no inverse calculations, either of kinematics or of dynamics. The system is set up so that the error signal is translated into output effects that make the error smaller; by suitable choices of controlled variables (including hierarchical relationships), this is sufficient to guarantee tight control and stability.

As an example, consider "Little Man Version 2," a model of limb control I have been working on for a couple of years. I took my design straight from neurophysiology, using a model of the muscle and of the tendon and stretch reflexes that is a simple and literal rendition of how the neuromotor system is actually wired. If there's anything clever in this design, give the credit to nature.

The lowest-level control loop is the tendon reflex, which senses the force applied by a muscle through a tendon to an attachment to a bone. The force is sensed directly as a neural signal which is compared with a reference signal, and the error is amplified to shorten or relax muscle fibers. Shortening stretches a series elastic component to produce the force that is sensed, closing that loop. This system causes the applied force to track the reference signal.

The second loop is a rate-control loop, and the third is a position- control loop. The position error determines the rate reference signal, and the rate error determines the force reference signal -- a three-level hierarchy, when decomposed this way. In the real system, the two higher loops are combined into one, with the rate control being combined with position control and determining the damping of the system.

What's interesting about this model is that stability is achieved without any special attention. With acceleration, velocity, and position all under specific and tight hierarchical control, the arm dynamics (which are included in the model of the system's environment) are almost perfectly compensated, without any need for inverse dynamic calculations, Jacobians, or all that stuff. Of course a real control engineer could probably work out how this design actually takes the required inverse calculations into account, but in fact the control system model itself does no such calculations. The dynamical equations appear only in computing the response of the arm to torques applied at its joints, which is really a model of the environment, not of the control system.

There is nothing in this model that could be construed as calculating the desired effect of an output and then producing the signals that will create that desired effect. The output calculations, in their entirety, consist of leaky integrations. What has struck some control engineers as interesting about this model is its extreme simplicity in comparison with other approaches. All six control systems involved in controlling three degrees of freedom in the arm are expressed as no more than 75 lines of rather loose C code, with no matrix operations or other shortcuts. The other 5000 lines of code are concerned with graphic presentation, facilities for parameter adjustments and test modes, and modeling the physics of the arm.

There is one sense in which it can be said that the Little Man model incorporates a full-time running model of the environment. As far as the control systems are concerned, the perceptual signals ARE the environment.

Each perceptual signal is a different function of environmental variables, and so represents the state of the variables as the variables in a running model represent the state of something else. We could say that the model is there -- but from the standpoint of the control system, it is simply "reality."

In the final analysis, PCTers and control engineers are really talking about the same thing. The main difference is in how one parses the system before modeling it. There's nothing wrong with the way control engineers go about their business, although I sometimes think that they have overcomplicated some problems, but the way they go about their business isn't the right approach for modeling organisms, which is what PCT is about. There is no friendly engineer standing by to tell an organism what it is really doing to its environment when it controls its own perceptions. Whatever an organism can figure out about the external world has to be done entirely through information carried in its perceptual signals, which are its feedback signals, and functions of those perceptions, and functions of those functions. It's as though the organism's brain lived in a control room lined with meters and levers, but the meters carry no labels showing what their readings represent and the levers have no labels describing what they do to the world outside the control room.

I have always hoped that real control engineers would lend their expertise to expanding the PCT approach. Now some of them are actually trying it out, and I hope to see more sophisticated analyses than I can produce. The first hurdle has always been to convince control engineers that there IS another approach, and even though it is mathematically equivalent to standard ones, it is also very different -- and simpler to implement -- in many regards. It's more important to understand the differences than to see the similarities.

Best, Bill P.

Date: Fri Jul 01, 1994 2:18 pm PST
Subject: Re: PCT models

From Tom Bourbon [940701.1601] >[Paul George 940701 10:00}

> A last comment on this thread.

I hope it isn't concluded. Much remains to be discussed and clarified.

For example, in my first post or two, I asked you how you go about assessing models and theories -- which criteria and procedures do you use before you decide a theory or model is "good" or that it "works?" For one thing, your answers to those questions would help me, and I'm sure many others on this net, to prepare appropriate replies to your messages. A general answer would help, but a specific example of how you used those criteria and procedures to assess Albus's model would be even better. I hope your announcement of the end of this thread does not preclude an answer to those questions.

> Thanks to all for the welcomes.

Thanks to you, for breaking your silence.

> From my point of view, what you are discussing in PCT is not particularly strange. The patterns are familiar, if uncommon in the psychological and medical arenas. There seems to be mostly a communication and terminology problem, at least with Tom. I am mildly amused when I am told I don't understand PCT, immediately followed in another post by an example of what I mentioned.

I hope there isn't a serious communication problem between us. My initial responses to you might have seemed abrupt -- perhaps that is my style when I am trying to catch up on a stack of accumulated mail, answering piece by piece, rather than waiting until I've read the entire pile. My intent was to suggest that your impressions about PCT, as it is named and discussed on this net, might not be accurate. I was not trying to imply either malice or stupidity, or anything akin to them, on your part. I _was_ reminding you of your own admission (which I respect) that you had not read material on PCT, hence, your ideas about it, and your attempts to compare it to other ideas (Albus's theory, for example) might suffer by virtue of your lack of exposure to PCT. I believe the replies you received from participants on this net _all_ have been civil and welcoming, and that all have emphasized the idea that an examination of the material on PCT might be important for you, if you wish to enter the discussion of the phenomenon of control and of PCT.

I'm puzzled by your statement that you are "mildly amused" to read that you are uninformed about PCT, as it is discussed on this net. You did say you had never read anything about it. In my case, I'll freely admit that I'm ignorant of (uninformed about) the actual contents of the Albus paper in 1991. I believe my ignorance and state of uninformedness on the paper disqualify me from making _any_ pronouncements about his theory or his model, circa 1991, although I do know his ideas circa 1979 and they do not seem to differ dramatically from what you have said about the more recent article. But I will make no further comments about his "model" or about your interpretations of it until after I have read it this weekend.

Uninformedness and ignorance are merely descriptors of a state of absence of knowledge about a particular topic. The terms do not imply a value judgment about the person who is ignorant or uninformed on that topic.

Later, Tom

Date: Fri Jul 01, 1994 4:03 pm PST
Subject: Re: PCT models

[Paul George 940701 15:00 EDT]

OK I'll bite. Principae Mechanica is not a prerequisite for catching a ball.

[From Bill Powers (940701.1100 MDT)]

> It's fascinating how one can read into expositions what one expects to find there.

> This is exactly the conventional control-of-output conception of control that we are deviating from. In the conventional engineering view, the problem is to cause some objective effect to appear in the environment

It seems to work on both sides. You conveniently ignore the later description of process control software and intermediary steps which match exactly your discussion.

> This means that when, as engineers, we design a system from the PCT standpoint, the first step is to create a perceptual signal that reflects the state of the external variable to be controlled. {snip}

> What's interesting about this model is that stability is achieved without any special attention. With acceleration, velocity, and position all under specific and tight hierarchical control, the arm dynamics (which are included in the model of the system's environment) are almost perfectly compensated, without any need for inverse dynamic calculations, Jacobians, or all that stuff. ... The dynamical equations appear only in computing the response of the arm to torques applied at its joints, which is really a model of the environment, not of the control system. {snip} There is nothing in this model that could be construed as calculating the desired effect of an output and then producing the signals that will create that desired effect. The output calculations, in their entirety, consist of leaky integrations. What has struck some control engineers as interesting about this model is its extreme simplicity in comparison with other approaches.

I said the control system defined a desired result - "move to particular coordinates". Such result is determined from the inputs. The control node has access to a 'perception' of the external world entities which are the mirrored objects (we call this a 'process centric view'. Process means something like 'make chemical with a certain recipe'. 'Self' is one of those entities). It is a set of instance variables. There also exist reference variables to which this object 'state' is compared (albeit by changing formulae and sets of 'controlled variables').

I said a 'mechanism' exists to generate outputs to bring the 'perceived' state in line with the 'reference' state. I said nothing about reverse kinematics, nor of feedback, nor of computing outputs _within the control node_, though this is a conventional mechanism for robot control. Your error correction method in the 'Little Man' simulation is just another (though more elegant) mechanism. I would like however to see it directly applied to a robot arm.

> In fact, what is perceived is what is controlled; if we want to make sure that some objective variable is controlled, we have to make sure that the perception accurately reflects its state.

Yes, as I said, the control signals (error signal) are sent to the mirrored object which reflects the real world entity. That object then has methods to affect the outside world through some chain of transducers and actuators. Measurement theory states that a measured value need to have a known relationship (not necessarily liner) to the thing being monitored, within a particular range, such that it is useful for prediction or control. A 'model' of sorts is needed to determine what output will cause a change of the controlled variable in the desired direction to diminish the error signal. The model _can_ of course be a simple mathematical function. It could however be a complex hierarchical finite state machine (which I think _could_ be implemented as a network of simple control loops).

> The other 5000 lines of code are concerned with graphic presentation, facilities for parameter adjustments and test modes, and modeling the physics of the arm.

These also must exist in a control system at above level 2 (0=AtoD conversion, 1 = PLC, 2=local net or process, 3=plant, 4=enterprise). At the level of nerve and muscle your mechanism works well, and probably does map to reality. However at higher levels _everything_ is virtual. You are usually 'controlling' lower level's reference variables and comparison algorithms (though you may have higher level controlled variables). They are also more focused on planning or predicting _future_ actions (or results) based upon extrapolation of current trends. The key idea is that granularity of data (input & output signals, reference & control variables) and time horizon increase as you go up the hierarchy. In my terms, your position, rate, and neuro/muscular loops are on the same level, though with a directional control flow. Note that in the (3D) Albus(perhaps giving him too much credit) architecture, nodes on the same level interact directly.

BTW - have you tried to cause the object being tracked in your simulation to change velocity in 3 space randomly? I suspect the little man would have trouble pointing at the target. You might find that an environmental model incorporating the physics of motion was needed. To make the problem harder, we observe that trying to 'catch' objects under water, or in lower gravity, or different coriolis force, is difficult and requires re-learning because things don't move as we expect. What happens if we make the little man observe in a mirror? How do you correct your output signals when the error signal doesn't diminish as expected? {My apologies if these are already dealt with, I haven't yet received Dag's package}

How about at the level of verbal communication? I construct some set of statements to influence your behavior. This is based both upon my perception of your response, my understanding of language, and my predictions of your reaction. This seems a bit beyond a simple error signal as there is considerable lag and indirection. It does seem based upon some set of predictive world models, or if you prefer, simulations. I compare the results of my model to my perceptions, and modify the model accordingly. (here control system = conscious & unconscious mind). However, I can view these models as control networks themselves encapsulated within a control system.

Hope there is a spark of something new here, If nothing else that different views or terminology can apply to the same thing. See any responses Tuesday

Paul George

Date: Fri Jul 01, 1994 7:39 pm PST
Subject: Re: PCT models

<[Bill Leach 940701.21:58 EST] >[Paul George 940701 15:00 EDT]

Paul, with Bill in the discussion I am reluctant to comment since his postings are so concise and usually outstanding... however, I will anyway if for no other reason than that typically I learn something.

Control system engineers should understand fundamental PCT quite rapidly. Understanding the deep implications of PCT might take quite a bit more time and effort.

For control systems engineers, it is important to remember that:

1. PCT is closed loop negative feedback control only. No open loop stuff.

2. The control system IS everything. There is no possible exact independent check of performance (at least for higher levels).

3. The control system IS adaptive.

4. The control system ALWAYS sets it's own references.

> BTW - have you tried to cause the object being tracked in your simulation to change velocity in 3 space randomly? I suspect the little man would have trouble pointing at the target.

No doubt however, you provide the answer to your own question:

> ..., we observe that trying to 'catch' objects under water, or in lower gravity, or different coriolis force, is difficult and requires re-learning because things don't move as we expect.

It is quite likely that humans don't do any "calculations" for issues of motion. That many have the ability to recognize certain types of motion is itself probably a learned ability that is independent of other control system functions.

-bill

Date: Sat Jul 02, 1994 4:44 am PST
Subject: Re: PCT models

[From Bill Powers (940702.0400 MDT)]

[Returned from conference in Wales] The internal clock is slowly creeping back into synch; good thing I don't have to go to work.

Paul George (940701.1500 EDT)--

We may slowly be converging to a common description of the control phenomenon in behavior. You should realize that when I began this work some 40 years ago, my target audience was in the life sciences, and I simply hoped to communicate what control engineers knew to people working (unknowingly) with living control systems. So I've always hoped to see real control-system engineers getting into this act. I would be quite alarmed to find that PCT and control engineering had nothing in common!

I have not really ignored your comments on process control. My problem is simply that in the language you use to describe control, you use terminology that is like what others have used in expressing a theory of control that is very different from PCT. It isn't easy to see through language to the underlying concepts behind it; the language may suggest divergences where there are none, but it can also suggest agreement where there is actually a divergence. So we just have to slog through this and make sure we are talking about the same things. Words don't always convey to a listener the meanings that the speaker had in mind.

Here is an example:

> I said the control system defined a desired result - "move to particular coordinates".

As I read this sentence, it says that the control system moves something to particular external spatial coordinates, in an objective coordinate system. This is not how we would describe the situation for an organism in PCT. What we would say is that the control system acts on the world in such a way that perceptions of position -- i.e., signals representing position in approximately orthogonal ways -- come to particular values. The measure of position is not a meter-stick reading obtained by an external observer in the environment of the control system, but a set of perceptual signals existing inside the control system. So even at this level (in fact, at every level), the world being controlled is, as you say, "virtual."

You then add

> Such a result is determined from the inputs.

Here's the seeming problem. As you describe the situation, there is first an objective position, produced by actions of the effectors. Then this position is represented in the system by signals arising from its input sensors. From this point of view, it is possible for the perceptual representation of position to be incorrect or distorted. But from the organism's point of view, the perceptions ARE the position; what is controlled is not the objective position, but the perception. Given that the perceptions in each coordinate are maintained at particular values, the objective position that corresponds to those values is simply the inverse input function of the perceptual signals. In general, when the reference signals for position are varied, the resulting position in external coordinates will change along curved lines, the curvature being determined by the nonlinearity of the sensing apparatus (for example, the "horopter" of visual three-dimensional space). So the space in which the organism controls position does not coincide with the Cartesian or polar coordinate systems that an external observer might use to characterize behavior.

You say

> The control node has access to a 'perception' of the external world entities which are the mirrored objects (we call this a 'process centric view'. Process means something like 'make chemical with a certain recipe'.

Good: the process-centric view is then what we refer to as control of perception. However, the strictly correct way to describe the process would then be 'make the perception of a chemical appear by manipulating perceptions of a certain recipe.' In the purely process-centric view, the laws of chemistry and the objective consequences of carrying out a given recipe are known to the controlling system only as perceptual signals; what is actually happening may be different by a small or large amount. From the standpoint of a controlling process, the objective world is known ONLY through representations in signals generated by its input transducers, and computations performed with those signals as arguments. Furthermore, what is controlled can be only the representations; the external correlates of those representations may not be what an external observer assumes.

I don't follow this: did you say what you mean?

> There also exist reference variables to which this object 'state' is compared (albeit by changing formulae and sets of 'controlled variables').

I understand comparing one signal against another, but how is such a comparison done by "changing formulas and sets of controlled variables?" Did you really mean that the _comparison_ is done by these means?

> I said a 'mechanism' exists to generate outputs to bring the 'perceived' state in line with the 'reference' state. I said nothing about reverse

kinematics, nor of feedback, nor of computing outputs _within the control node_, though this is a conventional mechanism for robot control.

Sorry if I misread you. What alerted me to a possible disagreement was the fact that others have proposed "mechanisms" designed to bring the objective variable, and thus the perception of it, TO a reference state rather than simply TOWARD it. This interpretation has turned up in many places, among them cognitive psychology and the motor control literature. The underlying model relies on assessing the situation, then computing what action is required to bring an external variable to a goal-state, and finally executing the program that will accomplish the pre-planned result. That is how the need for computing inverse kinematics and dynamics arises. Where feedback is involved, it occurs after the pre-planned movement, and serves to inform the system about the results, so that the next movement can be planned. This type of control loop might appear in some circumstances, but is far from adequate as a general model and is not a particularly effective mode of control. It is necessarily very slow because of all the calculations that need to be done to make the results come anywhere near the desired goal, and it relies on the absence of disturbances occurring between the initial assessment and the final execution of the plan.

I now take it that this is NOT the model you are speaking of.

> Your error correction method in the 'Little Man' simulation is just another (though more elegant) mechanism. I would like however to see it directly applied to a robot arm.

Thanks. I would like to see it applied that way, too. In fact, that application might impress grant-givers enough for them to supply the funds necessary to develop that application. Catch-22. In the meantime, the best I can do is a careful simulation.

> ... as I said, the control signals (error signal) are sent to the mirrored object which reflects the real world entity. That object then has methods to affect the outside world through some chain of transducers and actuators.

If I understand this rather terse statement as you intend it, this seems to be somewhat like the concept of hierarchical control in PCT. The error signal from a higher-level control system, after passing through an output function, fans out to enter a set of lower-level control systems each of which is concerned with controlling one dimension of the perceived world. However, it's not clear in what sense your "mirrored object" reflects the real world entity. In hierarchical PCT (HPCT), all that is represented explicitly, in the form of signals, is the value of a variable: the properties of the external world are not represented as signals. We have toyed with the idea of actual analog models of the external world being used by the control system, largely to fill in missing information, but so far have not come across a case that requires this (not to say that such cases will not be found).

The lower-level "object" in HPCT consists of a lower level of control systems, plus the environment with which they interact. In the absence of a literal model, that lower-level environment serves as its own model; that is, it responds to output signals from the control system by altering input signals in the control system. Any mirroring of the properties of the external world would be contained in the organization of the perceptual (input) and output functions of the control system -- for example, the appearance in the output function of the complex conjugate of external dynamics (to use a term from the old frequency-domain days).

You seem to agree with this when you say

> A 'model' of sorts is needed to determine what output will cause a change of the controlled variable in the desired direction to diminish the error signal.

I think that a "model of sorts" is an improvement over the more literal concept of a complete analog model of the environment, which some have proposed but doesn't seem practical to me.

Do I detect a little influence here from OOPS? One has to beware of letting programming languages dictate theories of behavior!

> These also must exist in a control system at above level 2 (0=AtoD conversion, 1 = PLC, 2=local net or process, 3=plant, 4=enterprise).

This sounds reminiscent of Albus. We use an entirely different set of conjectures about the nature of higher-level control processes. Albus' definitions of levels are not tied very well to real behavior and experience -- they're more like an engineer's assessment of an objective problem in plant control.

I've been concerned with trying to identify broad types of variables that people seem actually to control, relating them as nearly as possible to neurophysiology and direct experience. What is referred to above as "A-D conversion" occurs in the HPCT model at the 7th proposed level of control. I do not consider that the basic sensing process is an A-D conversion: it is a conversion from stimulus amplitude to impulse frequency, which is still an analog measure. The conversion from analog processes to digital processes, I have guessed, occurs where we create categories of perception out of the continuum of analog perceptions at lower levels. This is the level where symbols come into existence, representing either-or classifications rather than continuous processes. Above that level we begin to have what is usually described as a finite-state machine -- a machine that can carry out discrete sequences of acts and that reasons in terms of arithmetic, logic, grammar, and other such program-like rule-driven methods.

I'm sure that within the levels currently defined under HPCT, we could account for the things referred to in your Albus-like list above, but obviously we are slicing the pie along different planes.

> At the level of nerve and muscle your mechanism works well, and probably does map to reality. However at higher levels _everything_ is virtual.

Actually, in the HPCT model, everything is virtual at all levels, as I said earlier in this post. It isn't that the perceptual hierarchy "maps to reality" -- it _creates_ reality as we know it, out of the raw input of stimulus intensities at the lowest level. As a friend of mine said, the organism isn't the black box; the environment is. All that an organism can know of reality comes from seeing what happens when it emits output signals into the world beyond the senses, and observes the resulting changes in its own perceptual signals. In our modeling, we use physical models of the external world derived by the same means, and look for correspondences between that physical model and the world we experience as reality. These worlds are not very similar, although there are many points of contact. The ultimate criterion is not what the physical models say, but what experience tells us directly. Physics is valid only to the extent that it correctly predicts experience (as, I should add, it does exceedingly well).

> You are usually 'controlling' lower level's reference variables and comparison algorithms (though you may have higher level controlled variables).

I would say "varying" rather than "controlling" lower-level reference variables. What we control are input signals; we do this by _varying_ outputs (reference signals for lower level systems) as required by variations in environmental conditions (disturbances).

I should mention that the PCT model works strictly with scalar perceptual variables; vector inputs are reduced by higher-level input functions to the values of one-dimensional variables such as distance and angle. The result is actually consistent, pretty much, with models expressed in matrix notation, the main difference being that in the HPCT model all the matrices are presented as completely expanded. When you think about it, that has to happen in any real system anyway: the matrix notation is just a way of handling a lot of more detailed, and simpler, processes in a convenient way. The real system still has to do every individual operation implied by the matrix operations (just as a computer must).

The reason I mention this is that when control systems are seen as working with scalar perceptions only, there are no complex "comparison algorithms." Comparison is simply subtraction. All the complexity that would appear in other modeling approaches happens in the perceptual input functions at each level, which apply algorithms that extract, in parallel, many different scalar measures from the common set of lower-level perceptual signals. The general rule in HPCT is "one perception, one control system." Complex control is achieved by many such systems operating in parallel at the same level. I don't know how long this concept will survive; right now it seems to work pretty well in terms of experimental tests.

> They are also more focused on planning or predicting _future_ actions (or results) based upon extrapolation of current trends.

I agree that there are cases in which planning and extrapolating do occur, but generally the kind of control you get from this approach is poor. Real environments are too variable to permit much more than statistical success from this method. Where it's possible, and I think it is possible much more often than some people seem to think, a better approach to controlling is to get the component variables under individual control at a lower level, then simply control higher-level derived variables relative to specific reference levels: don't try to _predict_ the future, _control_ it. Where you can do that, it works a hell of a lot better than making predictions and plans. A compromise method is contingency planning, which is much more like real control. Instead of deciding what you are going to do in complete detail, you set up control systems to handle a range of possible occurrences, without trying to predict which of them will occur.

> The key idea is that granularity of data (input & output signals, reference & control variables) and time horizon increase as you go up the hierarchy.

Well, that may be true, but it's not a very useful concept for making models. It's sort of like Ashby's "Law of Requisite Variety," which says that the variety (possible output states) of a control system must at least equal the variety of its environment. A true statement, but you can't use it to design a system having the requisite variety. It's an after-the-fact description.

> Note that in the (3D) Albus (perhaps giving him too much credit) architecture, nodes on the same level interact directly.

While this idea could be added to the HPCT model, we haven't tried it, because so far we haven't done any experiments complicated enough to call for it. Some day, maybe.

> BTW - have you tried to cause the object being tracked in your simulation to change velocity in 3 space randomly? I suspect the little man would have trouble pointing at the target.

I didn't mention the visual control level in the Little Man, which uses binocular vision to generate x, y, and z signals for the visually-observed positions of the fingertip and a movable target. When you get the demo disk, you'll see. The Little Man tracks an arbitrarily moving target with its fingertip in real time. Version 1 does this better than Version 2 (the one with full dynamics, muscle model, reflex loops) because in version 2 the dynamics of the visual system are poorly implemented in relation to the arm dynamics. Still no inverse kinematics or dynamics, however, and they won't be needed even when I clean up the top level of the model (some day real soon now).

> You might find that an environmental model incorporating the physics of motion was needed.

Well, such a model exists in version 2, converting the torque outputs of the control system model into motions of the physical arm, with all the interactions, Coriolis forces, etc. . The model of the control system, however, does no such calculations. The program only does them to show how the signals cause the physical arm to move, turning the control system's outputs into perceptual inputs of joint angles. Only forward dynamics and kinematics are used to simulate a real arm.

> What happens if we make the little man observe in a mirror?

It would be necessary to switch a sign-inverter into the pathway carrying the x position signal (or the error signal) from each eye. Otherwise positive feedback would occur in the x direction. That would take another level of control.

> How do you correct your output signals when the error signal doesn't diminish as expected?

Since there's no prediction of how the error signal will diminish, there's no expectation (in the model, that is -- in me there is a definite one!). There's a provision in the model for switching gravity on and off: it makes hardly any difference in behavior, although you can see the muscle forces changing to compensate. The model handles disturbances in the usual way.

> How about at the level of verbal communication? I construct some set of statements to influence your behavior. This is based both upon my perception of your response, my understanding of language, and my predictions of your reaction. This seems a bit beyond a simple error signal as there is considerable lag and indirection.

We've talked about such things now and then on the net, but haven't reached any conclusions about how to model the situation. The big problem, of course, is that we don't understand how language works. That kind of modeling is a long way in the future.

> It does seem based upon some set of predictive world models, or if you prefer, simulations.

I like "simulations" better. Yes, this is true. It's on the agenda, for my grandchildren (or someone's). Lots of fertile ground there for research.

Best, Bill P.

Date: Tue Jul 05, 1994 10:46 am PST
Subject: Albus waits

From Tom Bourbon [940705.0924]

Last Friday I said that over the weekend I would remove my state of ignorance and uninformedness about Albus, 1991, but it turns out I was also ignorant and uninformed about the hours of the local university libraries during the Independence Day weekend. Albus 1991 must wait until this evening.

I did re-read the Albus series that ran Byte, 1979, in parallel with the series by Bill Powers. Now I see why I never took seriously anyone's suggestion to read Albus, in more recent years. There are a few points in common between Albus and Powers -- some of the _topics_ they discuss and some of the _words_ they use (including hierarchical perception and control [but of what and by what?], feedback [but of what, to what?], the analog nature of synaptic events in nervous systems, the idea that _all_ behavior [no matter how "cognitive" it may seem] is motor behavior, and so on) -- but the differences are glaring. I'll summarize my interpretation of those differences after I've read Albus, 1991; perhaps he made some radical changes in his "model" since 1979 although I doubt that is the case.

Later, Tom

Date: Tue Jul 05, 1994 10:52 am PST
Subject: Re: PCT models

[Paul George 940705 09:30 EDT] [Bill Powers (940702.0400 MDT)]

I am not really a control system engineer, though I have some experience with programming robot controllers. I only recently joined a control systems company, and the people I am working with are kind of on the leading edge of control theory. I may thus have a biased view of how control engineers think. By profession I am an engineering process and methods specialist (i.e human processes).

> As I read this sentence, it says that the control system moves something to particular external spatial coordinates, in an objective coordinate system.

There may be motion in the real universe, but the control system can never KNOW. It has an internal coordinate system it uses as a reference grid. It orients things it perceives in its environment in that system. But the system only detects the 'outside world' through it's senses.

>> There also exist reference variables to which this object 'state' is compared (albeit by changing formulae and sets of 'controlled variables').

> I understand comparing one signal against another, but how is such a comparison done by "changing formulas and sets of controlled variables?" Did you really mean that the _comparison_ is done by these means?

Yes. In the higher levels of distributed control systems what constitutes an error is a function of history and sets of control variables. Which inputs are important changes, as do the reference values. The system is using something like statistical process control(as you note later) to determine whether the environment appears to be changing in the desired directions, or is going out of control (a fuzzy concept). The system is trying to predict the future and take pro-active action.

Re actual application to a robot. I vaguely recall that there were some toy robot arms that are programmable. There are also some computer games where one programs a robot's behavior to interact with other robots. These would be a lot cheaper that an industrial quality robot. Sorry I can't give more solid references. BTW I don't know if you have seen it, but some research has successfully used chaos formulas to allow a multi-legged robot insect to walk without directly controlling the individual leg motions. The old joke about the centipede tripping when it paid attention to its feet appears to have a grain of truth.

> However, it's not clear in what sense your "mirrored object" reflects the real world entity. ..... We have toyed with the idea of actual analog models of the external world being used by the control system, largely to fill in missing information, but so far have not come across a case that requires this (not to say that such cases will not be found)

Objects were being used in the Object Oriented Analysis & Design sense. By convention we try to have objects in the control system software that correspond to the real world entities. They 'own' both the reference variables (in our terms instance variables) and the functions (methods) which may be used to influence the object. The object interacts with the available sensory data to 'set' the instance variables, often algorithmically (we can rarely directly measure what we are interested in). We must 'reverse engineer' the real entities' state based upon a chain of indirect information. A thermocouple produces a voltage, which is converted into a digital 'count'. This is input to the object and converted to a temperature value. However, often we will average the results of several sensors to infer the temperature of the material in the vessel. Further, we are usually interested in the temperature trend, which involves assumptions about the heat capacity of the material and the energy of the heaters.

> I do not consider that the basic sensing process is an A-D conversion: it is a conversion from stimulus amplitude to impulse frequency, which is still an analog measure.

Ack. A to D conversion is an artifact of computers and electronics. I would be surprised to find it (as a general mechanism) in living systems. My weak understanding of nerve dendrite interaction leads me to believe that the signals are fairly complex.

Note that my descriptions are of what plant control systems use. It may or may not apply to living systems, though it hints that data and modeling mechanisms _may_ be an efficient architecture at higher levels of control(particularly with multiple entity systems - e.g. social). OTOH, biological analogies are becoming fashionable in the control industry.

>> The key idea is that granularity of data (input & output signals, reference & control variables) and time horizon increase as you go up the hierarchy.

> Well, that may be true, but it's not a very useful concept for making models. It's sort of like Ashby's "Law of Requisite Variety," which says that the variety (possible output states) of a control system must at least equal the variety of its environment.

It is useful, and basic. The idea is that the sensing rate and the 'reaction time' slows as you go up the hierarchy. Higher levels have a wider 'span of control' (Management Science(?) concept), care only about 'summary information' and have a longer time horizon. Higher levels can sample the control variables of lower levels (among other things).

More directly to the point, one of the key design problems of safety critical systems is determining the complete state space of the environment, and insuring all unsafe states are detected. Fortunately, most states and events are of the 'don't care' variety. The key concept of a model(and a control system) is that it handles the 'important' behavior and variables of the real world system.

>> What happens if we make the little man observe in a mirror?

> It would be necessary to switch a sign-inverter into the pathway carrying the x position signal (or the error signal) from each eye. Otherwise positive feedback would occur in the x direction. That would take another level of control.

Precisely my point. How does the 'little man' tell that a mirror has been inserted? How does he (or more fairly a sophisticated control system like an organism) detect the positive feedback and correct the comparison algorithm(I look at another level of control as a modification of the control logic)? As I have heard it tossed around, the concept of 'reorganization' of the control system seems a bit like a Deus ex machina. There needs to be some mechanism to detect that the system is out of control (or trending that way) and that the control system needs to adapt. Something like a deductive mechanism would appear to be needed.

> It's on the agenda, for my grandchildren (or someone's). Lots of fertile ground there for research.

What?? You _can't_ give us all the knowledge of the world while standing on one leg? How disappointing :-)

(From Tom Bourbon [940701.1601])

> I asked you how you go about assessing models and theories -- which criteria and procedures do you use before you decide a theory or model is "good" or that it "works?"

I thought I had answered. We create a model when we can't handle the complexity of the real world, or if we need answers in faster than real time. The model should exhibit the same behavior as the modeled system or allow its computation, in terms of the factors of interest (aye, there's the rub). A model is good IF it is useful for the intended purpose. It need not be accurate, in the sense of using the exact mechanisms as the real world, unless that is its intent. It almost by definition cannot be complete in the sense of capturing all the real system's complexities, again unless that is its intent.

Models can be for several purposes. One variety is for understanding, as in an analogy, Einstienian thought experiment, or scientific theory. Such models allow us to reason about how the real system should react, and deal with currently 'unexplained' effects. Another type mimics a portion of the real system in order to allow analysis or temporal prediction, as in a simulation.

The model is good if it allows us to come to the 'same' conclusions as the real world system, or allows us to obtain our desired results. A theory should explain all observed behavior of interest. A simulation's outputs should be related in a known (deterministic) way to the real system's outputs, within a certain range of inputs and outputs.

Alchemy works to a certain extent, though quantum mechanics works better for explaining how chemical compounds will behave. Natural selection helps explain ecological adaptation and speciation. Fluid dynamics and aircraft simulations allow us to analyze airframes (and train pilots) without littering the countryside. The validation of the model is what happens with the real plane.

Albus' model is good if it allows us (my company) to build a more robust and easy to manage control architecture. It is partially based upon observation of natural and artificial systems. Nature has devised mechanisms that work, and if we can incorporate validated biological control architectures to avoid re-inventing the wheel. Whether it in fact is the general architecture for intelligence, remains to be seen.

With PCT, the question is the intent of the model. Is it to be a 'universal field theory' for biological systems, a biophysical model of a nervous system, or some other purpose? The simulations you have built appear to show that the model provides a mechanism which mimics living systems at the reactive level. It appears to be based upon what we observe (IMIO) at the neural level. It may come to pass that it can be used to predict or understand the behavior of social groups or cognition. However, because the model behaves the same as nature, it does not necessarily follow that the PCT mechanism is that which is used by nature at all levels. It might however be a very powerful technique for engineering PLC level industrial controls.

Perhaps PCT is the mechanism used at the biological level, and Albus' is that used at the cognitive or inter-personal level. Perhaps over the next few generations we will find out. It would surprise me if the exact same mechanism was used at all levels of control. Genetic drift and thermodynamics would seem to give examples that different mechanisms often provide similar functions.

Date: Tue Jul 05, 1994 7:49 pm PST
Subject: Re: PCT models

From Tom Bourbon [940705.1442] >[Paul George 940705 09:30 EDT]
. . .
>(From Tom Bourbon [940701.1601])

>> I asked you how you go about assessing models and theories -- which criteria and procedures do you use before you decide a theory or model is "good" or that it "works?"

> I thought I had answered.

Perhaps I didn't state my questions as clearly as I thought, either time. I asked how you assess models and theories _in general_ and then I asked if you would tell us specifically how you applied your preferred criteria to an assessment of the model of Albus. I am simply trying to gain an understanding of how and why you believe an examination of the Albus model might be of value to PCT modelers.

> We create a model when we can't handle the complexity of the real world, or if we need answers in faster than real time.

This sentence heightens my interest in learning specifically how you assessed the Albus model and found it good. You have described uses of models that are somewhat different from the way we use them in PCT. In "Models and their worlds," Bill Powers and I talked a little about the differences between uses such as you described here and our use of models in PCT. We use them in a way that is common in engineering -- to test the functionality of what we think is an explanation for an instance of control -- we devise a PCT model of the system we think is producing control in a particular situation, than run the model to determine if our ideas are "legitimate," or if they are all wet.

> The model should exhibit the same behavior as the modeled system or allow its computation, in terms of the factors of interest (aye, there's the rub).

Here, we agree, but I wonder how you test for the behavior of an Albus model? I ask this in earnest; I cannot imagine (my failure) how one would set up a working model of Albus's conjectures, then run it in simulations that yield bouts of continuous quantitative control behavior.

> A model is good IF it is useful for the intended purpose. It need not be accurate, in the sense of using the exact mechanisms as the real world, unless that is its intent. It almost by definition cannot be complete in the sense of capturing all the real system's complexities, again unless that is its intent.

If it is offered as a model for systems that produce control, then the model should recreate the control behavior of the modeled system(s). It should do that quantitatively and accurately, moment by moment. Shouldn't it? (If not, then we are not engaged in the same kinds of modeling activity.) If so, how does one go about setting up and testing an explicitly Albus-like model for control phenomena? I am being very concrete about this question, which may explain why I didn't recognize the fact that you thought you had answered me.

> Models can be for several purposes. One variety is for understanding, as in an analogy, Einstienian thought experiment, or scientific theory. Such models allow us to reason about how the real system should react, and deal with currently 'unexplained' effects.] Another type mimics a portion of the real system in order to allow analysis or temporal prediction, as in a simulation.

We may be talking from different sides of that distinction. I believe Albus said he offered his model as a speculation that would inspire empirical work by other people; we are trying to do the empirical work, using a different approach to building and testing a model that I see in Albus's work.

> The model is good if it allows us to come to the 'same' conclusions as the real world system, or allows us to obtain our desired results. A theory should explain all observed behavior of interest. A simulation's outputs should be related in a known (deterministic) way to the real system's outputs, within a certain range of inputs and outputs.

Fine. And which criteria would one use to determine how well a simulation's results are related to the real system's outputs? (Actually, in PCT we aren't just interested in the system's outputs; we are more interested in the _consequences_ of the outputs of either a model or the living system it represents.) We typically look for very high correlations between variables associated with the behavior of the model and of a person, or very low rms error between predicted and actual results, or some other strict quantitative agreement between predicted and actual results.

I'll defer on answering the remainder of your reply to me until after I read Albus, 1991. I finally have a copy.

Later, Tom

Date: Tue Jul 05, 1994 7:52 pm PST
Subject: Re: PCT Models

[From Rick Marken (940705.1340)] Paul George (940705 09:30 EDT)

> With PCT, the question is the intent of the model.

This is the question, indeed. The intent of the PCT model is to explain the phenomenon of control as it is manifested in the behavior of living systems. PCT is not a newer or snazzier version of control theory; it is just plain old control theory. What distinguishes PCT from all other models of behavior (such as Albus') is HOW it maps control theory to behavior. PCT begins with the observation of a phenomenon -- control. "Control" refers to the fact that organisms reliably produce consistent results in an inconsistent environment. People walk, talk, go to work, plant tress, cut down trees, eat, go to bed, etc etc -- and they do these things over and over again, always under slightly (and sometimes profoundly) different circumstances.

Powers recognized the existence of the phenomenon of control and saw that it was what had always been called "purposeful behavior". Control occurs because organisms are able to vary their effects on the world, as necessary, in order to produce the results they intended, all the while precisely counteracting unpredictable (and typically undetectable) disturbances to these results. Powers saw that no existing psychological theory explained how organisms could possibly do this -- ie. how they could control. Even control theory, as it had been applied in the life sciences, didn't explain the phenomenon of control; control theory (before PCT) had been used to describe the relationship between inputs; non-PCT control theory is really just S-R theory in the Laplace domain ;-).

Powers' figured out how to correctly map control theory to actual control behavior. When you apply control theory correctly to the phenomenon of control you find that living control systems control perceptual input variables and that their outputs depend largely disturbances and variations in the feedback function relating output to input -- factors that are independent of the living system itself.

It is impossible (I think) to appreciate the extraordinary significance of PCT without a thorough understanding of the phenomenon that it explains -- control. This phenomenon is so ubiquitous that it is almost invisible. But you can learn to see it -- sometimes by just looking at it and other times by actively testing for control via the introduction of disturbances to hypothetical controlled variables.

The life sciences, and the behavioral sciences in particular, have missed the phenomenon of control completely; all behavioral science research and modelling (to the extent that it occurs) is based on the assumption that living systems respond to stimuli or emit behavioral outputs; in fact they control. Thus PCT, which explains a phenomenon that doesn't even exist, from the point of view of conventional behavioral science, has been almost completely ignored.

PCT is not a sexy new model, like chaos, fuzzy logic, dynamical systems, etc, because the main goal of PCT is not to develop a complex theory but to understand a real phenomenon (you can probably guess what that phenomenon is by now). There are some fairly rich PCT models -- such as Bill Powers' hierarchical arm pointing system and my three level spreadsheet model with six systems controlling a different type of variable at each level -- but these are just demonstrations of what can be done with the hierarchical model. The most important work in PCT (I think) involves demonstrations of real control phenomena.

PCT is unique (and revolutionary -- unfortunately) not because it is a new theory but because it deals with a phenomenon that is not dealt with explicitly by any other theory -- and because it properly applies the existing engineering model of control to the controlling done by living systems.

Caveat emptor: Just because a theory has the word "control" in its name or a theorist talks about "control of this" and "control of that" does not -- repeat, DOES NOT -- mean that the theory or theorist is dealing with the phenomenon of control. In fact, I don't believe that there is ANY theory besides PCT that is an explicit attempt to explain control (but I would love to be surprised).

So accept no substitutes; if understanding control is your game then PCT is the name :-)

Best Rick

Date: Wed Jul 06, 1994 7:52 am PST
Subject: Re: PCT models

[Paul George 940706 11:00 EDT] >Tom Bourbon [940705.1442]

> I am simply trying to gain an understanding of how and why you believe an examination of the Albus model might be of value to PCT modelers.

It may not be of immediate use to PCT modelers (focusing on PCT control in the small), but It might be to PCT theorists focusing on higher levels of the hierarchy.

I found the overall architecture appealing aesthetically, and applied the framework to both control systems and organizations in thought experiments. It seemed to fit quite well as a technique that might work to rationalize both types of systems (note that my focus is on organizational and methodological technology for development projects). Its separation of sensor/actuator judgement, and a world model (perception/simulation) using autonomous object like structures seemed a good differentiation approach (division of labor & recognition of separable error sources). Organisms have different cell types and organs. Organizations have roles and departments.

Having now read Dag's data packet and Power's posts over the last few days, I think there are strong correspondences between Albus' architecture & PCT, particularly at the individual cognitive and organizational levels. In particular, both use planning, History/memory, goal seeking, configuration/influence and data passing of between lower and higher level nodes, an effective model or simulation of the putative 'real world'. If you put the PCT concepts of comparison in the value judgement node, and the concepts of controlled perception in the transducer and world model nodes, PCT adds 'mechanism' to the Albus architecture. The exact nature of the signals transmitted between nodes is almost in the noise range from an architectural standpoint.

I find it interesting that you appear to come from the standpoint that you should by default _not_ consider another's work, particularly from a 'rival' school. My approach is to examine any theory to see if it contains anything useful that might provide insight into others, if nothing else than another point of view. There is a tendency in the academic and methodological (a.k.a. consulting) worlds to focus on details for the purpose of differentiation. Brand names and ideological purity are the basis of recognition and reputation. One must deny any validity to another's approach in order to make one's own more prominent. "I cannot be RIGHT unless everybody else is WRONG" (or inferior, or at least different and incomparable). "Ignore them, they are not of the faithful".

> We use them in a way that is common in engineering -- to test the functionality of what we think is an explanation for an instance of control

Ah. This may be the basis of confusion. I was using model in the general sense of a theory - PCT _is_ a model of behavior. Executable models may be used to test mental models, or at least determine their feasibility. I would normally call this an experiment or simulation. It is however not the _only_ method of testing a theory. The experiments you have done (sorry still haven't gotten access to your book)

> Here, we agree, but I wonder how you test for the behavior of an Albus model? I ask this in earnest; I cannot imagine (my failure) how one would set up a working model of Albus's conjectures, then run it in simulations that yield bouts of continuous quantitative control behavior

Easily. Albus proposes an organizational framework, not a model in the PCT usage. We observe that control and information systems have been built with input and output handlers separate from the 'main' program logic. Similarly applications often separate the knowledge base and algorithmic processing. Distributes control systems (industrial) are hierarchical and show the time horizon behavior. Albus attempts to suggest a general framework abstracted from specific instances. Creating a control application using the 'object classes' identified in Albus' paper is no great technical challenge (we're doing it). The test is a system or organizational architecture that is simple and works. OTOH I may be unclear on your goal. " bouts of continuous quantitative control behavior" is not terribly meaningful to me, except at a limited scope - that of neurophysiology or device and PLC level automatic control.

> ....And which criteria would one use to determine how well a simulation's results are related to the real system's outputs?

If you want to test whether the model works for _human_ behavior, than you merely ;-} have to program the various object methods and attributes to correspond to your understanding of the specific variables and comparators, and sensory/communications distortion functions. Then you can see if it works like the real thing, as you later describe. But, you are testing your control logic and reference variables & values, _not_ the architecture.

Date: Wed Jul 06, 1994 12:00 pm PST
Subject: Re: PCT models

From Tom Bourbon [940706.1227]

Thanks, Paul, for you thoughtful reply to my previous post. I have read Albus, 1991 (and re-read Albus, 1979) and I am still trying to decide how to summarize that material -- my own notes on the articles add up to over 18 pages, so I _must_ try to identify the key points. Your post is helping me zero in on them.

>[Paul George 940706 11:00 EDT]

>>From Tom Bourbon [940705.1442]

>> I am simply trying to gain an understanding of how and why you believe an examination of the Albus model might be of value to PCT modelers.

> It may not be of immediate use to PCT modelers (focusing on PCT control in the small), but It might be to PCT theorists focusing on higher levels of the hierarchy.

Perhaps, but I'm not sure that will prove to be the case. When I decide how best to summarize Albus, I'll try to identify my reasons for that feeling or belief.

> I found the overall architecture appealing aesthetically, and applied the framework to both control systems and organizations in thought experiments. It seemed to fit quite well as a technique that might work to rationalize both types of systems (note that my focus is on organizational and methodological technology for development projects).

I believe this is one of the primary differences in our perspectives on Albus. Understandably, you are interested in designing organizations and systems that achieve your intended ends -- or those of your clients. Some of the uses for which you design systems would overwhelm any single person acting as a control system (a topic addressed nicely by Bill Powers (940706.0630 MDT). In that role, you can (and should) use any trick you can either dream up, or locate already in use somewhere. In his "model," Albus has done the same thing; he has put together an overwhelming array of bits and pieces from robotics, traditional control engineering, neurophysiology, psychology, popular literature and other sources, all in a way that would understandably have aesthetic appeal to someone who is looking for any conceivably useful tools he or she can find. I have no problem either with what Albus has done, or with anyone who uses his "model" as a source of ideas and techniques that will help in designing and building systems for clients.

I have no problem, that is, _unless_ Albus's "model" is presented as a model for living things. I am tempted to say that it is not such a model, but I will simply say that if it is intended as a model of living things, then it is not the same kind of model as the one in PCT. More on that idea when I put together my summary of Albus.

> Its separation of sensor/actuator judgement, and a world model (perception/simulation) using autonomous object like structures seemed a good differentiation approach (division of labor & recognition of separable error sources). Organisms have different cell types and organs. Organizations have roles and departments.

I agree that organisms have different kinds of cells and organs. I also agree that organizations have roles and departments. Let's call those two statements "facts." I do not see that the one fact has any significance for the other. My inability to see such significance does not imply a rejection by me of the idea that you, or anyone else, might see it. It is up to me (rather, it is up to PCT modelers) to demonstrate that our way of modeling living things can succeed without relying on the relationship of one such fact to the other.

> Having now read Dag's data packet and Power's posts over the last few days, I think there are strong correspondences between Albus' architecture & PCT, particularly at the individual cognitive and organizational levels. In particular, both use planning, History/memory, goal seeking, configuration/influence and data passing of between lower and higher level nodes, an effective model or simulation of the putative 'real world'. If you put the PCT concepts of comparison in the value judgement node, and the concepts of controlled perception in the transducer and world model nodes, PCT adds 'mechanism' to the Albus architecture. The exact nature of the signals transmitted between nodes is almost in the noise range from an architectural standpoint.

In some ways, I see what you are saying and I can agree. But in other ways -- major ways -- I believe there are significant differences between the two architectures. As you say, PCT includes a specific mechanism (the input function, comparator function, and output function, with their accompanying signals -- perceptual signal, reference signal and error signal), while Albus does not. But I do not believe one can so easily map PCT functions and signals onto the Albus architecture. For one thing, all of the components in the PCT model are "dumb," in that each of them (except perhaps the input functions) performs a relatively simple operation; in contrast, the elements and modules in Albus's "model" are often very "smart" indeed, with each one of them performing many different functions in many different ways. More on this in my summary.

> I find it interesting that you appear to come from the standpoint that you should by default _not_ consider another's work, particularly from a 'rival' school.

That is not at all the case. I must have overspoken in some of my previous posts. I had read Albus's earlier work and concluded that it would not help us in our attempts to model living systems, _as much as possible_ in their own terms and from their own perspectives. My decision was not made out of any default rejection of work by other people; I believed Albus was simply talking about a different kind of game -- one I can fully accept and respect, up to the limit I described above. Now that I have read Albus, 1991, my assessment has not changed.

> My approach is to examine any theory to see if it contains anything useful that might provide insight into others, if nothing else than another point of view.

I (and I believe Bill Powers and Rick Marken) did just that with Albus, and we independently came away with the idea that his work did not provide insights we could use. We could be wrong.

> There is a tendency in the academic and methodological (a.k.a. consulting) worlds to focus on details for the purpose of differentiation.

True. But there are also real differences between various theories and models. In science, people who observe such differences are obligated to make them known.

> Brand names and ideological purity are the basis of recognition and reputation. One must deny any validity to another's approach in order to make one's own more prominent. "I cannot be RIGHT unless everybody else is WRONG" (or inferior, or at least different and incomparable). "Ignore them, they are not of the faithful".

That may be true (to some degree or other) for people who play the game of Big Science. I do not knowingly or intentionally play that game. If PCT modelers cared about such things, we all would have opted for other "brands" long ago! ;-))

We _do_ say things like, "Here is an example of control and here is our model for the system that produces control; you are invited to show us any other model that you believe performs as well as or better than ours." As our reward, we have had editors and reviewers refer to that approach as "cute" and a "clever ploy."

>> We use them in a way that is common in engineering -- to test the functionality of what we think is an explanation for an instance of control

> Ah. This may be the basis of confusion.

I had begun to suspect that this might be a point where we were using similar words and talking past one another.

> I was using model in the general sense of a theory - PCT _is_ a model of behavior.

I thought so.

> Executable models may be used to test mental models, or at least determine their feasibility.

Yes.

> I would normally call this an experiment or simulation.

No problem; so would we, if editors and reviewers in behavioral science would allow it.

> It is however not the _only_ method of testing a theory.

Certainly it is not, but it is a rigorous one -- a strict one -- and we prefer to apply just such a method when we test our ideas about models for living control systems.

>> Here, we agree, but I wonder how you test for the behavior of an Albus model? I ask this in earnest; I cannot imagine (my failure) how one would set up a working model of Albus's conjectures, then run it in simulations that yield bouts of continuous quantitative control behavior

> Easily. Albus proposes an organizational framework, not a model in the PCT usage. We observe that control and information systems have been built with input and output handlers separate from the 'main' program logic. Similarly applications often separate the knowledge base and algorithmic processing. Distributes control systems (industrial) are hierarchical and show the time horizon behavior. Albus attempts to suggest a general framework abstracted from specific instances. Creating a control application using the 'object classes' identified in Albus' paper is no great technical challenge (we're doing it). The test is a system or organizational architecture that is simple and works. OTOH I may be

unclear on your goal. " bouts of continuous quantitative control behavior" is not terribly meaningful to me, except at a limited scope - that of neurophysiology or device and PLC level automatic control.

So the criteria by which you judge success or failure are different from those we elect to apply when we test our ideas. I have no problem with that. Incidentally, for those who have not read Albus's papers, you have just given a nice glimpse of how he describes a system and its workings; those who are familiar with PCT may begin to see why I am wrestling with how to characterize the similarities and differences between Albus and Powers! :-)

Later, Tom

Date: Wed Jul 06, 1994 5:24 pm PST
Subject: Re: PCT models and simulations

[From Bill Powers (940706.1210 MDT)] Paul George (940706.1100 EDT)

RE: Albus' model.

> I found the overall architecture appealing aesthetically, and applied the framework to both control systems and organizations in thought experiments. It seemed to fit quite well as a technique that might work to rationalize both types of systems...

Unless you are a very competent modeler, dabbling in thought experiments can be the basis of delusions more easily than truth. The problem with doing thought-experiments on aesthetically-pleasing patterns of system design is that you _think_ you know what a real system with that kind of design will do, but you can easily be completely wrong. The fact that Albus claims that a system build according to his design would behave in a certain way is no indication that such a system could even be built, or if built that it would actually do anything like what he thinks it would do.

At one point, some psychologists heard about positive feedback, and thinking that anything "positive" must be good, they starting musing about the effects of "deviation-amplifying responses" and such things, coming up with all sorts of fanciful tales about the advantages of positive feedback. Perhaps by amplifying deviations, a system could increase its variety and so better match the variety of the environment. They imagined that such systems would have useful properties, but they left out one critical step: actually analyzing, building, or simulating a positive-feedback system to see how it would really behave. Being, in fact, incompetent modelers, they didn't realize that an actual system organized in the way they proposed would have properties they hadn't anticipated, and would in fact blow itself up in a few seconds, or else do nothing at all of any interest.

If there's one thing any competent modeler knows, it's that any complex system design, unless it's thoroughly familiar and completely analyzed, is going to have important modes of operation that are different from what one thought they would be, and even more important modes of operation that are total surprises. Anybody can sit down and draw a fancy system diagram with all the parts labelled with function-names, with nice-looking symmetries and impressive patterns of repetition. Furthermore, anyone can claim that such a system would behave in certain ways. That sort of stuff is a dime a dozen; anyone with a fertile imagination can do it. But if that's all that is done, the result is just a fairy-tale.

The whole trick in successful modeling is to PROVE that a system designed in a certain way will actually behave as you imagine it will behave. There are only three ways I know of to do this. One is to do a complete mathematical analysis of the components and their connections and find analytical solutions of the system of equations. That is almost never possible for any system of even moderate complexity. The second way is to build a physical system using components having the specified properties and turn it on. That is usually too expensive an approach, especially considering that the system is going to require many revisions before it does anything like what you have imagined. And the third way is to design a simulation in which the properties of each component are accurately portrayed as elements in a working computational model, and run the simulation. Most of the time, that's the only practical way to test an idea.

The reason that Tom Bourbon (and Rick Marken and I) is skeptical about models that haven't been tested in simulation is that the human brain is very good at guessing wrong about what systems will do. We let what we WANT the system to do persuade us that our design will actually do what we want. If the system contains any reasonable amount of complexity, we will usually find that we have designed parasitic loops into it, or subtle self-contradictions, or have required some variable to do something physically impossible such as going from positive to negative values without crossing zero, or have built in requirements that assume a physical component capable of behaving with 18 decimal places of accuracy, forever. A competent modeler anticipates such problems and tries to take them into account. The neophyte just sails right in, drawing diagrams and claiming that they explain something.

When you sit down to produce an actual simulation, the most common thing that happens is that you come across something you have assumed to happen that the system you have actually designed can't do -- or that you have completely failed to provide for in the model. One of my arguments with Martin Taylor about the "alerting system" is based on this problem. It's easy to propose a system that will alert other systems to a problem. But when you sit down to simulate such a system, you discover all the machinery you were unknowingly assuming and that is essential to create the situation you had in mind. Providing this machinery very often turns out to be a far larger problem than the original design. And you can't get a simulation to run properly unless you include in it EVERYTHING it needs in order to run. You can wave your arms all you like; the system will still do only what you gave it the capability of doing. Also, like a computer program, it will do exactly what you designed it to do, not what you wanted it to do.

Rick Marken and Tom Bourbon and I have all tried to make system designs proposed by psychologists work. And I mean we have really tried. We have even supplied missing properties in our attempt to make such simulations produce the kind of behavior that the proponents of psychological theories claim their systems would produce. And despite the fact that we are fairly competent modelers, we have failed. Just try designing a simulation of an operant-conditioning experiment according to traditional explanations, one that will reproduce what is actually observed. You immediately discover that there is no specification for the effect of reinforcement on behavior that can be turned into a model. When you try to supply one, you find that translations of the usual verbal descriptions result in models that either don't do anything at all or settle into modes of behavior that are never observed in real systems. We have publicly challenged psychologists to supply their own designs for simulations, and have even offered to do the programming work ourselves if they don't know how. These challenges have, of course, disappeared into an echoing silence, because the psychologists involved don't have the vaguest idea of what we are talking about. They think that if they _say_ that reinforcement works the way they assume, such a system _will_ work as they say it does. The idea of actually testing these assertions has never crossed their minds.

One reason that PCT models tend to be simple is that we can't prove that any more complex models would actually behave as we imagine they would. All the models we talk about have been proven to work exactly as claimed, as well as matching real behavior with acceptable accuracy. The idea of jumping into extremely complex behavioral problems while we still can't simulate simpler behaviors is just ridiculous. We're building upward from a solid base. All the guesses as to what we will find as models of higher levels of control are just that, empty guesses. Whatever we guess right now is almost certain to be proven wrong; why waste time in airy-fairy-land? Why bluff when we can stick with what we know about?

Best, Bill P.

Date:     Thu Jul 07, 1994  6:10 am  PST
Subject:  Albus articles

From Tom Bourbon [940707.0813]

Paul George is the most recent person to suggest that PCT modelers might find useful ideas in the writings of James Albus. (Some people have also suggested that the model in PCT is a special, limited, case of a more general model of control developed by Albus.) For some years now, Albus has written about robots and artificial systems. He uses a particular interpretation of ideas from psychology and neurophysiology and combines them with ideas from work on artificial systems, with the intention of suggesting an "outline" or "framework" to be used in developing a general theory or model of behavior.

After seeing Paul's suggestion, I re-read four articles by Albus that ran in _Byte_ magazine, June-September 1979, concurrently with a series on control system theory by Bill Powers. I also read the more recent article suggested by Paul:

J. S. Albus (1991). Outline for a theory of intelligence. IEEE Transactions for Systems, Man, and Cybernetics, 21(3), 473-509.

I took more than 18 pages of notes on the articles. I cannot possibly summarize that much material here. All I present here is a brief description of a few of my conclusions. I use several quotes from Albus, thereby avoiding the need to present my own interpretations of him, which might be biased. I urge anyone who wants to challenge or check my conclusions to read the same articles and to perform the modeling I suggest at the end of this post.

The Elemental Albus Architecture:

Albus proposes the following architecture for his theory. I believe the diagram and a brief summary of his descriptions of the various elements in it clarify many differences between his theory and PCT.

```
                                     ____
     Perceived Situations           |    |      Plan Evaluations
   ,------------------------------->| VJ |------------------------------,
   |                                |____|                              |
   |                 Situation       |  |                               |
   |                 Evaluation      | /|\ Plan Results                 |
   |                                \|/ |                             \ | /
 __|__     Updates                   |__|       Plans                __|__
|     |---------------------------->|    |<---------------------------|     |
|  SP |                             | WM |                            |  BG |
|     |<----------------------------|    |--------------------------->|     |
|_____|    Predicted Input          |____|      States                |_____|
   |                                                                     |
  /|\                                                                    |
   |                                               Commanded Actions     |
   |  Observed Input                                                     |
   |______                                                      _______\ | /
|Sensor|                                                       |Actuator|
|______|                                                       |________|
   /|\                                                              |
    |                                ____________                   |
    |       Events                  |            |  Actions         |
    '-------------------------------|Environment |<-----------------'
                                    |____________|
```

The system portion of this busy looking elemental loop repeats in a hierarchical manner, much as the elemental loop in PCT repeats in a hierarchy.

DEFINITIONS AND DESCRIPTIONS OF ELEMENTS IN ALBUS'S ARCHITECTURE

Albus's sensors and actuators are essentially like the input and output functions in PCT. Very little else about his architecture is like anything in PCT. All of his other modules or elements are "intelligent" -- loaded with functions, powers and abilities unlike anything in the simple elements of a PCT model. I quote from Albus. Pay close attention to the various "system elements," and to the bewildering array of abilities, roles, powers, and functions he attributes to those elements and systems.

SP = "Sensory Processing. Perception takes place in a sensory processing system element that compares sensory observations with expectations generated by an internal world model. Sensory processing algorithms integrate similarities and differences between observations and expectations over time and space so as to detect events and recognize features, objects, and relationships in the world." p. 476.

"Perception is the establishment and maintenance of correspondence between the internal world model and the external real world." p. 493.

WM = "World Model. The world model is the intelligent system's best estimate of the state of the world. The world model includes a database of knowledge about the world, plus a database management system that stores and retrieves information. The world model also contains a simulation capability that generates expectations and predictions. The world model thus can provide answers to requests for information about the present, past, and future states of the world. The world model provides this information service to the behavior generation system element, so that it can make intelligent plans and behavioral choices, to the sensory processing system element, in order for it to perform correlation, model matching, and model based recognition of states, objects, and events, and to the value judgment system element, in order for it to compute values such as cost, benefit, risk, uncertainty, importance, attractiveness, etc. The world model is kept up-to-date by the sensory processing system element. p. 476.

VJ = "Value Judgment. The value judgment system element determines what is good and bad, rewarding and punishing, important and trivial, certain and improbable. The value judgment system evaluates both the observed state of the world and the predicted results of hypothesized plans. It compares costs, risks, and benefits, both of observed situations and of planned activities. It computes the probability of correctness and assigns believability and uncertainty parameters to state variables. It also assigns attractiveness, or repulsiveness to objects, events, regions of space, and other creatures. The value judgment system thus provides the basis for making decisions -- for choosing one action as opposed to another, or for pursuing one object and fleeing from another." pp. 476, 477.

BG = "Behavior Generation. Behavior results from a behavior generating system element that selects goals, and plans and executes tasks. Tasks are recursively decomposed into subtasks, and subtasks are sequenced so as to achieve goals. Goals are selected and plans generated by a looping interaction between behavior generation, world modeling, and value judgment elements. The behavior generating system hypothesizes plans, the world model predicts the results of those plans, and the value judgment element evaluates those results. The behavior generating system then selects the plans with the highest evaluation for execution. The behavior generating system element also monitors the execution of plans, and modifies existing plans whenever the situation requires." p. 477.

Albus says some of these "system elements" have subsystem elements -- as many as three or four of them. At all hierarchical levels of the SP and WM elements, information is represented in both iconic and symbolic forms. "At each level in the control hierarchy, the difference vector between planned commands and observed events is an error signal, that can be used by executor submodules for servo feedback control modules for evaluating success or failure." (p. 481) Differences between predicted and actual perceptions are error signals that are used to induce changes _inside the system_, perhaps similar to reorganization in PCT -- but I can't be sure about the degree of similarity. At each level, Kalman filters make predictions, using recent historical data to compute parameters that are then used to make extrapolations into the future. (p. 480)

Emotions: Albus uses an outdated and highly questionable interpretation of the anatomy and neurophysiology of the "limbic system" in the brain, as a source of emotions. In the _Byte_ series, he wrote about how emotions from the limbic system are input, along with "moderator variables," into the "will," a name he gave to the behavior-generating system element.

Reification from Observed Events: In all of the articles I read, Albus reified the observed results of behavior, then attributed many different "powers" to the reified concepts. A clear example is his definition of "intelligence," the subject of his article in 1991. He defined intelligence as, "the ability of a system to act appropriately in an uncertain environment" p.474. An observation that organisms _do_ act appropriately (intelligently) is explained by an assumed property (trait, power, ability, etc.) called "intelligence." Reification of this sort is common in the behavioral sciences and in behavioral neurophysiology and neuroscience. Once he assumes intelligence into existence, Albus says it (intelligence, which he described as an ability) requires further abilities -- "to sense the environment, to make decisions, to control action" p. 474. "From the viewpoint of control theory," says Albus, intelligence is an, ". . .integration of knowledge and feedback into a sensory-interactive goal-directed system that can make plans, and generate effective, purposeful action directed toward achieving them." From the viewpoint of psychology, he says, intelligence is a behavioral _strategy_. Once he brings intelligence into existence as reification from behavior, he says it is an ability, an integration, and a strategy. There is more on this theme, but I believe you get the idea.

General Conclusion:

Albus does not describe a generative model of the organization of living systems. Instead, he describes a highly complex flow chart of relationships (some empirical, many more conjectural), arranged in a hierarchy that includes (in conjecture) feedback loops, error signals, sensory processing with perceptions from lower levels combined to form more complex perceptions at higher levels, and a number of other features that, on first encounter, look similar to ideas in PCT, but he often uses those ideas differently than we do in PCT. I believe the "outline" or "framework" Albus presents may be useful to some people, in their efforts to develop engineered systems, but it cannot serve as a generative model for control by living systems.

More to the point, as Albus described his theory in the article in 1991, I do not believe it can be programmed as a working model of any system, living or not. I _do_ believe someone could pick various elements from his outline, which is like a giant buffet of elements and processes, and use some version of that subset to produce a working model, but that person would not be using "The Albus Model" in its fullness.

An Oft-repeated Invitation:

I admit that my inability to imagine how the Albus outline could be turned directly into a working model may be only that, my inability, but I believe there is more to it than that. Whatever I may believe about the Albus outline, I welcome simulations and demonstrations from anyone who wants to show that my conclusion is wrong. I will happily give any objecting person(s) copies of the computer code for the procedures Bill Powers and I used to generate the target and disturbance functions for our paper, "Models and Their Worlds." I encourage the objector(s) to use those procedures in the "environment" box of the Albus architecture, then implement a single-loop version of an Albus model, with all of the elements, connections, and processes that are shown in the figure I reproduced above. Then the person should run the model in simulation and report the results so that we can compare the results from the two models when both are applied to what many people think is a trivially simple, low-level, example of human performance. If the Albus and PCT models perform equivalently at that level, then we can begin to talk seriously about how the two might compare as models of higher levels in the control hierarchy.

I am not being "cute" or "judgmental" in this offer; I simply cannot think of another way to compare the Albus and Powers models, as working models of control behavior.

Later, Tom

Date: Thu Jul 07, 1994 10:07 am PST
Subject: Re: Albus articles

[From Rick Marken (940707.0900)]

I would like to thank Tom Bourbon (940707.0813) for posting the summary of the Albus model. How in the world did you do it, Tom? I started getting nauseous while studying the system diagram and nearly barfed on the keyboard after reading about the Sensory Processing, World Model, Value Judgment and Behavior Generation components of the "model". I hope you were able to get hold of an ice-cold glass of 7-Up after posting this stuff. You're a better man than I.

Now that I have recovered my sea legs I can see that the Albus model is an excellent example of what PCT is NOT. Albus throws in every misconception in the book, each one neatly summarized in the description of the function of each box in the model -- SP, WM, VJ, BG. It's tough to decide where to start dealing with all these juicy misconceptions -- each one part of the bread and butter (I've become partial to scones and clotted cream after the England trip) of current behavioral science dogma. But I think your suggestion is probably the best:

> I encourage the objector(s) to ... implement a single-loop version of an Albus model, with all of the elements, connections, and processes that are shown in the figure I reproduced above.

Of course, maybe the Albus model can't do anything as dumb as keep a cursor aligned with a target; maybe the only thing the Albus model can do is act "intelligently". That would sure show us, wouldn't it? ;-) (NB. I AM being cute and judgmental).

> I am not being "cute" or "judgmental" in this offer;

Sure. Sure. I know you. You're the killjoy who would keep Uri Geller from touching the fork before doing his great psychokinetic fork bending routine. People love the Albus-type stuff because it uses all the right buzzwords. Now you wanna go and spoil it with all that science stuff. You're just no fun (for psychological charlatans, anyway) but I can't wait to see you again in Durango ;-).

Best Rick

Date: Thu Jul 07, 1994 3:18 pm PST
Subject: Re: Albus articles. System design is easy!

[From Bill Powers (940707.1100 MDT)] Tom Bourbon (940707.0813)

Thanks for that review of Albus: you have done us a service.

How easy modeling is when you don't have to make anything actually work! That "world model" is very handy; it contains all the properties of the world as well as everything that can happen in it, and all that's needed to sense it is one little arrow that goes over to a box called sensory processing, with an "update" arrow going the other way. I wish somebody had told me that this is all there is to it. Out of this 40 year project I could have saved 39 years and 11 months.

Actually, this is just the sort of system diagram that certain kinds of executives like to draw up.

"What's your big problem?", they say, whipping out a sheet of newsprint and a dry-marker. "Watch, it's easy: just make a world model, connect it to a sensory processor, hook that up to a value judgement box with all sorts of good values and judgements in it, run the result over to the latest model of Behavior Generator, stick in some more connections to the world model and some ways of using it, and off you go! I don't see why _I_ have to do everything around here. I expect a progress report on my desk by Friday."

The poor engineer flunky, if he is wise, will say "Gee boss, that's a great idea. You've really wrapped it all up here, all right. We'll get right on it." Then later, after he has worked out a simple control system and got it to reproduce some real behavior, he'll bring the results back and say "Gosh, you were really right about that. Look at this schematic. Here's your sensory processor, and we've got a signal coming in here that's exactly what you described as a value, here here's the judgement that takes place just the way you said, and the output goes over here to your Behavior Generator that moves the control stick. It all works beautifully, just the way you said it would. Only thing is, we need some advice on that world model part -- we're getting good behavior out of this, but you're going to have to give us some guidance about what the world model does. Maybe you could say that this system has an _implicit_ world model in it, so it really works just like your diagram. Amazing!"

You can always bullshit a bullshitter.

Best, Bill P.

Date: Thu Jul 07, 1994 3:18 pm PST
Subject: Re: PCT models and simulations

[Paul George 940707 10:30]

>[From Bill Powers (940706.1210 MDT)]
Heavily edited, with many good points omitted.

> Unless you are a very competent modeler, dabbling in thought experiments can be the basis of delusions more easily than truth.

> Being, in fact, incompetent modelers, they {Psychologists} didn't realize that an actual system organized in the way they proposed would have properties they hadn't anticipated, and would in fact blow itself up in a few seconds, or else do nothing at all of any interest.

> If there's one thing any competent modeler knows, it's that any complex system design, unless it's thoroughly familiar and completely analyzed, is going to have important modes of operation that are different from what one thought they would be, and even more important modes of operation that are total surprises. Anybody can sit down and draw a fancy system diagram with all the parts labelled with function-names, with nice-looking symmetries and impressive patterns of repetition.

Amen. This is one reason why systems analysis is gravitating towards using executable models and simulations for validating requirements and design. In our world unforseen behavior can be deadly. Unfortunately, psychology has been limited by a tradition that experimentation on and vivisection of humans is unethical (Small minded of them ;-). There is no strong tradition of experimental proof of theory or models. Further, training in the soft sciences often contains little engineering, basic science, or mathematics. Statistics are of course grossly misused and (I hope) misunderstood. Instead it has essentially been a branch of philosophy (or theology ?? :-).

> One reason that PCT models tend to be simple is that we can't prove that any more complex models would actually behave as we imagine they would. All the models we talk about have been proven to work exactly as claimed, as well as matching real behavior with acceptable accuracy. The idea of jumping into extremely complex behavioral problems while we still can't simulate simpler behaviors is just ridiculous. We're building upward from a solid base.

Forgive what may be misunderstanding, but from Dag's materials Rick Marken's post yesterday, PCT derived from the concept that _People_ exhibit purposeful behavior and perceptual control, not insects. That would seem (to oversimplify) to focus of the functions of the cortex, not the thalmus and nervous system. Thus, while your approach is good science and engineering, you have validated only the PCT 'mechanism' (or one of them), not the PCT Theory. You (as a group) have demonstrated that simple positive feedback is sufficient for low level neuro-muscular functions. You have not demonstrated that it is necessary or sufficient for higher level behavior (to my extremely limited knowledge). Thus Tom Bourbon's statement in an earlier post (as I recall) that "...we have found no need for more complex mechanisms" is not completely accurate.

In my mind this state of knowledge in no way decreases the value or utility of the PCT theory, but fundamentally leaves it at the same level as other behavioral theories (including Albus) at the application level. It can be demonstrated experimentally that operant conditioning works, but not _how_ it works (i.e. the internal mechanism). Further we cannot generalize to say that _all_ behavior is just operant conditioning. The same kind of thing can be said for PCT behavioral control.

Paul George

Date: Fri Jul 08, 1994 9:32 am PST
Subject: Testing models

[From Bill Powers (940708.0905 MDT)] Paul George (940707.1030)

> Unfortunately, psychology has been limited by a tradition that experimentation on and vivisection of humans is unethical (Small minded of them ;-). There is no strong tradition of experimental proof of theory or models.

What I was talking about was not matching models against behavior, but showing that the model itself would behave as one claims it will behave. This is a much more serious problem. Look at Tom Bourbon's rendition of the Albus model. If you start at any arrowhead, for most of them you can trace half a dozen or more ways of going around the system and ending up back where you started. Each distinct path is a feedback loop, involving different combinations of components. So you really have a large number of feedback loops superimposed and intertwined in this diagram, and if you ever constructed a simulation having this connectivity, all those loops would start interacting. It's extremely unlikely that such a simulation would act in the way Albus imagines that it would act.

What Albus has done is to go from his conception of how people might act to a diagram that describes all the pathways he noticed. But in drawing the diagram, he created, probably without realizing it, a large number of OTHER pathways. Some of these pathways would be negative feedback loops, but other ways of following effects around the system might pass through an even number of sign-inversions and amount to positive feedback (unstable) loops. It thus becomes impossible to get from the properties of the model back to the same observations that led Albus to draw the model in the first place. If a model could actually be constructed, one might be able to show the correspondence with the real situation running in the direction from the observations to the model -- but when the model ran, it would not reproduce the original observations. In fact, it would probably go into violent oscillations or lock up in states having nothing to do with real behavior, even the behaviors that led Albus to draw this diagram.

This is the FIRST test that has to be applied to a model. Not a test to see if the model's behavior would match real behavior, but a test to see if the model would actually behave in the way you imagine it will behave, quite independently of the question of whether that behavior would match what a real system does.

All the models we use in experimental PCT actually behave in the ways we claim they will behave, without unwanted or unexpected modes of behavior. Any proposition that claims to be an actual model has to have this property. Very few models of behavior to be found in the literature do.

-------------------

> from Dag's materials Rick Marken's post yesterday, PCT derived from the concept that _People_ exhibit purposeful behavior and perceptual control, not insects.

The basic idea was that people exhibit _control_ behavior. From the model of control behavior (out of engineering control theory, by Weiner) it could be seen that the role of the reference signal fits the properties of what we informally call purposes, intentions, or goals.

> You (as a group) have demonstrated that simple positive feedback is sufficient for low level neuro-muscular functions.

Typo, I hope: simple _negative_ feedback. Positive feedback is the opposite of control.

> You have not demonstrated that it is necessary or sufficient for higher level behavior (to my extremely limited knowledge).

We have a few examples, but none instrumented or quantitatively modeled. There are many informal examples of higher-level control, but the problem with modeling it is that we don't know how to model the perceptual functions involved. We'll get there eventually.

> It can be demonstrated experimentally that operant conditioning works, but not _how_ it works (i.e. the internal mechanism).

I have modeled operant conditioning using control systems, and the model accurately reproduced some data obtained by others. One day I'll resurrect that model and post it. However, the model described by Skinner et. al. says only that reinforcement increases the probability of a behavior, increases the amount of behavior, or maintains a given behavior. When you try to turn those words into a working model, you end up with something that doesn't behave the way the real organisms do. The only model that does work treats the rate of reinforcement as a controlled variable, with a reference signal for the desired rate inside the organism, and the error adjusting the rate of behavior to bring the actual rate of reinforcement toward the desired rate. And from this model, we conclude that "operant conditioning" is mostly ordinary control behavior, misinterpreted. So it isn't operant conditioning that works; it's control that works.

---------------------------

In the realms where we haven't experimentally modeled the PCT organization, we have to treat the PCT model as an explanatory construct. Since that is what most other models are (and nothing more), we can then ask how realistic this explanation is, in terms of what we observe and what we experience, in comparison with what other models tell us. So far, those who have learned all the ins and outs of the PCT model say that this is the only model they have ever seen that represents their own behavior to themselves, and other people's behavior, in a believable form. Getting to that point of understanding takes rather a long time, a couple of years of asking questions and seeing how objections are met. I hope you stay the course.

Best, Bill P.

Date: Fri Jul 08, 1994 12:06 pm PST
Subject: Re: Testing models

From Tom Bourbon [940708.1205] >[Bill Powers (940708.0905 MDT)]

A taboo against vivisection of humans isn't the problem. In fact, where it has been performed on other species (in deafferentation studies, for example), the results have not helped increase our understanding of behavior. To the contrary, those studies have contributed to the great social danger that follows from most behavioral research and theory.

Bill P:

> What I was talking about was not matching models against behavior, but showing that the model itself would behave as one claims it . . . . do with real behavior, even the behaviors that led Albus to draw this diagram.

Yes, Bill. This was exactly what I saw when I began reading Albus again, and what I saw reminded me of why I had dismissed Albus back in 1979. It looks to me as though he noticed what, to him, seemed to be certain kinds of actions and results that people produced, then he reified those observations into "explanatory constructs" (X is intelligent behavior; intelligence causes intelligent behavior; intelligence causes X), then he strung those constructs together into a densely interconnected flow chart.

In 1979, while I was still a PCT neophyte, it was obvious that he had created a typical behavioral-life science flow chart, not a model that would behave, or if it would behave, there was no way he could possibly know in advance what it would do. In the post where I reviewed his papers, that thought was behind my remark that I could easily understand how people might carve out bits and pieces from the Albus densely-interconnected flow chart (the ADIFC) and use those pieces to build or program a working model, but their working model would in no way be a working ADIFC.

> This is the FIRST test that has to be applied to a model. Not a. . .

Yes. Yes. Yes. And that is one of the most difficult ideas for us to get across to our readers, editors, reviewers and critics. At this level, it is not a matter of PCT vs any other theory; it is a matter of whether _any_ particular theory or model would really produce the results its proponents claim. This is an issue that cuts right to, and through, the heart of nearly every single alleged theory and model in the behavioral-social-neuro-sciences. If those people tested their "models" in the way Bill is describing here, practically all of them would come up lacking and would either be modified (into PCT-like models, by the way -- control models), or they would be tossed aside. But how likely is it that practitioners of those "sciences" would willingly dismantle and close down their sciences? Don't count on it -- there are too many millions of dollars in research funds waiting for them -- yet another of those grave social dangers to which I allude from time to time.

Later, Tom

Date: Fri Jul 08, 1994 1:54 pm PST
Subject: Re: Testing Models

[From Rick Marken (949708.1310)] Tom Bourbon (940708.1205) --

> At this level [of seeing whether the model works as expected], it is not a matter of PCT vs any other theory; it is a matter of whether _any_ particular theory or model would really produce the results its proponents claim. This is an issue that cuts right to, and through, the heart of nearly every single alleged theory and model in the behavioral-social-neuro- sciences. If those people tested their "models" in the way Bill is describing here, practically all of them would come up lacking and would either be modified (into PCT-like models, by the way -- control models), or they would be tossed aside.

It's just this kind of extremist talk that alienates sober, well-established, prominent behavioral-social-neuro-scientists from PCT. Why, you're saying that the entire edifice of the behavioral-social-neuro-sciences is built on sand; that it's all a crock. Now who in the world is going to believe that (other than fools like you and me who have actually gone into the lab and found out for ourselves).

Excellent post, Tom!

Best Rick

Date: Fri Jul 08, 1994 3:44 pm PST
Subject: Re: PCT models

[Paul George 940709 11:30] >[Bill Leach 940707.23:04 EST(EDT)]

> Paul; Your discussions have been of value to me personally. The points you raise are resulting in some very illuminating responses.

Thanks. In truth I pointed out the Albus article in the spirit of throwing raw meat to wolves. The response of 'the masters':-) has been both informative and amusing. Defending and attacking expose what the proponents of a theory consider most important with amazing rapidity.

On the whole I have found PCT itself to be very solid, and capable of explaining a great deal. At its philosophical level I find it appealing, certainly in comparison to behaviorist purists (Though I have often thought Skinner was grossly misunderstood) or some of the humanist school (dates my studies, doesn't it).

But theory has two uses, understanding and application. As Jeff Vancouver is quoted as saying - Psychologists are trying to do something useful. I interpret this to mean apply theory to helping people, society, or personal life.

>From Tom Bourbon [940707.1729]

> ...one reason I say most claims to knowledge by psychology are dangerous is that innocent people might be tempted to believe those claims -- some might even go so far as to appeal to psychological "knowledge" when they attempt to establish new social policies, laws, and the like.

>From Tom Bourbon [940708.0816]

> All there is to support the majority of "theories" and "models" in the behavioral-social-neuro- sciences is their apparent plausibility -- they sound as though they could be true.

It goes beyond plausibility to the realm of utility. Behaviorists may not be able to provide a working model of the mechanism, but can give heuristics for how positive or negative reinforcement affect observed behavior and learning. Gestalt psychology has observed and defined certain stages of cognitive development, even if they don't provide the mechanical underpinnings. Priests and ministers can help people with life and make them more comfortable (and the reverse in spades!), even if metaphysics and theology are inherently un-provable. Alchemists were able to make chemical compounds even if the earth/air/fire/water element were not the correct model. A hot air balloon will work regardless whether we use gas density or the Aristotelian concept of a desire towards heaven as an explaination. The point is all these theories are or were useful to some extent. They could also cause people to come to wildly wrong conclusions under certain circumstances. It is always a question of fact. The question is not "is it right" but rather "does it work".

>[From Rick Marken (940707.1415)] "Out on a limb"
>From Tom Bourbon [940707.1729]

> Rick Once you demonstrate that a higher level behavior is a controlled variable, then you know you need a control model to explain it. A control

> model is necessary because you are modelling the phenomenon of control.... -- but when behavior involves control, you can only model it with a control model.

This I feel is an oversimplification. A control model is required to capture the _control component_ of the behavior. Further PCT has not (to my knowledge) that higher level behavior _is_ a "a controlled variable". I personally suspect it involves control of sets of controlled perceptual variables. It also IMHO involves sets of controlled reference variables, reference values, and some mechanism to determine from memory and/or extrapolation what variables of both types are or are going to be important.

{out of order}

> The answer, according to PCT, is that higher level behavior (like all behavior) is _a controlled result of action_.

I am not sure that control phenomena is all that exists. There are components that appear reactive. (I didn't fully follow the thread on alerting, other than to observe the apparent response that the behavior could be produced by a sufficiently fast control loop). When a hand is burned, the arm will jerk back. It is my understanding that this does not involve the brain, just neural reflex at the spine. I seem to recall something on the order that the optic nerve of a frog transmits no signal to the visual cortex until a moving object like a fly enters the field of vision.

Having a continuously active control loop monitoring every possible control variable seems very inefficient. Some inputs are discontinuous or periodic. Bill Powers statements recently notwithstanding, a control loop can stop, or at least idle. Computer systems often use something called and interrupt chain. If no sensory inputs (events) are currently active (one of the results of an error signal existing can be the generation of a discrete input to another loop) the control loop hibernates. When woken up, one kind of control loop determines what kind of event (or control variable) was populated or went out or range and activates the appropriate control loop to get things back in order.

> So the idea that PCT can handle simple behavior while the jury is still out on its ability to handle higher level behaviors really misses the point completely. Simple vs higher level is not the dimension on which PCT is discriminated from other theories of behavior; the dimension is "generated outputs vs control".

What I have noted in most discussions, particularly that about Albus, is the focus upon how "PCT is discriminated from other theories of behavior". You focus upon the differences, not the commonalities or how the approaches might be blended into a more powerful whole. It's kind of like "well, that's not what we're talking about, or it isn't the same point of view, or it doesn't address something we do. Now having disposed of that, we can go back to our own world and what we're doing."

The major difference I see between PCT proponents and others is that traditional psychology and sociology looks at behavior from the outside of the organism, while PCT focuses on the inside. Psychology has said 'We observe the following[...].. What can we say about predicting it or understanding the relationships between patterns of apparent events and observable actions of the organism?' PCT says 'can we identify something being controlled internally and a negative feedback{ack Rick} relationship? Can we simulate it?' (yes it's more involved than that). It is kind of like the difference between a 'top down' and 'bottom up' approach to analysis. Both have their pitfalls.

I would challenge you to try to design or understand a distributed real time system at the assembly language or register level. It _can_ be done, but not by many. But higher level constructs such as 'tasks', 'messages', data structures are useful for understanding and constructing the application, even if it does ultimately reduce to bit twiddling. Similarly, to understand thought using only the mechanics of dendrite firing is at best difficult.

To say "if they can't do the simple stuff why should we pay attention to anything you have to say at higher levels" is at best unfair and at worst disingenuous. Their models do not purport to address that level of detail. Frame of reference and scale of reference are important concepts to keep in mind.

Ultimately I suspect PCT will meet other approaches somewhere in the middle, providing both mechanism and greater understanding of the parameters of some kinds of human behavior. When you can build an executable control model showing how one puts on a sweater and/or adjusts a thermostat, and decides between the options, I will be convinced. A skin temperature control variable is not quite sufficient. Ah well, +sufficient unto the day be the troubles thereof'.

Date: Fri Jul 08, 1994 3:45 pm PST
Subject: Re: Testing models

[From Paul George 940708 15:00 EDT]
Can buy most of the post, but a couple nits:

[Bill Powers (940708.0905 MDT)]

> What Albus has done is to go from his conception of how people might act to a diagram that describes all the pathways he noticed. But in drawing the diagram, he created, probably without realizing it, a large number of OTHER pathways. Some of these pathways would be negative feedback loops, but other ways of following effects around the system might pass through an even number of sign-inversions and amount to positive feedback (unstable) loops.

This is true only if you model the diagram naively and permit all paths to be traversed. If you put a state machine(s) or PCT control node(s) within the nodes, and/or any discrimination node which can rout different signals to various control or sensory nodes, then all paths are not possible. Behavior could be rather complex, but so is organismic behavior. Unstable systems (like weather) may function within bounds (strange attractors). Besides, behavior (and life) is only a stable system within certain bounds.

A PCT control node must ignore all inputs that do not affect the controlled variable, or more accurately never receives them. How is it that a HPCT system assembles inputs, controlled variables, reference variables, and outputs, not to mention appropriate interconnections. All paths between nodes are not permissible in the HPCT model either (of course you define internal connection rules).

> And from this model, we conclude that "operant conditioning" is mostly ordinary control behavior, misinterpreted. So it isn't operant conditioning that works; it's control that works.

Close. Control is _a_ mechanism which _can_ implement operant conditioning, given certain assumptions. How is the control variable and negative feedback signal set up? (I have a little trouble with 'Unseen Hand' arguments)

The organism says 'gee, I liked that! Now what can I do to get more of it? The reinforcement must exist before its rate may be controlled. The organism must deduct a relationship between a behavior (or output of some sort) with the sensory input which is +reinforcement.+ Then a control loop must be set up which detects the rate (or more likely the presence). The control loop must essentially generate a constant maxed error signal that intermittently goes to zero when reinforcement arrives or it's rate increases. A rate based control variable would essentially be a trend computation based upon an intermittent binary signal (with a long period relative to the 'actions').

Several problems however: First, reinforcement is not a constant signal,and not even the same signal. Second, the desired behavior is not always rewarded by the trainer. While an increase in the rate may be desired by the organism, the rate is in fact random or constant, and ultimately declining. Third, the behavior persists after the reinforcement stops. The "...bring the actual rate of reinforcement toward the desired rate" doesn't occur. The problem is worse with negative reinforcement. There the control issue becomes 'what behavior must stop to reduce the unpleasant perception?". The concept of an error signal producing action becomes a little convoluted in the OC case, or that of Maslow's hierarchy. A 'happiness' or 'comfort' controlled variable is not much more convincing than a 'world model'..

Note: I don't really intend to have a lengthy discussion of the proper translation of OC to a control model. However, demonstrating learning behavior and OC using a 'dog training' type scenario would be an interesting test of the claim that 'all behavior is control'. You would have to come up with a mechanism for dynamic reconfiguration of control nodes and connections for adaptive purposes, which is a hard current computer science problem.

Not to worry, just because we don't know _how_ something works or we don't have the technology to reproduce it doesn't mean it can't work. This is a basic problem with simulation or modeling - the devil is in the details. We can't yet design and construct biological entities. PCTrs could be dismissed by neuro-physiologists because you can't describe how to build a PCT control node out of neurons and/or brain cells (I don't think).

>From Tom Bourbon [940708.1205]

> But how likely is it that practitioners of those "sciences" would willingly dismantle and close down their sciences? Don't count on it -- there are too many millions of dollars in research funds waiting for them -- yet another of those grave social dangers to which I allude from time to time.

.a tissue paper dog chasing an asbestus cat through hell. Of course the same can be said for any closely held belief, particularly one upon which one's reputation or world view is based.

Looking forward to next week. - Paul George

Date: Fri Jul 08, 1994 4:32 pm PST
Subject: Re: PCT models

[From Tom Bourbon 940708.1653]

Should I wait until Monday, so I can cool down a bit? Nah!

>[Paul George 940709 11:30]

>>[Bill Leach 940707.23:04 EST(EDT)]

>> Paul; Your discussions have been of value to me personally. The points you raise are resulting in some very illuminating responses.

> Thanks. In truth I pointed out the Albus article in the spirit of throwing raw meat to wolves. The response of 'the masters':-) has been both informative and amusing. Defending and attacking expose what the proponents of a theory consider most important with amazing rapidity.

Glad to know we could keep you laughing for a while. For our part, we don't usually suggest that people read and comment on published material unless we really mean it. I don't especially like to play games like that. Call it a silly bias of mine about how I approach people and deal with them on the net, or anywhere else, for that matter.

Would you let us in on what you found the most amusing?

And would you maybe let us in on your personal assessment of Albus, now that we know you weren't playing straight with us for the past week and a half? Whatever your assessment may be, I guess we don't need to wait to see any implementations of a genuine Albus model by you, do we?

> On the whole I have found PCT itself to be very solid, and capable of explaining a great deal. At its philosophical level I find it appealing, certainly in comparison to behaviorist purists (Though I have often thought Skinner was grossly misunderstood) or some of the humanist school (dates my studies, doesn't it).

I have a long-standing interest in Skinner and in the often-repeated claim that he was grossly misunderstood. Could you give a few specific examples of what you mean? (Or is there a hidden truth here, also? Are you just giving us some more raw meat, looking for a little more amusement?)

Paul, what you are seeing in my reply (as I would interpret it using PCT) is a little hint of how a hierarchical control system feels, and what it does, when it discovers that someone has deliberately jacked around with perceptions at the levels of system concepts and principles. You did that and I'm pretty hot.

> But theory has two uses, understanding and application. As Jeff Vancouver is quoted as saying - Psychologists are trying to do something useful. I interpret this to mean apply theory to helping people, society, or personal life.

That's what I'm talking about also, when I say the behavioral sciences (and their near relatives in the sciences) are dangerous. As in:

>> From Tom Bourbon [940707.1729]

>> ...one reason I say most claims to knowledge by psychology are dangerous is that innocent people might be tempted to believe those claims -- some might even go so far as to appeal to psychological "knowledge" when they attempt to establish new social policies, laws, and the like.

>>From Tom Bourbon [940708.0816]

>> All there is to support the majority of "theories" and "models" in the behavioral-social-neuro- sciences is their apparent plausibility -- they sound as though they could be true.

> It goes beyond plausibility to the realm of utility.

I'm open to specific examples.

> Behaviorists may not be able to provide a working model of the mechanism, but can give heuristics for how positive or negative reinforcement affect observed behavior and learning.

Fine. A technology and not a science. But that isn't where they stop with their own descriptions of what they do. Who was it that said behaviorism, more specifically radical behaviorism, _is_ _the_ science of behavior?

> Gestalt psychology has observed and defined certain stages of cognitive development, even if they don't provide the mechanical underpinnings.

Good enough, if we grant the veracity of those observations, which not all developmentalists will do. But I thought science was supposed to entail more than a few empirical observations.

> Priests and ministers can help people with life and make them more comfortable (and the reverse in spades!), even if metaphysics and theology are inherently un-provable.

Yes, but do they then pretend to the title of "behavioral scientist?"

> Alchemists were able to make chemical compounds even if the earth/air/fire/water element were not the correct model. A hot air balloon will work regardless whether we use gas density or the Aristotelian concept of a desire towards heaven as an explaination. The point is all these theories are or were useful to some extent. They could also cause people to come to wildly wrong conclusions under certain circumstances. It

is always a question of fact. The question is not "is it right" but rather "does it work".

All well and good, but all irrelevant to the discussion about a science of behavior. Oh, but I forgot: you weren't really serious about that discussion, were you?

>>[From Rick Marken (940707.1415)] "Out on a limb"
>> From Tom Bourbon [940707.1729]

>>Rick

>> Once you demonstrate that a higher level behavior is a controlled variable, then you know you need a control model to explain it. A control model is necessary because you are modelling the phenomenon of control.... -- but when behavior involves control, you can only model it with a control model.

> This I feel is an oversimplification. A control model is required to capture the _control component_ of the behavior.

And you can demonstrate that a control model has no applicability beyond that? Even Albus knew better than that. ("There is no such thing as _mere_ motor behavior. . .. The intellect is _not_ something distinct from behavior.")

> Further PCT has not (to my knowledge) that higher level behavior _is_ a "a controlled variable".

Something is missing from your post. Did you want to say that PCT has not _said_ that higher level behavior is a controlled variable? Or that PCT has not demonstrated ...? Or was it something else?

> I personally suspect it involves control of sets of controlled perceptual variables.

So do we, if by "it" you mean that what most people call higher level behavior, PCTers call behavior that controls higher level perceptions.

> It also IMHO involves sets of controlled reference variables, reference values, and some mechanism to determine from memory and/or extrapolation what variables of both types are or are going to be important.

Hmm. So you weren't _completely_ disingenuous when you appealed to Albus, after all. Time will tell, on this point. I would still like to know how you might program and run the model for this kind of performance. Albus did not provide that level of detail in his "outline." And I would like to know if you think such a model would be a representation of how living systems function.

> {out of order}

>> The answer, according to PCT, is that higher level behavior (like all behavior) is _ a controlled result of action_.

> I am not sure that control phenomena is all that exists.

Certainly not. There are many natural phenomena other than control. For example, there are Cause--->Effect phenomena. There just aren't many (any?) examples of those in the behavior of living things.

> There are components that appear reactive. (I didn't fully follow the thread on alerting, other than to observe the apparent response that the behavior could be produced by a sufficiently fast control loop). When a hand is burned, the arm will jerk back. It is my understanding that this does not involve the brain, just neural reflex at the spine. I seem to recall something on the order that the optic nerve of a frog transmits no signal to the visual cortex until a moving object like a fly enters the field of vision.

I was about to launch into a detailed reply to this point, but I forgot that you still haven't read any of the PCT literature. In that literature there are numerous discussions of examples in which a PCT analysis reveals that "components that appear reactive" rarely are purely reactive.

> Having a continuously active control loop monitoring every possible control variable seems very inefficient.

Not every variable -- every controlled perception.

> Some inputs are discontinuous or periodic. Bill Powers statements recently notwithstanding, a control loop can stop, or at least idle.

Or do they _seem_ to do that?

> Computer systems often use something called and interrupt chain. If no sensory inputs (events) are currently active (one of the results of an error signal existing can be the generation of a discrete input to another loop) the control loop hibernates.

Sure a computer system can use that procedure; but we are trying to study and model living things and they are a whole different breed of cat.

>> So the idea that PCT can handle simple behavior while the jury is still out on its ability to handle higher level behaviors really misses the point completely. Simple vs higher level is not the dimension on which PCT is discriminated from other theories of behavior; the dimension is "generated outputs vs control".

> What I have noted in most discussions, particularly that about Albus, is the focus upon how "PCT is discriminated from other theories of behavior". You focus upon the differences, not the commonalities  or how the approaches might be blended into a more powerful whole.

Please.  At the start of this post, you said you weren't really serious when you appeared on this net with the following post:

======================
P. George, 23 June 1994:

I have been lurking for a while listening to the PCT debates, and think you may be using too simplistic a view of a control system, which is biasing your discussion (in fairness I haven't had access to the books/papers on the subject, just posts on BPR_L by Dag).

Sophisticated control systems don't use reference variables (e.g. setpoints or alarm thresholds) they use reference _models_ (reflected in control logic). In a sense we have a continuously running simulation of the 'real' world to which we compare our perceptions (sensor inputs). As PCT stresses we also filter our perceptions based upon our current model, trying to separate signal from noise.

The basic components of this view of a control system are Sensory Processing (filter/transducer), Value Judgement (comparator), World Model/database, and Behavior Generator. Part of the model is some set of goals or desirable states we wish to maintain or approach. We modify the model to better reflect 'outside reality' (as perceived) so that we can better predict the results of our actions (actuator outputs).

For a good exposition of this model I suggest:

"Outline for a Theory of Intelligence" James S. Albus IEEE Transactions on Systems, Man, and Cybernetics, vol 21 #3, May/June 1991, p 473-509. IEEE log # 9042583 about the Albus article.

========================

It looks to me as though _you_ directed the focus to ""how "PCT is discriminated from other theories of behavior".""  Please don't tell us that _we_ pressed that focus upon you. You appeared making strong assertions of differences, which you now say were not really serious but were in fun, and now you say we have focused on the differences. Come on, Paul.

> It's kind of like "well, that's not what we're talking about, or it isn't the same point of view, or it doesn't address something we do. Now having disposed of that, we can go back to our own world and what we're doing."

Another of your little jokes, looking for some more laughs? I'm more than mildly amused, myself.

. . .

> I would challenge you to try to design or understand a distributed real time system at the assembly language or register level. It _can_ be done, but not by many.

Fine. And I repeat my invitation (not a challenge) for you to take a single-loop, fully-interconnected, Albus model and show that it duplicates the performance of a person and a PCT model on a simple control task. Let's start there, Paul.

> But higher level constructs such as 'tasks', 'messages', data structures are useful for understanding and constructing the application, even if it does ultimately reduce to bit twiddling. Similarly, to understand thought using only the mechanics of dendrite firing is at best difficult.

The simple stuff first, please. My invitation still stands.

> To say "if they can't do the simple stuff why should we pay attention to anything you have to say at higher levels" is at best unfair and at worst disingenuous.

Did someone here say that? I believe I said something to the effect that if the causal mechanisms assumed in a theory or model of behavior can't produce accurate and reliable results in simulations of the "mere" or "simple" kinds of behavior, then there is no possibility that that causal mechanism can explain and predict "higher" or "more complex" behavior. Maybe I'm wrong. Can you tell me where my reasoning has gone astray?

> Their models do not purport to address that level of detail. Frame of reference and scale of reference are important concepts to keep in mind.

Don't they, now? I seem to see a literature _filled_ with assertions about the details of how behavior is produced -- it is planned, commanded, controlled, programmed, and so on. Those untested and unworkable assertions come from people who purport to have a theory of complex behavior -- the important word is "behavior." I don't put that word into their writings, they do. If they speak of behavior, and further if they say they can explain complex behavior, then I expect to see at least _a little_ evidence that their assumed causal mechanisms can work. I guess there is very little I can do about the fact that this attitude of mine strikes you as disingenuous and amusing. Them's the breaks.

> Ultimately I suspect PCT will meet other approaches somewhere in the middle, providing both mechanism and greater understanding of the parameters of some kinds of human behavior.

Go ahead and suspect all you want. I _had_ hoped that you might actually join in and help with the project, but now I know you were only playing.

> When you can build an executable control model showing how one puts on a sweater and/or adjusts a thermostat, and decides between the options, I will be convinced. A skin temperature control variable is not quite sufficient. Ah well, +sufficient unto the day be the troubles thereof'.

Ah, well, since all we have ever written about and modeled is skin temperature, I guess we should simply leave the field to the real scientists -- the ones looking at the Really Big Questions. (Read some of the PCT work before you say such things about it, Paul. Don't keep speaking out of ignorance.)

I am happy that you suggested we, "build an executable control model showing how one . . . adjusts a thermostat." That happens to be one of the models I will demonstrate at the upcoming meeting of the CSG, in Colorado. It's nice to know there is at least one person in the world who might be interested in the results.

Oh, and by the way. Would you tell us about some of the people who have already built the models for putting on a sweater and adjusting a thermostat? Or has no one done that and you are just urging us to be the first? If they haven't already done those things, what _have_ they done that has convinced you they are on the right track?

Should I send this flaming piece? Probably not, but what the heck.

Later, Tom

Date: Fri Jul 08, 1994 5:43 pm PST
Subject: Look before you leap.

[From Bill Powers (940708.1830 MDT)] Paul George (940709.1130)

> Thanks. In truth I pointed out the Albus article in the spirit of throwing raw meat to wolves. The response of 'the masters':-) has been both informative and amusing. Defending and attacking expose what the proponents of a theory consider most important with amazing rapidity.

Sorry -- they expose what you THINK they consider most important, the exposure occurring more rapidly the more willing one is to jump to simplistic conclusions.

Paul, I really do urge you to get up to speed with the PCT literature. If you keep making comments of the kind you've been offering before pausing to find out what you're talking about, you will have generated such a backlog of embarrassing misinterpretations that you would have to be Jesus Christ to admit to them. Why not bite the bullet and learn what PCT has to say? Come on, it won't kill you.

Best, Bill P.

Date: Fri Jul 08, 1994 5:52 pm PST
Subject: Re: PCT models

[From Dag Forssell (940708 1800)]

On June 16, I posted on BPR-L: (Business Process Reengineering -L)
------------------------------

> Bill Powers writes regularly on the Control System Group List: CSG-L

> To subscribe and learn of your options, send

> Subscribe CSG-L Firstname Lastname
> help
> info refcard
> query csg-l

> to: LISTSERV@VMD.CSO.UIUC.EDU

> CSG-L volume is about 300 KB per week, against BPR-L 300 KB per month. Many posts are long and technical, but sincere questions, no matter how simple are welcome and get careful answers. Listen in, please read the literature first and introduce yourself when ready. See you there.

-----------------------------

Paul, you lurked for an entire week, before you emerged from the shadows without having read the literature. You still have not.

Now you say:

> In truth I pointed out the Albus article in the spirit of throwing raw meat to wolves. The response of 'the masters':-) has been both informative and amusing. Defending and attacking expose what the proponents of a theory consider most important with amazing rapidity.

You have asked insincere questions and received careful answers.

Now you say to Rick:

> This I feel is an oversimplification. A control model is required to capture the _control component_ of the behavior. Further PCT has not (to my knowledge) [shown??] that higher level behavior _is_ a "a controlled variable". I personally suspect it involves control of sets of controlled perceptual variables. It also IMHO involves sets of controlled reference variables, reference values, and some mechanism to determine from memory and/or extrapolation what variables of both types are or are going to be important.

and

> I am not sure that control phenomena is all that exists. There are components that appear reactive. (I didn't fully follow the thread on alerting, other than to observe the apparent response that the behavior could be produced by a sufficiently fast control loop). When a hand is burned, the arm will jerk back. It is my understanding that this does not involve the brain, just neural reflex at the spine. I seem to recall something on the order that the optic nerve of a frog transmits no signal to the visual cortex until a moving object like a fly enters the field of vision.

You are arguing from a standpoint of rather complete ignorance of PCT. PCT explains the reflexes you mention. You have been shown great courtesy in the form of sincere answers to your "throwing raw meat to wolves". I would like to perceive you returning the favor, by asking sincere questions. Sincere based on a real desire to understand.

> Looking forward to next week.

Please help make this mutual. Read the literature so you can learn from the exchanges, not just entertain yourself by jerking sincere PCT advocates around.

Best, Dag

Date: Fri Jul 08, 1994 9:35 pm PST
Subject: Re: PCT models

<[Bill Leach 940708.21:38 EST(EDT)] >[Paul George 940709 11:30]

Unfortunately both Ed Ford and Dag Forssell appear to have been very busy of late. Both of these people are "practicing" PCT in counseling or teaching. The both spend a great deal of time interacting with people that 1) have never heard of PCT and 2) could care less (at least initially).

I believe that Ed might tell you something like: "The very best counselors and therapists actually employ many principles of PCT in their work but would (and usually do) fervently deny doing so."

> It goes beyond plausibility to the realm of utility. Behaviorists may not be able to provide a working model of the mechanism, but can give heuristics for how positive or negative reinforcement affect observed behavior and learning.

The real root of the problem here is precisely what Tom is talking about:

Yes, INDEEDY, they provide heuristics and such for various behavior. They have names for about every "abnormal" behavior condition that exists and in many cases they have "treatments" designed to "correct" the behavior. The reason that this is actually dangerous to the world is that, if the basic PCT concept is right -- and there is compelling evidence that the basic concept is indeed right --, then these behaviouralists are treating symptoms and not causes. This can work (and indeed sometimes actually might work) but the problem in the realm of psychology (as predicted by PCT) is that any given set of symptoms could be caused by an almost infinite range of "causes".

This could be compared to how a medical doctor operates. The medical field is still almost a soft science itself but at least doctors actually test for real, verifiable physical conditions prior to responding to some symptom set. As I am sure that you are aware, failure to do so would lose a large percentage of patients (again, same symptoms can be produced by wildly varying causes, each of which requires a different treatment methodology).

Basically, this is how PCT and the other behavioral "sciences" are related. PCT recognizes "behavior" as a symptom. PCT recognizes that a statement such as 83.5% of the test subjects reacted "such and such way" to a stimulus, is only saying that 83.5% of the test subject were controlling the same reference the same way (and that actually should be worded ... roughly the same reference in roughly the same way -- since in most studies many conditions important to the study are neither controlled nor monitored). The question that PCT both raises and is concerned with is... what about the other 16.5%? Are they not real people too? An 80 or 90% correlation sounds good (and it is rarely actually that high in any psych studies even if the study says otherwise) but it is pretty lousey.

How would one feel if the laws governing hydraulics and mechanics (such as pertain to the coefficient of friction) only applied about 90% of the time. Thus, roughly every tenth time that you applied the brakes on your automobile they just did not work -- of course the same would be true for everyone else's brakes too (and we think that we have carnage on our hwys now!).

> Priests and ministers can help people with life and make them more comfortable (and the reverse in spades!), even if metaphysics ...

And the logical extensions to the theory of PCT can help to explain both results and why.

> Alchemists were able to make chemical compounds even if the earth/air/fire/water element were not the correct model.

Yes and it is quite likely that Alchemists killed a fair number of people too (including themselves). There is nothing wrong with doing something useful using a script even if the theory is incorrect. Where the error occurs is when one tries to generalize and apply the incorrect theory where the theory's error(s) are significant.

All of the sciences have had their "bad days" and rejected the correct approach in a big way (most do something similar in a small way all the time of course). The issue with the behavioral sciences right now is that they are only "making noise" concerning application of scientific method to their work. Behavioral scientists are not rigorous. They basically all run around saying that human behavior is too complex to be able to prove a theory. Having accepted that idea, they then refuse to either examine hard science data or have such principle based tests applied to their own theories.

This might not be so bad except that like the doctor, they are literally killing their patients and are providing "guidance" to policy makers, law makers, leaders and managers that is just plain wrong. Fortunately (or maybe that should be unfortunately) people seldom actually follow the advice fully (including the promoters themselves).

The behavioral sciences is a field full of inconsistent rules and exceptions or theories so vague that one can not actually conclude anything.

You can probably no more meld PCT to other theories than you can design one of your plant control systems without using negative feedback control (thinking about that, you actually can control some industrial processes without explicit negative feedback control though it tends to make a failure prone system).

> This I feel is an oversimplification. A control model is required to capture the _control component_ of the behavior. Further PCT has not (to my knowledge) that higher level behavior _is_ a "a controlled variable".

And this is the essence of misunderstanding of PCT (I believe). What PCT as a theory states is that BEHAVIOR IS THE CONTROL OF PERCEPTION. That's it, that's all of it and there is pretty sound reasoning (in addition to what actual testing has been performed) to be able to conclude that there need not necessarily be anything else.

> Having a continuously active control loop monitoring every possible control variable seems very inefficient. Some inputs are discontinuous or periodic.

I don't know that the theory says that this is the case for all situations but I doubt that the opinion that it "seems very inefficient" is a particularly compelling reason to suspect that it is not the situation.

I suspect that you are possibly confusing "conscious" control with control in general. In some of your plant systems, the control systems do not stop controlling when they do not have anything to do but rather they control for no action. In your distributed processing systems, it is quite common for controllers to operate for extended periods without reporting or requesting from their host and indeed in some such systems no report (other than maybe a periodic sanity check) is made unless there is either a control failure or "perception" that there may soon be one. This is still control and it is also control in the PCT sense too.

> What I have noted in most discussions, particularly that about Albus, is the focus upon how "PCT is discriminated from other theories of behavior". You focus upon the differences, not the commonalities or how the approaches might be blended into a more powerful whole.

But the difference IS significant. If behavior is "response to stimulus" (as it can often appear to be when the stimulus disturbs a controlled perception) as opposed to control of perception then PCT is utterly wrong and has nothing to offer the world about the operation of living organisms. I am sure that a great deal has been learned by the PCT folks in studying other behavioral scientists' work, indeed they often provide the "proof" of the failure of their theories within their own reports.

> The major difference I see between PCT proponents and others is that traditional psychology and sociology looks at behavior from the outside of the organism, while PCT focuses on the inside. Psychology has said 'We observe the following[...].. What can we say about predicting it or understanding the relationships between patterns of apparent events and observable actions of the organism?' PCT says 'can we identify something being controlled internally and a negative feedback{ack Rick} relationship? Can we simulate it?' (yes it's more involved than that).

Unfortunately, the problem with the observation is that if such reasoning is based upon a faulty model of the fundamental operation of living beings the conclusions will at best be true for only a very limited (and probably unspecified) set of conditions -- which is exactly the case now.

> It is kind of like the difference between a 'top down' and 'bottom up' approach to analysis. Both have their pitfalls.

This sounds like a valid argument but it is not (in my opinion anyway). "Top down" and "bottom up" both deal with exactly the same knowledge and exactly the same understanding of operational details.

"Top down" and "bottom up" are not equivalent to "Alchemy" and "Chemistry" nor is other "behavioral science" related to "PCT" that same way. If the PCT field were large enough then indeed there might be people working both approaches within PCT but I don't think that is the case right now.

> I would challenge you to try to design or understand a distributed real time system at the assembly language or register level. It _can_ be done, but not by many.

Again, the comparison is not in my mind fully valid. Someone somewhere did actually fully understand the details of hardware operation of such systems. The libraries (or OS access modules) for example are written by people that have intimate knowledge of the hardware. They precisely specify the calling sequence for their use and "hide" the gory details but the details are known.

-bill

Date: Sat Jul 09, 1994 12:12 am PST
Subject: Re: Testing models

<[Bill Leach 940708.23:01)] >[Paul George 940708 15:00]

I probably should do a little more controlling for accomplishing some work around here but seems I can't resist. <G>

> This is true only if you model the diagram naively and permit all paths to be traversed. If you put a state machine(s) or PCT control node(s) within the nodes, and/or any discrimination node which can rout different signals to various control or sensory nodes, then all paths are not possible.

This supports what has been said. Something would have to CONTROL the "model" so why bother with that model?

> A PCT control node must ignore all inputs that do not affect the controlled variable, or more accurately never receives them.

Interesting assertion. Why does this have to be true?

Pavlov's dog is explainable through PCT and pleasantly enough, the explaination is even capable of explaining why the dog does not always seek the reward.

> The organism says 'gee, I liked ...

You really do need to read some of the publications in PCT. My impression is that you have tried to define how the control system must operation and then challenge your own design. I am not qualified to really discuss in depth how this "reinforcement" process actually works but am willing to note that with all of the claims of its proponents there are (or were) eminent psychiatrists (Frankl comes to mind) that refuted the validity of the theories (though they did not have any replacement theories).

-bill

Date: Sat Jul 09, 1994 12:30 am PST
Subject: Re: PCT models

<[Bill Leach 940708.21:11 EST(EDT)] >Tom Bourbon [940708.0816]

I probably should have phrased that in more precise terms. PCT is not proven true in the sense that it can be taken as a physical "law" (such as like charges repel -- and even that is not a good example), ... yet.

The choice of the word "plausible" was poor also in that people (at least in the "soft" sciences and the general public) tend to view plausible as meaning "reasonable" rather than the idea that there is a actual body of physical evidence related through consistent laws or theories that can be examined objectively.

Basically, what I hear again and again on the CSG-L is that HPCT and the conjectures concerning it are hypothesis not theory. However, that hypothesis has the advantage over other behavioral "theories" in that it has not failed any serious challenge. The problem being that we don't KNOW that HPCT is capable of generating the observed human behavior but rather it is plausible that it could do so.

-bill

Date: Sat Jul 09, 1994 6:44 am PST
Subject: No control without feedback

[From Bill Powers (940709.0740 MDT)]

Bill Leach (940708.xxx) --

Thanks for the morale-building private post, Bill.

In trying to get through to Paul George you make a whole series of lovely observations. In some places you could make an even stronger case:

> You can probably no more meld PCT to other theories than you can design one of your plant control systems without using negative feedback control (thinking about that, you actually can control some industrial processes without explicit negative feedback control though it tends to make a failure prone system).

I think we can say flatly that there is no way to design a control system (one that creates consistent results by acting on and in a real environment) without feedback. To prove this, all you have to do is keep a sharp lookout for the feedback that is ALWAYS present when continuing regularities are seen.

To control anything literally without feedback, you would have to set up the situation so the person or machine doing the controlling could issue commands that affected the processes in the plant, BUT WAS NEVER ALLOWED TO OBSERVE THE RESULTS either while issuing the commands or at any later time. Just try to imagine how we could control ANYTHING under those conditions. You would remain blissfully unaware of failures, wouldn't you?

Or go even further: give a person or machine the ability to perform acts that affect the plant, and the ability to sense the resulting operation of the plant, but eliminate any conception of the desired operation of the plant: whatever it is seen to do is OK, no error implied, no preferred kind of result implied. The perceptual information is just for the entertainment of the observer, and does not call for any change in the action. Would anyone call this control?

Sometimes engineers think they are really building an open-loop system. They calculate every contingency, build everything to micron specifications, double and triple check every component and every physical theory involved, and eliminate every possible source of disturbance. Then they send the design to the production department. But do you think for an instant that such an engineer (if sane) would walk away without seeing what this system actually does? No, what you find is that the engineer is standing by anxiously with a twiddle-stick, and spends a lot of time trimming and adjusting and bending things a little until the system is actually producing the effects that the engineer -- the control system -- wanted to see. And even after that, no sane engineer would seal the system to protect it from further alterations and sign the guarantee that it will work forever. Far more likely, the engineer will leave behind a thick book of maintenance and calibration procedures, which he expects will be followed daily, preferably hourly, and certainly at least monthly, by the person using the system. The engineer knows that in the environment of the real Turing machine, there is a kid with a spraycan who like to spray dots onto any tapes he comes across.

There is ALWAYS a control system involved, even if it's only the person who is using the open-loop device. Without that control system, only an imitation of control is possible. If the control system does not act, sooner or later the open-loop system will wander blindly off into stupid actions that not only fail to accomplish the original intention, but actively produce results other than those intended.

The concept of an open-loop automaton is a vestige of 18th- and 19th-century engineering, from the time when people felt that the universe was a divine clockwork mechanism that, once God had set it in motion, would run forever with divine precision. In fact, natural processes wander randomly, never exactly repeating and always doing some sort of random walk toward no goal-state at all. Chaos prevails. Everything disturbs everything else. The more closely you look at natural processes, the more obviously the sense of orderliness is violated. The basins of attraction in nature are shallow and subject to disturbance.

In this universe only a feedback control system can create the same results over and over indefinitely, defying both chaos and probability. Only actions adjusted by purposes can create permanent regularity while the purpose lasts.

Best, Bill P.

Date: Sat Jul 09, 1994 9:58 am PST
Subject: The Nature of Control

[From Rick Marken (940709.1045)]

Paul George (940709 11:30) and (940708 15:00 EDT) --

I'll try to answer or ask questions about some of Paul's points in another post. But first let me suggest that Paul's basic problem -- the reason for all the apparently ridiculous statements in Paul's posts -- is a lack of understanding of the nature of the phenomenon of control. It looks like Paul does not intend to read the PCT literature or work with the computer demos (which is fine with me since doing so is no guarantee that a person will get the idea anyway) so before he makes any more (what I'm sure are unintentionally) silly statements, let me explain what we mean by "control" in PCT.

We begin with the observation that, in living organisms (with nervous systems), efferent neural signals cause actions (variations in muscle tension) that have effects on the environment. We can represent the situation with an equation:

(1) $q = f(o)$

where o is an action variable (or output) and q is a result of that output. The function f() represents the environmental laws that determine the nature of the relationship between output and result. (Of course, any output variable produces many results; I'm just focusing on one result here -- one that might be under control). An example of what is described by equation 1 is lifting a book; q is the position of the book and o is the muscle force(s) that influence the book's position. f() is the differential equation that transforms muscle force into book position (taking things like the mass of the book into account). But it is important to note that q can be ANY result of output, including very complex results. For example, q could be "amount read about control theory" , a very complex (and variable) result of many outputs and intermediate results of those outputs (like lifting a book); but a result of output, nevertheless.

The next observation is that any result of output is also a result of non-organism produced, environmental influences -- disturbances. Assuming that these influences add to the influence of the organism (output), the complete equation for the cause of any result of output is:

(2) $q = f(o) + g(d)$

In this equation, the function g() represents the differential equations that relate the disturbance variable(s), d, to the result, q. A disturbance to the position of the book, for example, is the gravitational force, which is basically constant; a variable disturbance might be a force on the same book exerted by someone trying to help you lift the book.

So, at any instant, the state of the result of an organism's outputs (the value of q) depends on the state of the organism's outputs as well as the state of (possibly many) environmental disturbances, d.

We say that a result of output is under control when q remains in a constant or variable reference state, q*: that is, when q = q*. This only happens when the influence of output on q -- f(o) -- precisely mirrors the influence of all other variables on q -- g(d). Letting q* be constant and zero (for simplicity), we say that there is control when:

(3) f(o) = - g(d) (approximately)

Remember that o, d and q are variables -- changing over time. It is, therefore, highly unlikely that f(o) would remain precisely opposed to g(d) for long simply by chance. Thus, equation (3) is the basis of The Test for control; we apply arbitrary influences to q (we vary g(d)) and look either for precise opposition from the system (-f(o)) or _no effect_ of variation in g(d) on q. If there is precise opposition to g(d) or no effect of g(d) on q then we say that q is a _controlled variable_ and it's reference state (which can vary over time) is q*.

The Test need not be applied quantitatively; we can get hints about controlled variables by simply "disturbing" them. For example, if we think a person is controlling the amount they read about control theory, then we can disturb this variable by talking about nothing but politics and religion on CSG-L; a person who wants to read about control theory might protest, question what is going on or try another list (the fact that there are many different ways to influence a controlled variable is why we usually monitor the effect of g(d) on the controlled variable itself, rather than on f(o), since we rarely know all the outputs that might be used to keep a variable under control).

Again, it is important to realize that a controlled variable, q, can be any _variable_ result of an organism's output -- it can be the position of a book (in x,y z space) or one's position on abortion (in political space?). It is also important to remember that all the results people produce, which means everything we call their "doings" -- lifting books, reading about control theory, going to get an abortion or preventing others from getting one -- all these results of output are a result of influences of _both_ outputs and disturbances (equation 2). To the extent that people produce these results reliably (produce consistent values of q) then these results must be under control -- that is, outputs must be varying appropriately to counter the influence of disturbances to these results (equation 3). This is what we mean when we say that all behavior -- all consistently produced results of action -- are under control. And we do have evidence -- quantitative evidence -- that people can control simple variables, like the position of a cursor, and "higher level" variables, like a sequence or program of numbers that is occurring on a computer screen.

But the evidence of "higher level" control is all around you; virtually every aspect of the world you live in is a result of people (including you) contriving to produce the results you see rather than the results that would have been produced by the operation of environmental disturbances (even if these were combined with carefully pre-programmed outputs,o). Higher level control is evidenced by the fact that there is a computer sitting on your desk rather than a pile of sand,

Now, don't you feel like starting all over again, Paul ;-).

Best Rick

Date: Sat Jul 09, 1994 8:29 pm PST
Subject: Re: No control without feedback

<[Bill Leach 940709.18:54 EST] >[Bill Powers (940709.0740 MDT)]

And thank you, maybe I am learning something after all :-)

>Bill Leach (940708.xxx):

> ... (thinking about that, you actually can control some industrial processes without explicit negative feedback control though it tends to make a failure prone system).

As is usual with such discussions, "levels" can be a problem. I can think of literally hundreds of actuators that operate with negative feedback control. In some cases there are automatic systems that will detect the error in a rather indirect way but in many cases about all the control system perceives is that there IS something wrong with neither the "knowledge" of what the problem might be nor the ability to correct the error.

So, actually what I was thinking of is that in some cases simple "loops" are sometimes created where an actuator operates blindly in the sense that there is no sensing of either the actuator's operation nor the results of the task performed by the actuator. This is still called control or at least is a part of control system though the use of the term in those cases is really a bit loose (of course if one specifically calls it "open loop control" then there would be less confusion).

Indeed, that was exactly that type of "control" that caused the TMI incident. The Power Operated Relief Valve was operated "blind" and failed to close after the control system opened it to reduce an over pressure condition. The open operation is actually blind also but there is rather obvious negative feedback in the sense that the controller calling for the valve open condition will respond to the resulting pressure drop if the valve does indeed open.

This is not intended to excuse the operators of the plant for failing to check the actual valve condition. It is and always has been common knowledge among competent steam plant operators that relief and safety valves should be expected to leak following operation.

An automobile power steering system is a good example of a hydraulic negative feedback control system (with mechanical over-ride). A hydraulic metal punch OTOH is often NOT a closed loop system unless the human operator is counted as a part of the loop. That is, when the valve opens to admit oil to the cylinder to drive the punch head it is assumed by design that the punch will travel the required distance and in a like manner when the cylinder is relieved (or the withdrawal cylinder is pressurized) it is expected that the punch head returns.

I don't think that there is necessarily anything wrong with designs that are open loop as long as the operator that is monitoring or using the equipment is aware of the fact and its' significance. For really complex systems however, such components are a poor idea.

-bill

Date: Sat Jul 09, 1994 8:56 pm PST
Subject: Re: open-loop "control"

[From Bill Powers (940709.2140 MDT)] Bill Leach (940709.1854 EDT)

> Indeed, that was exactly that type of "control" that caused the TMI incident. The Power Operated Relief Valve was operated "blind" and failed to close after the control system opened it to reduce an over pressure condition.

Yes, I've used that example, too. I believe there was a light on the control panel indicating "open" and "closed," but all it indicated was the state of the switch. It should have been connected to a flow sensor, shouldn't it? That's what you _really_ want to control. Even connecting it to the valve gate isn't enough: there might be somebody's overalls stuck in the orifice, or the water supply might have run dry (or, as at TMI), you may think the cooling water flow is shut off while it in process of draining dry. What you sense is what you control.

> A hydraulic metal punch OTOH is often NOT a closed loop system unless the human operator is counted as a part of the loop.

That was my point: the loop is _always_ closed, somehow, if there is to be control. When you build an open-loop device, it becomes part of the output function of a human being who watches the result to make sure it's what is wanted. Nobody trusts a purely open-loop system, and shouldn't. You don't press the actuate button on the hydraulic press without watching what the press does.

I've been watching Nasa Select for the last couple of days, and listening to the traffic between mission control and the spacecraft personnel. You can see the results of lessons learned over many years: everything is fed back. Sts bio, Houston. Houston, go ahead. I have a temperature reading from the biorack. Ready to copy. Incubator A reads niner five point 3. Copy niner five point three on A. That's affirmative, Houston.

Even when someone asks an astronaut to flip a switch, there are people on the ground monitoring the effect. A couple of missions ago, this led to tracing down a communications malfunction to a power supply switch that was flipped to standby instead of on. This time, it led to figuring out that the biorack power went off because an astronaut had caught a switch with his toe (special problems of working in zero-g!). And the astronaut didn't just flip it back on. He reported the event, and asked whether or not to turn the power back on. Then, given the go-ahead, he reported that he had switched the power back on, AND that the biorack was again operating. And Mission Control acknowledged: biorack operational again. These people understand exactly what control means.

> I don't think that there is necessarily anything wrong with designs that are open loop as long as the operator that is monitoring or using the equipment is aware of the fact and its' significance.

But if the operator is monitoring the results, the loop is closed. In a truly open-loop system, NOBODY is monitoring the results. We just don't speak of control when nobody is monitoring the results. We speak of effects or influences, not of control (or should). The air temperature at the north pole influences the freezing of the sea. But it doesn't control the freezing of the sea. If the temperature drops and the sea doesn't freeze, the temperature won't do a durn thing about it.

> For really complex systems however, such components are a poor idea.

For simple ones, too, unless the consequences are unimportant. Remember, for truly open-loop control, nobody gets to see whether the device operated, or whether its operation had the expected effect. All you really have to do to convince an engineer that all control is closed-loop is to keep him from seeing what happens when he turns on his "open-loop control system."

Best, Bill P.

Date: Sun Jul 10, 1994 1:59 pm PST
Subject: Re: open-loop "control"

<[Bill Leach 940710.17:31 EST] >[Bill Powers (940709.2140 MDT)]

Ok, but this is a "perspective" issue to an extent. At some level, probably no engineered control system is open-loop by design. Portions of it may however be open-loop in the sense that under normal conditions the operation of that portion is not monitored.

When designing a control system, normally the engineer considers the physically engineered parts as either open-loop or closed-loop based upon the control system behavior sans observer.

When designing something that is open-loop, a competent engineer will at least consider how the rest of the system will respond to a control failure in the open-loop part and how an observer might recognize such failures.

I am not sure that further discussion of this aspect is particularly important to PCT but recognize that viewing all systems that have human operators present with a perspective that they are a part of the system feedback is a good idea.

I think that much analysis of system failure modes fail to actually consider the operator in such a fashion.

-bill

Date: Sun Jul 10, 1994 4:15 pm PST
Subject: Replies to Paul George

[From Rick Marken (940710.1700)] >Paul George (940709 11:30) --

> (Though I have often thought Skinner was grossly misunderstood)

All Skinnerians feel that Skinner was grossly misunderstood. The fact is, however, that there is nothing to misunderstand about Skinner because he had no understanding to miss. Skinner did no modelling so he could change his fantasies about behavior faster than you could question his theoretical claims. Skinner could no more be misunderstood than could a Talmudic scholar; both know where their reasoning is going and they won't let data and models stand in their way. For example, I have experimental data that shows that operant behavior involves selection OF consequences (consistent with PCT), not selection BY consequences (as Skinner argued). Skinnerians have a nifty way of dealing with this data; rather than showing how a reinforcement model handles the data they simply say "you've got it wrong" and blythly press ahead. Quite an interesting "scientific" dialog.

> The question is not "is it right" but rather "does it work".

It's true that people can make things work without knowing the right (best) model of how it works; most human controlling works that way (we have no knowledge, for example, of the nature of f(o), the physical laws that relate our outputs to most controlled consequence of those outputs; nevertheless, we are able to control those consequences; this is just the way control works) but there are many things that would have been quite impossible to do without the right model (that is, without a detailed understanding of how our outputs influence the consequences of those outputs). For example, most high technology (space travel, computer chips) would have been quite unlikely to happen if people did not have the right models of physics and materials.

> A control model is required to capture the _control component_ of the behavior.

Correct. And that's a very big "component". Now that you have read my previous post on the nature of control, can you think of ANY behavior that does not involve control? That is, can you think of any result of action that is not brought to an intended level while being protected from the influence of disturbances? One behavior I can think of that has no "control component" is the downward velocity of a person who has jumped off a cliff (with no parachute or wing); not a very common behavior but a behavior that involves no control (after the jump, anyway). Can you think of any others?

> Further PCT has not (to my knowledge) [shown?] that higher level behavior _is_ a "a controlled variable".

Look back to last month's posts for a description of my program control study; I have a copy on HyperCard now. It shows unequivocally that a simple program of events ("if X then Y else Z") can be controlled (maintained against disturbance) by non-programmatic outputs.

> I personally suspect it involves control of sets of controlled perceptual variables.

Me too. It's called Hierarchical Control Theory. It's implemented as a real WORKING model. A nice description of a working, hierarchical control model is given in my book "Mind Readings", Ch. 6. See how useful reading can be:-).

> I am not sure that control phenomena is all that exists. There are components that appear reactive.

As Tom said, cause-effect phenomena exist, too. In fact, they are far better documented; they are what physicists study.

What appears "reactive" in a living control system, however, is part of the phenomenon of control; it is a reaction to sudden disturbances to a controlled variable. It looks like the disturbance causes the reaction but that's not quite how it works. The "reactive" component of behavior is disturbance - resistance that keeps the controlled variable at the reference level. The nature of the "reaction" to a disturbance depends (among other things) on the reference setting for the controlled variable -- which is determined, autonomously, by the organism itself. In my previous post I wrote the equation for disturbance resistance under the assumption that the reference for the controlled variable was fixed at zero. Here I rewrite the equation allowing for the possibility of a non-zero and variable reference, q*:

$f(o) = q^* - g(d)$

So the nature of the "reaction" (f(o)) of a control system to a disturbance depends on the reference setting for the controlled variable. What this means in real life is that the same disturbance will not always produce the same reaction. When someone takes a swing at your face (d) you usually duck (o) in order to maintain the amount of pain you experience (q) at a reference level of zero (q*). But if you want to prove how tough you are you might revise your reference for pain -- and the same disturbance (the swing) results in no reaction (you don't duck) and, therefore, you get hit -- bringing the controlled variable (pain) to its new, non-zero reference.

You say:

> When a hand is burned, the arm will jerk back.

Now you can see that this is only true if your reference for pain is zero. There is a great scene in the opening to "Lawrence of Arabia" where Lawrence holds his hand over a fire to demonstrate the absence of the "jerk reflex" when he changes his reference for pain. When asked if it hurts, I think Lawrence says something like "Of course it does; you just have to not mind". Lawrence of Arabia: control theorist.

> Bill Powers statements recently notwithstanding, a control loop can stop, or at least idle. Computer systems often use something called and interrupt chain.

The control loops we are talking about go through the environment, where variables don't behave like idealized digital circuit elements -- they are differentiable functions of time. A programming loop, for example, is not (usually) a control loop, in the control theory sense of control.

> You focus upon the differences, not the commonalities or how the approaches might be blended into a more powerful whole.

I wrote a paper called "The blind men and the elephant"; it is published in the CSG journal "Closed Loop". In that paper, I explain the "commonalities" between PCT and S-R, reinforcement and cognitive views of behavior. PCT shows that these three views of behavior capture three different aspects of control, but each, alone, misses the big picture. PCT shows why behavior has appeared to be a response to stimulation, selected by consequences or planned output.

In fact, behavior is none of these things, though it can look like each under the appropriate circumstances. Behavior is only one thing -- control.

> Ultimately I suspect PCT will meet other approaches somewhere in the middle,

I wish this were true. Unfortunately, reconciling PCT with other approaches is a bit like trying to reconcile Copernicus and Ptolemy. While it was eventually possible to see why Ptolemy would see things as he did, there was no "compromise" solution (like having the earth be close to the center of the planetary system -- where Venus is -- instead of three planets out). PCT shows why the "other approaches" see things as they do; unfortunately, it shows that what these approaches see is not what is. The "elephant" of control is not the "snake" of S-R, the "tree trunk" of reinforcement or the "wall" of cognition. And yet, they control;-)

> When you can build an executable control model showing how one puts on a sweater and/or adjusts a thermostat, and decides between the options, I will be convinced.

You should only be "convinced" of PCT when you are convinced that behavior IS control. This is an empirical -- not a theoretical -- question. If you are already convinced that psychologists know what they are talking about when they talk about "behavior" then PCT is irrelevant to you.

PCT can only be of serious interest to people who are convinced -- by OBSERVATION AND TEST -- that behavior IS the control of perceptual input variables.

If your selection of theories of behavior is just a "beauty context" -- to see which theory "looks best" then, even if PCT wins your favor, it is only a superficial victory and you will join the ranks of people (like Carver and Scheier, and other useless "adherents" of PCT) who like PCT but don't like what it is about -- THE PHENOMENON OF CONTROL. The first thing to learn and experience about PCT is the PHENOMENON that it explains.

I sometimes think that people should be required to test for controlled variables for a year before being allowed to read anything about the theory of control. Phenomena first!!

Paul George (940708 15:00 EDT) --

> All paths between nodes are not permissible in the HPCT model either (of course you define internal connection rules).

You seem to have answered your own point in parentheses. Powers explained how the HPCT model worked and how it could be implemented on a computer. Albus leaves most of the workings of the model to our imaginations -- which is fine but rather subjective, don't you think?

Powers says:

> And from this model, we conclude that "operant conditioning" is mostly ordinary control behavior, misinterpreted. So it isn't operant conditioning that works; it's control that works.

Paul:

> Close. Control is _a_ mechanism which _can_ implement operant conditioning, given certain assumptions. How is the control variable and negative feedback signal set up? (I have a little trouble with 'Unseen Hand' arguments)

In operant conditioning (once the animal has learned how to use the operant apparatus; that is, after "shaping" has occurred) the rate of reinforcement is kept under control; rate of reinforcement is the controlled variable, q. The so-called "schedule" that relates bar press rate to reinforcement rate is the "feedback" function, f(), that relates outputs to inputs. There are usually no disturbances introduced in operant experiments but they are quite precisely resisted when they are introduced; rate of reinforcement is maintained at a reference level, q*, despite disturbances or changes in the feedback function. A huge amount of operant data can be explained precisely by assuming that organisms control a perception of the rate of reinforcement by varying their output (bar press rate) appropriately to compensate for the changing circumstances (schedules mainly). So reinforcement doesn't "control" behavior ( now that you know what control means, wouldn't you find it amazing if it did?); behavior (the outputs of the organism) control reinforcement.

> Note: I don't really intend to have a lengthy discussion of the proper translation of OC to a control model.

That was clear from your brief description of what would be required to translate OC to a control model. I suggest that you read the description of a control model of operant conditioning starting on p. 67 of Powers "Living Control Systems" to see a proper translation of OC to a control model.

Best Rick

Date: Mon Jul 11, 1994 12:10 pm PST
Subject: Re: Replies to Paul George

[From Paul George (940711 ????)] >[Rick Marken (940710.1700)]

> All Skinnerians feel that Skinner was grossly misunderstood.

I'm not a Skinnerian. But I do think some of the ideas in B&D and Walden II were useful. There is wheat among the chaff. And you malign Talmudic scholars. OTOH philosophy is like that.

> And that's a very big "component". Now that you have read my previous post on the nature of control, can you think of ANY behavior that does not involve control?

I agree, and said so earlier. Never suggested that any behavior (above the reflex level) didn't involve control. Reflex can be overridden, but I am not yet convinced that reactive mechanisms don't exist in nature. As I said elsewhere, the term behavior covers a very wide semantic range. It does not always refer to the same thing.

> What appears "reactive" in a living control system, however, is part of the phenomenon of control; it is a reaction to sudden disturbances to a controlled variable.

Forgive a question probably answered in the literature. How do you know this for a fact? Can you empirically show each and every controlled variable, or just that they can exist? How do you identify a variable in tissue? Reactive behavior can be _modeled_ using control, but it does not necessarily follow that the map is the territory. Note that I do not disagree that behavior can be based upon perceptual control, just that it must be. That has not been proven.

> When someone takes a swing at your face (d) you usually duck (o) in order to maintain the amount of pain you experience (q) at a reference level of zero (q*).

This is the level where your model fails to convince me. There may be a gestalt phenomena, but the concept of ducking to keep a pain variable at 0 is a bit of a stretch. I don't disagree that the behavior involves control, I just have trouble with the described mechanism. The model requires considerable elaboration.

> I wrote a paper called "The blind men and the elephant"; it is published in the CSG journal "Closed Loop".

Any way to get an electronic copy? I can't get into anywhere the journal is carried.

> In operant conditioning (once the animal has learned how to use the operant apparatus; that is, after "shaping" has occurred) the rate of reinforcement is kept under control; rate of reinforcement is the controlled variable, q.

Sorry, I was refering to the likes of dog training, not the controlled experiments used to document the principle. Again, I am trying to apply the PCT mechanisms to reality (not to imply that they don't apply).I am focusing on design of control hierarchies to reflect actual behavior, i.e what would be the controlled variables, inputs, outputs, and control hierarchy. We should be able to define a HPCT system to cover known any behavior, whether or not the technology exists to actually implement it.

Date: Mon Jul 11, 1994 2:30 pm PST
Subject: Re: PCT models

[Paul George 940711 11:30 EDT]

>[From Tom Bourbon 940708.1653]

> Should I wait until Monday, so I can cool down a bit? Nah! Would you let us in on what you found the most amusing? And would you maybe let us in on your personal assessment of Albus, now that we know you weren't playing straight with us for the past week and a half? Whatever your assessment may be, I guess we don't need to wait to see any implementations of a genuine Albus model by you, do we?

Similar sentiments from [From Dag Forssell (940708 1800)]

Cool Down allready! Mighty thin skin for senior (I presume) researchers. Dry humor seems to fall on deaf eyes. I guess I have to make greater use of smilies, and be very careful about phrasing and images.

I wasn't playing with you. I suggested the article since it seemed to provide a sounding board for PCT. It _was_ a sincere question. I got much reaction and discussion that revealed much about PCT, which was my intent. As I said repeatedly, I didn't consider the Albus article to be a seminal work, just interesting. We are designing a control system using some of Albus' concepts, but that isn't the kind of implementation you are interested in. I am an engineer, not a scientist, even if that's not the way I planned it many years ago.

FYI I find discussion amusing, and a number of the comments and examples that have been made. Amusement is not the same as mockery. I will read the literature as soon as I can get my paws on it. So far I've had no luck in libraries, and I don't have access to a university that has a strong psychology or neurological department.

> I have a long-standing interest in Skinner and in the often-repeated claim that he was grossly misunderstood. Could you give a few specific examples of what you mean?

It has been a number of years since I read Skinner's books. Other schools often argued he felt that there was no such thing as intelligence or personality, just conditioned response. My reading was that he felt that one could not determine the existence of 'soul' and such empirically and so one should stick to what could be observed. He felt that OC was a useful technique for training humans to get more 'civilized' behavior. And he thought that treating people as anything more or less than animals was at best hubris. This enraged the Humanists in particular. I always thought they were saying much the same thing from different angles, but were too busy attacking one another to allow communication. Of course professorships and tenure were on the line.

> Further PCT has not (to my knowledge) that higher level behavior _is_ a "a controlled variable".

> Something is missing from your post.

Sorry, too much cutting and pasting. Should have read "demonstrated" as in experimentally. PCT makes the assertion that 'all behavior is control', where I would suggest that something like "all behavior involves control" is likely more accurate.

> I would still like to know how you might program and run the model for this kind of performance. Albus did not provide that level of detail in his "outline." And I would like to know if you think such a model would be a representation of how living systems function.

I suspect it describes how organizations and high level 'reasoning' type behavior works. I doubt it goes all the way down. At some point the concepts of value judgement and world model atrophy into a reference variable value. I also suspect that something like a PCT control system is involved in the internals. I can think of no way to prove such a model except by building a true artificial intelligence, and that would only prove that the model _can_ work, not that living systems are in fact is constructed that way.

> That happens to be one of the models I will demonstrate at the upcoming meeting of the CSG, in Colorado.

I would be interested in the particulars. The hints you and Bill Powers have given in other posts made it appear to be at a bit lower level than I was talking about. Part of the problem may be that the word 'behavior' is used to refer to everything from reflex to painting the Sistine Chapel.

> Oh, and by the way. Would you tell us about some of the people who have already built the models for putting on a sweater and adjusting a thermostat?

My point exactly. No one has yet built working models of high level behavior without gross oversimplification. When we do we will have created an artificial lifeform.

>[Bill Leach 940708.21:38 EST(EDT)]

Thanks for a lot cooler post.

> The reason that this is actually dangerous to the world is that, if the basic PCT concept is right -- and there is compelling evidence that the basic concept is indeed right --, then these behaviouralists are treating symptoms and not causes.

> Behavioral scientists are not rigorous. They basically all run around saying that human behavior is too complex to be able to prove a theory. Having accepted that idea, they then refuse to either examine hard science data or have such principle based tests applied to their own theories.

I concur, and that I why I find CSG and PCT so interesting. You have experimentally demonstrated the underpinnings, and generalizing the theory allows insight into higher level behaviors and clinical applications. There remains however a lot of work in the middle before we can say "all behavior is control" (or more accurately is the result of controlling perception). {Note the 'we'. I think you are on the right track, if occasionally locked into the worm's eye view}

>> It is kind of like the difference between a 'top down' and 'bottom up' approach to analysis. Both have their pitfalls.

> This sounds like a valid argument but it is not (in my opinion anyway). "Top down" and "bottom up" both deal with exactly the same knowledge and exactly the same understanding of operational details.

Using analogies is always risky. Actually In my area (software engineering) different levels of a model contain very different knowledge. Understanding of the operational details also varies widely (a severe problem). One of the most interesting things about working with OJT trained developers and users is the wildly varying assumptions they make about how things actually work in an application and how the hardware functions. Yes, "Someone somewhere did actually fully understand the details of hardware operation". But developer's often do not know the details, nor do they care.

I still think the analogy is valid. A top down approach takes the big picture and attempts to decompose it 'logically' (variously defined) into components. The bottom up attempts to create structures from existing components in order to meet goals. The two approaches can produce very different architectures. However, we have found that a better technique is to work from both directions meeting a third effort expanding from the middle (in reality you jump around).

The point is that different abstractions are useful at different levels of detail, and from different points of view. Asserting that only one theory is valid is dangerous, though occasionally right. I often find the parable of the blind men and the elephant enlightening.

Take care      Paul George

Date:     Mon Jul 11, 1994  3:48 pm  PST
Subject:  Re: The Nature of Control

[Paul George 940711 12:00] >[Marken (940709.1045)] (940710.1700)]

Thanks for the primers.

A number of you seem to be working under the misapprehension that I don't understand control from an engineering standpoint and that I disagree with PCT. I don't. I am not attacking you. I have no axes to grind, nor do I belong to any competing school of thought. Remember, _I_ am not a behavioral psychologist and have no vested interest in your being wrong.

Is there some reason you feel that PCT is not understandable? I know that you are apparently used to being misinterpreted by others in your field. Your examples make perfect sense to me and are rarely surprising. The higher level concepts flow logically from the basics. They largely match my own thinking. While I have more reading to do, this group has brought out the core of PCT quite well. I do have some questions about _how_ it scales up, and the mechanisms used to assemble and modify the controlled variables and control hierarchy. These are not defects, just areas for future work. I also tend to question phrasing that appears to 'fluff the wares'.

What you usually seem to read in my posts is rarely the point I am trying to make. Perhaps I use analogies with which you are unfamiliar, or terms that produce a reflex response. I try to feed back your examples from a different angle to produce discussion (I tend to play devil's advocate). The responses (from my point of view) seem to home in on side details that I didn't bother fleshing out or re-iterate points that were eloquently described earlier. We don't seem to be focusing on the same issues. This indicates that I am failing to communicate. But please be open to the idea that you actually transmit PCT concepts very clearly. Don't presume misunderstanding.

Date:     Mon Jul 11, 1994  3:49 pm  PST
Subject:  Re: Testing models

[Paul George 940711.15:00]      <[Bill Leach 940708.23:01 EST(EDT)]

>> A PCT control node must ignore all inputs that do not affect the controlled variable, or more accurately never receives them.

> Interesting assertion.  Why does this have to be true?

Perhaps I should clarify. This relates to the thread on open loop control. The universe provides a nearly infinite variety of inputs, or sense-able things. The body has a limited set of sensors. Their signals must routed to a control structure that interprets them in terms of some hierarchy of control variables. Sensory input is distinguishable from perception. A HPCT network has some set of controlled variables at a given time. Similarly the nodes 'generate' a given set of possible actions to attempt to control those variables. Part of the universe of possible perceptions are thus 'open loop' at a given time, as they are not being controlled. Many of the results of our actions are not perceived, as are the actions themselves.

One of the things I do not yet understand in PCT is how an organism 'decides' to control a perception and then discovers the actions that will affect it. It appears to be a bit of a chicken and egg problem. It even worse when a Higher node must set up a control relationship with lower nodes which actually do the sensing and acting. {I did read Bill Powers post on how organisms develop, and found it persuasive, if not complete}If just random chance was at work we would expect the control solutions to vary radically from individual to individual. OTOH we observe that people and other organisms seem to share control variables and reference values.

Best Paul

Date: Mon Jul 11, 1994 6:14 pm PST
Subject: Re: Testing models

[Avery.Andrews 940712.1022] (Paul George 40711.1500)

> One of the things I do not yet understand in PCT is how an organism 'decides' to control a perception and then discovers the actions that will affect it. It appears to be a bit of a chicken and egg problem.

It's either hard-wired, or stumbled across by reorganization. Cockroaches have some hard-wired circuitry that figures out from what direction a puff of air is directed, and sets a `direction of turn' parameter appropriately. On the other hand, people can alter the signs of feedback loops involved for controlling their behavior - when a sign is wrong, you get a very distinctive runaway effect. In fact, Bill Powers and Rick Marken discover that people spontaneously reverse signs during even highly practiced tracking tasks (in one of the Mind Reading articles, I believe), indicating, I think, a constant low level of reorganization going on in the background.

In between human flexibility and cockroach rigidity, there is a broad range of variation. In one of the chapters of the Gallistel's `Behavior, a New Synthesis', there is some discussion of the apparent fact that birds that use their feet to manipulate food can be taught to operate food-dispensing machinery with their feet, birds that use their beaks for this purpose can be taught to operate food-dispensers with their beaks, but not vice-versa. It is said to be impossible to condition a pigeon to operate a food-dispenser by hitting a lever with its wings. This suggests that there is some hard-wiring of the reorganization system to the effect that an unsatisfied reference level for some physiological variable connected to food intake causes reorganization in a specific area of the motor-control system, varying with the species.

The reptile-amphibian divide is also important here: if you cross the nerves operating an amphibian's limbs, the central activation patterns won't change, and they will for example scuttle backwards from food instead of toward it (I recall this from an old scientific american article - there must be some more recent work on this). Lizards on the other hand can, I think I remember, reorganize their way out of this inappropriate structure, as of course can mammals.

So the basic idea is that unsatisfied higher-level reference signals cause lower-level systems to reorganize on an essentially random basis (although there may be fixed constraints on where the reorganization caused by a particular kind of error happens), & reorganization continues until the higher-level errors stop. Obviously, there is a huge amount of detail left to be filled in here, and I would very much like to see a demo in which a simulated creature learned to do something interesting by reorganization (in fact, I'd like to write one, but haven't worked up to the point where that would be a sensible project to take on).

Avery.Andrews@anu.edu.au

Date: Mon Jul 11, 1994 7:32 pm PST
Subject: Re: Testing models

<[Bill Leach 940711.21:54 EST(EDT)] >[Paul George 940711.15:00]

> Perhaps I should clarify. This relates to the thread on open loop control. The universe provides a nearly infinite variety of inputs, or sense-able things.The body has a limited set of sensors. Their signals must routed to a control structure that interprets them in terms of some hierarchy of control variables. Sensory input is distinguishable from perception. A HPCT network has some set of controlled variables at a given time. Similarly the nodes 'generate' a given set of possible actions to attempt to control those variables. Part of the universe of possible perceptions are thus 'open loop' at a given time, as they are not being controlled.

Speaking of precision in use of terms... Is the use of "universe" the same in both instances?

One of the problems in understanding PCT, I suppose, is that perceptions do not necessarily have to be under control. That is not what PCT says, only that behavior is the result of controlling perceptions.

In this area, I may be "stepping out" beyond my expertise but will risk it:

It appears that all biological control systems are "unidirectional". By that I mean that that reference signal can only run from zero to some magnitude. Thus, in any given control loop there is no such thing as reversing the sign of control.

From a physical standpoint, this is actually rather obvious. If you want to turn your head to the left and then to the right, you don't reduce the "turn to the left reference until it changes sign" but rather actually evoke a different control system to make the change (portions of the systems may be common but not the portion that actually effects the rotation). Even in a "simple" example such as this one, there are many perceptions actually being controlled to accomplish the "single goal."

Many control loops must have their references set to zero -- there is no higher control loop that requires their operation.

> Many of the results of our actions are not perceived, as are the actions themselves.

I am not sure what the significance of this statement is suppose to be.

> If just random chance was at work we would expect the control solutions to vary radically from individual to individual. OTOH we observe that people and other organisms seem to share control variables and reference values.

This is again a reason to call for reading the literature. I will remark in one way though... have you studied the behavior of an infant much?

-bill

Date: Mon Jul 11, 1994 7:46 pm PST
Subject: Re: PCT models

<[Bill Leach 940711.21:27 EST] >[Paul George 940711 11:30 EDT]

> I think you are on the right track, if occasionally locked into the worm's eye view}

I think that I have addressed some of the source of this "view of PCT". As has been suggested, you really have to read some of the PCT literature to see just how much thought has gone into the "challenges" that you raise. In many cases, suggestions that have been made for changes to the PCT paradigm can be shown to be impossible requirement for living systems (using real data acquired from various sources).

> I still think the analogy is valid. A top down approach takes the big picture and attempts to decompose it 'logically' (variously defined) into components. The bottom up attempts to create structures from existing components in order to meet goals. The two approaches can produce very different architectures. However, we have found that a better technique is to work from both directions meeting a third effort expanding from the middle (in reality you jump around).

At issue here in this minor disagreement is that while both methods work, they are neither inconsistent with reality (when they both work). If when employing either method, one is using an incorrect fundamental concept set, the effort will fail.

The idea that behavior is ONLY the result of controlling for perceptions is rather unique view in behavioral science (yes disturbances affect behavior but ONLY because the disturbances cause changes in the perception).

HPCT claims (rather convincingly in my mind) that the paradigm of negative feedback control works equally well for the thought processes that I am experiencing while I write this to you. There are one hell of a lot of man years of thought and experiment involved in coming to that conclusion as well as many well thought out challenges to the idea.

-bill

Date: Mon Jul 11, 1994 9:02 pm PST
Subject: Re: The Nature of Control

<[Bill Leach 940711.21:42 EST(EDT)] >[Paul George 940711 12:00]

> I do have some questions about _how_ it scales up, and the mechanisms used to assemble and modify the controlled variables and control hierarchy. These are not defects, just areas for future work. I also tend to question phrasing that appears to 'fluff the wares'.

While PCT might well be the "richest" field as far as new work is concerned, I think you will be more than mildly surprised as how much of what you have been talking about is anything but new to PCT.

> Don't presume misunderstanding.

I don't believe that anyone has presumed misunderstanding. The responses to you postings have been quite genuine and honest.

To give you a little example of just how a PCTer might have to "beat someone over the head" to get them to understand the real issue involved in a topic, just go back a month of two in the CSG-L archives and take a look at my exchanges on "Society" with the net. Yes, sometimes I was undoubtedly misunderstood but I actually always was missing a point (it is just that it often was not the point that others thought that I misunderstood -- it was something else).

These guys tend toward extremes of precision, particularly with fundamental PCT/HPCT concepts and especially with a "new comer". I used to think myself that they got a bit "carried away" but now have come to appreciate that it is precisely this dogged insistence on getting down to specifics and locking onto exact meanings for terms that makes it possible for PCT to deal with the complexity of behavior in a consistent, reliable, reproducible manner.

Speaking for myself, it really can be a bit exasperating for someone new to experience but worth the trouble.

-bill

Date: Mon Jul 11, 1994 9:34 pm PST
Subject: Replies to Paul

[From Rick Marken (940711.2200)] Paul George (940711.?)

> I am not a Skinnerian. But I do think some of the ideas in B&D and Walden II were useful.

Not surprising; Skinner was heavy into "using" behavior to achieve his ends -- ie. controlling behavior. His ideas may be "useful" but PCT shows that they are also a recipe for intra- and inter-person conflict.

> There is wheat among the chaff.

Could you give just one example of Skinnerian wheat?

> Never suggested that any behavior (above the reflex level) didn't involve control.

Ok. But you'll have to throw the reflex level in, too, because reflexes are control processes themselves.

> Can you empirically show each and every controlled variable, or just that they can exist?

In principle, all controlled variables can be empirically shown to exist using The Test for the Controlled Variable. In practice, we just try to identify the variables controlled in examples of behavior that happen to be of interest at the moment -- like pointing at a target or navigating up a chemical gradient. The point of PCT is not to discover every variable controlled by every organism. The point is to show that organisms do control variables. Once you see that a variable, like reinforcement rate, is under control, your concept of "what is going on" with a particular behavior, like operant conditioning, changes completely. In operant conditioning, for example, you realize that you don't need a model of reinforcement because there is no such thing as reinforcement (events that strengthen responses); you need a model of what is actually occurring in operant conditioning -- control (by the organism, not by the environment).

> Reactive behavior can be _modeled_ using control, but it does not necessarily follow that the map is the territory.

"Reactive behavior" is simply an aspect of the phenomenon of control; when control is occurring, there is simply no way to model the reactive behavior associated with it except with a control model. If the reactive behavior is not associated with control (if, for example, it is like the "falling" behavior that is a reaction to gravity) then a non-control model (such as Newton's laws) will work fine.

Me:

> When someone takes a swing at your face (d) you usually duck (o) in order to maintain the amount of pain you experience (q) at a reference level of zero (q*).

You:

> This is the level where your model fails to convince me.

I wasn't trying to convince you about the model of control; I was trying to convince you of the existence of the phenomenon of control as it is manifested in the behavior of a living system.

> There may be a gestalt phenomena, but the concept of ducking to keep a pain variable at 0 is a bit of a stretch.

Why? What is stretched? What does "gestalt" have to do with anything? If, every time you swing (always from a different direction with a different speed) , I move my head appropriately so as to avoid the swing (and feel no pain), I am controlling pain; disturbances (swings) do not have their expected effect (producing pain); their effect is systematically resisted; this is control; where's the stretch? A perceptual variable (pain) is clearly being controlled; the output (head movement) is always exactly what is required to maintain a variable (pain) in a particular state. Control is indicated by the fact that the output is always exactly what is needed to compensate for the effect of disturbance, keeping a consequence of head movement (degree of pain) in a reference state (zero). This phenomenon (controlling pain -- keeping it at zero) could not be produced by a device that produces outputs in response to disturbance inputs; it can only be produced by a closed loop control system that is controlling a perceptual variable -- pain in this case.

> I don't disagree that the behavior involves control, I just have trouble with the described mechanism. The model requires considerable elaboration.

If you agree that the behavior [moving the head in reaction to a swing] involves control then I suggest we put your claim to the test; let's see the elaborations that you think are needed to make the model work.

> Any way to get an electronic copy [of "The blind men and the elephant"]?

It might be on the biome server. I actually don't know. It's one of the papers that I was never able to get published in a "real" psychology journal. I guess they were less interested in "reconciling" than I was.

> Sorry, I was refering to the likes of dog training, not the controlled experiments used to document the principle.

Same thing applies; the dog is controlling some variable-- attention, food, petting -- or trying to. During training, the dog is randomly varying it's outputs until the value of the controlled variable is close to its reference. If rover has to lift his paw when you say "shake" or bark when you say "speak" in order to get his controlled variables where he wants them, then that's what he does.

Best, Rick

Date: Tue Jul 12, 1994 8:29 am PST
Subject: Control

[From Rick Marken (940712.0830)]

Paul George (940711 12:00) re my "Nature of control" post

> A number of you seem to be working under the misapprehension that I don't understand control from an engineering standpoint and that I disagree with PCT.

Many of the things you said suggested a lack of understanding of the nature of control as it is manifested in the behavior of living systems. Knowledge of control engineering is no guarantee (we have found, to our great dismay) that one understands the controlling done by living systems . In fact, the controlling done by living systems is typically very hard to see. For example, disturbances and the compensatory responses to them are often invisible; behaviors (like lifting a book) just seem to happen. You rarely notice the different forces (actions) that are and MUST be used each time to produce consistent results. In fact, the only person who actually noticed this "simple" fact about behavior is W. T. Powers. That little observation --that what we call "behavior" consists of controlled results of action-- combined with the realization that what must be controlled is a perceptual representation of controlled results, is the basis for a monumental revolution in thinking about, studying and dealing with living systems, one for which most life scientists (and control engineers, for that matter) are clearly not ready. This may seem like an extreme claim but it helps me understand why there are only about five people in the world doing PCT research.

Best Rick

Date: Tue Jul 12, 1994 9:29 am PST
Subject: Re: Testing models

[Paul George 940712 10:30] <[Bill Leach 940711.21:54 EST(EDT)]

> Speaking of precision in use of terms... Is the use of "universe" the same in both instances?

In the mathematical sense (i.e. set theory) yes. The first 'universe' is the objective world, if any. The second is the set of possible perceptions based upon available sensory inputs.

> One of the problems in understanding PCT, I suppose, is that perceptions do not necessarily have to be under control.

No problem. It is an obvious conclusion, if not the essence. We are not controlling the universe or behavior, just _a set_ of perceptions vs _a set_ of reference values.

>> Many of the results of our actions are not perceived, as are the actions themselves.

> I am not sure what the significance of this statement is suppose to be.

Just a re-iteration of the PCT view that behavior is not directly perceived, not is the effect of the behavior on the universe. We perceive nothing that we are not controlling for.

> This is again a reason to call for reading the literature. I will remark in one way though... have you studied the behavior of an infant much?

Yup. Years of Kittens, a three year old, and one on the way. As Avery pointed out in his response, some things are hard wired. Some things are assembled. There are apparently rules or heuristics that guide development. Exploration is not precisely a drunkards walk search. Tuning of control systems has random elements, put still appears goal directed.

Date: Tue Jul 12, 1994 9:45 am PST
Subject: Re: Replies to Paul and Jeff

[Paul George 940712 10:30] >[Rick Marken (940711.2200)]

> Could you give just one example of Skinnerian wheat?

How about what you cited. Much of his work suggested that people could be guided away from conflicting behavior (as externally perceived) by setting up reinforcement structure that inclined their behavioral patterns towards 'better' patterns. Behavior (i.e. my perceptions of what you do) may be modified by altering the inputs to you. This causes your perceptions to be disturbed and perhaps adjust themselves so that my inputs resulting from your outputs are closer to my reference values. There is of course a kind of Heisenburg effect involved. Skinner didn't put it this way, but the concept is there. If you wish its 'truth' can be demonstrated by the fact that we can suggest a control model for the observed results. In my experience child raising works, and I am certainly trying to modify my son's behavior. Truly trying to control another in all things is very difficult and will cause conflicts, unless I can understand your control structure and control your inputs (e.g. brainwashing or indoctrination).

> In principle, all controlled variables can be empirically shown to exist using The Test for the Controlled Variable.

Eh, I'll hold off major comment until I can get hold of the literature. However, it is not clear to me that because I can model the system using a control variable (or network thereof) that it follows that it is the only mechanism that can work. {note that I am impressed with the resolve shown to use only this simple negative feedback mechanism to model all behavior; may you succeed} This appears similar to saying that since I can model control using a digital system that all life must be digital. Different functions can produce the same behavior. They are equivalent, not equal.

> What does "gestalt" have to do with anything?

A set of control nodes containing controlled variables produce the observed behavior, i.e it is the behavior of the system, not the parts. There _need_ not be a master variable. Distributed control systems do not necessarily have to have 'master' nodes.

> If you agree that the behavior [moving the head in reaction to a swing] involves control then I suggest we put your claim to the test; let's see the elaborations that you think are needed to make the model work.

The example you use is familiar, being a martial artist. At a high level there is an observed 'desire' to avoid harm. This doubtless involves avoiding pain. To avoid getting hit I can duck, sidestep, parry, or strike first interrupting the punch. Some network of control determines that a punch is coming, is likely to hit, and might cause harm. Some network of control 'selects' the appropriate 'pattern' of behavior, probably by activating other control loops 'programmed by long training. I am likely controlling many high level variables such as status (getting knocked down is losing face), form (pride in technique), principle of minimum force, etc. If I was just controlling for pain I would duck when anything hurt.

I hope you see that what I mean by elaboration of the model is an elaboration of the set of controlled variables, reference values, and the control network; not the model of perceptual control. Note that different sets could produce the same behavior. Much of martial arts training is simplifying or tuning the control structure to speed response and effectiveness. The simplest structure is not usually the one that forms 'naturally' or initially.

I don't argue that the phenomenon of control, particularly "produced by a closed loop control system that is controlling a perceptual variable " is most applicable. At worst I would change "a perceptual variable" to "a set of perceptual variables".

Date: Tue Jul 12, 1994 11:17 am PST
Subject: Testing models

[John Anderson (940712.1330)] > [Paul George 940712 10:30]

>> One of the problems in understanding PCT, I suppose, is that perceptions do not necessarily have to be under control.

> No problem. It is an obvious conclusion, if not the essence. We are not controlling the universe or behavior, just _a set_ of perceptions vs _a set_ of reference values.

>>> Many of the results of our actions are not perceived, as are the actions themselves. I am not sure what the significance of this statement is suppose to be.

> Just a re-iteration of the PCT view that behavior is not directly perceived, not is the effect of the behavior on the universe. We perceive nothing that we are not controlling for.

Paul, this seems contradictory. First, you appear to say that you agree with Bill Leach's statement that perceptions do not have to be under control, and then you turn around and say "We perceive nothing that we are not controlling for". Or do I misunderstand?

John E. Anderson Beckman Neuroscience Center

Date: Tue Jul 12, 1994 11:18 am PST
Subject: Re: top-down, bottom-up, etc.

[From Bill Powers (940712.0745 MDT)] Bill Leach (940711.2154)

> One of the problems in understanding PCT, I suppose, is that perceptions do not necessarily have to be under control. That is not what PCT says, only that behavior is the result of controlling perceptions.

Still not quite the way to put it. Behavior is the _means_ of controlling perceptions. As you say, not all of a particular organism's perceptions are systematically affected by that organism's behavior, so not all perceptions are under control at a given moment. But all behavior is produced by that organism as a means of controlling _some_ perception. Action is generated only in order to control some perception.

Actions have an array of effects on the world. To an outside observer it is not always obvious what the controlled effect is. To find it when in doubt, you apply the Test for the controlled variable, which will distinguish unintended effects of actions from intended ones.

Nice post.

-----------

Paul George (940711.1130)--

> Dry humor seems to fall on deaf eyes. I guess I have to make greater use of smilies, and be very careful about phrasing and images.

Dry humor and condescending amusement are not the same thing. Rather than choosing more careful phrasing and images, it might be more conducive to avoiding flames if you spoke whereof you know rather than whereof you guess about PCT.

> We are designing a control system using some of Albus' concepts, but that isn't the kind of implementation you are interested in.

I think you will find that we are very interested in any actual implementation of any model of behavior. One of our complaints about critics from other fields is that they dispute the PCT model without offering either any criticisms of it or any substitutes for it. If you have an alternative model, by all means lay it out to us.

RE: Skinner

> He felt that OC was a useful technique for training humans to get more 'civilized' behavior. And he thought that treating people as anything more or less than animals was at best hubris.

The most interesting aspect of B. F. Skinner as a theorist is that he recommended using one theory but lived by another. He had a goal of helping people toward more civilized behavior, a goal which existed inside himself. He tried to act in a way that would move his perceptions of how people do behave closer to the way he felt they should behave. He felt that hubris was not a desirable character trait, so he adjusted his actions, and recommended that others adjust theirs, so as to change from being hubristic to being less so. He had a low reference level for hubris, obviously. In the laboratory, he recommended getting in mind a firm conception of the behavior that an experimenter wanted to see an animal performing, and then issuing rewards whenever the actual behavior was observed (i.e., perceived) to change a little in the direction toward the desired behavior. He emphasized that this procedure was not just a rote sequence of actions, but had to be adjusted on the basis of what the animal was doing at every moment, in relation to the desired behavior. This perceptual control process he called shaping.

So B. F. SKinner himself was a negative feedback control system of the type described under PCT, and recommended that other experimenters behave that way, too (as if they could do otherwise). However, the experimental animals and people to whom he applied his methods were NOT negative feedback control systems. Instead, they were systems that blindly emitted behavior into the environment under the control of discriminating stimuli and the consequences of that behavior, which he labelled reinforcers if they were successful in maintaining the behavior. Skinner would vehemently have denied that he had goals "in mind", of course, saying "the environment made me say that."

RE: top-down, bottom-up

> A top down approach takes the big picture and attempts to decompose it 'logically' (variously defined) into components. The bottom up attempts to create structures from existing components in order to meet goals.

The top-down approach is a snare and a delusion. It sounds fine when put into words, because we have lots of abracadabra words to get us past the tight spots. "Decomposing" the big picture into components is actually impossible if done strictly from the top down. What you actually do is set up a trial set of components, see if they add up to the big picture you have in mind, and alter the components, if they don't add up, until the correct big picture is perceived. So all the judgements are being made in the bottom-up direction, not the top-down direction.

"Decompose" is one of those dormitive-principle words. It means whatever process will create a set of components that, viewed together, add up to a valid instance of the thing to be decomposed.

Suppose the Big Picture is "Wash the car." This task, according to top-down philosophers, would be analyzed into subtasks such as "get car out of garage," "find hose, bucket, soap, and sponge", and "apply washing method to car."

The problem is that this "decomposition" assumes that you already know what the parts will be. There is nothing in the Big Picture of washing the car that tells you what you need to do in order to achieve that result. The way you decompose the big picture depends on what you already perceive about the actual situation at the time the main task is to be carried out. If the car is in the street, you can't get it out of the garage. If the sponge can't be located you'll have to find a rag. And if circumstances are just right, you won't have to apply the washing method; your son might do it for you, or it might rain.

There are many different sets of tasks that would add up to the Big Picture of washing the car. You can't get from a statement of the top-level task to a description of a particular unique set of methods that happens to be workable at the time you decide to do the task. All you can really do is set the perceptual goal of seeing the top-level task being done, and if you don't already perceive it being done, send signals to more detailed systems which will try to supply the perceptions needed to match the top goal with an actual perception.

As the lower systems work on producing candidate possibilities, the top level is continually putting the perceived elements together into some Big Picture, comparing that result with the desired Big Picture, and saying "Keep going, keep going, keep going, STOP, that will do."

It's the perceptual side, not the output side, that determines whether a given set of subtasks adds up to accomplishment of the higher task. "Decomposition" as a top-down process is a myth.

> Asserting that only one theory is valid is dangerous, though occasionally right. I often find the parable of the blind men and the elephant enlightening.

Ah, then I take it you have obtained Rick Marken's book, "Mind readings," and have seen his article called "The blind men and the elephant," in which he shows how S-R theory and top-down cognitive theory capture different but limited aspects of the organism described by PCT.

Best to all, Bill P.

Date: Tue Jul 12, 1994 11:44 am PST
Subject: Comments from Mary

[from Mary Powers 940712]

In reply to Paul George "RE: Replies to Paul George"

Rick said: "can you think of ANY behavior that does not involve control?"

You said: "I agree" - and then excepted reflexes and reactive behavior. And went on to ask : "How do you know this for a fact?" And continued, having disagreed with the model: "note that I do not disagree that behavior can be based on perceptual control, just that it must be. That has not been proven".

It's a little difficult to know where you stand ;-)

Where PCT stands is here: we have a model of the organization of living systems. We have applied this model, or hypothesis if you will, as a means of analyzing a variety of facts scattered through the literature of human and animal behavior (bacterial chemotaxis, enzyme systems, neural physiology, reflexes, operant conditioning, tracking tasks, movement of people in crowds, and so on. The same model (which is far more elaborated and sophisticated than you have so far troubled yourself to find out) applies in all cases. Using the PCT model, various of these phenomena can also be simulated with high accuracy. Other models, such as S-R and planned-output, cannot be simulated - or if they are, do not produce the claimed results. In many cases, there are no models at all, just great steaming piles of facts.

The purpose of a model is to organize facts. Proving, or testing, a model, of course is necessary. There does come a point, however, where one can legitimately project one's model into untested territory with a high expectation that it will continue to hold. Aren't you pretty confident that if you let go of the next thing you hold in your hand, it will fall to the floor?

This is not to say that it's time for PCT to quit testing - huge amounts of experiments need to be done. And they will be done, poking and probing at gaps and flaws. But, if you don't mind, PCTers will conduct those tests without concluding in advance that the model doesn't apply. Simply saying off the top of your head that it does not is far too superficial. It's like refusing to look through a telescope because (gasp) you might see the moons of Jupiter.

While no one that we know of has trained dogs _consciously_ using PCT principles, that would certainly be a situation "in the wild" where the principles apply - however you go about training a dog. You exploit the dog's controlled variables, either giving him things he likes when he's good or doing things to him he hates when he's bad. Dogs have reference levels for things like milk bones (high) or being jerked on a choke chain (low), and will do what they must to maintain those levels. There's nothing new about the way dogs react to such things (part of the steaming pile of facts) - PCT models why they do, and how. There are lots of other theories as to why - PCT is, so far, the only plausible model of how - of the organization required.

More in the wild: there are a number of school teachers, principals, psychologists, and parents who are pretty happy with using PCT as an organizing principle to guide what they do and better understand the consequences, good and bad, of one or another way of doing things. They feel it beats cookbook education courses a mile (if this happens, do that). One of the main things they get out of it is the value of teaching children that they are control systems, and so are their classmates, and their teacher, and that everyone's interactions have to take that into account. It isn't a secret recipe for success in the classroom; it's fundamental knowledge to be shared.

You have quibbled about us jumping from a few paltry controlled experimental situations to conclusions about real life. You have now complained that we don't apply PCT to real life. Aside from that not being true, you can't have it both ways. PCTers do experiments _and_ look at real life. The roots of PCT lie in really close and detailed examination of real life. Any one of us could show you four levels of your own hierarchical control organization (including a well-controlled reflex) in about two minutes.

You said: "I am focussing on design of control hierarchies to reflect actual behavior ... we should be able to define a HPCT system to cover any known behavior, whether or not the technology exists to actually implement it".

Yes. We should. We are in the process of doing that very thing. But you keep bringing up examples that have already been taken care of, like reflexes. Why should we have to explain stuff on the net that's available in the literature?

The journal that has "the blind men and the elephant" is Closed Loop, the journal of the CSG. The same issue also has Bourbon's "Models and their worlds". A bargain at $6. Send your check to me at 73 Ridge Place, Durango CO 81301-8136.

Mary Powers

Date: Tue Jul 12, 1994 12:32 pm PST
Subject: Re: Replies to Paul and Jeff

From Tom Bourbon [940712.1341] >[Paul 940712 10:30] >>[Rick (940711.2200)]

>> Could you give just one example of Skinnerian wheat?

> How about what you cited. Much of his work suggested that people could be guided away from conflicting behavior (as externally perceived) by setting up reinforcement structure that inclined their behavioral patterns towards 'better' patterns. Behavior (i.e. my perceptions of what you do) may be modified by altering the inputs to you. This causes your perceptions to be disturbed and perhaps adjust themselves so that my inputs resulting from your outputs are closer to my reference values. There is of course a kind of Hisenberg effect involved. Skinner didn't put it this way, but the concept is there.If you wish its 'truth' can be demonstrated by the fact that we can suggest a control model for the observed results. In my experience child raising works, and I am certainly trying to modify my son's behavior. Truly trying to control another in all things is very difficult and will cause conflicts, unless I can understand your control structure and control your inputs (e.g. brainwashing or indoctrination).

But what you have described, Paul, is an example of one control system (A) disturbing a variable controlled by another (B) in such a way that A sees B's actions matching a pattern that A wants to see. (I have numerous working models of pairs of PCT systems that interact in that way.) In that case, for A to "control" the actions of B, A must not prevent B from controlling the variable that B intends to control -- in a very real sense, the actions of A are "controlled" by the fact that B is still in control of B's chosen variable. By the way, this situation sets up a beautiful opportunity for B to "counter-control" A's actions, while at the same time, A controls B's actions. It is indeed true that Skinner described similar instances of control and counter-control, but he never understood why they worked as they did -- he thought it was a matter of stimuli controlling actions. He missed by a mile.

Incidentally, you haven't really lived until you've exchanged control and counter-control with a PCT model running as your partner, in "real human time" in a laptop computer. Come to one of our meetings and I'll give you an opportunity to do that. Or maybe I can send you a program and let you try it at home.

>> In principle, all controlled variables can be empirically shown to exist using The Test for the Controlled Variable.

> Eh, I'll hold off major comment until I can get hold of the literature. However, it is not clear to me that because I can model the system using a control variable (or network thereof) that it follows that it is the only mechanism that can work.

I think what Rick was talking about was what you had asked about -- the idea that there are empirical techniques for demonstrating whether _or not_ control exists in a particular situation. If it does, there is no other model than a control theoretic model that can explain the observed phenomenon of control. (Notice I said control theoretic in the general sense -- PCT provides a control theoretic model that is set up in a particular way to emphasize the properties of living systems, whereas most engineering versions of control theory are set up with a different emphasis.)

> {note that I am impressed with the resolve shown to use only this simple negative feedback mechanism to model all behavior; may you succeed}

At least we will give it a good try.

> This appears similar to saying that since I can model control using a digital system that all life must be digital. Different functions can produce the same behavior. They are equivalent, not equal.

But that is not at all what we are saying -- we begin by saying that _if_ control is found, _then_ only a control system can produce it. Can you suggest to us another way control might occur? No S-R system or plan-driven system can produce reliable control in a variable environment, or can they work after all?

>> What does "gestalt" have to do with anything?

> A set of control nodes containing controlled variables produce the observed behavior, i.e it is the behavior of the system, not the parts. There _need_ not be a master variable. Distributed control systems do not necessarily have to have 'master' nodes.

As Rick, what does this have to do with a "gestalt?"

>> If you agree that the behavior [moving the head in reaction to a swing] involves control then I suggest we put your claim to the test; let's see the elaborations that you think are needed to make the model work.

> The example you use is familiar, being a martial artist. At a high level there is an observed 'desire' to avoid harm. This doubtless involves avoiding pain. To avoid getting hit I can duck, sidestep, parry, or strike first interrupting the punch. Some network of control determines that a punch is coming, is likely to hit, and might cause harm. Some network of control 'selects' the appropriate 'pattern' of behavior, probably by activating other control loops 'programmed by long training. I am likely controlling many high level variables such as status (getting knocked down is losing face), form (pride in technique), principle of minimum force, etc. If I was just controlling for pain I would duck when anything hurt.

Paul, you are describing control as though it were a lineal process of:

input-->process ("determine")-->select appropriate preprogrammed action-->act

I believe you left off that last step -- act. The system you described would not control anything; it would simply act in a pre-programmed way that might be suitable for the average value of similar input conditions it had encountered in the past, but I'll guarantee you that such a system would never win the heavy-weight boxing title! I doubt that it's head would even be in place after the first punch thrown by a challenger.

Also, we have been trying to tell you from the very start that the hierarchical PCT model includes many -- very many -- control loops running simultaneously in a system that is hierarchical and richly parallel. What's with this idea you keep going back to -- the idea that there must be more than one reference signal at a time. Of course that is the case, but sometimes only one reference signal is needed in a working model. When more are needed, they are available.

Later, Tom

Date: Tue Jul 12, 1994 2:34 pm PST
Subject: Fact vs theory of control

[From Rick Marken (940712.1330)] Paul George (940712 10:30)

> Much of his [Skinner's] work suggested that people could be guided away from conflicting behavior (as externally perceived) by setting up reinforcement structure that inclined their behavioral patterns towards 'better' patterns.

And this is just one of Skinner's observations that PCT shows to be a crock. 'Better' behavioral patterns means behaviors that are closer to the reference specifications of the observer (like Skinner). The behavior patterns that meet the reference specifications of an observer are unlikely to produce perceptual results that are "better' for the organism itself.

> Behavior (i.e. my perceptions of what you do) may be modified by altering the inputs to you. This causes your perceptions to be disturbed and perhaps adjust themselves so that my inputs resulting from your outputs are closer to my reference values.

That's one very hopeful "perhaps".

Me:

>> In principle, all controlled variables can be empirically shown to exist using The Test for the Controlled Variable.

You:

> it is not clear to me that because I can model the system using a control variable (or network thereof) that it follows that it is the only mechanism that can work.

This is not the point of what I said . You are having a problem, I think, distinguishing the FACT of control from any THEORY that might be proposed to explain it. The Test for controlled variables establishes the FACT that a variable is under control. In our tracking experiments, for example, the fact that a person's actions precisely counteract disturbances to the position of the cursor demonstrates the FACT that the position of the cursor is a controlled variable. The Test can be used to establish that any other variable is under control as well; reinforcement rate, for example. The existence of controlled variables is a FACT; there is no theory involved.

We are not saying that PCT is a good theory because we can "model the system using a control variable". We are saying -- HEY LOOK, THERE ARE CONTROLLED VARIABLES ALL OVER THE PLACE AND LIVING SYSTEMS ARE RESPONSIBLE FOR THEIR EXISTENCE. This is a fact to be explained, just as it is a fact that objects accelerate to earth at 32ft/sec^2. The existence of controlled variables is a fact that was not observed until Powers pointed it out -- just as linear acceleration is a fact that was not observed until Galileo pointed it out.

Newton provided the theory that explained Galileo's (and Kepler's) findings; the theory that explains Powers' findings was already in existence (and Powers knew it); control theory.

> {note that I am impressed with the resolve shown to use only this simple negative feedback mechanism to model all behavior;

So now you see that this is not quite a correct statement; we use negative feedback control to model the FACT of control behavior. "Simple" feedback control is the only model we know of that explains the fact of controlled variables. If you (or anyone) knows of another model that explains the same fact (control) then, by all means, show us.

PCT doesn't have a problem of acceptance because people don't like the THEORY; PCT has a problem of acceptance because people don't know what FACT HPCT has been developed to explain -- the FACT of control.

Moreover, psychologists have no idea that CAUSING something and CONTROLLING it are two completely different things. I don't think there are many psychologists, for example, who would recognize that a statement like "reinforcement controls behavior" is demonstrably ridiculous. Reinforcements won't act to resist disturbances that move behavior from its reference state.

Me:

>> If you agree that the behavior [moving the head in reaction to a swing] involves control then I suggest we put your claim to the test; let's see the elaborations that you think are needed to make the model work.

You:

> I hope you see that what I mean by elaboration of the model is an elaboration of the set of controlled variables , reference values, and the control network; not the model of perceptual control.

Then I agree that these elaborations are necessary; they are already part of the HPCT model. I know that control of almost any variable really involves control of a hierarchy of controlled variables; control of reinforcement rate involves control of (perceptions of) muscle tensions, limb configurations, transitions between configurations, sequences of transitions, and many other variables that result in a particular rate of reinforcement. But, for simplicity, we can often lump all these variables together to produce a simple model of control of one variable -- rate of reinforcement. We do not imagine that a rat, for example, is just a furry thermostat with one sensor and one output device. If you have this impression of PCT, I think it would be quickly dispelled when you _read_ "Behavior: The control of perception". HPCT is one very rich model; "simple feedback system" doesn't really capture it.

> At worst I would change "a perceptual variable" to "a set of perceptual variables".

It's a deal! Even the simplest organisms control thousands (really) of perceptual variables simultaneously; and many of these perceptual variables are controlled as the means of controlling other perceptual variables. That's hierarchical perceptual control theory. It looks kinda like an Albus model on paper; the main differences are 1) HPCT is explicitly a model of control (the fact of the existence of controlled variables) and 2) the variables that are controlled in PCT are PERCEPTUAL variables.

Best Rick

Date: Tue Jul 12, 1994 4:04 pm PST
Subject: Re: Replies to Paul and Jeff

[Paul George 940712 15:40] >Tom Bourbon [940712.1341]

> But what you have described, Paul, is an example of one control system (A) disturbing a variable controlled by another (B) in such a way that A sees B's actions matching a pattern that A wants to see....

Deliberately. I attempted to describe in PCT terms (or a reasonable facsimile ;-) a useful idea from Skinner. Hoped I might demonstrate that I have some inkling of what you are talking about.

> It is indeed true that Skinner described similar instances of control and counter-control, but he never understood why they worked as they did -- he thought it was a matter of stimuli controlling actions. He missed by a mile.

That he didn't 'properly understand' the mechanism is to me of little importance - that is the chaff. I'm just not a purist or true believer by nature. And as I commented to Bill P. today, I think he meant that all you could observe was stimuli apparently producing actions, and that was sufficient for behavioral modification.

>> A set of control nodes containing controlled variables produce the observed behavior, i.e it is the behavior of the system, not the parts. There _need_ not be a master variable. Distributed control systems do not necessarily have to have 'master' nodes.

> As Rick, what does this have to do with a "gestalt?"

We seem to have very different definitions for the term. I mean something that does not appear until a 'critical mass' of components exist. Systems show group behavior, and it is sometimes a step function of complexity. (N-1) components won't do it, and there must be the right N components. The little man doesn't work until all three control functions exist, the sensory inputs exist for feedback, the output mechanisms exist, and all are properly interconnected.

> Paul, you are describing control as though it were a lineal process of:

> input--process ("determine")--select appropriate preprogrammed action--act

> I believe you left off that last step -- act.

Nope, while action and input may be lineal, control is usually continuous or at least periodic. (I vaguely seem to recall some medical research that indicated that sensory inputs such as vision had some kind of sense>process cycle in the brain with a detectable period). I considered the 'selection' the same (from the control system's point of view) as acting. The act is usually doing something to allow another series of nodes to 'take action'.

> What's with this idea you keep going back to -- the idea that there must be more than one reference signal at a time. Of course that is the case, but sometimes only one reference signal is needed in a working model. When more are needed, they are available.

It had been stated as the general model, but every time an example was given it was in terms of 'a controlled variable' as if each loop had only one, usually refering to something complex. When I suggested a more complex set, I was told it was not needed. This produced confusion. I am relieved that we were saying the same thing (?) and that my initial understanding was correct.

Hope all these replies aren't eating up too much bandwidth.

Paul

Date: Tue Jul 12, 1994 4:07 pm PST
Subject: Re: Fact vs theory of control

[Paul George 940712 15:50] >[Rick Marken (940712.1330)]

Most of my comment is in my previous post to Tom on the same subject.

> 'Better' behavioral patterns means behaviors that are closer to the reference specifications of the observer (like Skinner). The behavior patterns that meet the reference specifications of an observer are unlikely to produce perceptual results that are "better' for the organism itself.

If that were purely true neither society nor parenting would be possible. Since a fair amount of our perceptual variables are hardwired, there are commonalities, and there can be complementary goals.

> This is not the point of what I said . You are having a problem, I think, distinguishing the FACT of control from any THEORY that might be proposed to explain it. The Test for controlled variables establishes the FACT that a variable is under control..... We are saying -- HEY LOOK, THERE ARE CONTROLLED VARIABLES ALL OVER THE PLACE AND LIVING SYSTEMS ARE RESPONSIBLE FOR THEIR EXISTENCE. This is a fact to be explained....

I don't think I dispute the fact. Sometimes I question the theory or its modeling. The capitalized point is the basis of my interest in PCT, I too observed it a long time ago (but I don't publish, no credentials ;-)), though I think in terms of perceptual models (i.e a control network or set of variables)) The only question I have is whether postulating and modeling a particular variable(or set thereof) establishes _in all cases_ that it is in fact the correct one used by an organism apparently demonstrating the modeled behavior. You (Rick) seem to be a little more dogmatic on the question than most others. I have interpreted most of your examples of "the test" as 'proving' a particular control variable, not the existence of control. In your example the cursor position could be the variable, or it could affect the variable. The cursor could be perceived in a number of ways, and 'position' is not a simple concept. At any rate this is a bit of a fine distinction, and adds little insight. It is fine to look at position as if it were the variable. It works to describe the control behavior.

Paul

Date: Tue Jul 12, 1994 4:33 pm PST
Subject: Powers Responses.

[Paul George 940712 17:00] >[Bill Powers (940712.0745 MDT)]

> Skinner would vehemently have denied that he had goals "in mind", of course, saying "the environment made me say that."

Umm, It's your field not mine, and it has been a long time since I've read him. I interpreted him to be saying that we couldn't determine what was happening within the organism, and so should focus on what was without as we could observe it. And in a sense the environment does 'make us do it' as that is what we perceive and indirectly attempt to control. Proponents of a school sometimes make extreme pronouncements to 'stake out the territory'. Their other statements and actions sometimes reveal they really don't believe them.

> The problem is that this "decomposition" assumes that you already know what the parts will be. There is nothing in the Big Picture of washing the car that tells you what you need to do in order to achieve that result....The way you decompose the big picture depends on what you already perceive about the actual situation at the time the main task is to be carried out.

When decomposing behavior perhaps. When we do it we are trying to determine what we are talking about(general to specific) and what the components need to be. It is a 'divide and conquer' strategy. When we get to a set of concrete components or features we usually need to totally reorganize the hierarchy.

OTOH when I do process (a.k.a. activity) modeling, tasks have preconditions that are evaluated to determine if the task needs to be performed.Upon completion postconditions to affect the activation and processing of other tasks. The model tries to capture all potential tasks. Executing the model involves traversing appropriate paths. I have no idea how this might apply to PCT, though it could be a mechanism for dynamically reorganizing or activating components (heresy I know) of a control network.The evaluation of conditions could be viewed as an intermittent or reactive control loop.

> Ah, then I take it you have obtained Rick Marken's book, "Mind readings," and have seen his article called "The blind men and the elephant,"

Nope, the parable predates it somewhat :-), and applies to a lot of what I do. I do however look forward to reading it.

>[from Mary Powers 940712]

> It's a little difficult to know where you stand ;-)

In answer, at this point I don't stand I just raise questions and examples from my experience. I consider most of my points or observations to be minor, usually on how things have been phrased, or the exact contents of a given control model. Occasionally I question the logical necessity of a conclusion or extrapolation. I try to feed back my perceptions of what was said so I can evaluate my interpretation. When I get hold of the published literature and have the oportunity to study it in detail, then I'll take a stand. Y'all project a lot more certanty upon me than I have, and assume a lot stronger criticism than is there (apparently because of historical attacks with similar patterns). I use words like 'may', 'appear', 'seem', and 'not sure' quite carefully. As the old saw goes "I know you think you understand what you thought I said, but you may not realize that what you heard was not what I meant". Communication is a tricky thing, particularly without consistent sets of semantics.

> You have quibbled about us jumping from a few paltry controlled experimental situations to conclusions about real life. You have now complained that we don't apply PCT to real life. Aside from that not being true, you can't have it both ways.

I said that they were minor factors, at worst logical leaps. "Not Proven" is not the same as rejected. I have acknowledeged repeatedly that PCT is and has been applied to real life, and derived from observations about it. In my view that is the beauty of PCT, that it has both high level application and low level (demonstrable) mechanisms. My 'quibble' is a repetition of the frequent statements, including yours, that both ends haven't _yet_ met in the middle in terms of executable models. That simply is what I consider the set of interesting problems to be addressed.

> But you keep bringing up examples that have already been taken care of, like reflexes.

Based upon what I have seen (my fault not yours) you have not 'taken care of it', you have addressed it. You have built models that demonstrate that it is possible for control systems to demonstrate equivalent behavior to an organism. A true accomplishment. But it still doesn't follow that it is in fact the only possible model, or the mechanism in fact used by living systems at all levels. The fact that the other current behavioral models are even less validated (or not at all) doesn't to my mind change things. I may accept your hypothesis or conclusions without agreeing that they are proven or inescapable.

> Why should we have to explain stuff on the net that's available in the literature?

Because the literature is not readily available in public libraries or bookstores. As you are an unpopular school of thought, they are not even present in many university libraries (which in any event not everyone has access to). You don't (according to the FAQ) provide Closed Loop or other papers at the ftp site. Until I can arrange to borrow or afford to buy the material, I must rely on what Bill Powers, Marken, and Tom Bourbon say in their posts. And BTW they are very clear writers and likely successful educators. I have learned a lot about their thinking and technical aspects of PCT. Their posts are informative, if not always directed towards what I thought I was bringing up.

I am trying to get BCP via inter-library loan, and will order CL from you shortly (though the issue number might be helpful).

Thanks, Paul George

Date: Tue Jul 12, 1994 4:59 pm PST
Subject: Re: Testing models

[Paul George 940712 15:00] >[John Anderson (940712.1330)]

> Paul, this seems contradictory. First, you appear to say that you agree with Bill Leach's statement that perceptions do not have to be under control, and then you turn around and say "We perceive nothing that we are not controlling for". Or do I misunderstand?

May be a terminology problem. I think I agree with Bill. If I understand the PCT basics correctly, we do not perceive anything but controlled variables (though I suppose a noop control loop is conceivable). The nature of our perceptions change only as a function of changing variables or associated reference values (or possibly comparison algorithms). At a given point in time a given variable may be set within reference bounds or no sensory inputs have caused it to perturb. At such times no actions are being generated and so the control loop could be considered as not actively controlling the perception. {I think this matches [From Bill Powers (940712.0745 MDT)]} I can't say if Bill L would consider the perception as being controlled or not in this situation. If we allow a control loop to hibernate or in some way prevent error signals from being acted upon, then too we are not controlling the perception.

Another interpretation is possible is if you are allowed to distinguish 'perception' variables from 'controlled' variables. Either perceptual variables are processed in some way to 'compute' the values of controlled variables (e,g, a trend, upper/lower limit, etc.), or the latter are a subset of the former. In the first case the perceptions are controlled indirectly, and in the second not all perceptions are controlled at a given time. Can't say how formal PCT judges the situation. I'm not sure either elaboration is needed.

Paul

Date: Tue Jul 12, 1994 7:42 pm PST
Subject: Say what???

[From Rick Marken (940712.2030)] Paul George (940712 15:50)

> The only question I have is whether postulating and modeling a particular variable (or set thereof) establishes _in all cases_ that it is in fact the correct one used by an organism apparently demonstrating the modeled behavior.

What do you mean by "in all cases"? I have no idea what you could possibly mean. We establish by Test that an organism is controlling a particular variable in a particular situation; if The Test reveals that reinforcement rate is a controlled variable then it is -- IN THAT CASE. The next step is to build a model that can also control that variable. What in the world is your concept of a controlled variable??

> You (Rick) seem to be a little more dogmatic on the question than most others.

On what question?? Actually I think Bill and Tom can match me dogma for dogma. I'm just a little stricter (high gain control) than they are -- though Tom is making great strides ;-)

> I have interpreted most of your examples of "the test" as 'proving' a particular control variable, not the existence of control.

I would feel a lot more comfortable about your interpretations of my examples if you would call a controlled variable a "controlled variable" rather than a "control variable". The term "control variable" implies that a variable is doing the controlling. In control, one of several different variables, (by virtue of its position in a negative feedback loop) is _controlled_; there is no single variable that could be called the "control variable".

And how in the world could one demonstrate the existence of a controllED variable and not be demonstrating the existence of control?

> In your example the cursor position could be the variable, or it could affect the variable. The cursor could be perceived in a number of ways, and 'position' is not a simple concept.

Position may not be a simple concept in some way or other, but its pretty easy to measure. In compensatory tracking tasks it's just the horizontal position of the cursor on the screen, in pixels. Call this variable x. In our experiments x = h + d where all variables vary over time. Position, x, is at any instant the sum of a disturbance variable, d and the position of a "control" handle, h. If, over 2000 or so samples of each variable, the variance of x is virtually zero because h = -d (approximately) then the x variable is a controlled variable. It's true that the actual controlled variable (the perceptual representation of x) could be any monotonic function of x, but that would have to be determined by other test. There is no question that x is a VERY GOOD approximation to the actual controlled variable. The nearly perfect relationship between h and d is proof that x (or some monotonic function thereof) is under control -- it is a controlled variable.

Am I missing something?

Best Rick

Date: Tue Jul 12, 1994 7:53 pm PST
Subject: Re: Fact vs theory of control

[Avery Andrews 940713.1335] (Paul George 940712 15:50)

> The only question I have is whether postulating and modeling a particular variable(or set thereof) establishes _in all cases_ that it is in fact the correct one used by an organism apparently demonstrating the modeled behavior.

I'd say that it doesn't, but that it leaves you with a hypothesis that can be further challenged and investigated. In this respect I don't think PCT is different from any other empirical field of inquiry. We can be wrong about what variables an organism is controlling, and also about how apparent control is achieved. For example, perhaps one of the `transactionist AI/Alife' people (Beer, Brooks, Horswill, etc.) will come up with an architecture that produces what appears to be control by some means completely different from that proposed in PCT. Then we can look for evidence as to whether living systems employ this architecture or not.

Fowler and Turvey, somewhere (I think in a book edited by Stelmach in the late 70's) noted that respiration rate looked like it might be under feedback control, but argued that it wasn't. But an important point of their argument is that it doesn't really act like a control system, but like a dynamic system with a whole series of attractor basins, so that a sufficiently large disturbance will push it from one to another.

It might in fact be difficult to distinguish a multilevel control system from some random kind of dynamic system - since lower level reference signals are derived from higher-level error signals (plus perhaps perceptions), you might get attractor-like behavior as a system shuffled around between different lower level goals in an attempt to satisfy a higher level ones. But then PCT predicts that there will be a higher-level goal that you might succeed in identifying, whereas dynamic systems theory says nothing, and so is, I think, inherently dead-ended.

Another problem in evaluating human performance is that there can be several control systems operating simultaneously, at different time-scales. E.g. I can be controlling for getting my sailboat from Manchester Harbor to Singing Beach, but doing it at the moment by controlling for heading towards some island, but then I discover that it's the wrong island, so this reference level is replaced by some other one. Sorting out what's really going on is obviously not going to be easy, but I think that this is a difficulty inherent in the subject matter (living systems), rather than a nasty feature of the approach (PCT).

Avery

Date: Tue Jul 12, 1994 10:07 pm PST
Subject: Re: Fact vs theory of control

<[Bill Leach 940712.20:44 EST(EDT)] >[Paul George 940712 15:50]

Paul, I'm not sure if you picked up on the most significant thing that Rick was saying there:

Negative feedback control theory is the ONLY known explaination for behavior. That is the point and that is the reason for PCT.

You mentioned yourself (several times I believe) that "control theory could explain a certain behavior" and added that "you were not convinced that control theory could be the only explaination".

Turn that around a bit for a minute...

S-R (stimulus-response) actually does appear to explain some behavior (at least under some conditions).

PCT predicts that S-R will appear to work under the "right" conditions but where the two differ is that PCT works to explain the "anomalies" that appear in S-R testing instead for just dismissing them.

An additional thought is that some behavior quite clearly can only be explained by control theory unless one is willing to postulate a "theory" that requires different explanations for every change in behavior.

When one recognizes that living systems do indeed exhibit control system operation in testable situations and that control theory is the only theory that can explain the manner in which behavior is actually seen to vary as environmental conditions change AND that there is an ever mounting body of physical evidence to support that there are actual physical control structures present in living beings then it seems reasonable to me to contend that until proven otherwise, PCT is the proper approach to understanding.

Much of the difficulty that you are presently experiencing on the net will "go away" when you have had a chance to read some of the literature. This is especially true because you will see how most of what we tend to believe as "characteristics" of humans is rather learned traits.

Also, you will see where many things that humans do that are thought to be extremely complex could in fact be quite simple if performed by a control system with certain references set.

The literature will show that a great deal of consideration has gone into understanding "general human nature" already and that many of the conclusions are surprising.

-bill

Date: Tue Jul 12, 1994 10:16 pm PST
Subject: Re: top-down, bottom-up, etc.

<[Bill Leach 940712.20:24 EST] >[Bill Powers (940712.0745 MDT)]

Not to belabor a point but...

Is the statement; "Behavior results from the control of perception." imprecise? Is there really a reason for not stating it in such a manner?

Or for that matter, again what specifically is wrong with saying; "Behaviors is the result of controlling perceptions."?

I am not trying to be contrary here but rather I really do want to know the reason the above is not a desirable way of stating the case.

-bill

Date: Wed Jul 13, 1994 7:27 am PST
Subject: Re: Fact vs theory of control

[Dan Miller (940713)]

To Rick Marken:

Hi, again, Rick. In your tutorial with Paul George you made the following assertion:

> The existence of controlled variables is a FACT; there is no theory involved.

There you go again with FACTS. Do you mean to say that facts have an existence without theory, hypotheses, or epistemology? If this is true, then all we scientists have to do is reveal the facts of the universe.

I do not disagree with your perceptions of FACTS, but didn't those before Galileo and Kepler think that the Earth was the center of the universe. Weren't their observations FACTS? Now we understand that their model was invalid (it didn't fit all the observations). Of course their fate will not be ours. I would hesitate to say that we have "immaculate perceptions." However, an element of doubt might not hurt.

As for FACTS. If facts are perceptions (shared and reproducible), then they must be attached to theory (broadly considered). As conscious perceptions we see them (feel, etc.) and think of them in certain ways - temporally and often spatially separated from the situated perceptions themselves.

Later, Dan Miller

Date: Wed Jul 13, 1994 7:55 am PST
Subject: Re: Say what???

Paul George (940713 0930)

I think we have a common understanding of control and testing for it.

>[Rick Marken (940712.2030)]

> What do you mean by "in all cases"?

Always x itself. The variable I test must be the actual variable. Sometimes it is x, and sometimes a function of x, as you say below. Sometimes (though not in your examples) x could be a function of the true variable(s). We can but try to find a 'metric' that lets us determine something useful about the perceptual function.

> And how in the world could one demonstrate the existence of a controllED variable and not be demonstrating the existence of control?

You couldn't, and I didn't intend to ever suggest that you could or would want to.

> I would feel a lot more comfortable about your interpretations of my examples if you would call a controlled variable a "controlled variable" rather than a "control variable".

No problem, though that is the term process control system designers use. Ack on the semantic distinction. Much of the interaction of the last couple of weeks is reaching a common set of terminology.

> Am I missing something?

Don't think so. Don't think I am either, at least on your particular area of focus. I'll grant that 'strict' or 'rigorous' is a better characterization than dogmatic.

Please note that I don't usually comment on things I agree with in another's post in the interest of saving bandwidth. Silence is acknowledgement or agreement. My view of a forum is constructing an artifact in cyberspace is the product of all posts.

Later, Paul George.

Date: Wed Jul 13, 1994 7:55 am PST
Subject: Artificial

Paul George 940713 10:00 >[Avery Andrews 940713.1335]

> For example, perhaps one of the `transactionist AI/Alife' people (Beer, Brooks, Horswill, etc.) will come up with an architecture that produces what appears to be control by some means completely different from that proposed in PCT.

Has anyone in the CSG taken a close look recently on the work on Artificial Life, Core Wars, or Genetic Algorithm software?

I haven't looked at it in detail for a while, and never from the standpoint of formal control theory. However, these things do seem to demonstrate behavior similar to ecologies. Simple strategy has the effect of goal seeking, without it in fact being programmed. The observer or programmer intends to maximize survival, but the program does not explicitly have such a goal. Intent and effect need not be the same thing, and a successful survival strategy will tend to persist. This corresponds to an attempt to model in software ethology from the behavioral science domain.

Comments?

{This is not criticism, nor challenge, just a question}

Paul George

Date: Wed Jul 13, 1994 9:07 am PST
Subject: Paul George 940713 10:00

Thomas Baines ???

Reference the use of genetic algorithms in control, look at the stiles and Glickson article in IBM JOURNAL OF RESEARCH & DEVELOPMENT, Vol. 38, No. 2, p. 157.

Although they are dealing with a "Highly parallelizable rout planner based on cellular automata algorithms", the translation to search and control is not a biggie.

Date: Wed Jul 13, 1994 9:20 am PST
Subject: Controlling Skinner

[From Rick Marken (940713.0900)]

I said:

> 'Better' behavioral patterns means behaviors that are closer to the reference specifications of the observer (like Skinner). The behavior patterns that meet the reference specifications of an observer are unlikely to produce perceptual results that are "better' for the organism itself.

Paul George (940712 15:50) --

> If that were purely true neither society nor parenting would be possible.

It is purely true, and society and parenting are still possible (and actual) because neither is based on the control of other people's behavior. When people seriously apply Skinnerian methods to building societies or to parenting children the result is always disaster. People have always tried to use Skinnerian methods -- even before Skinner -- but they can't apply them for long because they almost always place the "applier" of these methods in a conflict - - with him or herself and/or with the controllee. My wife and I successfully (to our minds) parented two kids without ever resorting to Skinnerian nonsense (the down side is that I now I enjoy eating all those M&Ms that we didn't need to use as reinforcers;-))

PCT shows that it was natural for Skinner to have wanted to be in control; all people are controllers; controlling is a good thing. The problem is that you run into trouble when you try to control other controllers -- especially other controllers who can perceive and control the same variables that you can. Skinner just never figured out that other people are just like him -- they are controllers. Skinner was so wrapped up in his controlling that he never realized that what he was doing (controlling) was the interesting phenomenon; not the results of his controlling (like pigeons playing ping pong).

Best Rick

Date: Wed Jul 13, 1994 11:47 am PST
Subject: Facts, Theories & Illusions

[From Rick Marken (940713.1030)] Dan Miller (940713)

> There you go again with FACTS. Do you mean to say that facts have an existence without theory, hypotheses, or epistemology?

I think I'll have to say "yes".

> If this is true, then all we scientists have to do is reveal the facts of the universe.

I don't see why that's true. I think that one aspect of science is "revealing" or describing facts; another is trying to make sense of them -- to figure out why these facts exist. That's where theory comes in. I think.

> I do not disagree with your perceptions of FACTS, but didn't those before Galileo and Kepler think that the Earth was the center of the universe. Weren't their observations FACTS?

I don't think that they "observed" the earth to be in the center of the universe; they observed (as you still can) that the earth is stationary, that the sun, moon, stars and planets move across the sky at different rates, and so on. A universe with the earth at the center was their explanation (theory) of why they saw what they saw; it was an explanation of the facts. Their observations were (and still are) facts.

I don't think theory determines facts; it determines how we interpret the facts. For example, I personally believe that the sun is the center of the solar system, and that the earth rotates on an axis as it moves around the sun. But I still observe the same facts that the ancients observed; I perceive the earth as stationary with the sun, moon, stars and planets moving across the sky at different rates. I just think that the reason why I see these facts is different than the reason assumed by the ancients.

> Of course their fate will not be ours. I would hesitate to say that we have "immaculate perceptions." However, an element of doubt might not hurt.

Our model will ultimately be replaced -- as was theirs. But the facts will remain (as did theirs). Our perceptions (aided or unaided) are the same as those of the ancients -- it's our explanation of the cause of those perceptions -- our theory -- that will get replaced.

> As for FACTS. If facts are perceptions (shared and reproducible), then they must be attached to theory (broadly considered).

I guess I disagree with this. If facts are perceptions, then I don't see why they MUST be "attached" to a theory. In fact, the idea of theorizing to explain facts is rather new -- starting in the 1500s or so. Perceptions are just perceptions -- some can be controlled; some can't. It's a relatively weird breed of cat (a scientist cat) who finds it necessary to build models (theories) to explain the reality that is presumed to lie beyond what we perceive. Western science proves that such models can provide a powerful tool for controlling perceptions; but, as Paul George pointed out, we are able to control most perceptual variables without having any understanding of how we are able to do it.

By the way, PCT does not deny that certain facts exist; the statistical results observed in the behavioral sciences still count as observations (though they are often so unreliable that we don't like to call them "facts"). Still, the fact that we do observe statistical relationships (across subjects) between variables like the color of the room and the productivity of the workers in the room (to pick a weird "fact" off the top of my head) is not denied by PCT; PCT just says that the REASON for this fact (observation) differs from what has seemed like the obvious reason -- that there is a causal relationship between wall color and productivity.

I think theory influences how we interpret the facts; it also influences what facts we consider important and what facts we attend to; but I don't think theory influences what the facts are. At least, I don't think it should.

I agree that I am not always clear in how I deal with the relationship between fact and theory. For example, I often refer to the S-R relationships observed in the behavior of living organisms as an "illusion". This suggests that these relationships (in light of the PCT model) don't exist; that the theory (PCT) revokes what was thought of as a fact (S-R relationships). But the S-R relationships are facts; PCT doesn't change that. The "illusion" is not the S-R relationship itself -- that is a fact. The illusion is that there is a causal relationship between S and R mediated by the organism; the illusion is not that S-R is a fact; the illusion is the theory typically assumed to explain that fact. Saying that S-R is an illusion is like saying that the sun moving across the sky is an illusion. The perception of the sun moving across the sky is NOT an illusion; it is a fact. The illusion is the theory -- that it is the sun that is moving. The right theory (currently) is that the earth is rotating and that the sun just appears to move as we rotate past it. So maybe the word "illusion" is inappropriate when referring to incorrect theories of the facts. Mea culpa.

Best Rick

Date: Wed Jul 13, 1994 1:11 pm PST
Subject: 6=+3g~},w_/_/{w{{^{o

[From Bill Powers (940713.1230 MDT)] Paul George (940713)

Paul, I think some of our semantic difficulties need clearing up by way of a diagram. Below is our standard diagram of a control system with verbal definitions of the components and signals. Would you attach to this diagram the labels you are accustomed to using for the parts of a control system? It seems to me that a process-control engineer uses the term "control variable" to mean a variable that causes control actions, while we use "controlled variable" to mean a variable that is altered by actions and, via the closed loop, is controlled by them. There may be other problems; for example it is customary in control engineering to speak of the reference signal (in our diagrams) as the "input."

I will present this diagram on its side (as compared to the way we usually draw it), so it will look more like standard engineering diagrams. Higher levels would then be to the left, and systems operating in parallel at the same level would be stacked vertically. You might want to insert blank lines here so the following will appear on one page.

```
                          :
                          :
               B      ======
           =========>|| F2 ||---C----------D
          ||          ======               |
          ^^              :                v
          ||              :              =====
       ======             :             || F3 ||
A ===>|| F1 ||            :              ======
       ======             :                |                 E
         ^^               :                |                 |
         ||     H      ======              v       ======    |
          ====<======|| F4 ||<-----G------ F <----|| F5 ||<---
                       ======                      ======
                          :
                          :
```

The vertical dotted line separates the control system (to the left) from the physical (observable) environment outside it (to the right).

Let's start with the environment part, to the right of the vertical dotted line.

C = path by which action of the system (the output of F2) affects
D = an observable physical variable. F2 is an output transducer.

The state of D has an influence on
F = another physical variable, via
F3 = a function representing physical laws connecting D to F.

There is in general a physical variable

E = variable that independently contributes to the state of F via
F5 = a function representing physical laws connecting E to F. (There may be many independent variables like E, contributing to the state of F via different functions like F5).

Note that F = F3(D) + F5(E). Also, C and G have no special names in PCT; they are pathways.

Now we come to the control system.

At the input interface, there is a function F4 which is affected via
G = a pathway through which the variable F physically affects the system's sensors, and produces an output
H = a signal standing for the state of F.

F4 can contain not only sensors, but computations on the values of outputs from sensors.

The signal H enters the function
F1 = a function which subtracts H from
A = a signal that has a state equal to the desired state of H, and produces an output
B = a signal that is a measure of the discrepancy between A and H.

The signal B enters
F2 = an output transducer which converts the signal B into
D = the physical variable already mentioned, via
C = a pathway connecting the output of F2 to the variable D.

These are the functions, signals and variables that comprise the basic control system as seen in PCT. Engineering diagrams generally leave out one more of these elements, or lump several together, but they are always there in real closed-loop control systems.

A comment on the environment part of the model is needed. In fact, the variable D might be multiple physical variables, affected by many pathways C; there might be many functions F3 connecting the effects of the system to many variables F affected by many independent variables E. The variables F might affect many sensors through different paths at the input to F4.

The output of F4, however, is always a scalar variable, the value of the function F4(F[i]) for all i. In a real system there may be many functions F4, each producing a signal H that is a different function of the same set of sensor inputs. Each F4 would be part of a different control loop acting in parallel and sensing a different aspect of the collection of physical variables F. Each such loop would be drawn as a separate control system and placed vertically above or below the system shown above. A given F4 determines what function of the detailed set of environmental variables F is represented by H.

-----------------------------------

OK, now what we need from you is a set of labels to go with A..H, and F1..F5. Each signal (inside the system), variable (outside the system) and function has a specific name in PCT. These names are always used in the same way to mean the same part of a control system or its environment. If we can set up a translation table, it should be possible to translate a PCT description of control into a description in terms of any other engineering subculture. How about giving it a try?

Best, Bill P.

Date: Wed Jul 13, 1994 1:35 pm PST
Subject: Facts

[From Bill Powers (940713.1300 MDT)] Dan Miller (940713) --

Butting in:

> There you go again with FACTS. Do you mean to say that facts have an existence without theory, hypotheses, or epistemology? If this is true, then all we scientists have to do is reveal the facts of the universe.

The answer isn't a simple yes or no. There's a progression from simply recording the occurrence of an experience (brightness exists now) to reading all kinds of inferences into such experiences (somebody is shining a flashlight into my window).

In the old two-psychoanalyst joke (How are you --- what did he mean by that?), the joke has to do with an unwarranted model-based perception of implications that take a theory even to detect. But if the second psychoanalyst said "I heard him say 'How are you'" there would be no joke, and indeed no point. It is a fact, as we think of facts, that the first psychoanalyst said "How are you?". The inter-rater reliability of observers asked whether this sentence was spoken with an interrogative inflection (one of them) would be extremely high. Even a person who spoke only a different language might be able to identify that sentence among tape-recordings of different sentences and declare that recording 17 is the one that was spoken.

The relationships we observe in PCT which identify the existence of control are, as nearly as possible, incontrovertible observations. We don't often go through all the details because we're so familiar with them, but we could. I apply a force to an arm. I ask the person who owns the arm, "Am I applying a force to your arm?" The answer would be "Yes." I say "I feel resistance to my push; do you feel yourself resisting it?" The answer would be "yes." I then inject curare into the person's muscles, and try again. This time there is no resistance to my push. So I say that what I am observing meets the criteria for the phenomenon I call control, or at least some of the main ones.

The Test for the controlled variable, while it is used in PCT, is not theory-based in the sense of depending on control theory to work. The theories it depends upon are those of physics, and they are among the oldest and best-supported theories known to human beings. A force applied to an object will cause that object to accelerate. If one force is applied, but the object does not accelerate according to A = F/M, then there must be a second force being applied in the opposite direction so that the net force is zero. When we find the second force, we observe that it is equal in magnitude and opposite in direction to the first force. So we accept this aspect of control as having been demonstrated as a fact.

There is one kind of fact that is, as far as I can see, absolutely reliable. If I see a purple dog floating three feet off the floor, all kinds of disputable statements can be made about this experience, such as "there isn't (or is) really a dog there" or "the dog that is floating there isn't really purple." But there is one indisputable statement that I and I alone can make: I see a purple dog floating three feet off the floor.

That to me is a ground-zero bedrock Fact. It is a description of experience devoid of opinion about it. The moment you get away from that level of Fact -- the moment you say anything ABOUT that fact other than simply reporting its presence -- interpretation, opinion, and theory get into the act.

The aim of experimental science is to provide facts that are as near as possible to bare reports of experience. A true experimental scientist would never report that an animal presses a bar at x presses per minute and gets y reinforcements per minute. It does not get reinforcements; it gets little brown food pellets. This scientist would never say that 10 reinforcements per minute is enough to maintain 150 bar-presses per minute. The report would merely say that when the rate of delivery of little brown food pellets was 10 per minute, the rate of bar-pressing was 150 per minute. "Maintaining" is not a fact, but an interpretation.

Best, Bill P.

Date: Wed Jul 13, 1994 2:02 pm PST
Subject: Behavior and control

From Tom Bourbon [940713.1512]

>[Bill Leach 940712.20:24 EST] >>[Bill Powers (940712.0745 MDT)]

> Not to belabor a point but...

> Is the statement; "Behavior results from the control of perception." imprecise? Is there really a reason for not stating it in such a manner?

> Or for that matter, again what specifically is wrong with saying; "Behaviors is the result of controlling perceptions."?

Bill L., I think what Bill P. is getting at is that behavior is not something that (merely or just) "results from" the control of perception. If that were the case, then we would need to identify what it is, other than behavior, than controls perception and that simultaneously produces behavior as an unintended result. In that light, either of the alternative renderings you suggested says something different from Bill P's original statement -- Behavior: (is) the control of perception. Behavior is the means by which perception is controlled; behavior is not a result of something, other than behavior, by means of which perception is controlled.

Of course, we then go on to say that many of the outward appearances of specific _actions_ we see a person make at a particular time are probably unintended side effects of the person's control of perception. In that case, we are emphasizing the fact that, while behavior controls perception, the specific actions required to establish and maintain control must vary, moment-by-moment, any way necessary to eliminate the effects of disturbances that act on the person's controlled variable(s). When actions vary, so do many side effects of the actions -- many of the "outward appearances" that often catch the eye of an uninformed observer, who then arrives at a wrong conclusion about what the person is doing.

Specific actions are a result of any momentary changes in the person's intentions (modeled as reference signals) _and_ in the effects of disturbances that affect the person's controlled variable(s). Those actions are the means by which the person controls perceptions.

If this isn't sufficiently confusing, just let me know.   ;-)

Later,   Tom

Date:     Wed Jul 13, 1994  3:51 pm  PST
Subject:  Re: Diagram Terms o

[Paul George 940713 17:00]

{The original subject line was scrambled}

>[Bill Powers (940713.1230 MDT)]

>   Below is our standard diagram of a control system with verbal  definitions of the components and signals. Would you attach to  this diagram the labels you are accustomed to using for the  parts of a control system?

{"I reserve the right to revise and extend my remarks....."}

```
                          :
                          :
               B      ======
           =========>|| F2 ||---C----------D
          ||          ======               |
          ||              :                v
          ||              :              =====
       ======             :            || F3 ||
A ===>|| F1 ||            :              ======
       ======             :                |               E
         ^^               :                |               |
         ||     H     ======               v       ======  |
          ====<======|| F4 ||<-----G------ F <----|| F5 ||<---
                      ======                       ======
                          :
                          :
```

Note: Because of the way physical devices work, we can't (always) use just negative feedback control. Two electric motors opposing each other like muscles would burn out. Many devices like switches, pumps, and valves are discrete (on/off, forward/stop/reverse). And, as (I recall) Bill Leach observed, sensors and actuators are expensive and so kept to a minimum. We also usually just simulate continuous control because we are using electronics and digital computers which (often) operate in a discrete fashion. Biological systems have certain advantages.

Inside:

A = reference value or setpoint. Usually the desired state of H, but sometimes a limit - i.e. action should be taken when H exceeds the limit until H falls below. Sometimes A = f(H), and corresponds to an estimated state of E or F

B = Control output, output variable, or sometimes 'command'. Note that we usually keep it around and use it in computations (F1).

H = Control input, input variable, or input data (precise aren't we). Same apparent semantics as in PCT controlled variable. In our case H can be a scalar or vector array, but that is an artifact of design. We could construct things your way with one loop per variable.

F1 = Control logic or function Block. Usually not a simple, 2 variable difference. Usually a number of H values and B values and (sometimes) A values go through a set of algorithms to determine the desired value for B. We are usually forced to be pro-active due to latencies in the system and the process under control.

{This _could_ have implications for HPCT at levels where there is a considerable delay between an action and any perceived 'result'. Back in school we had some software in the Human Factors lab similar to the 'little man' where the user had to try to track and object with a cursor. The software inserted all kinds of linear and non-linear delays and vectors to the 'actions'. Made the task real difficult, and frustrated people to no end}

F2 = Output point. An output transducer. Performs D to A conversion in some cases. Basically translates B to whatever the actuator instrument requires. In some cases performs computations or sampling on B. Depends on the cycle time of F1 vs that of the instrument.

F4 = Input Point. An input transducer. Usually performs A to D conversion. Usually samples or averages F

Note: Input and output points are referred to collectively as 'I/O Points' within the control system. They are often grouped into related sets (for practical reasons) referred to as I/O Blocks, which are packaged as a device.

Environment:

We would view the outside structure a little differently, because we view the Process (environment), instruments (nerves, muscles, eyes), and Control system (brain) as being different things. In other words the system boundary is a little different. It doesn't really affect anything from the control system's point of view. I'll just re-diagram the input side

```
   :                         |
 ======         =======      v       ======
|| F4 ||<--G--|| F4.1 ||---- F <----|| F5 ||<---
 ======         =======              ======
   :
```

F4.1 = Sensor Instrument. Translates the sensed variable F into some kind of a reading or metric. A pressure sensor turns a pressure into a voltage. Sometimes combines some set of physical sensors into a reading or set of readings. Note that I may measure the temperature of a vessel and translate it into the temperature of the substance inside the vessel. Sometimes samples or averages a continuous reading.

I think this has the same semantics as your G, but we consider it a function, since it is a device. The channel is separate. Some advanced instruments are control systems in their own right. They may have both sensors and actuators, and may transmit an error signal.

The F4.1>F4 chain is analogous to the process of detecting light in separate rods and cones, converting it to an optical nerve signal, and then into an 'retinal image' (I'm fuzzy on the exact physical process). The image then needs to be processed to extract objects and motion which are 'perceived'. Some of this can map to control nodes, but having programmed such for radar systems it is far from simple and involves a lot of computation and prediction.

G = Input signal or measurement. We would usually have a G' which is the channel (wire), but the distinction is not important to the control system. As noted above, G in some cases is an error signal.

F - An attribute or characteristic of the process or of the equipment being controlled . A temperature, pressure, etc.

F3, F5 = As in yours, though we would think of it as some characteristic or part of the process system under control. We usually 'mirror' this function within the control system for various practical reasons.

C - Could be a channel as in yours, but again the 'output' signal or 'output data' is usually distinguished from the wire.

D - not real clear on its usage in PCT. The action of the actuator usually affects the process or equipment being controlled, frequently indirectly

F2.1 Actuator Instrument. Translates a command into an action. E.g trips a switch. Sometimes involves a sequence of actions which affect F. In a sense has some of the semantics of D

Hope this hasn't munged things up too badly and has met your intent. I think this view can be translated readily into HPCT terminology as we appear to simply be making slightly different groupings and distinctions. The only real difference seems to be the definition of the 'Boundary' of the control system and the 'Environment'. For simple PCT, environment appears to be everything outside of A, B, and H. For HPTC it is a point of view or (apparently) arbitrary scale {not a problem for any system modeler}. Process control just uses a kind of Mind/Body distinction that I haven't seen HPCT use explicitly. The distinction between a HPCT and multiple interacting PCSs or HPCSs is unclear, and probably unimportant.

The only likely difference I can see is that HPCT nodes can rout C of one node directly into G of another without intermediary processing. In the degenerate case B and H can be the same variable.

I am still thinking about how to directly use HPCT to modify process control system architectures, but suspect I don't understand it well enough yet. I think there is some real possibility of striking gold, but it is just an intuitive reaction. First theory, then application.

Later Paul

Date: Wed Jul 13, 1994 4:41 pm PST
Subject: Re: Behavior and control

[Paul George 940713 17:30] >Tom Bourbon [940713.1512]

> Behavior is the means by which perception is controlled; behavior is not a result of something, other than behavior, by means of which perception is controlled.

> Of course, we then go on to say that many of the outward appearances of specific _actions_ we see a person make at a particular time are probably unintended side effects of the person's control of perception.

Could you clarify the distinction between Behavior and Action? I think it may be the basis of much confusion.

I usually look at behavior as being a pattern or set of actions, usually within the context of some environmental pattern of events or inputs. Tracking is behavior, a given arm motion is action.

You also seem to be saying that behavior directly causes other behavior. This makes sense in the context of muscles (unobservable) moving an arm (observable) or a arm motion propelling a ball. Is that all you meant?

Date: Wed Jul 13, 1994 4:42 pm PST
Subject: Re: Replies to Paul

From Tom Bourbon [940713.1538] >[Paul 940712 15:40] >>Tom [940712.1341]

>> But what you have described, Paul, is an example of one control system (A) disturbing a variable controlled by another (B) in such a way that A sees B's actions matching a pattern that A wants to see....

> Deliberately. I attempted to describe in PCT terms (or a reasonable facsimile ;-) a useful idea from Skinner. Hoped I might demonstrate that I have some inkling of what you are talking about.

Good enough.

>> It is indeed true that Skinner described similar instances of control and counter-control, but he never understood why they worked as they did -- he thought it was a matter of stimuli controlling actions. He missed by a mile.

> That he didn't 'properly understand' the mechanism is to me of little importance - that is the chaff. I'm just not a purist or true believer by nature.

But what you call "chaff" is the entire social and scientific establishment called "radical behaviorism." The quality of Skinner's "understanding" is not something I dismiss as readily as you. He touted that understanding as _the_ very science of behavior, not as a personal understanding. And from that mistaken understanding (that behavior must be described as originating from environmental stimuli), he drew up his guidelines for research, and for re-engineering society. His "chaff" is still used to justify operant techniques and interpretations in many quarters of society. (Including the courts in the United States, where the idea that "the environment made me do it" is currently very popular. However, in a demonstration of true eclecticism, the major players in the same legal system appear to be enamored of the idea that individuals are "victims" of their own genes and brains. In our courts, you can have it either way, or both ways at once -- either way, it's pure lineal thinking -- all cause-->effect.)

> And as I commented to Bill P today, I think he meant that all you could observe was stimuli apparently producing actions, and that was sufficient for behavioral modification.

I think you are right; I think Skinner _did_ mean that all _he_ could observe were stimuli apparently producing behavior. He went so far as to say (in so many words) that behaviorism is the science of that which can be observed. But all that proves is that Skinner wasn't a very astute observer of behavior. Control _by_ organisms can be observed; it is an observable, not a conjecture. Skinner didn't observe it; therefore Skinner did not develop the science of that which can be observed. Instead, he developed a technology, which he called a science, in which he treated all behavior as though it were a result of environmental stimulation.

>>> A set of control nodes containing controlled variables produce the observed behavior, i.e it is the behavior of the system, not the parts. There _need_ not be a master variable. Distributed control systems do not necessarily have to have 'master' nodes.

>> As Rick (asked), what does this have to do with a "gestalt?"

> We seem to have very different definitions for the term. I mean something that does not appear until a 'critical mass' of components exist. Systems show group behavior, and it is sometimes a step function of complexity. (N-1) components won't do it, and there must be the right N components. The little man doesn't work until all three control functions exist, the sensory inputs exist for feedback, the output mechanisms exist, and all are properly interconnected.

I never thought of a control system as something that functions because the number of its parts is at least one greater than some threshold number of parts. I think of a pile of parts as a pile of parts, not as a sub-threshold system. And I always thought that systems "show system behavior," with each system sort of "doing what comes naturally" for its particular organization. On that reading, a PCT system with missing parts isn't a PCT system with critical-1 parts; it's simply not a PST system -- it is something else.

I'm still left wondering what this might have to do with "gestalt."

In these exchanges, we haven't even begun to explore the differences between a hierarchical PCT model and the kinds of distributed control systems to which you allude. In hierarchical PCT models, there are no distributed nodes, each with its controlled variable(s). We seem to be talking about different kinds of systems.

>> Paul, you are describing control as though it were a lineal process of:

>> input--process ("determine")--select appropriate preprogrammed action--act

>> I believe you left off that last step -- act.

> Nope, while action and input may be lineal, control is usually continuous or at least periodic.

But you didn't describe it that way. You described a lineal process that includes the production of pre-programmed outputs, selected to match present-time inputs. Or did I misread you then, and in this reply? It has been our experience that lineal systems operating this way cannot control variables in disturbance-filled world.

> . . I considered the 'selection' the same (from the control system's point of view) as acting. The act is usually doing something to allow another series of nodes to 'take action'.

In a reply to Bill Powers [Paul George 940712 17:00], you said more on this topic:

> OTOH when I do process (a.k.a. activity) modeling, tasks have preconditions that are evaluated to determine if the task needs to be performed.Upon completion postconditions to affect the activation and processing of other tasks. The model tries to capture all potential tasks. Executing the model involves traversing appropriate paths.

I think these two replies (to Bill and me) are the clearest indications yet that you are talking about a different kind of system from us. A system, incidentally, that is very similar to the ones described by Albus -- you must think of him as more than just someone to use as raw meat to cast before the PCT modelers. ;-)) You seem to envision a "system" as a kind of "meta-assemblage" of independent systems, each of which accomplishes an assigned task, then passes off the result to the next system(s) in the chain (or net) which uses the received result as fodder for its own processes, and so on. Then, in the language of Albus, the "behavior" of the entire assemblage "traces a trajectory" in space and time. I don't think that kind of system has much in common with a hierarchical PCT model. Let us know what you think, after you have a chance to read some of our work.

By the way, we are acutely aware of the problem you mentioned -- part of our work is published in our own somewhat inaccessible ghetto press -- but a might fine ghetto it is!

>> What's with this idea you keep going back to -- the idea that there must be more than one reference signal at a time. Of course that is the case, but sometimes only one reference signal is needed in a working model. When more are needed, they are available.

> It had been stated as the general model, but every time an example was given it was in terms of 'a controlled variable' as if each loop had only one, usually refering to something complex. When I suggested a more complex set, I was told it was not needed. This produced confusion. I am relieved that we were saying the same thing (?) and that my initial understanding was correct.

I think Rick and Bill P. have had their shots at this topic with you, but let me try.

At the risk of losing readers whose interests are limited to the seemingly "big" topics, I'll use the example of a person running a compensatory tracking task. (No surprise there. I always hope that people interested in the big questions will come to realize this is not a trivially simple example of human behavior.) The person watches a computer screen and uses a control handle (h) to keep a cursor (c) aligned with a stationary target (t). A random disturbance (d) also affects the position of the cursor.

Paul, I know you will probably think this is a simplistic example, but I hope you will bear with me; I am concerned over the fact that you and we seem to be talking about different kinds of models and that you express confusion at our bouncing you back and forth between talk about simple PCT models and complex ones. I hope to help clear up some of the problems.

The momentary position of the cursor is determined by two present-time variables: the position of the handle and the magnitude of the disturbance.

$$c = h + d$$

That equation completely describes the environmental variables in the tracking task. It is also the "environment equation" in the simplest PCT model for the person performing the task. Of course, the environment is really more complex than is shown in the equation: there is the room, the computer, the power system, gravity, the air conditioning (the run occurs in Houston, in the summer), perhaps other people, and so on. But for now we will use the simple equation.

Now we know that the person is a rather complex device comprising what we call organs, cells, systems, sub-systems, processes and so on. How will we represent all of that in a PCT model? The simplest possible model includes three functions (input, comparison and output) and three signals (perceptual, reference, and error). We assume the person establishes a reference signal (p*) (pronounced p-star) that specifies an intended relationship between cursor and target (c-t). In this case, let's assume

$p^* = 0$, which means the intended difference of $c-t = 0$.

We assume the person compares p* with present-time perceptions (p) of the relationship, c-t. The comparison is a simple subtraction (p - p*) and the difference is the error signal. $e = p - p^*$

How can we plug that system into the loop with the environment equation? We need for the model of the person to do more than set p*, p and e; we need for it to behave. We assume that the person can be modeled as a negative feedback control system in which error signals drive the movements of its "handle." How is error converted into movement, hence into handle position? In the simplest possible PCT model, we use the equation (or program statement):

$$h = h - k(p - p^*),$$

where h is momentary handle position and k is the integration factor, a coefficient which represents the movement of the handle (in units/sec/sec) for one unit of the error signal. That is the "person equation" in the PCT model. The complete model is:

$$c = h + d$$

$$h = h - k(p - p^*)$$

That's it.

Now we know this is a pretty simple representation of a person and an environment, but it works. For tracking tasks, it duplicates the person's handle movements and the resulting c-t relationships to a very high degree of precision -- correlations of .996 and higher between thousands of pairs of modeled and actual positions of the handle, and similarly for positions of the cursor. All other descriptive statistics -- and plain old eye-balling -- all confirm the model duplicates and even predicts (as long as five years ahead of time) the person's actions and their results.

With correlations of .996 and above, the simplest possible PCT model accounts for over 99% of the variance in the "simple" tracking task, but it can be improved. For example, the model could be given an arm, like the one in Bill Powers's "Little Man" program, and equipped with the Little Man's visual system. (By the way, "Little Man" is correct -- the modeled arm is programmed to resemble many physical features of Bill's arm.) We could explore many other ways to complicate the model. Some changes might degrade its performance, others might improve it. If they did improve it, the improvement would be in chipping away at the 1% or less of the variance that is not accounted for by the simplest possible model.

So, we say the simplest PCT model is often good enough, but that we are well aware it is inadequate as a complete model of the entire person and the complex environment.

I hope this treatise will help clear up some of your confusion over our way of talking.

> Hope all these replies aren't eating up too much bandwidth.

Isn't that what bandwidth is for? Munch away! ;-)

Later, Tom

Date: Wed Jul 13, 1994 7:25 pm PST
Subject: $64,000 Question

[From Rick Marken (940713.1550)]

Re: Bill Power's diagram:

```
                          :
                          :
               B      ======
           =========>|| F2 ||---C----------D
          ||          ======               |
          vv              :                v
          ||              :              =====
       ======             :             || F3 ||
A ===>|| F1 ||            :              ======
       ======             :                |              E
         ^^               :                |              |
         ||     H     ======               v       ======   |
          ====<======|| F4 ||<-----G------ F <----|| F5 ||<---
                      ======                       ======
                          :
                          :
```

Paul George (940713 17:00) provides the following identifications

> A = reference value or setpoint.
> B = Control output, output variable, or sometimes 'command'.
> H = Control input, input variable, or input data.
> F1 = Control logic or function Block.
> F2 = Output point.
> F4 = Input Point.
> F4.1 = Sensor Instrument.
> G = Input signal or measurement.
> C - Could be a channel
> F2.1 Actuator Instrument.

Thanks. That helps. But I have a few questions.

1) I don't see the term "control variable" in your list. Wasn't it shown in the diagram? What did you mean by "control variable"; what did you think I meant by "controlled variable"?

2) I didn't find a definition of E. Don't you have a name for it? I would think that it would be a very important variable in control engineering.

3) Here's the $64,000 question: What variable(s) is controlled in this loop? Please explain your answer.

Thanks Rick

Date: Wed Jul 13, 1994 8:24 pm PST
Subject: Re: Diagram Terms o

<[Bill Leach 940713.22:11 EST(EDT)] >[Paul George 940713 17:00]

> Note: Because of the way physical devices work, we can't (always) use just negative feedback control. Two electric motors opposing each other like muscles would burn out.

Which is a significant statement. From PCT one can rather easily conclude that when two people (two independent control systems) attempt to control the same perceptual signals to different reference values, conflict will occur and one of the control systems might even "burn out" or at least "reorganize" until the error level is within acceptable limits.

In terms of engineered control systems, control loops may have positive feedback paths, "feed-forward" and the like but if the overall control loop is ever net positive feedback then you have what Dr. Armstrong used to make Radio really practical... an oscillator.

> Many devices like switches, pumps, and valves are discrete (on/off, forward/stop/reverse). And, as (I recall) Bill Leach observed, sensors and actuators are expensive and so kept to a minimum. We also usually just simulate continuous control because we are using electronics and digital computers which (often) operate in a discrete fashion. Biological systems have certain advantages.

Muscles also have "discrete" operation I believe. A particular cell either is or is not exerting force and the amount of force that a particular cell exerts when activated is not an available control function (though the number of cells activated is and for how long or how often is).

Digital computers always operate in "discrete fashion". Even the famous D/A converter is a quantized device. OTOH, what we call analog electronics is considered to be continuous over whatever its operating range may be (that is it is not considered quantized unless one is talking about individual electrons. Even there I don't believe that any theory maintains that the energy of a free electron is quantized).

> Inside:

> A = reference value or setpoint. Usually the desired state of H, but sometimes a limit - i.e. action should be taken when H exceeds the limit until H falls below. Sometimes A = f(H), and corresponds to an estimated state of E or F

The "but sometimes a limit" IS a reference. A control system does not have to be continuously driving some perception to an exact value. For instance, often we will allow ourselves to get rather cool before we "reach a reference limit" that results in our perceiving ourselves as feeling warmer (and maybe putting on a sweater).

> B = Control output, output variable, or sometimes 'command'. Note that we usually keep it around and use it in computations (F1).

I don't believe that human physiology indicates any sort of "direct path" between perceptual sensors associated with output devices (ie: muscle tension sensors) but there is clearly such control loops in existence. It is very probably (I think) that when we consciously perceive a physical control problem, it is likely that the perception comes from error signals "working their way up the chain" or from other sensors such as the trembling and shaking that can result when one stresses muscles for an extended period of time.

For many, as I understand it, demonstrable reasons, it is highly unlikely that perceptual sensors of output functions are used in high level control loops such as is done in engineered systems. One obvious reason is the response speed differences that are often measurable.

> F1 = Control logic or function Block. Usually not a simple, 2 variable difference. Usually a number of H values and B ...

This is not a problem for PCT on the conceptual level. That is, most behavior that we are likely to talk about ALWAYS involves a large number of reference signals and their associated "control logic". However, from a human physiology standpoint, it is quite likely that a single "control block" will never have more than one reference nor more than one perceptual signal at any one time. It may have a number of "control" or "bias" lines connected however.

The point here is that a series of control loops all connected together to control a single high level reference can be discussed as though it were a single monolithic control loop. It is valid to do so as long as one recognizes that it is really multiple loops. In the example often given of someone pushing against someone else's hand while that second person is trying to control the positions of their hand, the number of actual control loops involved is staggering.

> Environment:

I believe that there is a serious "mangling" of control system theory going on there. While in engineered systems, much of the environment IS "controllable" in the design, it is only controllable to an extent. This is the whole reasons why engineered control systems exist... to produce consistent results (goals) under varying and even unplanned environmental conditions (disturbances).

The mechanics of the design (digital computer with A/D converters, Analog Operational amplifiers or an old steam turbine's flyball governor) are only important to the specifics of a particular application. They are all doing (or attempting to do) the same thing functionally. Again, they are attempting to keep some measured parameter (perception) at a desired value (reference) by acting on the environment. It is the recognition that the Steam Turbine's flyball governor is functionally the same as a billion logic gate control system that is important.

-bill

Date: Wed Jul 13, 1994 9:03 pm PST
Subject: Re: Diagram Terms o

[Avery Andrews 940714.1400] (Bill Leach 940713.22:11 EST(EDT))

> I don't believe that human physiology indicates any sort of "direct path" between perceptual sensors associated with output devices (ie: muscle tension sensors)

?? My recollection is that the spinal reflex loops involve two such connections, from organs sensing muscle tension near the tendons, and also from spindles, roughly measuring muscle length. (I used to actually know this stuff last year, I'm horrified to discover how much I've forgotten in the last eight months, but there are plenty of books floating around in which it is all laid out). Or maybe I'm just misunderstanding something, such as what a `direct path' is supposed to be?

Avery.Andrews@anu.edu.au

Date: Wed Jul 13, 1994 10:19 pm PST
Subject: Re: Diagram Terms o

<[Bill Leach 940714.00:31 EST(EDT)] >[Avery Andrews 940714.1400]

Oh hell! I thought that maybe you "miss quoted" my message (by leaving out part of the statement) but then I went back and read the posting!

That should have read something along the lines of:

I don't believe that human physiology indicates any sort of "direct path" between perceptual sensors associated with output devices (ie: muscle tension sensors) -- and the high level perceptual control functions.

What I was trying to say is that measurement of output function operation is likely not used by high level control functions directly. Normally, (always?) only if an uncorrected error exists in a lower level control loop will a higher level control system "be concerned" with measured (perceived) output function performance (if then).

The so called "reflex loops" are complete control loops of their own (with their reference coming in as an output signal from a higher level control system.

> Or maybe I'm just misunderstanding something, such as what a 'direct path' is supposed to be?

The reference to "direct path' was to compare the concept of how engineered control systems often actually sense an measure the operation of an output device for use in calculation of system output signal values. I believe that such functionality does exist in living control systems but that it is distributed in the same fashion as the control hierarchy. That is muscle tension sensors are generally inputs to a control loop directly associated with the operation of the particular muscle (or maybe group) but are not an input to the "top" of the control loop that has for its reference the act of picking up a book (for example).

Obviously I "blew it" in trying to express that idea so thanks for the challenge to what I had actually stated.

-bill

Date: Thu Jul 14, 1994 6:52 am PST
Subject: Re: Facts, Theories & Illusions

[Dan Miller (940714)] Rick Marken:

Happy Bastille Day! - a day marked by good intentions, but the jail was virtually empty. Still, we should celebrate revolutions, and maybe we can bring about our own.

Rick, thanks for the post. I don't really disagree with most of it. I haven't been very clear myself when I have discussed facts. Let me try again.

As I said, I see facts as perceptions that have been made significant. That is, the perceiver has detached the perceptions from the situation (temporally and, perhaps, spatially). S/he speaks of them, thinks about them, tries to make sense of them, shows them to others (thus reproducible), and if lucky gets others to agree about those perceptions. In this sense facts are symbolic and social objects (as in the object of attention and action). Also, facts as significant social objects are used to support theories, models, hypotheses, hunches, etc. - or to negate others.

What I want to say is that the act of transforming the perceptions into facts is itself based on some model or hypothesis about what the facts mean (or how the facts relate to other observations). So, not only do we use theories to interpret facts, but we also use some underlying theory to see and make sense of perceptions as facts. Somewhere in the higher reaches of the hierarchy (language and above) we organize, objectify, and relate our perceptions (sensations, intensities, durations, etc.).

You note (continually) that control is a fact. I agree. I am trying to figure our what a fact is, how we discuss them, and how we use them to construct theories, destruct theories, (destruct? - sorry, I must have had some serious interference} that is destroy (or negate) theories, and a whole assorted set of problems centering around shared perceptions and concerted social action. I am afraid this post is degenerating into drivel.

More later, Dan Miller

Date: Thu Jul 14, 1994 6:53 am PST
Subject: Re: Facts

[Dan Miller (940714)] Bill Powers:

Thanks for the post. I just replied to Rick Marken about what I am thinking about facts. Any comments are greatly appreciated.

I agree that the job of experimental scientists is to get as close as possible to the bare facts (objective descriptions of perceptions?). However, I am not so sure how close we can get to this sense data as humans. My thinking is that those upper levels of the hierarchy really get in the way. Also they allow us to do this kind of work.

Lots of others, particularly behaviorists and new cognitive scientists, do not see our facts as we do. They don't get their significance. My excursion into this topic is as much about how they can be so blind to what are obviously very elementary facts to us.

I'm on-line and getting interference, so I'll post more later.

Dan Miller

Date: Thu Jul 14, 1994 12:53 pm PST
Subject: Re: Replies to Paul

[Paul George 940714 11:00] >From Tom Bourbon [940713.1538]

> But what you call "chaff" is the entire social and scientific establishment called "radical behaviorism." The quality of Skinner's "understanding" is not something I dismiss as readily as you. He touted that understanding as _the_ very science of behavior, not as a personal understanding.

That is what I call 'fluffing the ware's' (legal & economic phrase). People, particularly scientists, scholars, and methodologists control for recognition or status (I hope I phrase that correctly). This often leads to advertising for product differentiation and declaring that one has the one true solution. I concur with your conclusions about the effect of radical behaviorism.

> Instead, he developed a technology, which he called a science, in which he treated all behavior as though it were a result of environmental stimulation.

Just for clarification, the distinction between technology and science is...??

There is a bit of a point of view problem. Clearly the environmental stimuli (though that is a loaded term, I prefer 'influence' or 'event') affects behavior because stimuli is what can be perceived. When I have a continuous or rapidly cycling system (i.e. environment+organism) with what I deem cause and what I deem effect is largely an artifact of at what point I arbitrarily start to follow the chain of interaction. It is a classic chicken and egg situation. As I said earlier, I think PCTrs look at the situation from the organism's point of view while others look at from a theoretical external observer's point of view. However, as Hiesenburg pointed out, the observer and observed affect one another whether they want to or not. PCT thus gives a clearer view of what is going on as it recognizes interacting control systems.

> I never thought of a control system as something that functions because the number of its parts is at least one greater than some threshold number of parts.

I believe it has been demonstrated that there is behavior whose complexity is a step function of the number of components. It is certainly the case that certain concepts or theories cannot be understood until one has acquired a certain minimal set of facts (or other theories) - the pieces just won't fit together. I can't control a process (a.k.a control a perceptual pattern) until I have the right set of perceptions and actions to control it. I can't track the cursor or target unless I can recognize azimuth and elevation. Azimuth and range won't cut it. I also can't perform the tracking task unless I have a stick and the ability to move it in 2 directions.

Perhaps you should give me your definition of 'gestalt'.

> In hierarchical PCT models, there are no distributed nodes, each with its controlled variable(s). We seem to be talking about different kinds of systems.

Say what??? Then how would you describe a HPCS?? Must all of a control network occupy the same locus? Sight and motor control are the same parts of the brain and use the same variables? A HPCS can't involve more than one individual? I think we may be using different meanings for 'distributed'.

Me

>> Nope, while action and input may be lineal, control is usually continuous or at least periodic.

> But you didn't describe it that way. You described a lineal process that includes the production of pre-programmed outputs, selected to match present-time inputs.

Unfortunately language is lineal. Inputs arrive relatively constantly, evaluation occurs relatively constantly, and a given 'action programs' run concurantly, though usually for a period of time. When I decide to throw a punch or execute a block, I do not have to concentrate on the micro movements. That task is spawned into to what is probably an autonomous control loop or 'sub process'. It is a predefined pattern of actions that may be activated at need. I am not constantly prepared to punch or block at all times. I have a little trouble with the idea of a multitude of control loops monitoring at all times every existing controlled variable waiting for an positive error signal to occur. A cascade system seems much simpler and more efficient, but I am contaminated by my experience as an engineer.

I guess the empirical question is whether all action in an organism is the result of a continuously active control loop, or whether some are cyclical or activated on demand (i.e. triggered by a discontinuous event). Note that if the period of a cycle is sufficiently rapid relative to the rate of change of inputs and outputs it appears continuous. This is why we are able to use digital computers for control.

Tracking a cursor, as you eloquently explained below, certainly is a situation where continuous negative feedback loop applies. You are in fact continuously controlling something. It is not clear in my mind how or why the concept of triggering a learned action pattern that goes away when the task is completed is inappropriate. Perhaps an extension of the 'little man' problem to two little men playing the computer game 'pong' would be enlightening. The problem is probably my lack of understanding of HPCT.

> I think these two replies (to Bill and me) are the clearest indications yet that you are talking about a different kind of system from us. A system, incidentally, that is very similar to the ones described by Albus -- you must think of him as more than just someone to use as raw meat to cast before the PCT modelers. ;-)) You seem to envision a "system" as a

kind of "meta-assemblage" of independent systems, each of which accomplishes an assigned task, then passes off the result to the next system(s) in the chain (or net) which uses the received result as fodder for its own processes, and so on.

Sorry for the confusion, I use process in two contexts. The first is where we are trying to provide a control system for a physical process, like a paper plant. The second is when I am modeling human processes for the purposes of process improvement (SEI CMM or BPR or TQM style). It also applies for scheduling and planning projects. It could apply to the situation where organisms plan a future activity, like 'controlling for hunger' triggering a hunting behavior that involves a lot of other behavior patterns which are activated by the recognition of certain environmental conditions along the way.

Your description of a system I think accurately describes both human systems like companies, teams, and societies. It also applies to factories, herds of animals, and the 'environment' in general. It may apply to the control systems within the skin of an organism. The model pre-dated my reading of Albus. As I and (I think it was Rick) commented, Albus' model had correspondence with real systems and a passing resemblance to HPCT structures. We may be talking about different kinds of systems, but I think they can be blended either by combination or translation. I'll let you know what I think when I understand more about the more complex PCT applications and concepts.

Once more around the loop :-)

Paul

Date: Thu Jul 14, 1994 1:24 pm PST
Subject: PCT: A New Way to Look at Behavior

[From Rick Marken (940714.1045)] Paul George (940713 10:00)

> Has anyone in the CSG taken a close look recently on the work on Artificial Life, Core Wars, or Genetic Algorithm software?

I'm on the Santa Fe Institute's Artificial Life mailing list. The folks in Santa Fe are comfortably (and profitably, apparently) ignorant of the nature of behavior as control. Like those in the rest of the life sciences, they are busy trying to explain how organisms generate output (even if they call these outputs "goals"), not how they control.

Dan Miller (940714)--

> Lots of others, particularly behaviorists and new cognitive scientists, do not see our facts as we do. They don't get their significance.

This is an excellent point. And I think I see what you mean about "facts" and I agree with you. From a PCT perspective, any fact (perception) is experienced at many levels simultaneously. The facts we deal with in PCT have to do with perceptions of "behavior" (whatever that is). We experience behavior at many different perceptual levels at once; in terms of intensities, sensations, transitions (actions), events, etc (you know the drill). So a rat pressing a bar in a Skinner box is perceived as many simultaneous perceptual "facts"; color and shape of the rat, the relationship between paw and lever, the pressing event, etc. How one perceives these "facts" must depend on the perceptual functions that have developed in one's brain. Skinner developed perceptual functions that produced a perception of the environment shaping behavior; he probably actually perceived (at some level) food pellets selecting or strengthening the behavior of the rat; "reinforcement" was a "fact" for Skinner as much as "control of reinforcement" is a fact for me. He SAW the rat being controlled by the reinforcement; I SEE the rat controlling the reinforcement.

Skinner would probably agree with a PCTer about the lower level facts of what is going on in operant conditioning -- that the rat is white and that it's pressing the bar with its paw and that this results in the bar being pressed and a food pellet appearing. But Skinner's higher level perception of the situation would differ from that of a PCTer; the PCTer would experience the "behavior" that Skinner sees as "pellet selecting bar press" as "rat selecting food pellet".

So I think you make a very good point Dan -- "facts" (as higher level perceptions) must, to a large extent, depend on the perceptual functions people have built up and use to experience these facts. The perception of the same behavior that Skinner has and that the PCTer has are still "facts" but I think one fact is "better" than another. The PCT "fact" is consistent with other perceptions (like disturbance resistance) while the Skinner "fact" is not.

This provides a nice sequel into my proposal for a new agenda for PCT. I think we should stop pushing PCT as a new theory of behavior because everyone assumes that they know what the "behavior" is that PCT is a new theory of -- and they are virtually always wrong. I suggest that we spend more time publicizing PCT as a NEW WAY OF LOOKING AT THE FACTS OF BEHAVIOR. We have to teach people how to SEE CONTROL; we have to help people build new perceptual functions that will let them see what I call (perhaps incorrectly) the "fact" of control. In the operant situation, for example, we can do this by showing how a perceptual aspect of the pellets (like rate in pellets/sec or amount in volume/ sec) is maintained against disturbances (changes in schedule, size of pellets, etc). This can be observed in real time if done properly. We have to develop tons of demonstrations so that whenever a person sees a behavior, they are able to see that there is a perception affected by actions and disturbances and that the disturbances have no effect on the perception; we have to show that this is happening when people push on buttons, type e-mail messages, set up psychology experiments, go to the library, brush their teeth, write a program, kiss their wife, pick up their daughter after work, make a barbecue, explain why alternative models of behavior are horseshit, eat lunch, make a phone call etc etc. People have to learn to perceive control. Until people can see what we mean when we say that "behavior IS control" PCT will just be another, rather unglamorous model of how people "behave".

I plan to suggest a number of "portable" demonstrations of control over the next few days; I hope other people will have some suggestions for such demonstrations too. I'm sure there will be disagreements about whether what is being demonstrated is "really" control. But I think we should nail down the presumed "facts" that PCT is designed to explain before getting too far into the details of the PCT model itself.

I really think that the most important thing about PCT is that it gives us a new way of looking at what we call "behavior". Maybe a paper on "Watching people control" would get the attention of behavioral scientists -- maybe.

Oh, and we should also be clear that we know that not all behavior involves control; sometimes people are flailing away, learning how to do something; in this case the "something" being learned is not under control until the person learns to control it.

Best Rick

Date: Fri Jul 15, 1994 2:28 pm PST
Subject: PCT and process control; modeling spatial perception

[From Bill Powers (940714.0900 MDT)] Paul George (940713.1700)

Excellent replies to my request for labels on the diagram. I'm going to concentrate on main differences/agreements and not comment on side- issues, about which we have no _important_ disagreements.

```
                          :
                          :
               B      ======
           =========>|| F2 ||---C----------D
          ||          ======               |
          ^^              :                v
          ||              :              =====
       ======             :             || F3 ||
A ===>|| F1 ||            :              ======
       ======             :                |                  E
         ^^               :                |                  |
         ||     H     ======               v       ======     |
          ====<======|| F4 ||<-----G------ F <----|| F5 ||<---
                      ======                       ======
                          :
                          :
```

> A = reference value or setpoint. Usually the desired state of H, but sometimes a limit - i.e. action should be taken when H exceeds the limit until H falls below.

Right. We call A the reference _signal_, as it is variable. We call a "limit" system a "one-way" control system; if controls only when H exceeds (or is less than) A.

> B = Control output, output variable, or sometimes 'command'. Note that we usually keep it around and use it in computations (F1).

This is what we call the error signal. It represents the difference between A and H. We wouldn't use the word "control" here; see below.

> H = Control input, input variable, or input data (precise aren't we). Same apparent semantics as in PCT controlled variable. In our case H can be a scalar or vector array, but that is an artifact of design. We could construct things your way with one loop per variable.

Here there is some divergence. In PCT, H is the perceptual signal. It is _not_ what we call the controlled variable, because we use the term "variable" to refer to observables in the environment (what an external observer could see). A "signal" is a variable inside the control system. It is true that H is a controlled signal, in that the action of the whole loop is such as to make H match A, whatever A may be. I don't much like "control input", because this implies that the signal H is capable of controlling something -- i.e., maintaining it in a specified state.

Yes, you could construct things my way with one loop per variable; that's the current preferred model in PCT, but of course there's nothing carved in stone there. Thinking of one loop per variable makes it easier to understand what a system is doing. Sometimes, I think, it makes a system easier to design, too.

> F1 = Control logic or function Block. Usually not a simple, 2 variable difference. Usually a number of H values and B values and (sometimes) A values go through a set of algorithms to determine the desired value for B. We are usually forced to be pro-active due to latencies in the system and the process under control.

Again, I consider the term "control" unfortunate, because by itself this block can do no controlling. Control is what the whole loop does. We call F1 the comparator, and treat it as a simple subtraction. This is part of seeing how far we can go with the simple PCT model, as Tom Bourbon noted. Neurally, subtracting one signal from another to get a difference signal is very simple: one excitatory input and one inhibitory input do the trick. Of course this means that to get two-way subtraction you have to have two comparators, because neural signals can't reverse sign. We usually ignore this problem.

As to multiplicity of inputs to a comparator, I think that is taken care of by the "one-variable-one-loop" architecture. Also I think that allowing too many arbitrary connections in an ad-hoc way leads to messy designs, and probably unlikely designs for a living system. And much of what you might try to do with a complex comparator is taken care of in a simpler way by a hierarchy of control.

> F2 = Output point. An output transducer. Performs D to A conversion in some cases. Basically translates B to whatever the actuator instrument requires. In some cases performs computations or sampling on B. Depends on the cycle time of F1 vs that of the instrument.

Agree. We call F2 the Output Function. It converts the error signal into some immediate influence on the physical environment, unidirectionally (disturbing its output does not alter its input).

> F4 = Input Point. An input transducer. Usually performs A to D conversion. Usually samples or averages F.

We call F4 the Input Function or Perceptual Function. It converts one or more inputs from variables like F into a signal H that represents some aspect of the part of the environment where F is. The kind of function performed depends on the level in the hierarchy. For a system in the middle of the hierarchy, the inputs F are really copies of lower-level perceptual signals, some of which are under direct control and some of which simply represent states of the lower-level world. The first-level systems receive inputs only from the environment. Part of the modeling problem is to decide what function F4 performs. It could be sampling, averaging, or anything else.

Since the output of F4, or H, is what is really controlled (the ultimate criterion is that H must match A), it is the form of F4 that determines what the external observer will see being controlled in the environment of the control system. Thus in PCT, we do NOT begin by knowing what aspect of the external world is under control, as the engineer would. We have to deduce what is under control by hypothesizing forms of F4, then testing to see whether the environment, viewed through the hypothetical function, is resistant to disturbance because of the actions of the system. That's the basis for the Test for the controlled variable. Engineers don't have to use that test; they know what aspect of the environment is supposed to be under control.

> F4.1 = Sensor Instrument. Translates the sensed variable F into some kind of a reading or metric. A pressure sensor turns a pressure into a voltage. Sometimes combines some set of physical sensors into a reading or set of readings. Note that I may measure the temperature of a vessel and translate it into the temperature of the substance inside the vessel. Sometimes samples or averages a continuous reading. I think this has the same semantics as your G, but we consider it a function, since it is a device. The channel is separate. Some advanced instruments are control systems in their own right. They may have both sensors and actuators, and may transmit an error signal.

No problem here. In the nervous system, all input translations are either from a physical variable to a neural signal (considered as a continuously-variable frequency at the lower levels) or from a signal into a signal. No A/D conversions as such, of course.

You are simply expressing the hierarchical nature of perception. At the lowest level, where stimulus intensity is converted to neural frequency, there are some control systems (like the spinal reflexes or the iris reflex) which directly control intensity signals, and thus indirectly control the physical variable (F) of which the signal is a quantitative analog. Copies of the controlled signal also pass upward to higher levels, the same levels from which the reference signals for the spinal control systems arise. So a higher system may send its error signal (B) to the reference inputs of a number of spinal control systems, and receive its input information (F) in the form of copies of the controlled perceptual signals H from a number of intensity-control systems (as well as from uncontrolled sources). Thus the set of intensity-control systems becomes, on the output side, part of the output function (F2) of the higher system, and also provides input data (F) to the input function (F4) of the higher system. The input function generally is a many-to-one function.

When we model just the higher system, we lump all the lower systems into the higher system's output function, and treat the upgoing signals entering the input function as if they were aspects of the environment. So all control systems at all levels operate through the environment, although at different levels of abstraction. We can experiment with higher levels of control without necessarily understanding the lower levels that are involved.

So OK on your F4.1->F4 chain. In HPCT we actually have an F4.1...F4.11 chain! With, of course, control systems at each level.

> G = Input signal or measurement. We would usually have a G' which is the channel (wire), but the distinction is not important to the control system. As noted above, G in some cases is an error signal.

I meant for G to refer to physical processes at too low a level to be of interest in modeling (like the jiggle of electrons in a wire). It's whatever conveys a physical quantity like temperature measured near the skin to the surface of a cell like a temperature-sensor.

In PCT G would _never_ be properly called an error signal. The function F4 might construct a perception representing the difference between two environmental variables -- for example, the difference in position between a fingertip and a target, as seen. But that is just another perceptual signal, representing the current state of some variable in the environment. Zero difference in position might constitute an error, if the reference signal A is set to specify some non-zero difference that is to be maintained. In the Little Man, non-zero reference signals for controlling target-finger separation in X and Y are used to make the finger draw a circle around the target.

> F - An attribute or characteristic of the process or of the equipment being controlled . A temperature, pressure, etc.

Yes. That, in HPCT, would be a first-level controlled variable, with Ht standing for the temperature, Hp for pressure, etc. With both T and P under control, a higher-level system might receive the signals Hp and Ht as inputs at a second level, in which case those inputs become F for the higher system. Then the higher system could continually compute T/P, represent the ratio as a second-level perceptual signal H2, and control the ratio of T to P (or any other function of the two variables) by varying reference signals entering the two lower systems controlling T and P. Only the lowest level of system sends its output directly into the environment via transducers. All higher systems must act by varying the reference signals for lower-level systems.

> F3, F5 = As in yours, though we would think of it as some characteristic or part of the process system under control. We usually 'mirror' this function within the control system for various practical reasons.

In PCT we distinguish these functions for a good reason. The function F3 lies in the path between the visible physical output of the control system (D) and the visible controlled variable (F). The function F5 does not; it conveys the effects of independent _disturbances_ (E) to the controlled variable by a path that is independent of the system's actions. In general, both E and the form of F5 are unknown to the control system and cannot be deduced from the behavior of signals in the control loop. An example is the effect on your car's direction of a crosswind. The velocity of the crosswind would be E. The aerodynamic laws that convert the vector velocity into a lateral force on the car would be F5. You keep your car on the road without being able to sense E and without knowing how to compute F5.

As you can no doubt see immediately, if F is under control, there will be a necessary relationship between D and E. This is the apparent relationship on which stimulus-response theory was built, with the existence of controlled variables like F being unsuspected.

> C - Could be a channel as in yours, but again the 'output' signal or 'output data' is usually distinguished from the wire.

OK. A detail.

> D - not real clear on its usage in PCT. The action of the actuator usually affects the process or equipment being controlled, frequently indirectly.

Another critical point in PCT. D is the immediate effect of the actuator on the environment. In a motor, it would be the torque generated by the motor on the end of the shaft at the motor. This is rarely the variable that is to be put under control (F). What is usually to be controlled is some rather remote effect of D on the environment, such as the position of a load being wound up on a pulley at the other end of the shaft. The function F3 expresses the physical link between the actuator output D and the variable to be controlled, F. Because of the function F3, it is possible for other influences in the environment to affect the controlled variable F independently of the output D of the control system. The state of the controlled variable is really given by F = F3(D) + F5(E) (in the simplest case). This is why, if F is being actively controlled at zero, we get the interesting relationship that Rick Marken often talks about using other symbols: F3(D) = -F5(E). That's what gives the appearance of responses to stimuli.

The definition of D depends on the level of control you're talking about. It is always defined so that D is very difficult for the environment to affect, so D depends ONLY on the error signal in the control system.

F2.1: OK, just part of F3. Note that it's not part of F2 unless we know that the environment can't disturb it. The environment can actually insert disturbances anywhere between D and F; we represent all disturbances, however, as an equivalent disturbance of the kind shown.

> Process control just uses a kind of Mind/Body distinction that I haven't seen HPCT use explicitly.

Right. The distinction in PCT is between the nervous system and all that is not nervous system, with sensors and actuators lying in the boundary surface. This allows us to speak of "the control system" as an entity of relatively fixed organization, while "the environment" can be continuously changing. The same control system can operate in many environments. This is seldom of interest in process control, because the engineered control system is bolted down and wired to the process it is to control -- it can't wander around and encounter environments with different properties (F3, F5) the way most organisms can.

> The only likely difference I can see is that HPCT nodes can rout C of one node directly into G of another without intermediary processing. In the degenerate case B and H can be the same variable.

Right. Usually, however, we do use an output function (F2) even in higher-level systems, because a given control system will act by adjusting reference signals for many systems of the next lower level. The least computation we need is to set the signs of the outputs correctly so there is negative feedback in the path involving each lower-level system. I have been trying to avoid introducing complex output functions, because one of the nifty feature of control systems is their ability to determine what is to be controlled strictly in terms of the input function, F4. You don't need to get the output function just so in order to get good control; it just has to have the right sign and about the right form. Proportional or integral output will usually do, with perhaps some dynamic trimming. There's nothing fundamental to say that output functions can't be complex, but my instincts tell me that we should first explore what can be done with input functions, leaving complex output functions as a last resort. Some amazing things can be done with the output signals weighted only by 1, 0, or -1.

See my Byte articles, particularly Part 3:

Powers, W.T. (1979) The nature of robots: Part 3: A close look at human behavior. Byte,_4_, No. 8, 94-116.

> I am still thinking about how to directly use HPCT to modify process control system architectures, but suspect I don't understand it well enough yet. I think there is some real possibility of striking gold, but it is just an intuitive reaction.

Your intuition and mine agree.

Misc:

The _action_ of a system is D. D can in general affect many environmental variables in addition to F. So without finding which of these affected variables is under control, it can be difficult to decide what effect of D should be called the organism's "behavior." We mean by behavior the effect is that is being controlled, not the action by which is it controlled. The action, D, varies as disturbances E come and go, and as internal properties of F3 change. The real behavioral variable F changes only as the reference signal A changes. This is why S-R laws are so unreliable and have to be deduced statistically. They express relationships between D and E, but leave out the effect of A: what the organism wants, which is variable.

What the organism is "really doing" can be seen only by looking at F. Of course in another sense, you see what the organism is doing by looking at D. That's why we say that you can't tell what an organism is doing [F] just by looking at what it's doing [D].

Best, Bill P.

Date: Thu Jul 14, 1994 1:37 pm PST
Subject: Re: Diagram labels

[Paul George 940714 14:30] >[Bill Powers (940714.0900 MDT)]

Thanks for a Very enlightening post. It cleared up a lot of minor confusion. I agree that we don't seem to have any major differences.

The only real difference in terminology is kind of computer related. For you, and probably within living systems, signal and variable are equivalent. For us a variable is a static storage area whose value is changed by a signal. Anything which is interested in the signal reads the value of the variable. The Variable represents or records the input or output. The distinction is not significant until you try to model a living system with a computer or electronics

> In PCT G would _never_ be properly called an error signal. The function F4 might construct a perception representing the difference between two environmental variables

I meant to say, as you indicate above, that F might be a B of another loop. Sometimes it is the magnitude of the error signal that is of interest, not the 'action' (C or D) which is the transduced output of F2. For simplicity in your model you bundle a lot of transforms into F2 & F4 that we break out for practical reasons.

> In PCT we distinguish these functions for a good reason. The function F3 lies in the path between the visible physical output of the control system (D) and the visible controlled variable (F). The function F5 does not.

The distinction is very important, and critical for the _design_ of the process control system and the control algorithms. It is not important for the operation of the system. Of course as you point out, an engineer knows the real process while an organism tries to divine it from its perceptions.

> I have been trying to avoid introducing complex output functions, because one of the nifty feature of control systems is their ability to determine what is to be controlled strictly in terms of the input function, F4.... There's nothing fundamental to say that output functions can't be complex, but my instincts tell me that we should first explore what can be done with input functions, leaving complex output functions as a last resort. Some amazing things can be done with the output signals weighted only by 1, 0, or -1.

Ack, but it might make a more complete explanation of something like driving a car or performing a learned response simpler. Starting your way is probably best for Rick Marken who is doing low level modeling of physiology (it appears). However the complex output might be more useful for Dag in applying HPCT to individuals or organizations.

Thanks again for the clarification, particularly the 'Misc'. A really slick observation.

Paul George

Date: Thu Jul 14, 1994 1:41 pm PST
Subject: Re: $64,000 Question

Paul George 940714, 11:40 >[Rick Marken (940713.1550)]

> 1) I don't see the term "control variable" in your list. Wasn't it shown in the diagram? What did you mean by "control variable"; what did you think I meant by "controlled variable"?

In our world 'control variable' means a data structure used in control, which contains a value. Thus it would apply to both B & H. The way I use it in this forum is to refer to H. It is that input that the control system would like to get to match A through the change in (I guess) D.

> 2) I didn't find a definition of E. Don't you have a name for it?

Sorry, when I generated the post a second time (letting it simmer changed how I structured my response) it got left off. It is some facet of the process that we cannot directly sense, but must rather infer via F. I know of no formal name for it. Since we often cannot sense the perception we really want to control, we often have to provide an F5 and/or F3 inverse within F4 or F1 (or more usually between, this is the 'world model' or 'mirrored object'.

> 3) Here's the $64,000 question: What variable(s) is controlled in this loop?

From your point of view (and mine) H. By convention and practice the system is seen as controlling D in order to provide control of F (the 'process centric view'). However, the control system writ small (essentially F1) only cares about H & B. Recall that in our system we are often triggering an industrial machine (often with control capabilities) that may do a fairly complex series of things to the process under control. They respond to a fixed set of signal's that are usually interpreted as commands or instructions. An F3 could represent such a machine.Sometimes B may be a program which is downloaded to replace an F3. Similarly a F5 could represent a machine not under the direct control of the control system (I know the phraseology flies in the face of the face of the PCT concept of what is controlled).

Hope this helps

Paul

Date: Thu Jul 14, 1994 1:42 pm PST
Subject: Re: Erratum

Paul George 940714 12:00 >[Bill Powers (940713.1700 MDT)]

> A comment on the environment part of the model is needed. In fact, the variable H might be multiple physical variables, H should be changed to D.

In that case there may be a disconnect if the statement is not true (for D it certainly is).

In practice several Hs (and often B) must be input as some F1.n function must convert them to a value H' that may be compared to A. It is my understanding that allocating each H to a separate control loop usually doesn't work because the value of a given A depends on the values of the other H's (not to say that it is absolutely impossible). See also my response to Rick's $64k question.

BTW It slipped my mind but F1, or more accurately {A,B,H,F1}, would often be referred to as a 'Controller'. It could also refer to a HPCT network. It's all a question of the physical architecture of the process control system.

Paul

Date: Thu Jul 14, 1994 1:46 pm PST
Subject: Re: Diagram Terms o

[Paul George 940714 15:00] >[Bill Leach 940713.22:11 EST(EDT)]

> The "but sometimes a limit" IS a reference. A control system does not have to be continuously driving some perception to an exact value. For instance, often we will allow ourselves to get rather cool before we "reach a reference limit" that results in our perceiving ourselves as feeling warmer (and maybe putting on a sweater).

What I meant is that there might several H's representing upper and lower control limits SPC style. The error signal B could be {-1,0,1}, or a larger set, resulting in different D's depending on which limit was crossed. Slowing a car and slamming on the brakes can be a considered a different action, rather than just a difference in the gain of F2.

Paul

Date: Thu Jul 14, 1994 4:09 pm PST
Subject: We have a winner!

[From Rick Marken (940714.1445)]

Me:

> 3) Here's the $64,000 question: What variable(s) is controlled in this loop?

Paul George (940714, 11:40) --
> From your point of view (and mine) H.

YES! You WIN!! Congratulations. You are the first control engineer on CSG-L to correctly identify the perceptual signal as the variable (OK "signal") that is controlled in a control loop! The behavior of a control loop is the control of perception (seems like I've read that somewhere before).

A check for $64,000 will be in the mail to you soon -- I hope;-).

> By convention and practice the system is seen as controlling D in order to provide control of F

Oops :-( The only variable in the loop that is controlled (kept at a specified level against disturbance) is H. But you don't have to send the check back (if you get it); getting H right is VERY impressive. Really. You wouldn't believe how tough it is to get control engineers to accept this incredibly simple fact.

> However, the control system writ small (essentially F1) only cares about H & B.

Well, it really only cares about (that is, controls) H; the system cares what value H is -- it "wants" H to equal A and it will do what it can to make that happen. The system doesn't care what value B is; B just varies around as necessary (depending on disturbances -- E -- and changes in functional relationships in the loop) to keep H matching A.

In PCT we like to say that a control system doesn't care what it does (variations in B and D), just what it perceives (variations in H).

Paul George (940714 11:00) --

> I believe it has been demonstrated that there is behavior whose complexity is a step function of the number of components.

What is behavior? What are its components? It sounds to me like the behavior of which you speak consists of the visible actions and results produced by an organism. Anything that has been "demonstrated" about this behavior would be quite irrelevant, would it not? After all, we know that organisms control their perceptions, right?

Best Rick

Date: Thu Jul 14, 1994 8:10 pm PST
Subject: Re: Behavior and control

From Tom Bourbon [940714.1740] Replies to Paul George and Bill Leach.

>[Paul George 940713 17:30]

>>From Tom Bourbon [940713.1512]

>> Behavior is the means by which perception is controlled; behavior is not a result of something, other than behavior, by means of which perception is controlled.

>> Of course, we then go on to say that many of the outward appearances of specific _actions_ we see a person make at a particular time are probably unintended side effects of the person's control of perception.

> Could you clarify the distinction between Behavior and Action? I think it may be the basis of much confusion.

> I usually look at behavior as being a pattern or set of actions, usually within the context of some environmental pattern of events or inputs. Tracking is behavior, a given arm motion is action.

At least on first reading, that doesn't look terribly far from what I believe is the PCT-modeling interpretation. I said not terribly far -- there are some differences, the magnitudes of which we can explore. Most people, including most behavioral scientists, use the word "behavior" to refer to their own observations of what another creature is doing -- he is walking, she is talking, they are assembling, the rats ran along that path to a new source of food, and so on. In that usage, behavior is interpreted in terms of the observer's perceptual units and the labels are really names for _results_ the observer notices when the other creature acts. In that usage, actions and behavior are often equated or used interchangeably.

In PCT modeling, we have found it useful, often necessary, to think of the behavior of a control system as what it is doing _from its own perspective_. If, after performing the Test for Controlled Variables, we concluded that a system is controlling its perceptions of the position of an automobile on the highway, we speak of that as its behavior -- that is "what it is doing." It does what it is doing (which is controlling a particular perception in a desired way) by acting on the environment; the actions by which it achieves control of perception are not "what it is doing." The system's actions must be "out of control," in that the must vary any way necessary (for example, due to environmental disturbances) in order for the system to do what it is really doing -- controlling a particular perception or set of perceptions.

In the PCT interpretation, many, if not all, of the things an observer sees a system "doing" may well be outward appearances that are irrelevant to the observed control system -- in most cases, what an observer sees does not even exist for the observed control system -- it does not know that it is seen as doing what the observer sees.

This difference in "points of view" concerning what the observed system is doing is behind many misunderstandings between people -- "Why do you keep doing that? Don't you know it drives me crazy?" "Stop doing what?" "Damn it, you know what I mean! You keep making that stupid noise through your nose every time you type at the word processor." "What noise?" And so on. And in the behavioral and cognitive sciences, we have scientists who observe people "doing things" and then conclude that they (the scientists) know what those people are doing. Then they ask the people what they are doing. The observed persons reply from their own perspectives; they describe what they are controlling, whereupon the scientists (and a few famous "neurophilosophers" I could name) shout in unison,"Aha! Do you see? We _told you_ that people are unconscious of what they are doing. They don't really know what they are doing." In contrast to nearly all other behavioral scientists (and those very famous neurophilosophers) a PCT modeler is more likely to take people at their word when they say, "I was making the contrast on the picture just right," or "I was going to the store to buy flour to make cookies for Aunt Sally." The difference in our interpretation comes from our knowing that the actions of the person are not "what the person is doing;" the person is "doing" his or her own controlled perceptions.

> You also seem to be saying that behavior directly causes other behavior. This makes sense in the context of muscles (unobservable) moving an arm (observable) or a arm motion propelling a ball. Is that all you meant?

I don't think I was saying exactly that, but it is true that the actions of a hierarchical control system are "nested" in something like the manner you describe. (As an aside, are you perhaps alluding to the old "behavior cannot cause behavior" song, from radical behaviorism?)

><[Bill Leach 940713.21:10 EST(EDT)]
>>From Tom Bourbon [940713.1512]

Bill L. says:

> Let me try to see if I can express this in a cogent fashion... In common terms, "behavior" is a label for actions of people. In its' common usage, the term is pretty vague but that is how most people see the meaning. Behavior is what we observe another person do, that is their observed actions are behavior.

Bill, in my reply to Paul I was actually replying to your questions as well. Did I come close to addressing some of the questions you asked in your post? I agree that common usage is pretty vague on "behavior" and "actions" and even more vague or silent on "unintended side effects of actions." In PCT modeling, it is mandatory that we maintain some kind of clear distinction among the moment-by-moment products of the output function, the many environmental consequences of those outputs, and the perceptual signal(s) the system is controlling. In the modeling, the names of those variables are unimportant -- they can be assigned any symbolic label that is acceptable to the programming language one is using -- it is their _values_ that matter, when the PCT model is run in simulations. Of course, once we leave the confines of the little worlds we create when we are modeling, we are back in the world of people and words; that's when the fun begins. This is why it is so crucial that anyone who wants to understand PCT look beyond the words they see on this net, or in our sometimes hard-to-locate publications. A deep understanding of PCT can only come from grasping the significance, if not the computations, of the quantitative modeling. The PCT model is not about the words; the words are always inadequate for expressing the quantitative model, and the necessarily linear structure of language can never convey the simultaneity and continuity of the model in action.

. . .

> A typical "formal" definition of the term "behavior" is: The term "behavior" may be understood to embrace both the expressed and potential capacity for activity in the physical, mental and social spheres of life.

> <That looks pretty useless to me but that is what many will see as the meaning of the term.>

> The phrase "Behavior is the control of perception." is a definition of the term "behavior". I don't believe that I maintained an incorrect perception of the term as used in PCT but also recognize that most people will not see that "behavior" is a output and nothing more (I say output because as I understand PCT, my "conclusions" when thinking are behavior even if they are not observable to any outsider).

Do my comments above come close to addressing these ideas?

> It is also far more complete than what I said (this I see in retrospection). When I say something like "The control of perception results in what we call behavior." it is possible that the person hearing that might think that what is observed (so called behavior) is what is controlled (though I still think that such an interpretation is incorrect).

Agreed. You are speaking of the (easily understood by PCTers) problem an observer has if the observer is not aware of the phenomenon of control. To that observer, it seems extraordinarily easy to identify what the observed system "is doing" and never even suspect that the conclusion is wrong. (An aside: I believe this point is relevant to the current discussion on csg-l about "facts," but I won't have time to join in that discussion before I vanish for the wedding.)

> I will have to think about this a bit more. I still feel like taking the step to saying something like "The control of perception results in what we call behavior." and then going on to explain B:CP is not necessarily such a bad idea.

I think you can see now why I said, in the earlier post, that some of us would chose to disagree with you -- but good naturedly :-) -- if you were to say that.

. . .

Later, Tom

Date: Thu Jul 14, 1994 11:51 pm PST
Subject: Re: Replies to Paul and Jeff

From Tom Bourbon [940714.1750] >[Paul 940714 11:00] >>Tom [940713.1538]

>> But what you call "chaff" is the entire social and scientific establishment called "radical behaviorism." The quality of Skinner's "understanding" is not something I dismiss as readily as you. He touted that understanding as _the_ very science of behavior, not as a personal understanding.

> That is what I call 'fluffing the ware's' (legal & economic phrase). People, particularly scientists, scholars, and methodologists control for recognition or status (I hope I phrase that correctly). This often leads to advertising for product differentiation and declaring that one has the one true solution. I concur with your conclusions about the effect of radical behaviorism.

It looks as though we ar approaching something of an agreement on Skinner, at least in terms of how he did his fluffing. That's good. :-)

>> Instead, he developed a technology, which he called a science, in which he treated all behavior as though it were a result of environmental stimulation.

> Just for clarification, the distinction between technology and science is...??

I was using Skinner's term. For some audiences, he said behaviorism was only a technology, not a science; for others, he claimed behaviorism was _the_ science of behavior (TB: "the only" was understood). I assume that when he said it was a technology, he was trying to evade challenges to the quality of behaviorism as science -- "You don't think behaviorism measures up as science and those are the reasons you give? Not to worry! Behaviorism is only a technology -- a set of handy tools for controlling behavior." Notice the neat dance to the side? What a nifty way to stay out of arguments, one of those mythical feats for which Skinner was noted during his life. Of course, the fact that there was more to behavior than radical behaviorists allowed to meet their eyes was ignored and behaviorists-as-THE-behavioral-scientists could go on about their "Science" undisturbed. Come to think of it, Skinner was being a damned effective control system, wasn't he?

> There is a bit of a point of view problem. Clearly the environmental stimuli (though that is a loaded term, I prefer 'influence' or 'event') affects behavior because stimuli is what can be perceived.

Agreed, up to a point. In the right conditions, it is possible to see a contingent temporal sequence of stimulus followed by response. (See Rick's paper on the blind men and the elephant for a PCT analysis of that S-R interpretation -- oh, by the way Paul, Rick and the rest of us know the story was around long before Rick used that as his title; that's why he picked the title.) But the fact that I can see that sequence does not mean I have seen what the observed system "is doing:" controlling specific perceptions. Once I see _that_ fact, then an infinite array of specific stimulus-response contingencies can be predicted with great precision and they can be explained and modeled as though the system acts to eliminate the effects of disturbances that affect variables related to its (the system's) controlled perceptions.

> When I have a continuous or rapidly cycling system (i.e. environment+organism) with what I deem cause and what I deem effect is largely an artifact of at what point I arbitrarily start to follow the chain of interaction.

I guess that could occur, were you observing a truly discrete system-environment interaction sequence. Can you identify some interactions in which that is really what occurs? Those aren't the kinds of things you see when control systems interact with their environments, but there _are_ circumstances when that can _appear to be_ the case. Again, see Rick's paper, and "Models and their worlds," by Bill Powers and me, and my paper on "Mimicry and repetition." Easy for me to say that : all three papers are in our ghetto journal, _Closed Loop_. Mary Powers can supply copies, at cost.

> It is a classic chicken and egg situation.

Not in a closed loop.

> As I said earlier, I think PCTrs look at the situation from the organism's point of view while others look at from a theoretical external observer's point of view. However, as Hiesenburg pointed out, the observer and observed affect one another whether they want to or not. PCT

thus gives a clearer view of what is going on as it recognizes interacting control systems.

Heisenburg wasn't talking about the behavior of organisms, so I'm not certain we can assume that his comments will automatically apply to our work, but that caveat aside, I agree.

>> I never thought of a control system as something that functions because the number of its parts is at least one greater than some threshold number of parts.

> I believe it has been demonstrated that there is behavior whose complexity is a step function of the number of components. It is certainly the case that certain concepts or theories cannot be understood until one has acquired a certain minimal set of facts (or other theories) - the pieces just won't fit together.

But PCT is about an empirical, nontheoretical phenomenon -- control, and about the minimal organization of a system that can achieve control. It is not about how many _pieces_ there must be, but about the specific functions that must occur, whether they occur in a person, or a pigeon, or a bacterium, or a chemical reaction. A control system -- a person or a rat for example -- can loose many identifiable parts and still function as a control system. What matters is not the absolute number of parts, but the specific organization of the ones that remain.

Incidentally, in Wales, Pedro Mendes gave an exciting presentation on his work in which he applies the PCT model to biochemical reactions. There is no nervous system involved there, or in the modeling of bacteria by Bill P. and Rick. What is important is that the requisite functions occur. Their occurrence is of course dependent upon enough "pieces" being present, whatever the substance of the pieces might be, but the number of pieces alone has nothing to do with what happens when they are connected and put in motion. It is the specific organization of the system that matters. The organization represented in the PCT model seems to have wide applicability as a model for control -- we are willing to start suggesting that it will apply to control by all living things.

> I can't control a process (a.k.a control a perceptual pattern) until I have the right set of perceptions and actions to control it. I can't track the cursor or target unless I can recognize azimuth and elevation. Azimuth and range won't cut it. I also can't perform the tracking task unless I have a stick and the ability to move it in 2 directions.

And having said all of those true things (ignoring for now the possibility that we might need to explore the meaning of "recognize azimuth and X"), you (generic you) have said exactly nothing about how it comes to pass that a person can use a control stick to keep a cursor aligned with a target. That explanation requires a model for the organization of the person, not for the number of parts in the person.

> Perhaps you should give me your definition of 'gestalt'.

Ah, it was _you_ who introduced the term into the discussion. I've called _your_ hand. Show your cards! :-)

>> In hierarchical PCT models, there are no distributed nodes, each with its controlled variable(s). We seem to be talking about different kinds of systems.

> Say what??? Then how would you describe a HPCS?? Must all of a control network occupy the same locus? Sight and motor control are the same parts of the brain and use the same variables? A HPCS can't involve more than one individual? I think we may be using different meanings for 'distributed'.

See my post earlier today -- was it to Bill Leach? I was speaking of what I saw (perhaps incorrectly) as an appeal by you to the kind of "distributed system" that is in Albus's papers. The PCT model, whether in the single-loop or hierarchical version, is not like the "models" in the Albus papers. You still have not told us whet your own model would look like, if not like Albus's.

>Me

>>> Nope, while action and input may be lineal, control is usually continuous or at least periodic.

>> But you didn't describe it that way. You described a lineal process that includes the production of pre-programmed outputs, selected to match present-time inputs.

> Unfortunately language is lineal.

Yes, it is. But you also included a diagram, and it was clearly a diagram of a lineal system for input-process-output.

> Inputs arrive relatively constantly, evaluation occurs relatively constantly, and a given 'action programs' run concurantly, though usually for a period of time.

But living control systems don't work "relatively" constantly; they work constantly -- continuously. And the world works constantly, not intermittently.

> When I decide to throw a punch or execute a block, I do not have to concentrate on the micro movements.

Yes. Isn't it amazing?! That's the nature of perceptual control -- the actions that produce your behavior -- your controlled perceptions -- seem to "just happen" as though by magic.

> That task is spawned into to what is probably an autonomous control loop or 'sub process'.

Oops! I won't even start to reply to this until after you have a chance to read some things about PCT. If you still want to say this after you have done the reading, then we can go at it in earnest. ;-)

> It is a predefined pattern of actions that may be activated at need.

Ouch! In spite of what I said a sentence earlier, I can't hold back from saying that this is completely outside the understanding of behavior in PCT. I believe we have developed some pretty convincing arguments that nothing like this happens in control behavior. Nothing.

> I am not constantly prepared to punch or block at all times. I have a little trouble with the idea of a multitude of control loops monitoring at all times every existing controlled variable waiting for an positive error signal to occur.

Read about it. I think you will enjoy it. And I think you will find that the characterization you just gave is not an accurate one of the hierarchical PCT model. (That is not meant as a criticism or put down, but as an acknowledgement that you still have not read the literature or hung around the net long enough to know what we say about such things.)

> A cascade system seems much simpler and more efficient, but I am contaminated by my experience as an engineer.

That's OK. We'll forgive you for the bad company you fell in with, earlier in your life. ;-)

> I guess the empirical question is whether all action in an organism is the result of a continuously active control loop, or whether some are cyclical or activated on demand (i.e. triggered by a discontinuous event).

The question, the $64,000,000 question (I've raised thee ten fold, Rick) is is there a "controlled variable." The next question is, "If there is a controlled variable, what kind of system could control it?" The answer will always be, "A control system could control it." And many other kinds of systems -- including all lineal systems -- could not.

> Note that if the period of a cycle is sufficiently rapid relative to the rate of change of inputs and outputs it appears continuous. This is why we are able to use digital computers for control.

Fine, especially for people who want to make a living using digital computers for achieving control. But that does not necessarily have anything to do with how living systems achieve control -- in fact, it probably has _nothing_ to do with how living systems control.

> Tracking a cursor, as you eloquently explained below, certainly is a situation where continuous negative feedback loop applies. You are in fact continuously controlling something. It is not clear in my mind how or why the concept of triggering a learned action pattern that goes away when the task is completed is inappropriate.

Read "Models and their worlds," by Bill P and me. If after that you still think a triggered action pattern can explain tracking, I'll be pretty disappointed with our writing skills. (I'll stop short of saying, "I'll eat my control handles.")

> Perhaps an extension of the 'little man' problem to two little men playing the computer game 'pong' would be enlightening. The problem is probably my lack of understanding of HPCT.

I don't have exactly that program, but I do have a lot of programs with two people, two hands, two models, and a model and a person, all interacting in about five or six identifiable kinds of interactions. And I'm still overcoming my profound ignorance about the amazing phenomenon of control.

>> I think these two replies (to Bill and me) are the clearest indications yet that you are talking about a different kind of system from us. A system, incidentally, that is very similar to the ones described by Albus -- you must think of him as more than just someone to use as raw meat to cast before the PCT modelers. ;-)) You seem to envision a "system" as a kind of "meta-assemblage" of independent systems, each of which accomplishes an assigned task, then passes off the result to the next system(s) in the chain (or net) which uses the received result as fodder for its own processes, and so on.

> Sorry for the confusion,

No apology needed. I'm just trying to understand your point of view -- and you certainly are up against it, trying to understand ours.

Time to shut down and vanish for a protracted weekend, marrying off the "baby" in the family. If this post is filled with typos, I apologize -- it was done on the fly.

I can't imagine how I'll even read all of the mail I know will come in by Monday!

> Once more around the loop :-)

Keep chomping on that bandwidth!

Later, Tom

Date: Fri Jul 15, 1994 12:19 am PST
Subject: Re: Replies to Paul and Jeff

<[Bill Leach 940714.20:27 EST(EDT)] >Tom Bourbon [940714.0945]

> I didn't say what I wanted to say, or if I did it wasn't very clear.

Not a problem. Indeed, I am almost aghast at the grammar and sometimes incomplete sentences that I seem to be using of late. I think that there is still a serious "internal control conflict" present here. I have may things that I should be doing rather than participating so heavily in the CSG-L (of course, I would rather participate than not...).

Yes, I accept the idea that pulse rate is probably continuously variable over some range that represents the range of some signal and that such signals are really analog not digital.

All the best for the wedding and see you when you "recover". :-)

-bill

Date: Fri Jul 15, 1994 12:55 am PST
Subject: Re: questions

[From Bill Powers (940714.1645 MDT)] Paul George (940714.1140)

I'm glad Rick asked about E, what we call "the disturbance." The disturbance represents the class of independent environmental variables that can act directly on the quantity you're trying to control, independently of the outputs of your control system. Like someone pushing on your elbow while you're aiming a pistol.

Your reply to Rick didn't clear up the point:

> It is some facet of the process that we cannot directly sense, but must rather infer via F.

This would partly fit, except that in general there is no way to work backward from the state of F (actually, H) to the cause of a disturbance, that is, to E. All the control system knows is that H changed when the reference signal A did not change. In a properly designed control system, the change will be strongly resisted because of the closed-loop action. It isn't necessary to know E. But a complete model of a control process has to include E; otherwise you'll be designing for a disturbance-free environment instead of a real one.

> I know of no formal name for it.

Don't tell me I have invented something. Actually, I doubt it. I learned about disturbances, after all, from reading control-system texts. Maybe the subject has been dropped since the 1950s. If so, I might be dubious about trusting my safety to any control system designed since then.

One of the sources of misunderstanding we have had with control engineers (especially concerning "open-loop control") is that many of them seem to assume, in their designs, that all sources of disturbance can be accounted for, so they can be incorporated into a "world model." But that is a very unrealistic approach to designing control systems, unless they are so simple and transparent that you can anticipate everything the environment might do to them. Well, that point aside, it is certainly not true of organisms that they can identify all possible causes of disturbances and prepare to meet them. MOST disturbances are known only in terms of an unexpected and persistent change in the controlled variable. It takes a closed-loop system to maintain control when that happens.

So am I to take it that in your process-control models, you do NOT make an explicit provision for unpredictable disturbances to affect the controlled variable?

> Since we often cannot sense the perception we really want to control, we often have to provide an F5 and/or F3 inverse within F4 or F1 (or more usually between, this is the 'world model' or 'mirrored object'.

We never assume that either F3 or F5 (function linking a disturbing variable to the controlled variable F) is known to the control system; the organic system has to be able to maintain good control without knowing the form of either function.

It sounds odd to me to say that "we often cannot sense the perception we really want to control." But then I realize that you mean we, the engineers, don't have available a sensor that can be put into the control system to let IT sense the variable we, the engineers, really want to control.

I should think, though, that it would often be possible to sense component variables and compute the state of the variable you really want to control. Isn't that one solution? That counts as sensing it, in PCT.

This is not a problem in modeling organisms. There is no engineer. If the organism doesn't sense a variable, it can't be controlled by that organism.

> By convention and practice the system is seen as controlling D in order to provide control of F (the 'process centric view').

Maybe I'll accept that. But I'd have to believe that the conventional view is a hierarchical control view, which would surprise me. A more accurate way of expressing this idea as it is usually carried out is to say that the control system VARIES D in order to control F. To control something means to bring it to a reference state and keep it there, in my dictionary. If you're driving a car, you can't simultaneously control the steering wheel angle and control the car's position on the road. Not, that is, unless you mean this hierarchically: you send a VARYING reference signal to the wheel-angle control system, which controls wheel angle to make it match the reference signal, and thus you control the position of the car.

The reason you can't speak of controlling D is that D is determined just as much by disturbances E as it is by the reference signal A. You can't bring D to a predetermined state, because that may be the wrong state for counteracting a disturbance. The output action of the control system must NOT be controlled in some preferred state; it has to be free to vary as required to counteract the effects of disturbances. Just imagine trying to drive while holding the steering wheel at your favorite angle.

You can see why disturbances are considered important in PCT.

> Recall that in our system we are often triggering an industrial machine (often with control capabilities) that may do a fairly complex series of things to the process under control. They respond to a fixed set of signal's that are usually interpreted as commands or instructions.

"Our system?" Are you speaking of the organism, or about the process control systems you-all design? If you're speaking about the human system, I would dispute the claim that ANYTHING is done open-loop. There are always disturbances. But in industrial control systems, the system design is whatever the engineers decide it should be. If they're confident that the process will always produce exactly the result that the command specifies, without any feedback to check that what happened is what was wanted, and without any means of altering the command if the process strays off track, then they are braver souls than I am. But it's their necks.

Don't tell me. Management says that sensors are too expensive.

I notice that NASA systems are quite variable in this regard. Sometimes they are fully fed-back so that the result of every command, even every switch-transition, is immediately observed, and the command can be changed or canceled if the wrong result even starts to occur (or if the contacts don't close). Others are designed by engineers who seem never to have heard the term "disturbance." When they send commands to those systems, Mission Control sends an astronaut to stand by and observe the actual result, and report it back immediately. Good thing, too.

> Sometimes B may be a program which is downloaded to replace an F3. Similarly a F5 could represent a machine not under the direct control of the control system (I know the phraseology flies in the face of the PCT concept of what is controlled).

I don't care about phraseology as long as I can figure out what you're talking about. I suppose that a control system could be built that senses which program is running, which program should be running, and if there is an error downloads a replacement program to a lower-level system. So far we haven't done any experiments complicated enough to call for that model. But there's a place reserved for it, at the program and principle levels in HPCT.

Actually, since engineers aren't constrained by any evolutionary or survival considerations, they can design systems any way they want, as long as they do the job at hand. They aren't designing general-purpose control systems like organisms. They don't have to worry about ad-hoc patches to an existing design in terms of how it is going to affect the whole system in a different situation. If it works, screw the cover down and ship it.

I've been intermittently seeing the TEMPUS downlink video from STS-65. This experiment includes a levitation device for "containerless processing" of small spherical alloy samples. Maybe there's some basic problem about stabilizing the sample spheres, but after five days of watching I'm going crazy from seeing that sample jiggling and bouncing around with every little disturbance from the spacecraft, or from heating or from evaporation or from who knows what. Jiggle, jiggle, jiggle. At least three samples that I know of have hit the cage while molten and stuck to it, ruined. And then the PI says "Thanks, Columbia, it's nice and stable now." Jiggle, jiggle, jiggle. Jiggle, jiggle, jiggle. I'm probably doing the designer a terrible injustice; the control problem might be all but unsolvable. But I want to grab the phone and yell at the PI, "For God's sake, haven't you ever seen a REAL control system?" I have enough hubris to think that if someone asked me to stabilize a levitated sphere, they'd think it was nailed in place, like the parrot's feet. Sometimes being old and retired is hell.

Best, Bill P.

Date: Fri Jul 15, 1994 3:19 am PST
Subject: Re: Diagram Terms o

<[Bill Leach 940714.20:34 EST(EDT)] >>[Paul George 940714 15:00]

Indeed, Paul, as you have mentioned (and others before you, including myself), how you choose to view a particular control system example influences how you might "parse" the parts.

In the simplest biological control loops, it appears that a reference signal can only have magnitude (including zero) but may not change signs or "go negative". However, it also appears that there are "complementary control loops" where such bidirectional control is necessary (most all physical movement can be viewed that way but I am not sure that even that view is really correct).

> Slowing a car and slamming on the brakes can be a considered a different action, rather than just a difference in the gain of F2.

Yes they could be but probably are not. Then again... there might actually be two types of situations. I am thinking of the person that appears to really have two modes of operation of the brakes; Variable pressure control to slow the vehicle at some controlled rate is one and the other is "panic" mode where they just "press for all that they are worth" (hopefully at least on the brake but often not!).

-bill

Date: Fri Jul 15, 1994 8:15 am PST
Subject: Re: Behavior and control

[From Rick Marken (940715.0730)] Tom Bourbon (940714.1740)

Wow! What a nice post, Tom. You've managed to say, very clearly and concisely, what I've been struggling to say all week. Here are some of my favorites:

> Most people, including most behavioral scientists, use the word "behavior" to refer to their own observations of what another creature is doing -- he is walking, she is talking, they are assembling, the rats ran along that path to a new source of food, and so on. In that usage, behavior is interpreted in terms of the observer's perceptual units and the labels are really names for _results_ the observer notices when the other creature acts. In that usage, actions and behavior are often equated or used interchangeably.

> In PCT modeling, we have found it useful, often necessary, to think of the behavior of a control system as what it is doing _from its own perspective_...It does what it is doing (which is controlling a particular perception in a desired way) by acting on the environment; the actions by which it achieves control of perception are not "what it is doing." The system's actions must be "out of control," in that the must vary any way necessary (for example, due to environmental disturbances) in order for the system to do what it is really doing -- controlling a particular perception or set of perceptions.

> In the PCT interpretation, many, if not all, of the things an observer sees a system "doing" may well be outward appearances that are irrelevant to the observed control system -- in most cases, what an observer sees does not even exist for the observed control system -- it does not know that it is seen as doing what the observer sees.

> This difference in "points of view" concerning what the observed system is doing is behind many misunderstandings between people.

And the grand finale:

> The difference in our interpretation comes from our knowing that the actions of the person are not "what the person is doing;" the person is "doing" his or her own controlled perceptions.

Now THAT'S PCT!!!

You've done well, Tom. Now you can go off and enjoy your daughter's wedding; and don't worry about how much it costs; the grants will be pouring in now that you've explained with crystal clarity what PCT is about.

Mazel Tov Rick

Date: Fri Jul 15, 1994 11:11 am PST
Subject: Gestalt

[Paul George 940715 10:30] >From Tom Bourbon [940714.1750]

Me:

>> Perhaps you should give me your definition of 'gestalt'.

> Ah, it was _you_ who introduced the term into the discussion. I've called _your_ hand. Show your cards! :-)

I've been trying. And from your discussion I think we are saying much the same thing. One more time.

To perceive something which exists in the environment, one needs a set of inputs. The blind men couldn't correctly perceive the elephant because they didn't touch enough of it and couldn't see the whole. Similarly, I can't clap with only one hand. It isn't the number of components (def differs between examples) that matters, it's just that if you don't have the complete set, nothing happens. At some point you have enough knowledge, or control loops, or perception, or instrumentality to affect the environment in a way that allows you to control your perceptions effectively (hope that was bad enough wording :-). Adding more may increase your level (effectiveness) of control, or may not. But less means you can't control at all. At some point there is a critical mass.

A Baby can't act very successfully because it hasn't learned enough about what the sensory stimuli 'means'. I can't form complete sentences until I know enough words and grammar. Children don't comprehend conservation of area or volume until a certain age (exactly why is not apparently known). We may presume their HPC network is not sufficiently complex. In general there is some correlation between the number of neurons & connections and 'levels of intelligence' (speaking of fuzzy terms). I don't think it is a smooth curve (perhaps Rick can tell us).

Hope these examples clarify what I am talking about.

Paul

Date: Fri Jul 15, 1994 11:13 am PST
Subject: Chapman and Agre

[Paul George 940715 10:05]

Ok, I just gotta ask. {Putting on asbestos suit to avoid fire & brimstone}

What is it that the heretics :-) Chapman and Agre say that differs from 'pure' PCT? They seem to catch a lot of flack in this forum.

Date: Fri Jul 15, 1994 12:39 pm PST
Subject: Re: Behavior and control

[Paul George 940715 09:30] >Tom Bourbon [940714.1740]

Me:

>> You also seem to be saying that behavior directly causes other behavior. This makes sense in the context of muscles (unobservable) moving an arm (observable) or a arm motion propelling a ball. Is that all you meant?

> I don't think I was saying exactly that, but it is true that the actions of a hierarchical control system are "nested" in something like the manner you describe. (As an aside, are you perhaps alluding to the old "behavior cannot cause behavior" song, from radical behaviorism?)

I have no problem with behavior causing behavior (and don't remember the song). I just wondered if you intended something more than the simple interpretation. However, given your definition of behavior above (what the observer perceives) I'm no longer sure I understood what you meant.

Have fun at the wedding (Free at last! Free at last.... :-)

Paul

Date: Fri Jul 15, 1994 12:39 pm PST
Subject: Re: Diagram Terms o

[Paul George 940715 12:00] >[Bill Leach 940714.20:34)]

> Then again... there might actually be two types of situations. I am thinking of the person that appears to really have two modes of operation of the brakes; Variable pressure control to slow the vehicle at some controlled rate is one and the other is "panic" mode where they just

"press for all that they are worth" (hopefully at least on the brake but often not!).

That is the situation I had in mind. People often react very differently to gaining on another vehicle, and seeing a truck stopped in front of them.

Paul

Date: Fri Jul 15, 1994 2:27 pm PST
Subject: Misc

[From Dag Forssell (940715 0930)] >[Paul George 940708 15:00]

> Looking forward to next week.

It has been a good week. A handful more from BPR-L are lurking.

Oldtimers on CSG-L are committed to PCT and try to be very careful in their answers, no matter who asks; no matter what the question, because there are many lurking, watching the discussion and learning from it. I learn too. I think it has been abundantly clear from discussions this past week that PCT is a serious physical science, not an "off the top of your head" word model.

--------------------------------

I have found the discussion on FACTS very useful. Thanks, Dan.

Best, Dag

Date: Fri Jul 15, 1994 2:27 pm PST
Subject: Re: questions

[Paul George 940715 11:45] >[Bill Powers (940714.1645 MDT)]

> Don't tell me I have invented something. Actually, I doubt it. I learned about disturbances, after all, from reading control-system texts. Maybe the subject has been dropped since the 1950s. If so, I might be dubious about trusting my safety to any control system designed since then.

There probably is a formal name, the problem is that I am a software engineer by training, not a control engineer. I work on development processes and techniques not process control software. Unfortunately, most people developing control systems these days aren't trained in control theory. My company used to have a kind of apprentice program for all engineers for that purpose, but it was dropped about 15 years ago. Fear and trembling is appropriate. Read the RISKS forum (ACM SIGSOFT Transactions or comp.risks. is also a mailing list if you are interested).

> One of the sources of misunderstanding we have had with control engineers (especially concerning "open-loop control") is that many of them seem to assume, in their designs, that all sources of disturbance can be accounted for, so they can be incorporated into a "world model."... So am I to take it that in your process-control models, you do NOT make an explicit provision for unpredictable disturbances to affect the controlled variable?

Yup. We are forced to try to anticipate all possible problems and deal with them. In fact for safety critical systems we are legally required to. That is why requirements analysis and design are so difficult. If something unanticipated does occur and something bad happens, we can be held legally accountable if we 'knew or should have known' that it could occur. The big problem is that if an E (disturbance) occurs we usually must have a D to affect it (making it another F). Disturbances (E) which may be corrected via D are not a problem. Our task is to insure we actually have all the Fs and D's that are required to keep the process under control. Fortunately the task is eased by having a human operator as part of the System. They are a bit more flexible and don't just 'fly IFR' (instrument flight rules).

> But then I realize that you mean we, the engineers, don't have available a sensor that can be put into the control system to let IT sense the variable we, the engineers, really want to control.

Or more likely, the customer wouldn't pay for it, or there is no place to put it. Sensors and wiring are expensive and kept to a minimum. One of the reasons that continuous control isn't used is that message passing networks are used for communication between controllers and instruments to minimize cabling. Point to point wire connections, nerve style, would be prohibitively expensive. There are also transmission problems with analog control signals and long cables. Practicality raises its ugly head.

>> By convention and practice the system is seen as controlling D in order to provide control of F (the 'process centric view').

> Maybe I'll accept that. But I'd have to believe that the conventional view is a hierarchical control view, which would surprise me.

These days it is, at least in companies who build plant control systems. Control engineers working at the PLC loop level is probably what you are familiar with. Hierarchical systems emerged in the 70's. The car example is exactly how we usually design it.

> "Our system?" Are you speaking of the organism, or about the process control systems you-all design?

My whole post dealt with process control systems, or more accurately hierarchical distributed plant control systems. Organisms work differently, though system architectures are patterned on various understandings of how organisms function.

> Don't tell me Management says that sensors are too expensive. I notice that NASA systems are quite variable in this regard.

NASA doesn't _sell_ control systems. Their contractors (system developers) are in the business of selling engineering hours. In the industrial world, system proposals are evaluated mostly on price. The number of I/O points and cable lengths are the major cost drivers.

> I suppose that a control system could be built that senses which program is running, which program should be running, and if there is an error downloads a replacement program to a lower-level system.

That is actually a normal situation. If a processor crashes, you often need to reload the software. But I meant something more akin to learning. You effectively replace the old control node (or HPCS) with a new one that is better for handling the current situation, or is more efficient (say a different F2 or F4). In addition, since (due to bad design) the H values are often hard coded into the software, to change it you need to download a new version.

Later Paul

Date: Fri Jul 15, 1994 5:38 pm PST
Subject: Chapman and Agre

[From Bill Powers (940715.1350 MDT)] Paul George (940715.1005)

> What is it that the heretics :-) Chapman and Agre say that differs from 'pure' PCT? They seem to catch a lot of flack in this forum.

Just from me, and you can put it down to prejudice. They seem to make up a lot of peculiar terms for what turn out to be simple things, like their proposal that behavior is "situated." This means that behavior takes place in an environment, as near as I can figure. For AI types this may be a major discovery, but not for anyone else. They're probably doing no harm; I just don't go in for high-falutin' generalizations. There are lots of true statements one can make, the only polite response to which is "My, that _is_ interesting." I have a feeling that they belabor me with earnest urgings to grasp simple facts, as if they see something in them of, unfortunately, inexpressible importance.

So now you know that I'm not very impressed by their work. Won't I look silly when they turn out to be right?

Best, Bill P.

Date: Fri Jul 15, 1994 6:53 pm PST
Subject: Re: We have a winner!

[Paul George 940715 10:00] >[Rick Marken (940714.1445)]

Me:

>> By convention and practice the system is seen as controlling D in order to provide control of F

> Oops :-( The only variable in the loop that is controlled (kept at a specified level against disturbance) is H.

Ack. I just meant that that is the way process engineers usually talk about it. And while PCT usually just works within the context of the control loop, In process control, the Instruments and the equipment performing the process are viewed as part of the 'System'. The system boundary is in a slightly different place. Also from a terminology standpoint (backward from PCT) the thing acted upon is seen as being controlled, the sensors just tell you what happened (you hope :-) {sound familiar?}.

>> However, the control system writ small (essentially F1) only cares about H & B.

> Well, it really only cares about (that is, controls) H; the system cares what value H is -- it "wants" H to equal A and it will do what it can to make that happen. The system doesn't care what value B is; B just varies around as necessary (depending on disturbances -- E -- and changes in functional relationships in the loop) to keep H matching A.

I understand that is the way it is usually modeled. But is it necessarily impossible that B is also 'perceived'? To anthromorphisize a bit, If I push on a door lightly and don't perceive it moving, I then push harder. To increase the force I must have some idea how hard I was pushing originally. Now, in this particular case I have sensors that tell me how hard my hand is pressing, or how hard my muscles are contracting. But the concept of direct feedback when I can't detect (indirectly) C or D doesn't seem invalid on the face. The signal exists, and the control loop could 'wire itself' that way.

Date: Fri Jul 15, 1994 6:54 pm PST
Subject: Heretics, etc

[From Rick Marken (940715.1300)] Paul George (940715 10:05)

> What is it that the heretics :-) Chapman and Agre say that differs from 'pure' PCT?

Chapman and Agre are not heretics because they were never in the cult (of PCT that is). Like most behavioral scientists I imagine that they are either blissfully unaware of PCT or have no idea what it is about.

> They seem to catch a lot of flack in this forum.

No more than any other psychologists who are clueless about the nature of control- - ie. all of them. They were once mentioned (long ago) as people whose work was (unbeknownst to us or them) compatible with PCT because they talk about "situated action" and seem to be aware of the fact that actions have to adjust to environmental circumstances (disturbances) if people are to achieve their ends. But, like all other psychologists, they have no idea that this means that it is perception, not observable action or behavior, that is "achieved". So they are busy building superficial models of superficial behavior. And, of course, ignoring PCT.

The only PCT "heretics" I know of are Carver and Scheier and Hyland (there are a few other names as well but I can't remember them). These people are "heretics" only because they talk in terms of Powers' PCT model and refer to Powers a lot but they are applying the PCT model to the wrong phenomenon -- to observable actions and behavior (which are usually, as Tom noted, irrelevant side effects of control) rather than to control.

Paul George (940715 10:30)--

> In general there is some correlation between the number of neurons & connections and 'levels of intelligence' (speaking of fuzzy terms). I don't think it is a smooth curve (perhaps Rick can tell us).

I'm not a physiologist; I study control at the same levels as Tom and Bill P. ie. all levels. But I will say that I don't think that a correlation that is less that .99 is of much interest in PCT. What's your guess at the correlation between number of neurons & connections and 'levels of intelligence'?

Rick

Date: Fri Jul 15, 1994 7:32 pm PST
Subject: Control of output? No such thing

[From Rick Marken (940715.1530)]

Me

>> Well, it really only cares about (that is, controls) H; the system cares what value H is -- it "wants" H to equal A and it will do what it can to make that happen. The system doesn't care what value B is; B just varies around as necessary (depending on disturbances -- E -- and changes in functional relationships in the loop) to keep H matching A.

Paul George (940715 10:00) --

> I understand that is the way it is usually modeled. But is it necessarily impossible that B is also 'perceived'?

Why don't we go back to PCT terminology now. B is the error signal in a control loop. It would make no sense to have it be perceived by the control loop itself (perhaps re-entering the loop through the perceptual function - - F4 in Bill's diagram, or added directly to the perceptual signal -- H in the diagram -- I don't know what you had in mind) though the error signal could become the perceptual input to another control loop.

> To anthromorphisize a bit, If I push on a door lightly and don't perceive it moving, I then push harder.

Right.

> To increase the force I must have some idea how hard I was pushing originally.

Not at all, if by "some idea" you mean that the control system must have or would be helped by a perceptual representation of the force that is being exerted on the door (variables C or D in the diagram). All the control loop needs is a perceptual representation of the variable under control -- the position of the door in this case -- which is variable F in Bill's diagram; it is called the controlled variable in PCT.

> Now, in this particular case I have sensors that tell me how hard my hand is pressing, or how hard my muscles are contracting.

You only need those sensors if you are controlling those variables. It looks like you've let psychologists sell you a bill of goods about how control systems work. A control system only needs to perceive the variable it is controlling; it does not need to (and typically CANNOT) perceive its own actions (outputs) , the actions of the environment (disturbances) or how those actions (outputs and disturbances) affect the controlled variable.

> But the concept of direct feedback when I can't detect (indirectly) C or D doesn't seem invalid on the face.

Just one of the many surprising results of analyzing the operation of a control system WITHOUT PRECONCEPTIONS;-)

A simple control loop can detect neither its own output (C in Bill's diagram) or the effect of this output on the controlled variable; nor can it detect any disturbance(s), E, or the effect of any disturbance on the controlled variable.

Is PCT starting to smell a little revolutionary yet?

> The signal exists, and the control loop could 'wire itself' that way.

But there would be absolutely no reason for a control system to wire itself "that way". I can't think of any way you could add a "perception" of the system's own error signal or output that wouldn't be either superfluous or a major hinderance. Maybe you could draw a diagram of what you had in mind?

Best Rick

Date: Fri Jul 15, 1994 9:02 pm PST
Subject: Re: Diagram Terms o

<[Bill Leach 940715.22:32 EST(EDT)] >[Paul George 940715 12:00]

> That is the situation I had in mind. People often react very differently to gaining on another vehicle, and seeing a truck ...

Yes but the difficulty is that this is something that we often observe and directly tells us nothing about what the person is controlling. We can probably surmise that the control system is either making a discontinuous change in reference signals or a perception involved in a loop with exceptionally high gain suddenly deviated from its' reference.

Even when just thinking rather subjectively about such a situation one can easily come up with between a few dozen and a few hundred possible controlled perceptions nearly all of which can influence the act of driving a car.

I can't help but think that when driving a car (as with many other situations), there are possible perceptions associated with continued existence, lack of injury or pain, etc. that normally are quite well within their respective control references. Coming around a curve at 60 MPH and seeing a truck across the road combined with the perceptions that 1) the car will not stop before reaching the location of the stopped truck, 2) I have seen many mangled cars, 3) I have seen accident victims, 4) I have seen and known what WERE accident victims but now they are classified as dead, etc., is likely to provide a disturbance to the perceptions that, among other things include, the idea that I will enjoy the rest of the day.

I suspect that an error signal caused by a perception that maybe I am about to die, is likely to cause related control systems to engage in most vigorous control activity. I do not think that "learned behavior" (experience) necessarily "goes out the window" but many people have had little experience in dealing with life threatening situations. In such cases, their "behavior" probably should be similar to the infant.

-bill

Date: Sat Jul 16, 1994 2:37 pm PST
Subject: Disturbances

[From Bill Powers (940717.0620 MDT)] Paul George (940715.1145)

Your description of the constraints under which process control engineers work makes the whole thing sound either (a) challenging or (b) infuriating. It's probably both.

On disturbances:

I'm very puzzled about your remarks concerning disturbances of controlled variables. We don't seem to be talking about quite the same thing. I said

>> So am I to take it that in your process-control models, you do NOT make an explicit provision for unpredictable disturbances to affect the controlled variable?

And you replied,

> Yup. We are forced to try to anticipate all possible problems and deal with them. In fact for safety critical systems we are legally required to. That is why requirements analysis and design are so difficult. If something unanticipated does occur and something bad happens, we can be held legally accountable if we 'knew or should have known' that it could occur.

By "disturbance" I don't refer to "problems," but just normal effects that the environment would have on controlled variables if there were no control. In our Figure, the disturbance is shown explicitly to allow for such effects. I'm hesitating here, because if you already know what I'm about to say, this whole comment will seem silly and patronizing, but if you don't it's very important to explain. I can't tell from your words which is the case.

Suppose you're controlling temperature (qi) of a sample of something. This is done by sensing (Fi) the temperature to produce a temperature signal (p), comparing (C) that signal with a reference-signal (r), and converting the error (e) via an amplifier (Fo) to send power output (qo) to the heater coil. The output of the heater coil affects the sample according to the sample's heat capacity and mass (Ff includes the heater coil and thermal constants of the sample). Also affecting the temperature of the sample is the temperature of the surroundings (d), which drains heat from the sample according to the conductivity, radiation properties, and convective properties (Fd) of the intervening medium; this affects the sample temperature according to the same thermal properties of the sample.

So we have identified all the variables, signals, and functions in our standard Figure.

OK, the control system will vary qo (power to the heater) until the sensed temperature matches the reference temperature. At this point there is an actively-maintained balance between the heat input from qo and heat losses to the sample's surroundings. This is a continuous control system, not a bang-bang home thermostat.

Note that we can't even make this system work without the disturbance. The disturbance in this case consists of heat losses to the surroundings of the sample. Without those, the control system couldn't cool the sample when it is too hot; it can only heat the sample.

The disturbance d is, in general, a variable; the temperature of the sample's surroundings can vary unpredictably. But we don't need to predict those variations or even know what is causing them, because we have a temperature control system that acts directly on the sample temperature. If the surroundings warm up, the resulting error will immediately reduce the heater output, preventing the temperature from changing more than a minute amount.

What you call a "problem" would be, for example, a situation in which the surroundings came to a temperature higher than the set-point of the temperature controller. The control system is unable to handle this situation. If that is at all likely to happen, the control system would have to be redesigned with an output function capable of heating or cooling the sample according to the sign of the error signal. Then the only "problem" that could occur would be for the ambient temperature to become so extreme that the maximum heating or cooling capacity was exceeded.

A "problem," as I'm defining it here, is a situation that goes outside the range of control of the control system. All closed-loop control systems are limited as to the amount of output they can produce and the speed with which they can vary the output between maximum and minimum. That defines the universe of disturbances that the control system can handle automatically, simply out of its basic design. Any disturbance that goes outside that space-time envelope constitutes a "problem," because control will fail as long as the disturbance remains outside the envelope.

A "disturbance" is not a problem if it does not call for more output or faster changes in output than the control system can produce. When I speak of disturbances, I'm talking about normal operation of the system. That's how I define a "normal" disturbance: one that does not exceed the capacities of the control system.

It's not necessary to understand what is causing normal disturbances. Neither we nor the control system needs to know that. If the comparator generates an error signal for ANY reason, whether an external disturbance or a change in the reference signal, the control loop will immediately correct it by matching the perceptual signal to the reference signal (language requires us to describe this process sequentially; it is really a simultaneous rebalancing of all variables in the loop).

So: where do we stand on the subject of disturbances now?

Best, Bill P.

Date: Sun Jul 17, 1994 9:20 am PST
Subject: Re: Control of output? No such thing

<[Bill Leach 940717.11:27 EST(EDT)] >[Rick Marken (940715.1530)]

> ... All the control loop needs is a perceptual representation of the variable under control -- the position of the door in this case -- which is variable F in Bill's diagram; it is called the controlled variable in PCT.

Rick; I think that this statement could lead to some confusion. What you are saying is (I believe absolutely correct) but it is predicated upon the simple idea that "position of door" is all that is being controlled and I suspect such is rarely the case.

For various different reasons, we likely DO perceive the force that we apply to operate a door and have an active control loop for just that perception. It is probable that the reference is a moderately loose reference (that is very low gain until the applied force rises to near some reference limit).

I suspect that the reference is set based upon experience with similar (or previous, if the same door) attempts to open the door.

This whole thing can get so complex when used as an example even though the example appears so simple. For example, when one pushes upon a door to have the perception for the door match the reference (in this case door open) and the door does not move there are literally a multitude of perceptions that might be affected.

One perception is that of "balance" as a result that the applied force will as per Newton's laws be equal in both directions and since the door did not move...

Now then, while I might appear to be "challenging" what you are saying, I don't believe that I really am. This "knowing the force applied" has nothing to do with the actual perception "door open" that is desired (except that the lower level control loops associated with actually operating the muscles deal with such matters of course) but rather it is associated with a whole host of other perceptions that are also under control. This would include even such perceptions that one could break something if too much force is applied (say like opening the plastic door on a piece of equipment).

OTOH, this does appear to me to be an excellent discussion of the concept of the TEST. A control system DOES NOT need to know the force exerted to control the perception "door open". If one observes someone "opening a door"; 1- the door does not open and 2- the person does not exert the maximum physical force that they can to achieve the an "open door condition" then (as I understand it) the TEST tells us that there are other perceptions under control here that are related to the one that we are testing.

-bill

Date: Sun Jul 17, 1994 11:10 am PST
Subject: Mary on B:CP

from Mary Powers 940717 To Bill Leach:

Despite Tom Bourbon's statement, that behavior is the means by which perceptions are controlled, you still, 940715, want to maintain your position that "behavior results from the control of perception".

I hope this is a semantic rather than a conceptual difference. But from the PCT point of view, once again, behavior is not a result, it is a means. Behavior results _in_ perceptual control, and is not a consequence of it.

I can appreciate your desire to take as a starting point something that is a little easier to understand for those who are new to PCT. But I believe your formulation is way, way down a slippery slope. It is the formulation that Carver & Scheier and other "self-regulation" psychologists are promulgating, and it misses the point entirely.

If behavior results from anything, it is from the discrepancy between perceived and desired/intended states. That is the part of the loop immediately preceding action. However, even this is probably an excessively lineal and sequential way of looking at a process which is essentially simultaneous all around the loop.

There really are no baby steps to take between behavior as outcome, consequence or result, and behavior as the control of perception. One of the big difficulties PCT has in making its way in the world is that you can't shape the understanding of it incrementally. Either you force PCT data to fit your world view, or you make the jump, and the world never looks quite the same again. Nothing in between.

Mary P.

Date: Sun Jul 17, 1994 12:26 pm PST
Subject: Re: Mary on B:CP

<[Bill Leach 940717.15:42 EST(EDT)] >Mary Powers 940717

> There really are no baby steps to take between behavior as outcome, consequence or result, and behavior as the control of perception.

I am beginning to believe the truth of that statement. It seems to me that most of the differences in opinion between you, Bill, Rick, Tom and virtually anyone else (where such difference occur of course) rest squarely upon a disagreement with the truth of the following:

Behavior as a function of some environmental condition can not exist at all unless there is both a perception for that environmental condition and a reference FOR THAT PERCEPTION. This is likely also true for "unobservable" behavior such as thinking.

No change in behavior can occur unless either the perception changes or the reference changes.

Now there are other matters such as reorganization are maybe even physical damage to the control system but such matters are usually mentioned in any post where applicable.

I admit to being a bit "dense" now and then or should that be... all the time? :-) . . . I realize that you did not raise this question but only my assertion that "behavior results from the control of perception" and I guess that I am trying to say that I am beginning to see were such "minor" differences are not so minor after all... they are rather the root of misunderstanding.

-bill

Date: Mon Jul 18, 1994 10:21 am PST
Subject: Re: Disturbances

[Paul George 940718 0930] [From Bill Powers (940717.0620 MDT)]

> So: where do we stand on the subject of disturbances now?

Just fine. I share your distinction between 'problem' and 'disturbance'.

The engineer is usually interested in 1) keeping some controlled variable within a range (normal operation), 2) designing the system so that it has the necessary sensors and instrumentality to keep the process under control, and 3) anticipating and detecting any circumstance that indicate the process is getting out of control (i.e. a defect in 2). The latter is why control systems have 'alarms' to detect potentially dangerous situations in areas not under adaptive control. 'Alerting' (ducking to avoid bricks from Rick Marken) is a big part of industrial control systems, of course they are alerting the human part of the control system. The operator may then take mitigating action by opening a valve, shutting off power, calling for evacuating a tri state area :-}, etc.

Date: Mon Jul 18, 1994 10:22 am PST
Subject: Re: Heretics, etc

[Paul George 940718 10:00] >[Rick Marken (940715.1300)]

> What's your guess at the correlation between number of neurons & connections and 'levels of intelligence'?

It's not my field of expertise, so my opinion isn't worth very much. However, complexity of behavior in species appears to be something of a step function (possible multiple curves or sawteeth rather than right angel steps) associated with brain size. A certain critical mass seems to be needed to support a given level of complexity.

Date: Tue Jul 19, 1994 1:40 am PST
Subject: Re: Diagram Terms o

From Tom Bourbon [940718.1739]

Back from the wedding weekend with my credit cards melted down from overuse -- the true role of the father of the bride!

What a flood of mail on the net!

>[Paul George 940715 12:00]

>>[Bill Leach 940714.20:34 EST(EDT)]

>> Then again... there might actually be two types of situations. I am thinking of the person that appears to really have two modes of operation of the brakes; Variable pressure control to slow the vehicle at some controlled rate is one and the other is "panic" mode where they just "press for all that they are worth" (hopefully at least on the brake but often not!).

> That is the situation I had in mind. People often react very differently to gaining on another vehicle, and seeing a truck stopped in front of them.

Why would those two scenarios call for two different "modes" of control? If perceived rate of closure is being controlled, with a reference value of zero, or something very small, then think of how the error signals would be changing in the two scenarios, prior to the person applying the brakes. Even if the person's gain did not change, or the "mode" of control remained the same, the error signal would be much bigger and growing more rapidly in the one scenario than in the other. Perhaps?

Later, Tom

Date: Tue Jul 19, 1994 1:51 am PST
Subject: Disturbing alerts

[From Rick Marken (940718.1430)] Paul George (940718 0930)

> The engineer is usually interested in 1) keeping some controlled variable within a range (normal operation), 2) designing the system so that it has the necessary sensors and instrumentality to keep the process under control, and 3) anticipating and detecting any circumstance that indicate the process is getting out of control (i.e. a defect in 2).

I don't understand point 3). Are you trying to anticipate deviations of a controlled variable from its reference state or are you trying to anticipate disturbances to the controlled variable?

> 'Alerting' (ducking to avoid bricks from Rick Marken) is a big part of industrial control systems

Satellite control systems, too. Of course, a signal is only "alerting" if the operator has learned to treat the signal that way; "alertingness" is determined by the operator, not the signal.

Rick

Date: Wed Jul 20, 1994 1:05 am PST
Subject: Re: Disturbing alerts

[Paul George 940719 11:30] >[Rick Marken (940718.1430)]

>> 3) anticipating and detecting any circumstance that indicate the process is getting out of control (i.e. a defect in 2).

> I don't understand point 3). Are you trying to anticipate deviations of a controlled variable from its reference state or are you trying to anticipate disturbances to the controlled variable?

I didn't phrase it very clearly. The system normally perceives more than it can affect. In design I must anticipate different kinds of disturbances that the system must detect in order to control the process (sorry;-). Sometimes that 'controlled variable' can be controlled (i.e. affected by an action of a control system) sometimes it can't. I must anticipate accident, ignorance, and malice. Ideally I will design the control system so that the process _cannot_ run away and no hazardous situation can occur. Otherwise I must detect the condition and inform the operator, who (hopefully) can do something to correct the situation, or prevent disaster, or at least mitigate the effects.

Kodan lieutenant: "We are trapped in the moon's gravitational pull, what do we do!!"

Kodan commander: "We Die."

{From the movie 'The Last Starfighter'}

Paul

PS: I finally got hold of B:CP, YEA!!

Date: Wed Jul 20, 1994 1:34 am PST
Subject: Re: Disturbing alerts

<[Bill Leach 940719.19:12 EST(EDT)] >[Rick Marken (940718.1430)]

>>Paul George (940718 0930) --

>> 3) anticipating and detecting any circumstance that indicate the process is getting out of control (i.e. a defect in 2).

>Rick:

> I don't understand point 3). Are you trying to anticipate deviations of a controlled variable from its reference state or are you trying to anticipate disturbances to the controlled variable?

Among other things, in engineered control systems, there are often environmental conditions whose status is quite important to the control process but is not controlled by the engineered system hardware.

An example that I can think of right off the top of my head would be a batch process tank that is manually made up by an operator. An engineered control system would often be designed to at least estimate when a new batch of chemicals might be needed.

A different sort of example is the situation where a sensor has a lifetime and it is difficult to determine the sensor's status. The control system then may attempt to create "lifetime statistics" for the sensor and even attempt to derive failure clues from experience with the actual operating process.

I would say that living control systems do the same sort of thing but 1) we are often probably not aware of it even when we are doing it ourselves and 2) we are so damned adaptive at times that even a careful observer might not notice a change that we make to continue to control a perception when something "goes wrong" with whatever method that we were using.

-bill

Date: Wed Jul 20, 1994 1:36 am PST
Subject: Cause systems

[From Rick Marken (940719.1400)] Paul George (940719 11:30) --

> In design I must anticipate different kinds of disturbances that the system must detect in order to control the process (sorry;-).

I think you're going to have to diagram one of these controllers of yours. In a real control system, there is no need to detect or anticipate disturbances. I thought that you agreed with Bill's definition of disturbances?

> Ideally I will design the control system so that the process _cannot_ run away and no hazardous situation can occur.

I think we need a diagram of what you call a control system. Could you give us a diagram of one of the "control systems" you build at work, identifying controlled variables, disturbances, etc.

I'm getting the distinct impression that what you call a "control system" is not a control system . A lot of people you the word "control" as a synonym for "cause"; could this be what's going on in "process control"?

Best Rick

Date: Wed Jul 20, 1994 1:38 am PST
Subject: Re: Gestalt

From Tom Bourbon [940715.1722]

I'm way behind with my replies, but here goes a start at trying to catch up.

>[Paul George 940715 10:30] >>Tom Bourbon [940714.1750]

>Paul:

>>> Perhaps you should give me your definition of 'gestalt'.

Tom:

>> Ah, it was _you_ who introduced the term into the discussion. I've called _your_ hand. Show your cards! :-)

Paul:

> I've been trying. And from your discussion I think we are saying much the same thing. One more time.

> To perceive something which exists in the environment, one needs a set of inputs. The blind men couldn't correctly perceive the elephant because they didn't touch enough of it and couldn't see the whole. Similarly, I can't clap with only one hand. It isn't the number of components (def differs between examples) that matters, it's just that if you don't have the complete set, nothing happens. At some point you have enough knowledge, or control loops, or perception, or instrumentality to affect the environment in a way that allows you to control your perceptions effectively (hope that was bad enough wording :-). Adding more may increase your level (effectiveness) of control, or may not. But less means you can't control at all. At some point there is a critical mass.

I Believe I'm beginning to see why you thought you were telling me what you mean by a "gestalt" and I didn't realize that was the case. I think of "gestalt" in terms of the historic Gestalt "school" in psychology. Its adherents urged that perceptual experience is _always_ "whole" and perception is different from elementary sensations -- an idea that in some inexplicable way became the textbook chestnut, "the whole is greater than the sum of the parts."

They also made a big thing of the idea that perception and neurological events are "isomorphic" and that both neural events and perception are entirely "contemporaneous:" they are simultaneous and in the immediate present -- nothing from the past or future can enter into present neurological and perceptual events. (I guess that would rule out "feedforward," wouldn't it? ;-)

In light of that tradition, the story of the blind men and the elephant is not a story about a failure of three people to have their individual "gestalts." On the traditional Gestalt reading, each of the blind men had the only perceptions he could have: those that were isomorphic with the organization and functional state of his nervous system, _whatever_ those might happen to be. Each blind man had "whole" perceptions -- "gestalts," if you like -- not incomplete ones. Only a differently organized person (say one who was not blind) could have whole perceptions that were different from those of each of the blind men. And who is to say the sighted person sees all there "really is?" Different organizations, different activity, different perceptions -- and all of them "whole."

> A Baby can't act very successfully because it hasn't learned enough about what the sensory stimuli 'means'.

And the organization isn't there to support "mature" actions; but babies function as whole systems, with whole perceptions and actions. (I like this more accurate reading of Gestalt theory -- it flies in the face of more familiar and popular ideas that infants are "incomplete" adults and that they must "develop" into adult "finished products.")

> I can't form complete sentences until I know enough words and grammar.

But you do whole things other than form complete sentences . . . And so on.

> . . . In general there is some correlation between the number of neurons & connections and 'levels of intelligence' (speaking of fuzzy terms).

Perhaps, in a very loose sense. That puts whales far ahead of us, I guess, and . . . Hmm. Maybe I won't even stand on my "perhaps." Whatever the case, the Gestaltists would have come back with the idea that no species is a partial realization of any other species; each is complete and its experiences and actions are "whole" for its particular organization.

> Hope these examples clarify what I am talking about.

I think they did. I hope my reply clarifies why your use of "gestalt" didn't look familiar to me. ;-)

Later, Tom

Date: Wed Jul 20, 1994 1:44 am PST
Subject: Modes of Control

[Paul George 940719 10:40] >Tom Bourbon [940718.1739]

> Why would those two scenarios call for two different "modes" of control? If perceived rate of closure is being controlled, with a reference value of zero, or something very small, then think of how the error signals would be changing in the two scenarios, prior to the person applying the brakes. Even if the person's gain did not change, or the "mode" of control remained the same, the error signal would be much bigger and growing more rapidly in the one scenario than in the other. Perhaps?

Perhaps. Sometimes the magnitude or rate of change of the error signal may be sufficient. I guess it depends how sophisticated you allow the output function (F2) to be. If it can select between different actions based upon thresholds for the error signal, fine. The idea is that different actions may be required in response to different amounts or rates of disturbance. To deal with a punch, sometimes you block, sometimes you duck. Depends on how fast it is coming, and where your arms are.

I think the more common structure might be having another control loop monitoring the error signal, treating it as a controlled variable. That loop's output function would then activate the appropriate loop for the needed kind of action (i.e an 'emergency' loop). This presupposes that you have multiple loops capable of controlling the same perception, but have their reference variables adjusted so only one actually generates an error signal at a given time. Possible?

Paul

Date: Wed Jul 20, 1994 5:12 pm PST
Subject: Process control systems

[Paul George 940720 11:00 ] >[Rick Marken (940719.1400)]

Me:

>> In design I must anticipate different kinds of disturbances that the system must detect in order to control the process (sorry;-).

>> Ideally I will design the control system so that the process _cannot_ run away and no hazardous situation can occur.

> I'm getting the distinct impression that what you call a "control system" is not a control system . A lot of people you the word "control" as a synonym for "cause"; could this be what's going on in "process control"?

I think this is just mis-communication. Having worked with satellite control systems at Loral Aerospace, they seem much the same as process control - from the control system's standpoint.

The concept is that when we _design_ a control system we must determine what needs to be controlled and what can go wrong. For a satellite, if we consider a loss of carrier or orientation to be significant, we had better have sensors to detect it. Further, it had best have some mechanism for station keeping and re-orientation. Then, we has better consider how the attitude control jets or controls could screw up (say lock on, vent the hydrazene, overheat fuel,.....) and design in preventive measures or sensors to detect the event.

Paul

Subject: Re: Disturbance Control, Political Steering

Paul George 940720 17:30

>[From Rick Marken (940720.0845)] >>Bill Leach (940719.19:12) --

>> Among other things, in engineered control systems, there are often environmental conditions whose status is quite important to the control process but is not controlled by the engineered system hardware.

> I'd say that that's true of ALL control systems; the environmental conditions of which you speak are called disturbances; they are the reason for building a control system in the first place -- to protect variables from the unpredictable effects of these conditions.

They are called disturbances _after_ the system is constructed. During analysis or design there are no controlled variables _yet_. That is the purpose of the design - to determine what input functions, output functions, and hierarchical levels of control are needed to meet the purpose of keeping the process stable and predictable.

> This estimate is not really part of the control process; it is just another variable in the environment

Here is the major terminology conflict. To us 'control system' is more than the controller. It includes the equipment, sensors, and multiple controllers. (so is a 'satellite control system') It even sometimes includes the operator via interaction with displays as well as instruments or controls not hooked to controller input or output functions. In biology one can distinguish between the nervous system and everything else. You could also draw a system boundary at the skin. A process control system is more like a work party than a person.

Paul

Date: Wed Jul 20, 1994 7:24 pm PST
Subject: Comments to Paul George (from Mary)

[from Mary Powers 940720 8:45] Paul George:

> Sometimes the magnitude or rate of change of the error signal may be sufficient. i guess it depends on how sophisticated you allow the output function (F2) to be. if it can select between different actions based upon thresholds for the error signal, fine. the idea is that different actions may be required in response to different amounts or rates of disturbance. to deal with a punch, sometimes you block, sometimes you duck. Depends on how fast it is coming, and where your arms are. I think the more common structure might be having another control loop monitoring the error signal, treating it as a controlled variable. that loop's output function would then activate the appropriate loop for the needed kind of action (i.e. an 'emergency' loop. this presupposes that you have multiple loops capable of controlling the same perception, but have their reference variables adjusted so only one actually generates an error signal at a given time. Possible?

Possible, maybe, if you are designing such a system. But remember PCT is under the constraint of being neurologically plausible, and also under the constraint of keeping things as simple as possible unless evidence demands otherwise. Therefore the output is not allowed to be very sophisticated.

Part of the problem here is that you are looking at the canonical, single loop model, and trying to see how it can do the work of what PCT conceives of as a hierarchical arrangement of a multitude of control systems, none of which is engaged in monitoring the size of the error signal (aside from the postulated reorganization function which has as input a state of large and chronic error).

In PCT, error signals are outputs to lower levels, where they function as reference signals. Only at the lowest level is output a signal for muscles to contract. There is no function selecting actions: the actions are dependent on numerous output signals to various muscles, which in turn are dependent on the comparison of the desired state of affairs to the perceived environmental situation. In the case of ducking or blocking a punch, the state of the environment is the perception of how fast the punch is coming, PLUS, as you correctly pointed out, where your arms are - which is as much a part of the environment of the system as the approaching punch. If your arms are in a particular place, you've got to duck because it's too late to block, etc.

In the heat of a fight, this is all going on at pretty low levels, and your success at blocking or ducking is not a high- level intellectual exercise - it's a matter of much practice and training to have lower-level loops - perceptions, comparison, outputs - be sufficient and correct; for the lower systems to duck rather than block because of experience with what can be done or not done when the arms are in various positions. The same applies to dealing with braking in emergencies. Hitting the brakes hard, which in normal situations is the right thing to do, is a poor idea on an icy slope, which is why incomers from Texas and California end up in ditches around here in the wintertime more often than the natives.

Mary P.

Date: Wed Jul 20, 1994 9:56 pm PST
Subject: B:CP impressions - 1st cut

[Paul George 940720 17:00]

Having read through the first 8 chapters, I begin to see the source of the misunderstanding of my posts, as most discussion seem to focus on the first three levels (orders) of the CNS. PCT does agree with my mental model, there were just some scope and terminology issues. All in all a great piece of work.

I guess it's lack of acceptance is due to perceptual filtering, or more likely that 'other schools' never actually _read_ (with understanding) the book. I found Bill Power's post on the e. coli paper review committee revealing as to the nature of 'peer review'. However you should be careful that you don't follow the same pattern in judging other's work. It is important to evaluate a piece of work using the _author's_ frame of reference, not your own. Translation is your job when you are not using terminology in a 'standard' way or when using 'non-standard' concepts.

May I suggest that the 'FAQ'/intro add a summary of the nature of neural currents, the types of neural circuits, and the distinction between first order and higher order systems. Since B:CP is relatively hard to find (lacking money to order it from the Powers), this would avoid a lot of confusion. {I had to get my copy via inter-library loan from a Cincinnati public library (I live in Cleveland). It apparently wasn't available from the universities in the area, of which there are many)

The essence of PCT is the functioning of organisms (a normal focus of psychology as opposed to sociology). It can be summarized thusly: All action and sensation is produced via the interaction of neurons and muscles. Further, the central nervous system is composed of neurons. It therefor follows that any perception or behavior, regardless of complexity, must be producible via the interaction of neurons through neural currents (in the absence of some other mechanism - 'a ghost in the machine'). We can demonstrate that negative feedback control is the mechanism used for first order interaction with the environment, and at the level of spinal reflex. We can use models to show it can work for higher order behavior. Thus, by default we presume it is the mechanism used at all levels.

When I talk about a system or control system I am usually looking at more than one entity. A biological analogy to a process control system would be a supervisor monitoring 2 operators, who are each operating a set of tools involved in a manufacturing process. This is similar to HPCT, except that neural currents are not the _only_ mechanism for interaction of the 'nodes'. I don't think that this really changes anything, except by adding 1st level input and output function errors and adding the possibility of non pulse coded signals.

Please bear with me on the following as I am using a 20 year old work. There may have been elaborations in the mean time, but I presume there would be a '2nd edition' if there were major changes.

An observation: error signals don't theoretically have to only have positive values from an information standpoint. You can cheat by 'biasing' the signal. If the signal may vary from 0-30 ppm, I can set the comparator and output function so that 15 = no error (a logical 0). I could then set an 'upper threshold' at 25 and a lower threshold at 5. The output function could use this information to 'select' the proper action or magnitude of action. This could allow a stepped sawtooth output function rather than a continuous one. I'm not saying that this _ever_ happens in biology, or that it can't be implemented with a network of simple nodes, just that it would work. This might be useful for designing automata using PCT.

Question: why is it illegal to pass an error signal (directly or through a 'repeater' output function) to another node as a 'sensation'? While the analogy may not bear close examination, pain might be something of this type. The output is just passed up to tell a higher level controller that something is wrong, and how badly. It can be used to indicate that 'control' is not working and another strategy must be applied. It is up to the higher level to set reference levels elsewhere in the network so that the error signal is mitigated. It's not very useful at the lowest couple of orders, but might be useful at the 'cognitive' levels of control. This also might correspond to 'alerting'. It is just another input signal, but generated by a comparator rather than a lower order input function. Again, I don't assert that it actually exists in nervous systems, or that it is necessary. OTOH in engineering we sometimes find that a more complex structure is more efficient than the equivalent constructed from simple structures. If nothing else you have propagation and processing lags. Biology has a limitation in that it developed by 'growing like topsy'. New levels and nodes were added to existing ones that worked.

As your thinking has evolved over 20 years, do you have any problem with having 'neural logic' as a part of input, output, or comparator functions? Is there any theoretical problem with a single node (or subsystem) being at different levels of the hierarchy with respect to other nodes or subsystems {Higher apparently means 'sets another's reference level'} at the same time? My mental model can envision a network rather than a true hierarchy

On another minor matter, I am not sure that nodes interact only through neural currents in biological systems. Some output functions result in the release of hormones, neurochemicals or substances like adrenaline (blanking on the name). These are certainly sensed by other nodes, and not necessarily hierarchically (i.e 4th level affecting 4th or other level). I guess this could be viewed as an 'environmental' interaction of physical laws, but seems to me more like a 2nd or higher order interaction. I'm not sure it really matters if a signal is transmitted in terms of neural current frequency or in terms of a chemical concentration which must be translated via an input or reference signal transducer function.  Again, it may be just where you draw the system boundary.

Re Tom Bourbon [940719.1202]

I don't think creating a taxomy of types of signals or types of control nodes is silly in and of itself, any more than distinguishing between orders of control systems or sections of the brain. Yea it is all 'just control' or 'just neural currents', but the distinctions are sometimes useful. We wouldn't get very far in medicine or physiology if we fixated on the idea that 'cells are just cells' and ignored their differentiations and groupings (e.g. organs). We may be able (and have) to define standard structures or patterns used for various 'types' of perception and control. The question is which groupings or distinctions make sense.

Enough digital diarrhea for today,  Paul

Date:     Wed Jul 20, 1994 10:02 pm  PST
Subject:  Re: B:CP impressions - 1st cut

<[Bill Leach 940720.23:22 EST(EDT)]   >[Paul George 940720 17:00]

> Having read through the first 8 chapters, I begin to see the source of the misunderstanding of my posts, as most discussion seem to focus on the first three levels (orders) of the CNS.PCT does agree with my mental model, there were just some scope and terminology issues. All in all a great piece of work.

Delighted to see that you have started into B:CP. Your difficulty in obtaining it is certainly a sign of at least part of the problem.

> I guess it's lack of acceptance is due to perceptual filtering, or more likely that 'other schools' never actually _read_ (with understanding) the book.

I am about willing to even make a bet, that few if any in a position of authority have read the book at all much less study the concepts.

> It is important to evaluate a piece of work using the _author's_ frame of reference, not your own.

This is hardly correct. It is essential to attempt to understand the author's meaning(s) but at some point you must evaluate the author's work based upon some standards of your own. The idea that a technical book is accepted based upon the author's remaining consistent his own "frame of reference" is absurd.

In the first place, the author may state goals for his work that are not met even though he maintains steadfastly that they were. The stated goals may not be the true purpose of the work. And the work may rely upon principles or truths that are not. All of this must be considered when doing a scholarly reading.

Admittedly, it is easy to "rationalize" and not do an honest evaluation of what the author actually has said. I believe in the case of PCT the matter is really much simpler in that so few people actually understand the phenomenon of control that they quickly reveal errors when discussing control.

> Translation is your job when you are not using terminology in a 'standard' way or when using 'non-standard' concepts.

This is true and it is also true that it is the "job" of anyone trying to understand the work of another to attempt to come to terms. A particular difficulty in this area for PCT is that almost NO ONE ELSE uses precise meanings for relevant terms. In the behavioral sciences the statement is true almost without exception but unfortunately even in the engineered control systems field terms are rather loose.

> The essence of PCT is the functioning of organisms (a normal focus of psychology as opposed to sociology).

Well, if sociologists had some concept that they are dealing with living control system and the implication of such a condition, they would have a great deal of interest in the matter.

> It can be summarized thusly: All action and sensation is produced via the interaction of neurons ...
> ... Thus, by default we presume it is the mechanism used at all levels.

I think "them's fighten' words" :-) PCT is about the idea that Behavior is the control of perception. It does not matter that it actually appears that neural comparators do not allow sign switching, it does not matter if references and their perceptual signals might be biased or might not. PCT does is not only not "bothered" by the idea that many signals may be chemical in nature (as opposed to electrical) but the theory even helps explain the fundamental operation of some of them and further just about requires the existence of some sort of chemical type system for reorganization to occur in the fashion that the theory predicts. The real point is that whatever it is, there is overwhelmingly strong evidence that it is a closed loop negative feedback system and that is what PCT is about.

> When I talk about a system or control system I am usually looking at more than one entity.

PCT does not have a problem in dealing with multiple control systems.

> Question: why is it illegal to pass an error signal (directly or through a 'repeater' output function) to another node as a 'sensation'?

I would suggest that it is probably NOT illegal in any sense of the word but rather it does not happen because it makes no sense from a control system standpoint to do so. When controlling, the error is so close to zero to be useless for any other purposes (it is in the noise). The perception itself could also be an input to another perceptual control function that could handle problems associated with loss of control by the first system. Additionally, there does appear to be a system for sensing loss of control.

> While the analogy may not bear close examination, pain might be something of this type.

If I am not misunderstanding you here, Pain is clearly not a good example. Pain happens to be one of the areas where a great deal of medical research has been conducted. Though certainly not all pain has been studied this way but many studies indicate that mechanics of pain is the result of sensors sending signals up through the nervous system to the brain. There is no indication that there is any connection between these signals and the output of any other system.

> OTOH in engineering we sometimes find that a more complex structure is more efficient than the equivalent constructed from simple structures.

I should like an example of this. This sounds a bit like the argument that it is impossible to write a program in assembler that is as efficient and effective as one written using an optimizing high level compiler. Such a statement is of course pure bunk. Attempting such a thing would be inefficient, would require one hell of a programmer, be impossible to maintain, etc. BUT it has to be doable since the optimizing compiler is using the very same "simple structures" that are available to the assembler.

> If nothing else you have propagation and processing lags.

What do you mean by this? Higher level engineered control structures don't have propagation and processing lags?

I would like to suggest that in using the high level development tools that you have at your disposal for designing control systems, what is happening is that 1) you have a great deal more knowledge available to you about the capability of the elements of the system and you have the ability to easily control perceptions that you would not likely attempt to control if you were trying to use only simpler controllers for the job.

> On another minor matter, I am not sure that nodes interact only through neural currents in biological systems. Some output functions result in the release of hormones, neurochemicals or substances ...

Keep reading :-)

-bill

Date: Thu Jul 21, 1994 11:29 am PST
Subject: Re: Comments to Paul George (from Mary)

[Paul George 940721 10:00] [Mary Powers 940720 8:45]

> In PCT, error signals are outputs to lower levels, where they function as reference signals. Only at the lowest level is output a signal for muscles to contract. There is no function selecting actions: the actions are dependent on numerous output signals to various muscles, which in turn are dependent on the comparison of the desired state of affairs to the perceived environmental situation.

Sorry, sloppy use of the term 'action'. What I meant was a higher level loop using 'neural logic' or some kind of rule set to select _which_ of a set of lower loops should have its reference levels adjusted to what degree, 'causing' them to control (generate outputs). Ultimately action occurs which hopefully causes the original error signal to diminish.

See also my post yesterday on B:CP

Thanks for the response Paul George

Date: Thu Jul 21, 1994 11:38 am PST
Subject: Misc comments

[From Bill Powers (940721.0815)] Paul George (940720.1100)

You (and Bill Leach) are making some good points about the differences between process control designs and modeling people. A lot of what goes on in the design of artificial control systems doesn't translate into PCT. The real affinity between the approaches is at a deeper level; the phenomena of closed-loop systems, basic ideas like perceptual representation, comparison, action, external feedback paths. One difference is that the process control engineer knows too much; he/she knows what is causing disturbances, what might cause disturbances, and what sorts of countermeasures can be taken before the disturbances ever happen. To an organism, everything comes as a surprise (other than the very few behavioral systems that are in something like working condition when we're born).

I'm delighted that you're making your way through BCP. It is hard to get; distributors don't stock it. Aldine-deGruyter has done a new printing (don't know how large), so it's available from them for, I fear, about $43. Total sales so far must be somewhere between 5000 and 7000 copies.

I haven't done a revision, though it could use one. Just haven't got the inspiration to try. The internet absorbs all my creative writing these days.

> I found Bill Power's post on the e. coli paper review committee revealing as to the nature of 'peer review'. However you should be careful that you don't follow the same pattern in judging other's work.

Right you are. I always make a pretty serious attempt to grasp what the other person's model is before saying anything about it. But you can go only so far in following another person's reasoning, especially when the person doesn't really understand what's required of a model.

The biggest problem I have usually comes right at the beginning, where the speaker wants us to accept certain vital premises for the sake of the rest of the argument. If I have no knowledge or opinions about the premises, I'll play along, but it gets really difficult for me when a premise is something I just flat can't believe. Sometimes the premise is right there in the title: "Control of responses by discriminative stimuli as a function of frequency of reinforcement." How am I going to get past that? I just can't. I have the same problem with complex mathematical developments. I find them difficult enough, but when they start out by assuming things about which I am very dubious I just can't get up the energy to play out the game, even assuming that the derivations and lemmas and theorems are free of mistakes. This has irritated some people; they say "Well can't you just grant that for the sake of the argument, and see what we can derive from it?" I say, "I'll grant it when you give me some reason to," and there goes a beautiful relationship down the drain.

That's one reason I'm so big on demos and experiments. That's how you establish that your premises are reasonable. Psychological theorists don't seem to pay much attention to that sort of thing.

Glad you liked the section on neural currents and computations; not many people read that chapter. At least I can claim that I was thinking about analog neural nets before 1973.

> All action and sensation is produced via the interaction of neurons and muscles.

OK, if you'll add " ... through effects of the muscles on the environment."

> ... is similar to HPCT, except that neural currents are not the _only_ mechanism for interaction of the 'nodes'.

A realistic model of nervous-system operation would have to include the biochemistry of neurons. That is implicit in the notion of neural computing functions, but not spelled out in PCT. Too detailed a level of analysis for our current state of understanding.

We are commencing some work on purely biochemical control systems. Stay tuned.

> An observation: error signals don't theoretically have to only have positive values from an information standpoint. You can cheat by 'biasing' the signal.

Right. The model we use is a "canonical" model, meaning that there are many other forms that would be equivalent to it in function. What you're suggesting are variations that could well exist, but which would be equivalent, in the end, to the canonical model. In the brain stem, the reference signals appear to enter the output function along with the perceptual signals (at least in the Red Nucleus), and there is no distinct comparator. But the actual arrangement can be reduced to the canonical model by redefining some functions and constants.

> Question: why is it illegal to pass an error signal (directly or through a 'repeater' output function) to another node as a 'sensation'?

When you start talking about rerouting error signals to other systems, then the architecture can't be reduced to the canonical model. That may be perfectly OK, but first you have to demonstrate how a system organized that way would actually work, and show that there's a behavioral phenomenon that needs it. There are lots of possible variations on the current model, but until they've been simulated and applied as quantitative explanations of real behavior, they have to remain "unofficial." There are lots of unofficial ideas floating around, some from me, but they're just possibilities so far. Until we're forced to explore them more seriously (because the existing model runs up against a phenomenon it can't explain and that we're willing to do the work on), they will continue to float. There are plenty of simple phenomena that the canonical model, alone or in hierarchies, handles very well. We're still at the point where we're very happy to do simple experiments that work out. We don't accept changes to the model that aren't backed up by some pretty solid evidence.

As you have guessed, most of what we say about higher levels of control is in the realms of entertainment. We can't really model those levels of organization yet.

As a system designer, you have to watch out for designing systems as opposed to modeling them. Any clever designer can think of six different designs to accomplish a given result before breakfast, especially when working without any factual constraints. What we're trying to do is figure out how the REAL system accomplishes what it does, and in that case there could be six designs that would work and all of them are wrong. This is not something that control-system engineers normally have to worry about.

> As your thinking has evolved over 20 years, do you have any problem with having 'neural logic' as a part of input, output, or comparator functions?

No, and I never did. You'll be running across a "program" level of control, which I envision specifically as using symbolic computations, whether logical or of any other kind (like calculus). I see the lower levels as exclusively analog in nature, but if some phenomenon arises that calls for digital logic in a low-level system, there's nothing in principle against it. All anyone has to do is demonstrate that it's needed.

> Is there any theoretical problem with a single node (or subsystem) being at different levels of the hierarchy with respect to other nodes or subsystems {Higher apparently means 'sets another's reference level'} at the same time?

Yes, because I see the hierarchy as physical, not conceptual. The lower- level control systems are located physically near the periphery of the nervous system (the spinal and brain-stem reflexes, for example). The pathways that connect different levels are fairly well known up to a point, and they resemble those in the hierarchical model (not by accident). There's nothing conceptually to prevent the output of, say, a first-level control system from becoming the reference input of a sixth- level system -- except that the outputs of first-level systems actually go to muscles, not to higher systems. There are neuroanatomical constraints on the model. They aren't obvious in the final product, but I did do quite a lot of study of that literature while putting the model together, looking for constraints that would narrow the possibilities. There isn't as much useful information there as I had hoped for, but there is some.

> My mental model can envision a network rather than a true hierarchy.

Yes, that thought has nagged at me for a long time. I keep thinking of cases where what I call a third-level control systems seems to operate via a fifth-level control system, and so on. So far I've been able to resolve most of these possibilities, but the question is still open whether we should think of these "levels" as "dimensions" instead. This possibility was put to me by my friend Kirk Sattley, a lurker on this net, in about 1954, on the very day when something else he said resulted in the picture of the control hierarchy suddenly falling into place. I opted for the hierarchy, but have wondered ever since whether there isn't something to the other view, too.

Whatever the case eventually proves to be, there will be constraints on the solution. You can mess around with network ideas and come up with lots of interesting-looking designs, but "interesting" doesn't mean "right."

> Some output functions result in the release of hormones, neurochemicals or substances like adrenaline (blanking on the name).

There are many hormonal control systems at the level of organs, and below that many control systems of more detailed nature, all the way down to RNA and DNA. I see these as part of the environment in which the behavioral (neural) systems live. At the level of the hypothalamus, there are neural signals which go into the pituitary, where they seem to serve as reference signals for many hormonal control loops (control of circulating thyroxin is one well-known control system of this kind). There are also many neural sensory signals originating among these organ and biochemical systems. So to the neural hierarchy, these systems are simply another part of the environment that can be sensed and affected via neural signals. Of course we give different names to the perceptions arising from inside the body; feelings and emotions and other proprioceptive things. I think of these control systems as being below the first level of control in the neural hierarchy. Not much has been done with them, by me, although others have studied isolated systems at these physiological levels.

> I'm not sure it really matters if a signal is transmitted in terms of neural current frequency or in terms of a chemical concentration which must be translated via an input or reference signal transducer function.

The neural currents interact chemically with neural cell bodies at the synapses. Inside the cell bodies, the chemical influences of many simultaneous incoming neural signals, represented by neurotransmitters, interact according to laws of chemistry and diffusion to produce new signals at the output that can be complex functions of the input signals. A single neuron can be quite a complex analog computer.

The literal _transmission_ of neural signals is an electrochemical phenomenon, and the all-or-nothing nature of the impulse says that information is most likely to be carried by frequency variations. The concentrations of neurotransmitters at synaptic junctions, and even more so inside the receiving cell body, vary on a longer time scale than a single impulse, so the natural relationship between the chemical processes in the cell body and the neural signals is one of frequency- to-concentration conversion. The physical chemistry in the nerve cell is really the temporal bottleneck that says an individual impulse has no significance. There are others who disagree with me, claiming that neural signals can be multiplexed into a single fiber (and distinguished from each other afterward!), and that information might be transmitted by small variations in amplitude of spike as well as frequency. But the simplest view is that neural signals are measured as continuous variables in terms of frequency. Maybe this isn't true everywhere; there are all kinds of synaptic effects other than my basic model, although they are not common. I don't really think we need to carry the model to that level just now.

I am very pleased to see the serious attempt to understand that you are bringing to my book. I thank you for the compliment.

--------------------------------

Bill Leach (940720.1925) --

Good observations on process control.

> ... engineered control systems don't really have goals in the same sense that living control systems do (indeed, their goals ARE the implementation of the designers goals).

Language difficulties again. In common parlance, a goal is something external to a system; If I say I have the goal of going to the movies, the movie theater is taken to be my goal. Another usage is to treat a goal as the final state in which you want something else to be, out there in the environment. You use the latter sense above: the goal of the control system is the designer's goal for what the control system should do.

What these usages overlook is that for any goal, the real goal is to _perceive_ that something has happened or is in a particular state. The control-system designer wouldn't be very satisfied if the control system actually did what he had in mind, but he was unable to find out what it actually did.

The concept of the reference signal gives us a general-purpose definition of a goal that makes sense in all circumstances. The reference signal is an example of a perceptual signal when it is in a particular state. So the reference signal is, physically, the goal.

That definition fits both the designer and the control system. If the control system has a reference signal, it has a goal for one of its perceptions (sensor signals). If the designer has a reference-signal specifying what the control system is to be perceived doing, then the designer has a goal, too. In both cases, the goal is located physically inside the controlling system, not in its environment.

Accomplishing this change of understanding, unfortunately, requires going all the way in absorbing PCT. It's all perception. When you look at a spot in the service court across the net, you can see where you want the ball to go: the goal is right there, in the far left corner. So it really seems that the goal is outside you. But as a PCTer, your understanding kicks in and says "Aha, the far left corner is a perception, represented by a neural signal in my brain. That's how the neural signals representing the far left corner look. So my goal is to have a perception of a ball bouncing right there where that perception is; I'm imagining a ball doing that (with or without vivid video)."

By that time, of course, the umpire has called you for delay of game.

We have to live with the fact that language incorporates a lot of theories that we would really not want to support any more, but becomes awkward when we try to speak correctly. It's just a lot easier to say "Look at the far left corner" than it is to say "Bring the far left corner to the center of your perceptual field."

-------------------------------

RE: your comments on disturbances (see Mary's post, appended at end, for a nifty example inspired by your mention of baking a cake):

One experiment I have done with the tracking experiments is to actually display a measure of the disturbance being applied to the cursor. Normally, the disturbance is simply added, numerically, to the number representing handle position and the sum determines cursor position. There's no way to separate out the effect of the disturbance. But in this experiment, I added another "cursor" to the screen which moved up and down in proportion to the amount of disturbance being applied to the cursor position. If knowledge about the _cause_ of a disturbance was helpful in controlling, then people should control better when this accurate information about the disturbance was displayed.

Of course they didn't. They controlled worse, unless they deliberately ignored this added information, in which case they controlled as well as usual. The problem is that the visual angle of precise vision is quite narrow, only a couple of degrees, and having to look at the disturbance indication required, if not moving the eyes, at least attending a little off center. Or at least it required some additional perceptual processing. Whatever the reason, the display of the disturbance amplitude turned out to be a distraction, not a help, and control became worse.

This experiment could probably be refined so the information about the disturbance could be obtained from a place located ON the target or the cursor, somehow. But I would still predict no better performance than when the information is missing. A control system doesn't use information about the disturbance if it has information about the controlled variable. I suspect that the cases where sensing the disturbance proves useful are those in which control is pretty poor to begin with.

Best, Bill P.

Mary's post follows:
----------------------------------
[from Mary Powers 940721}

Tom's tracking experiments do distinguish disturbances from variables affected by the disturbance. Here's another example, besides cats and dogs, from real life: cake baking.

At 6890 feet (where we live) atmospheric pressure is a lot less than at sea level. This is the disturbance, which we are not equipped to sense (I suppose we could, but what kitchen is equipped to do so?) OR control (unless we install the kitchen in a hyperbaric chamber :-)). Low pressure does strange things to cakes, which we can control, rather elaborately, by a) decreasing the amount of baking powder and sugar in the recipe, b) increasing the amount of flour and liquid (which includes eggs) and c) upping the oven temperature by 25 degrees. None of these changes controls the disturbance of low atmospheric pressure, but rather the effect of it on the physics and chemistry of the ingredients, which in turn affect the variables actually being controlled: doneness, texture, and so forth.

Mary P.

Date: Thu Jul 21, 1994 2:48 pm PST
Subject: Re: Disturbance Control, Political Steering

From Tom Bourbon [940721.0853]

>Paul George 940720 17:30

>>[Rick Marken (940720.0845)] >>>Bill Leach (940719.19:12)

Bill L.

>>> Among other things, in engineered control systems, there are often environmental conditions whose status is quite important to the control process but is not controlled by the engineered system hardware.

Rick

>> I'd say that that's true of ALL control systems; the environmental conditions of which you speak are called disturbances; they are the reason for building a control system in the first place -- to protect variables from the unpredictable effects of these conditions.

Paul

> They are called disturbances _after_ the system is constructed. During analysis or design there are no controlled variables _yet_. That is the purpose of the design - to determine what input functions, output functions, and hierarchical levels of control are needed to meet the purpose of keeping the process stable and predictable.

Rick:

>> This estimate [TB: of the disturbances an engineered system might encounter] is not really part of the control process; it is just another variable in the environment

Paul:

> Here is the major terminology conflict. To us 'control system' is more than the controller. It includes the equipment, sensors, and multiple controllers. (so is a 'satellite control system') It even sometimes includes the operator via interaction with displays as well as instruments or controls not hooked to controller input or output functions. In biology one can distinguish between the nervous system and everything else.

Yes.

> You could also draw a system boundary at the skin.

Yes, but that might not always be the best place to draw the line -- see my post (From Tom Bourbon [940720.1653]) in which I talked about cats, dogs and pursuit tracking. The skin is just something that gets dragged along when the skeleton is moved by skeletal muscles, all as part of the incidental environmental consequences of a nervous system controlling its own perceptual signals.

> A process control system is more like a work party than a person.

Agreed, Paul. That is the point I was trying to make when you appeared on csg-l. I didn't do a very good job, though. When you first appeared, you suggested that the work of Albus might be informative to PCT modelers, in that Albus described systems far more sophisticated than PCT models as you understood them then. What I tried to say in reply was that the systems Albus described were not individual living control systems, but were widely distributed groups of systems of some other kind(s). I'm happy to see that we have all worked our ways past some of the early "heat" in the conversations to some mutual acknowledgements that Albus-like process control is quite different from perceptual control by individual living control systems. :-))

Later, Tom

Date: Thu Jul 21, 1994 2:53 pm PST
Subject: Re: B:CP impressions - 1st cut

[Paul George 940721 11:00] >[Bill Leach 940720.23:22 EST(EDT)]

Me:

>> It can be summarized thusly: All action and sensation is produced via the interaction of neurons ... ... Thus, by default we presume it is the mechanism used at all levels.

> I think "them's fighten' words" :-) PCT is about the idea that Behavior is the control of perception.

I _knew_ I was going to get zapped for that :-}.

To me there are 2 significant factors in PCT, while y'all seem to focus on one. The first is that you can demonstrate that behavior is a hierarchical control process and that a neural mechanism can produce it. The second is that perception is being controlled rather than the environment or actions upon the environment. I took the latter en passant, as should anyone with any familiarity with epistemology.

However note that even control engineers do not view that which is being corrected as that which is being controlled. From _their point of view_ the controller is controlling the process, i.e. trying to control the environment. Of course the only way we know what is going on in the environment or with our actions (which you normally consider part of it), is through our senses which are abstracted up into perceptions through a hierarchy (or series, or network) of input functions. Engineers would grant that the control loop and hierarchy functions as described in PCT.

I grant that this control distinction is a major sticking point between you and other psychologists, but I am not sure that the first point is not the most significant in terms of being able to analyze human and social behavior. The second point seems to be more one of terminology or point of view than of real overriding significance. But then I'm not a psychologist or behavioral researcher :-).

>> When I talk about a system or control system I am usually looking at more than one entity.

> PCT does not have a problem in dealing with multiple control systems.

I meant the same thing as you did in your post on process control. We tend to draw the boundary of 'a system' in a different place that PCT'rs usually do. The theory of course has no problem with distribution of control nodes or systems, other than the mechanism of communication would possibly vary from the physiological model.

> I should like an example of this. This sounds a bit like the argument that it is impossible to write a program in assembler that is as efficient and effective as one written using an optimizing high level compiler.

I intended it as more the other way around. It is more efficient to use high level graphic design systems than to hack code, which is much more efficient than bit twiddling. (this is after all my field of expertise). I am reminded of a 'programmer' on one project who always hard coded incrementing loop counters, because he had never heard of a fortran 'do loop' (and never thought to look in a manual). I know there are examples in electronics of more sophisticated circuits replacing sets of simpler ones, but can't think of a specific on the spur of the moment.

> What do you mean by this? Higher level engineered control structures don't have propagation and processing lags?

No, but intra-node communication can be faster than inter-node, due to physical distance if nothing else. An IC based CPU is a might faster (and more reliable) than the equivalent made of tubes or transistors. As a team grows in size the increased communication and coordination overhead makes the increase in capacity diminish on the order of np^2 (a.k.a an NP complete problem). Centralizing computation or logic minimizes communication at the expense of complexity. All I am saying is that in a given situation a 'logic chain' and single control loop may be more efficient than a hierarchy of control loops.

Date: Thu Jul 21, 1994 11:28 pm PST
Subject: Re: B:CP impressions - 1st cut

<[Bill Leach 940721.20:16 EST(EDT)] >[Paul George 940721 11:00]

> I _knew_ I was going to get zapped for that :-}.

Then why d'ja do it? Huh, Huh? :-)

> To me there are 2 significant factors in PCT, while y'all seem to focus on one. The first is that you can demonstrate that behavior is a hierarchical control process and that a neural mechanism can produce it. The second is that perception is being controlled rather than the environment or actions upon the environment. I took the latter en passant, as should anyone with any familiarity with epistemology.

There are really three that you have listed. The necessity for demonstrating that the neural structure for control exists is really "someone elses bag." OTOH, if what you mean by "and that a neural mechanism can produce it." is that "we" can show that biological beings exhibit control system behavior then that is another matter altogether.

The real question there is "Can anyone show any other method of operation for living entities?" PCT can an does demonstrate that it can model to a high degree of accuracy the operation of living systems in limited spheres. What is important here with the phrase "limited" is that the evidence points to the limited ability to model rather than limits upon a particular model match (human in this case) behavior.

At risk of being slapped into line by Bill P., I am going to propose some overall "driving concepts" for PCT:

1. The macro view of the activities of mankind over as long a history as detailed information is available clearly demonstrate "closed loop negative feedback control system operation".

   This single observation, I gather, was one of the overwhelming influences upon Bill Powers concerning the possible "nature of man." You can probably add the scientist/engineers' disdain for the "mushy" opinion based existing behavioral "sciences".

2. Actual attempts at modeling aspects of human behavior were found to correlate well with real human behavior in the conduct of simple experiments. (and if I remember the story correctly, an error in prediction made by those discussing PCT was discovered but BOTH the human and the model behaved in the same fashion -- thus interpretation of the significance of control theory was what was in error, not the theory itself).

3. There seems to be physical evidence of the existence of actual control systems in living systems (including, of course, people). The discovery of actual control system loops is "encouraging" but is not central to the theory. If the biologists had not found such structures, the rest of the behavioral "sciences" would still have the problem of explaining closed loop control system behavior some other way and none has been demonstrated to work.

> I grant that this control distinction is a major sticking point between you and other psychologists, but I am not sure that the first point is not the most significant in terms of being able to analyze human and social behavior. The second point seems to be more one of terminology or point of view than of real overriding significance. But then I'm not a psychologist or behavioral researcher :-).

If you reword that first line to "... is THE major ...", I believe that you will have phrased it correctly.

As I tried to indicate above, the physical evidence of structure is not at issue and would not be even if it appeared to be conflicting. If the biological evidence indicated something other than closed loop control systems, their would still be the problem of explaining how something that is not a closed loop control system behaves as though it is one.

>> PCT does not have a problem in dealing with multiple control systems.

> I meant the same thing as you did in your post on process control. We tend to draw the boundary of 'a system' in a different place that PCT'rs usually do. The theory of course has no problem with distribution of control nodes or systems, other than the mechanism of communication would possibly vary from the physiological model.

Indeed, if I have learned anything from Tom Bourbon at all, it is that "social aspects" of human behavior are explainable by analyzing the dynamics of control systems operating in a common environment.

>> I should like an example of this. This sounds a bit like the argument that it is impossible to write a program in assembler that is as efficient and effective as one written using an optimizing high level compiler.

> I intended it as more the other way around. It is more efficient to use high level graphic design systems than to hack code, which is much more efficient than bit twiddling. (this is after all my field of expertise). I am reminded of a 'programmer' on one project who always hard coded

incrementing loop counters, because he had never heard of a fortran 'do loop' (and never thought to look in a manual). I know there are examples in electronics of more sophisticated circuits replacing sets of simpler ones, but can't think of a specific on the spur of the moment.

It is (generally) a more efficient use of resources including time, to program in a high level language. However, in anything but a trivial program it is always possible that assembly code could be written that would be more efficient as a program than that generated by a compiler. That programs have reached levels of complexity that would make it not only impractical to make the attempt but would be physically impossible for even a medium sized programming team to do so.

The statement is true because of the very nature of programming and logic processing. Indeed, high level languages are nothing more than programs written to convert tokens and symbols to binary code but such conversions are somewhat general. You might want to take a look at some of the comparisons that have been made between code size and function as a function of the method of generation.

In terms of efficiency, computers have moved away from efficiency at a remarkable rate... it is just that processing power has more than kept pace (most of the time).

I am "overly belaboring" the point here except that I think that it is worthwhile to consider that what we often perceive as an improvement in engineered control systems is in fact a system with vastly greater processing power that is only able to make marginal improvements in control system capability. The "real" improvement is in the change to the environment that the DESIGNER is working in. People do things in control system design today as a matter of course that would not have been much more than dreamed of even as recently as 10 years ago. The reason that they would not have done these things was not because they were impossible but rather because they were impractical (I will admit, of course, that there are things that are done today that were, at least perceived, to be impossible at one time).

>> What do you mean by this? Higher level engineered control structures don't have propagation and processing lags?

> No, but intra-node communication can be faster than inter-node, due to physical distance if nothing else.

For the technology in actual use. Indeed, we have deployed communications systems whose throughput would "crush" the sort of controllers that you are talking about but in general it is not practical to use such communications systems for your application.

> An IC based CPU is a might faster (and more reliable) than the equivalent made of tubes or transistors.

I'm not sure of the significance of this statement... is your company still using tubes? :-)

> As a team grows in size the increased communication and coordination overhead makes the increase in capacity diminish on the order of np^2 (a.k.a an NP complete problem). Centralizing computation or logic minimizes communication at the expense of complexity. All I am saying is that in a given situation a 'logic chain' and single control loop may be more efficient than a hierarchy of control loops.

This is true but its' relevance is dependent upon how control is structured. A biological control hierarchy seems to be an example in the extreme for what you are saying is NOT a good control methodology. If you think you have ever seen a "large" distributed control system design, just ponder for a moment the sheer number of "controllers" that a human might have.

The "trick" there is that each level also "removes" raw detail in a perceptual signal so that each "neural signal" in successively higher levels represents a greater amount of information. Likewise, each high level reference "decomposes" into many lower level references on the way down. It is sort of like the idea that a node in a binary tree represents all of the nodes below it (in its' chain). The datum is (at least close) to the same size and duration as datum "at the skin". Thus, no NP problem).

-bill

Date: Fri, 22 Jul 1994 07:14:12 -0600
Subject: control of perception

[From Bill Powers (940722.0510 MDT)] Paul George (940721.1100)--

> However note that even control engineers do not view that which is being corrected as that which is being controlled. From _their point of view_ the controller is controlling the process, i.e. trying to control the environment.

I am beginning to think that this is how _naive_ engineers view the process. Of course their task is to be sure that the environment is controlled, because that is what the customer wants. But the control system itself, the one they build and that must operate on its own, can control only what it perceives of the environment. If the perceptual system is set up to perceive the wrong thing, the engineer may believe that the environment is being properly controlled, when in fact the control system is controlling some other aspects of it which only happens, for now, to keep the variables the engineer is interested in within the required limits. But a change in circumstances can leave the control system controlling quite successfully while the environment goes to pot.

Suppose the engineer wants to control the temperature of a bath. He gives the control system a thermocouple to register the temperature, constructs the rest of the loop, and all is well. But the thermocouple, as it happens, is sensitive to radiated energy as well as conducted energy. During the tests, it happened to be shielded from radiated energy, or there were no radiant energy sources to worry about. So everything worked. But when the control system is moved into the factory where there are powerful lights overhead, or sunlight coming in through windows, it will start controlling for the sum of the input energies from ALL sources, not just the intended one. When the lights go on, the temperature in the controlled bath will go down!

Worse yet, when the engineer is summoned to see this problem, he goes right up to the bath and stands there, and sees that everything is operating normally. When he goes away, removing his shadow from the thermocouple, down goes the bath temperature again.

Engineers may intend to control variables in the environment, but the control systems they build can control only their own sensory signals -- no matter how the engineers think about it. Smart engineers know this,

if only intuitively; this is why they specify periodic calibrations of the sensors. When a sensor drifts in calibration, the control system continues to hold the sensor's output signal at the set point, but this results in the external variable's changing. This is why they make sure that a sensor which nominally senses one variable, like light intensity, is not also affected by other variables, such as magnetic fields or temperature or humidity.

A control system controls its own sensory signal. This is not a matter of interpretation; it's how control works.

Best to all, Bill P.

Date: Fri, 22 Jul 1994 11:11:15 EDT
Subject: Re: control of perception

[Paul George 940722 11:20] >[Bill Powers (940722.0510 MDT)]

> I am beginning to think that this is how _naive_ engineers view the process.

Note I meant that there is no disagreement on how a control loop actually works, just on focus on what is spoken of being 'under control'. However, I think something more like 'blissfully ignorant' would describe a frighteningly large number of them. Unfortunately a sense of infallibility and of order in the universe is a common attribute of engineers. Hazard analysis is rarely performed, and there are a lot of naive assumptions about the environment and equipment.

> If the perceptual system is set up to perceive the wrong thing, the engineer may believe that the environment is being properly controlled, ... But a change in circumstances can leave the control system controlling quite successfully while the environment goes to pot.

That is indeed the challenge in process control system design; figuring out what needs to be sensed and how to sense it.

> Suppose the engineer wants to control the temperature of a bath. He gives the control system a thermocouple to register the temperature, constructs the rest of the loop, and all is well. But the thermocouple, as it happens, is sensitive to radiated energy as well as conducted energy. During the tests, it happened to be shielded from radiated energy, or there were no radiant energy sources to worry about. So everything worked. But when the control system is moved into the factory where there are powerful lights overhead, or sunlight coming in through windows, it will start controlling for the sum of the input energies from ALL sources, not just the intended one. When the lights go on, the temperature in the controlled bath will go down! Worse yet, when the engineer is summoned to see this problem, he goes right up to the bath and stands there, and sees that everything is operating normally. When he goes away, removing his shadow from the thermocouple, down goes the bath temperature again.

I wonder if you have any idea how realistic this example is? (though RF, magnetic, or electrical interference is more common) Unfortunately a lot of systems are designed this way. Problems like this often show up in the field. And funny thing, when the equipment goes back to the shop, there is no problem found.

"...And he examined the test data and saw that some of it was good, and some not so good. He therefor divided the good data from the bad and he called one 'results' and the other he called 'spurious anomalies'+ ;-)

Paul

Date: Fri, 22 Jul 1994 10:28:06 EDT
Subject: Re: B:CP impressions - 1st cut

[Paul George 940722 10:20] >[Bill Leach 940721.20:16)]

> OTOH, if what you mean by "and that a neural mechanism can produce it." is that "we" can show that biological beings exhibit control system behavior then that is another matter altogether.

Kind of both, if the B:CP discussion of spinal reflexes is correct (which to my level of neurological knowledge it is). The latter point is what I meant by:

>> The first is that you can demonstrate that behavior is a hierarchical control process...

> If you reword that first line to "... is THE major ...", I believe that you will have phrased it correctly.

:-) Actually, I did phrase it that way the first time and later toned it down out of politeness.

> As I tried to indicate above, the physical evidence of structure is not at issue and would not be even if it appeared to be conflicting. If the illogical evidence indicated something other than closed loop control systems, their would still be the problem of explaining how something that is not a closed loop control system behaves as though it is one.

Want to run that one by me again? You are saying that if I could prove that the CNS anatomically did not use closed loop control structures that the burden of proof would be on me to show how and why it worked?? From my point of view it would simply show PCT to be incorrect in asserting that closed loop control is the mechanism for organismic behavior. It would still be useful for analyzing or predicting it. Two different mathematical functions that give the same results within a given range are equivalent within that range.

> A biological control hierarchy seems to be an example in the extreme for what you are saying is NOT a good control methodology........ It is sort of like the idea that a node in a binary tree represents all of the nodes below it (in its' chain). The datum is (at least close) to the same size and duration as datum "at the skin". Thus, no NP problem).

I'm not saying it is bad, just not always the most efficient or simple (in terms of number of components). The situation I am talking about is where there are a large number of nodes on the same level which must interact in order to 'evaluate' (not necessarily control) enough perceptions to 'make decisions', 'recognize patterns', or make plans. Computers use memory and logic structures other than trees (particularly in AI) for a reason. The levels 1-3 of a biological system are clearly hierarchical anatomically . Hierarchical structures to filter, summarize, or compress information are common. However, the hierarchy is often in terms of 'strata' rather than levels. As Bill B indicated in his response to my post:

>[Bill Powers (940721.0815 MDT)]

> I keep thinking of cases where what I call a third-level control systems seems to operate via a fifth-level control system, and so on. So far I've been able to resolve most of these possibilities, but the question is still open whether we should think of these "levels" as "dimensions" instead. ... I opted for the hierarchy, but have wondered ever since whether there isn't something to the other view, too.

At some point the 'canonical model' will likely break down due to the explosion of the number of nodes required. It takes a lot of control loops to control something like writing this post. When you get to the level of the sensory processors and upper cortex functions such as language,I strongly suspect that network rather than hierarchical structures predominate in order to generate the constructs we usually term 'perceptions' or 'ideas'. There will be more 'horizontal' connections than 'vertical'. I also suspect that control functions will involve more 'algorithmic' processing. However, I am not sure how you could tell empirically, though the stuff being done with PET scanners might help.

Paul

Date: Sat, 23 Jul 1994 17:07:49 -0400
Subject: Re: B:CP impressions - 1st cut

[Bill Leach 940723.16:23 EST(EDT)] >[Paul George 940722 10:20]

>> As I tried to indicate above, the physical evidence of structure is not

> Want to run that one by me again? You are saying that if I could prove that the CNS anatomically did not use closed loop control structures that the burden of proof would be on me to show how and why it worked??

First, let me say that this was my opinion and not necessarily shared by anyone else here. In particular, I know that Bill P. HAS studied anatomy to try to look for evidence in either direction and I'm sure that he cares very much what sort of new evidence shows up.

In a sense yes. The issue with PCT as I understand it IS NOT the representation of the structure, the possible levels and hierarchy. The issue is closed loop negative feedback control system is the only known mode of operation of living systems that explains what is actually seen.

That PCT explains S-R behavior is not even the point though the fact that it explains exceptions to S-R behavior is getting closer.

If you "come up with proof" that a person's neurological structure CAN NOT be a control system, then yes you indeed have a great burden of proof on your hands -- how does a non-control system consistently exhibit control system behavior?

> From my point of view it would simply show PCT to be incorrect in asserting that closed loop control is the mechanism for organismic behavior. It would still be useful for analyzing or predicting it. Two different mathematical functions that give the same results within a given range are equivalent within that range.

I don't challenge the last sentence of your statement there at all. Show me the function that provides the equivalent results. Short of metaphysics there is no known phenomenon that can make a non-control system behave exactly like a control system. If you do find such a phenomenon maybe the second sentence would still be true.

> I'm not saying it is bad, just not always the most efficient or simple (in terms of number of components).

I agree and will even add that when enough is known about the limits of environmental disturbances possible, even the choice of controlled perceptions may be poor.

> At some point the 'canonical model' will likely break down due to the explosion of the number of nodes required.

No, I don't think so. Even the discussions and models currently existing in the PCT world do major consolidation of control loops and even hierarchy. Where such "would break down" is in any attempt to exactly model the structure of the living being as opposed to modeling its behavior for specific instances of control.

-bill

Date: Sat, 23 Jul 1994 16:53:04 -0600
Subject: Levels and Misc

[From Bill Powers (940723.1540)] Bill Leach (940723.14:41)

(your note to Paul George):

> The issue with PCT as I understand it IS NOT the representation of the structure, the possible levels and hierarchy. The issue is closed loop negative feedback control system is the only known mode of operation of living systems that explains what is actually seen.

There's no way to separate these issues. A conceptually neat model that violates known neuroanatomy would be just as useless as a neuranatomically correct model that blew up when you tried to run it. What we must have is a workable model that recreates the behavior we actually observe, and fits what we know of the nervous system and muscles as well -- and works in the world described by physics and chemistry. This is supposed to be a model of a real system, not an intellectual exercise.

--------------------------------------

RE: boundaries.

The basic rule is: put yourself inside the system you're talking about. If you're the dog, your perceptual field contains a cat moving unpredictably around, as well as your nose and paws -- but no dog. If you're the cat, there's a dog in your perceptual field along with your nose and paws, but no cat. As an external observer, you are looking out of the wrong boundary. You can see a cat AND a dog, and the noses you see (other than your own) are seen in profile or in front, not from behind them. That's the incorrect view for understanding either system.

Best, Bill P.

Date: Sun, 24 Jul 1994 00:39:55 -0400
Subject: Re: Levels and Misc

[Bill Leach 940724.00:17)] >[Bill Powers (940723.1540)]

> (your note to Paul George):

>> The issue with PCT as I understand it IS NOT the representation of the structure, ...

> There's no way to separate these issues. A conceptually neat model that violates known neuroanatomy would be just as useless as a ...

I agree except that I also take the position that a "known" neuroanatomy that denies control loop operation would itself be wrong. I guess, in essence I am taking the stand that "so called" proof that the system is not a control system would be ok if and only if, it could also show that the control system behavior that is actually observed could also be explained.

> What we must have is a workable model that recreates the behavior we actually observe, and fits what we know of the nervous system and muscles as well -- and works in the world described by physics and chemistry. This is supposed to be a model of a real system, not an intellectual exercise.

I agree with this but what I am saying is that any other assertion about the structure MUST also account for the behavior. What "we know" about the nervous system is pretty tentative at best and a "generally accepted" theory of the nervous system structure that was used to assert that control system behavior had to be an illusion would not be particularly convincing to me at this point. Science has refuted a correct theory more than once in the past.

> RE: boundaries.

Ok, I think that Tom straightened me out on that one. I really will try to be more "reference oriented" to the control system itself in the future unless specifically stating otherwise.

-bill

Date: Thu, 4 Aug 1994 16:14:14 EDT
Subject: B:CP Reactions - Cut 2

[Paul George 940804 1600]

{Directed mostly to Bill Powers, but comments welcome from anywhere}

Having finished B:CP, I think it is a good piece of work. Too bad so much of it seems to be ignored in most of the research and the discussions here, which seem focused on the 'worm's eye view' of PCT (1st & 2nd order control). Perhaps because the other concepts are deemed 'uninteresting'? :-)

I do have a few questions & observations:

In terms of the reorganization system, why must the reference variables be _intrinsic_ and physiological? (I equate 'intrinsic' with hardwired or inborn) Some clearly would be, for example detecting life threatening situations or physiological conditions (Thirst, hunger, various levels of stress or pain, danger, 'instincts'). However, I don't see anything which would preclude having learned or otherwise generated reference levels as well. Recognizing the need to adapt would seem to be a characteristic of higher levels of intelligence. A general 'error level' or 'discontent' intrinsic would help to 'awaken' the reorganization function, but seems to me unsatisfying for directing it. Non innate values would also help explain how some people seem to change their behavior in self destructive ways. Their hierarchy of reorg reference variables have placed some artificial(??) goal at a higher priority than the natural or intrinsic ones. In other cases they may be trying to reconcile two incompatible goals. (I'm not sure that goal is the appropriate term, but is the one I'm most familiar with).

Note: I have not explicitly seen the concept of priorities between reference variables, perceptions, or control loops, except through the hierarchy and the actual constructs in the 'little man'. Levels of the hierarchy would appear to be more than one loop deep. I am not at all sure that such a concept is truly required, but at some level some disturbances are more important than others.

Re the close of chapter 14 on Platt's work. What experimental work has been done in the last 20 years on reorganization? Personally the application of control theory to adaptation, learning, and communication was what sucked me into this group, combined with it's effects on interpersonal relationships. The book has a number of references to work that needed to be done, but it is unclear to me as to what actually has happened.

I personally think you do your 'cause' a bit of harm by splitting off the reorganization hierarchy from the central control hierarchy, and then ignoring it in most of your posts and articles. I would publish figure 14.1 a little more widely. Similarly, just pushing the simple canonical control diagram as the definitive model instead of the figure 15.3 'final form' is misleading and I think hurts acceptance of your ideas. You don't highlight that the canonical model is a simplification of the 'true' form which you use because it is sufficient for your experimental purposes.

A large part of my initial impressions of weakness or oversimplicity in HPCT was because you (particularly Rick & Tom) appeared to be saying more elaborate structures were unneeded. I recognize that most PCT research is focused on levels 1-3 of the perceptual control hierarchy, partially due to amenability to modeling, but creates the impression that it is all HPCT consists of. The simple perceptual control loop does not explain all of what is commonly called behavior, much less psychology. The reorganization control hierarchy, memory,and the perceptual & memory switches are necessary constructs to deal with common concepts of planning, imagination, etc. Further the levels of the PC hierarchy above 3 are rarely explained, other than by occasional asides - e.g. "The lower six levels are concerned with control of intensities, sensations, configurations, transitions, events, and relationships. "

[Bill Powers (940803.1510 MDT)]

It is not intuitively obvious that these constructs derive from the simple concept of control. The terms are also subject to variable interpretation.

As the [From Bill Powers (940803.1510 MDT)] exchange with comp.ai indicates, your critics and the uninitiated dismiss you because you do not publicize the 'non-core' ideas, except through casual asides (in my experience), and so they presume such concepts are not incorporated in PCT.

Please accept a suggestion from a newbee who recently went through trying to get up to speed with PCT, even given a fairly strong understanding of control. Given the paucity of 'public domain' sources, I would recommend writing a little white paper and incorporating it in the monthly post (as well as archiving it (you could really use a FAQ). Few people are willing to shell out a significant amount of money or expend a lot of effort to investigate an apparently fringe or 'crackpot' (ITHO) theory.

I would summarize the levels of control discussed in chapters 7-13 from a role/responsibility standpoint at about a paragraph apiece. The introductory pages and chapter summaries could provide most of the meat. Introduce the reorg hierarchy as well, possibly with an elaboration of figure 14.1 (I think this hides the concept of their being two 'planes' of control hierarchies: x/z=PC,y/z=RC). This heads off objections concerning higher level behavior (from the standpoint of internal perceptions) and learning.

Next provide a paragraph for each of the neural gate constructs for the underpinnings of the feedback control mechanism and logic functions used elsewhere. Use the spinal reflex model (say figs 7.3 or 7.4) to discuss the basic control phenomena.(this is the part that is well covered by the current intro material, but it's derivation from physiological constructs is not highlighted). I would deem this a key point for acceptance. Use 'the parable of the rubber band' experiment (16.3), particularly the coin extension (Brilliant!!), to further illustrate the power of the concept.

Next briefly explain the final form diagram (15.3) with the memory and the 4 'control modes' (I think this was my favorite part of the book. What a neat model!).

The organization could perhaps use work, but I think it can be done clearly in less than 10 pages{Shorter than many of your posts ;-)} Also, consider putting your response to Bruce Buchanan (or perhaps both posts) into the intro materials. I wish I had read it much earlier as it would have avoided much misinterpretation as to what PCT encompassed. I am happy to see that PCT does in fact cover the areas I had anticipated it should.

Continuing to enjoy the discussion, Paul

Date: Fri, 5 Aug 1994 13:50:09 EST
Subject: worm's eye view

[Avery Andrews 940805.1347] (Paul George 940804 1600)

> Having finished B:CP, I think it is a good piece of work. Too bad so much of it seems to be ignored in most of the research and the discussions here, which seem focused on the 'worm's eye view' of PCT (1st & 2nd order control). Perhaps because the other concepts are deemed 'uninteresting'? :-)

I won't speak for the others, but I've been focussing on the lower levels (so far) because we have a long tradition of investigation of the upper levels (going back to Aristotle and the Stoics, if not further), and large numbers of clever and well-funded people still beavering away at them, but it all is and will remain very up in the air until we understand how the upper levels cash out as actual activity, which requires connecting them to lower-level systems, which requires reaching some minimal level of understanding of how these work.

Plus it seems to me the that actual properties of even simple closed-loop systems are often seriously misunderstood by people (such as Fowler and Turvey, or Abbs and Winstein) who one would really expect to understand them properly, so there is obviously some kind of problem there.

There is also the fact that the behavior of control systems is often wildly counter to even well-informed intuition, so that if you can't model, you run a big risk of producing bullshit. But our ability to model is limited by, among other things, our ability to construct perceptual functions, which is, especially at the higher levels, pretty minimal.

Avery.Andrews@anu.edu.au

Date: Fri, 5 Aug 1994 09:51:52 -0700
Subject: Appealing to Complexity Worshippers

[From Rick Marken (940508.0945)] Paul George (940804 1600)

> Having finished B:CP, I think it is a good piece of work.

To me, this is a bit like calling Newton's "Principia" a "good piece of work". But I guess I have to agree with your basic assessment.

> Too bad so much of it seems to be ignored in most of the research and the discussions here, which seem focused on the 'worm's eye view' of PCT (1st & 2nd order control). Perhaps because the other concepts are deemed 'uninteresting'? :-)

I think Avery Andrews (940805.1347) gave an EXCELLENT reply to this. In particular, I like his observation that the "actual properties of even simple closed-loop systems are often seriously misunderstood by people". I find it amusing when people storm off to model the higher levels before they grasp even the most basic concepts of PCT, such as the nature of control, the control of perception, testing for controlled variables, etc. By ignoring the basics, these people fail to see that a great deal of what might SEEM like "higher level" behavior is just the side-effects of very simple controlling (for example, consider the behavior of the simple control systems in the CROWD program or of the interacting humans in Tom's experiments on cooperation and conflict).

> A large part of my initial impressions of weakness or oversimplicity in HPCT was because you (particularly Rick & Tom) appeared to be saying more elaborate structures were unneeded.

In PCT, we try to explain phenomena with as simple a model as possible. It seems to me that one of the diseases of modern behavioral science is the fascination with complexity per se. I'll call it " complexity worship". It seems that many behavioral scientists think a model is weak and/or oversimplified if it is not complex (this is a big change from the scientific goals of Newton and Einstein). These behavioral scientists seem to be controlling for a high level of perceived complexity as a means of perceiving themselves as understanding behavior.

The modern disease of "complexity worship" seems to stem from three problems that are endemic to the behavioral sciences:

1) The first is an addition to the superficial; behavioral scientists are fascinated (as we all are) by the visible side effects of controlling, which can look quite complex.

2) The second is an addition to the verbal labels used to refer to these side effects. Thus, if a collection of observable side effects is called "optimal trajectory selection" it seems like one is dealing with something a great deal more complex and "high level" than "pressing a bar". Many PCT studies that had seemed simple and "low level" suddenly become "important and relevant" when they are given names like "helping", "cooperation", "conflict", "leadership", "learning", etc.

3) The third is an apparent aversion to experimental test and the acceptance of incredibly poor data when such tests are performed. There seems to be a growing interest in theory qua theory; there is very little testing of such theories against data. And the theories that are developed are not really attempts to account for real data; they are attempts to account for verbal _descriptions_ of data. So we see people trying to come up with theories of "path planning", "alerting", "intelligent search", etc. That is, people are trying to draw diagrams that they imagine will produce the kind of behavior that they imagine people would describe with these words.

People who enjoy PCT tend to be those who find that their understanding of human nature is controlled by building working models (the simpler, the better) that produce behavior that exactly duplicates real data. People who enjoy PCT tend to be "phenomena freaks"; they are interested in understanding a real phenomenon that is really interesting that they can really experience in themselves and others - - CONTROL. People who enjoy PCT tend to enjoy simple explanations (probably because we a simple minded) that are powerful (they explain a LOT of data); basic PCT explains a LOT of data.

I don't think there is much that PCT can do for behavioral scientists who suffer from the three addictions I listed about -- to the superficial, to the verbal, to the theoretical (undisciplined by the phenomenal -- i.e. data). But don't blame PCT for the fact that most of the behavioral science community is clueless. PCT only seems to be "simple minded" if you approach it from the perspective of clueless behavioral science.

In fact, the basic PCT model has a great deal to offer people interested in high level, "real life" problems; the discussions at the recent CSG meeting were evidence of that. Once you understand the basic PCT model (including the notion of hierarchy and reorganization) you can cut through a lot of the apparent complexity of behavior to see what people are actually doing (controlling perceptions) and why they often have trouble doing what they want (conflict).

In order to "sell" PCT to the "complexity worshippers" in the behavioral sciences, we would have to make it seem like PCT gives them what they want -- descriptions of the superficial complexities of behavior, understanding through verbalization and theoretical complexity for its own sake. I, personally, would rather just stay in the "simple minded" PCT ghetto, with simple-minded friends (like Tom, Bill and Avery) doing simple minded research and publishing it in simple minded books like "Simple MINDed READINGS".

Simply, Rick

Date: Fri, 5 Aug 1994 13:29:27 -0600
Subject: Re: B:CP evaluation, part 2

[From Bill Powers (940805.1150)] Paul George (940804.1600)

> Having finished B:CP, I think it is a good piece of work. Too bad so much of it seems to be ignored in most of the research and the discussions here, which seem focused on the 'worm's eye view' of PCT (1st & 2nd order control). Perhaps because the other concepts are deemed 'uninteresting'? :-)

Having finished the book, you can now speak with more assurance about what PCT is than others who haven't -- as you now realize.

Avery Andrews explained the position well. There are lots of people guessing about how the higher levels work, but there's no unifying principle to tie such work into the whole system. The biggest problem is in designing research and doing simulations without understanding how higher-level perceptions work, or even exactly what the higher levels are. So we work where we're more sure of our ground, trying to develop methods that will apply generally but not going too far beyond what we can demonstrate. It will be up to smart people like you to push the boundaries further upward in the hierarchy.

> In terms of the reorganization system, why must the reference variables be _intrinsic_ and physiological? (I equate 'intrinsic' with hardwired or inborn).

The main reason is that I needed something that _could_ be inborn, to provide the basis for the development of the hierarchy from infancy (or before) onward. I see the hierarchy in the neonate as a set of possible control systems, with the necessary sorts of computations provided for at each level but no wiring and no organization yet. We have to learn to perceive everything, at every level, and also how to control what we perceive. As I said in the book, this is a worst-case model, where there is an absolute minimum of inborn organization. The reorganizing system has to work under these worst-case conditions, where even the _concept_ of an algorithm or systematic strategy hasn't been developed yet.

If everything from sensations to systems concepts has to be learned, we obviously can't count on anything in those categories as an aid to learning. The process of reorganization, in short, can't be intelligent, because it has to work properly before any intelligence is developed. So I asked myself what the basis for reorganization could possibly be with those rather severe restrictions, and came up with the answer that it has to be concerned with the status of the organism itself. That's where the idea of intrinsic variables came from (that, and from W. Ross Ashby, who called these "critical variables" and postulated a random switching device which was the direct ancestor of my idea). There's no real need to guess at what the intrinsic variables are; all we need to specify is that they be affected by the way the organism behaves in an environment, and that no knowledge of the external world or its laws be involved.

Ashby assumed that there were built-in upper and lower limits on critical variables, exceeding which started the random switching process. I substituted the idea of reference signals, but that's a minor difference (two-way instead of one-way control). He also saw that such a system would be "superstable" in that it would keep randomly switching until the critical variables were once again within limits -- regardless of _why_ they came back within limits. I simply adopted that principle; it's still the basic principle of reorganization in PCT.

I was always somewhat unsure whether random switching could really be efficient enough to accomplish real reorganization, until I heard of Koshland's work on chemotaxis with E. coli, and set up some simulations to test his idea. The results astonished me. This turns out to be an unbelievably effective way of controlling things in a few dimensions at once. As long as reorganization is applied to small parts of the system at any one time, it can very quickly lead to local optimizations. The random nature of the reorganizing process even provides for getting out of local minima if they aren't too deep. I have shown that this method can be used to solve 50 equations in 50 unknowns in half an hour or so.

> However, I don't see anything which would preclude having learned or otherwise generated reference levels as well. Recognizing the need to adapt would seem to be a characteristic of higher levels of intelligence.

But you've put your finger right on the problem, haven't you? Where do those higher levels of intelligence come from? They have to be learned, too. And if they are to be learned, they must be the product of some learning process that predates them. Once each new level is organized, many control problems that could formerly be solved only by random reorganization can now be solved by a systematic control process, which is far faster and more accurate, and can keep errors so small that intrinsic errors are never seen -- at least for reasons having to do with that level of control. It's even possible that the reorganizing system creates (at the logical level, probably) specific strategies for learning, although it's learning of a different kind from the random trial-and-error of reorganization.

> A general 'error level' or 'discontent' intrinsic would help to 'awaken' the reorganization function, but seems to me unsatisfying for directing it.

Yes, general error level in the hierarchy (along with other types of variables inside the organism) is considered an intrinsic variable, with a built-in reference level of zero. When you say that reorganization is unsatisfying for "directing" the results, however, you're not recognizing the surprising property that random reorganization CAN be directed toward achieving a specific end-state. Whatever the variable under control, reorganization is driven by the difference between that variable and a corresponding reference state. This is the "intrinsic error" signal. Specifically, the interval between reorganizations depends on the rate of change of intrinsic error times the magnitude of that error -- which amounts to the first derivative of the square of the error. That is the model that seems to work the best.

What's hardest to grasp here is that the reorganizing system doesn't care what organization it produces as it acts on the system being reorganized. The ONLY thing it is concerned with is the intrinsic error. It will stop reorganizing only when the intrinsic error gets small enough. There many be better organizations, or the intrinsic error might drop to zero because of the arbitrary action of an external agency. None of that matters to the reorganizing system. It just wants its own intrinsic error to be zero. If it is zero, or small enough, reorganization stops, leaving _whatever_ organization exists at that point to go on operating.

It is this extreme pragmatism of the reorganizing system that makes it so effective. It does not need to know anything about the environment, or the organism, or the brain, or anything else besides the states of its own input variables in relation to their respective reference signals. It will try literally anything, and at random, until intrinsic error disappears, and then it will stop acting. How, during this process, it has changed the relationship between the behaving system and its environment is completely irrelevant to it. It can't either know or care about that. That is why it can produce a workable organization in any kind of environment. Its very stupidity is what makes it so powerful.

RE: experimental work on reorganization.

Very little has been done with real people, although Dick Robertson has done some experiments in which reorganization appears to occur, and Frans Plooij has recorded data on infant learning in chimpanzees and humans that seem to show definite periods of reorganization. We have done a number of simulations, one of them being a control system model that matches itself to behavior of a real person by treating the difference between the model's and the person's behavior as an intrinsic variable and randomly reorganizing the integration factor of the model's output function to reach a minimum in the difference. As mentioned, I've tested some models involving multiple control systems controlling a shared world of multiple variables, and reorganizing until the system of simultaneous equations is solved for independent control by each system of a different aspect of the shared environment. That's not a lot, but such things take time to work out and at least it's something.

RE: presenting the theory

> I personally think you do your 'cause' a bit of harm by splitting off the reorganization hierarchy from the central control hierarchy, and then ignoring it in most of your posts and articles.

Well, sheesh, how much can we talk about at once? We talk about the things we have done the most work with. Actually, if I had responded to your initial inquiries with a complete summary of everything that is in the PCT model, you would have thought I was some kind of nut. Haven't you ever seen real nut mail? If somebody you had never heard of sent you a two-inch-thick packet of drawings and text with the title, "A complete theory of everything that all organisms do under all conditions," wouldn't you alert the nearest looney bin? Most people who first hear about PCT have the impression that I thought it all up last week, so they're surprised that I haven't cited someone who said something somewhat similar last month. The idea that there's been this whole long development going on for 40 years, and that they've never heard of it, never occurs to them (and why should it?). And when I can't explain it all in one breath, they start fidgeting.

> Please accept a suggestion from a newbee who recently went through trying to get up to speed with PCT, even given a fairly strong understanding of control. Given the paucity of 'public domain' sources, I would recommend writing a little white paper and incorporating it in the monthly post (as well as archiving it (you could really use a FAQ).

I'd love to have an FAQ. How much do they cost, and can I get a used one cheaper? By the way, what is an FAQ?

I'm still leery of trying to cram the whole theory into the monthly intro, for the reasons cited above. Would you like to try it, being fresh from the learning experience? You seem to have a pretty clear view of what might make the whole thing more quickly understandable. Why not post a first try and let some of the other newcomers chime in. It might turn out to be a really effective document. "How I finally got the idea of PCT," or something. In fact, your suggested outline sounds pretty good to me.

Best, Bill P.

Date: Fri, 5 Aug 1994 15:57:20 EDT
Subject: Re: Appealing to Complexity Worshippers

[Paul George 940805 16:00]

>[Rick Marken (940508.0945)], also Avery Andrews (940805.1347)

> I find it amusing when people storm off to model the higher levels before they grasp even the most basic concepts of PCT, such as the nature of control, the control of perception, testing for controlled variables, etc. By ignoring the basics, these people fail to see that a great deal of what might SEEM like "higher level" behavior is just the side-effects of very simple controlling.

I share your reaction, and acknowledge that a little knowledge is a dangerous thing. I intend to "wait for fullness". But I wish that _you people_ had explored more in those areas, given that you presumably do understand.(Of course I haven't read all you have written) From one standpoint all you have done for 20 years is elaborate the evidence to support the models and experiments described in B:CP (not to suggest that the effort was wasted or trivial). Of course a major reason for 'slow' progress may have been that affordable computers with sufficient power only became available in the last 5-10 years.

I don't think I suffer (much) from the Three Addictions (particularly the third). I have no real issue with what most PCTers work with. Simple models are best, allowing for Einstein's caveat of "...as simple as possible, but no simpler". I just personally find the concepts of memory as reference value, the switching construct, and reorganization fascinating. Forgive me if I rely on _your_ work and conclusions modeling the basics :-). I'm mostly interested in conclusions, implications, and applications. Of course my job deals with designing processes and technology to allow people to adapt and communicate in order to deal with changing environments and problem solving.

> In order to "sell" PCT to the "complexity worshippers" in the behavioral sciences, we would have to make it seem like PCT gives them what they want -- descriptions of the superficial complexities of behavior, understanding through verbalization and theoretical complexity for its own sake. I, personally, would rather just stay in the "simple minded" PCT ghetto, with simple-minded friends (like Tom, Bill and Avery) doing simple minded research and publishing it in simple minded books like "Simple MINDed READINGS".

However, if propagation of truth or knowledge is your goal, rather than a cloistered purity of thought and purpose, you need to communicate with the unenlightened. IMHO you need to first communicate that the theory can encompass the complexities that many psychologists and commoners ;-) observe in human behavior, and then focus them towards the basics. Then they can be amazed at how much so called complex behavior simple hierarchical control can model. But first you have to catch the mule's attention.

If you make it easy for them to classify you as monomaniacal simpletons, then they will. It is much easier to dismiss you than to consider changing their way of thinking or studying a new approach. Why should they waste time investigating something that "obviously can't" scale up or map to reality?

Being "despised and rejected of men" may produce a kind of group pride or feeling of belongingness, but is not very effective for propagating your mnemes. The history of science deals with the mnemes that spread widely enough and survived long enough to make an impression. I have no doubt that university (and monastery) archives are littered with seminal dissertations that never again saw the light of day.

Keep up the good work, Paul

Date: Fri, 5 Aug 1994 17:56:35 CST
Subject: Re: B:CP Reactions - Cut 2

Tom Bourbon [940805.1719] >[Paul George 940804 1600]

> {Directed mostly to Bill Powers, but comments welcome from anywhere}

How can I refuse an invitation like that? Time is short today, so I will only reply briefly to one or two points.

> Having finished B:CP, I think it is a good piece of work. Too bad so much of it seems to be ignored in most of the research and the discussions here, which seem focused on the 'worm's eye view' of PCT (1st & 2nd order control). Perhaps because the other concepts are deemed 'uninteresting'? :-)

Paul, it is nice to see that you are reading something about PCT. That is a significant step from the time when you appeared on this net, all filled with thunder and lightning, but with no history of having read any of the material. :-)

Like Rick Marken (940508.0945), I think you understate the case. I would place Powers's B:CP, and his theory of behavior in the same league as Newton's "Principia," or, as I said in the "foreword" to the second volume of Bill's _Living Control Systems_, I think of B:CP in the same class as Aristotle's ideas about "final cause" and William James's ideas about purpose and intention -- only Bill's work is better -- he turned *ideas* about intention into a *science* of intention. But, then, I am a biased lover of only the lowest levels of behavior and perception. ;-)

. . .

> A large part of my initial impressions of weakness or oversimplicity in HPCT was because you (particularly Rick & Tom) appeared to be saying more elaborate structures were unneeded.

Ah, but that was the impression you formed back when you had read nothing, but burst on the scene telling us were using a model that was (or that might be) overly simple. Didn't you want us to reply to you on that topic? Should we have ignored you, or assented without protest to your claims?

> I recognize that most PCT research is focused on levels 1-3 of the perceptual control hierarchy, partially due to amenability to modeling, but creates the impression that it is all HPCT consists of.

Sorry, but have we read the same book? And which net are you reading from? What you say about the levels we focus on is a pretty serious distortion of the truth. You have repeated that assertion many times, even though it is patently false. Why do you do that? I don't understand. Can you clue me in? Remember, though, I probably won't understand your answer if it is more complex than a second- or third-level explanation. ;-))

> The simple perceptual control loop does not explain all of what is commonly called behavior, much less psychology.

I'm glad you agree with us on that!

> The reorganization control hierarchy, memory,and the perceptual & memory switches are necessary constructs to deal with common concepts of planning, imagination, etc.

The PCT and HPCT models aren't intended to "deal with concepts." They are intended to explain the phenomenon of control.

> Further the levels of the PC hierarchy above 3 are rarely explained, other than by occasional asides - e.g. "The lower six levels are concerned with control of intensities, sensations, configurations, transitions, events, and relationships. "

Hmm. That' interesting. Are we reading different nets, again? Even the simplest tracking task reaches the sixth level, at a minimum. And I have shown how easily a tracking task can be modeled as requiring a program level -- that's about the eighth level in a ten or eleven level hierarchy. Are we talking about the same model, Paul?

> It is not intuitively obvious that these constructs derive from the simple concept of control. The terms are also subject to variable interpretation.

Yes. That's why we keep telling you the theory and model are not about the words. The only way we know for a person to "get it" at the gut level -- to get it "in your bones" -- is to play with the model. Have you run any of the demonstrations?

> As the [From Bill Powers (940803.1510 MDT)] exchange with comp.ai indicates, your critics and the uninitiated dismiss you because you do not publicize the 'non-core' ideas, except through casual asides (in my experience), and so they presume such concepts are not incorporated in PCT.

"In my experience." That is an important qualifier, Paul; your experience to now is very limited -- not a criticism, but a restatement of what you have told us. Keep reading -- and run some demonstrations. :-))

While you read, please remember that until *very* recently, only three people were working (most often on their own time and at home) to seriously test the PCT model in research and simulations. Only three, and I was one. *No wonder* we have accomplished so little of what needs to be done!

Later, Tom

Date: Fri, 5 Aug 1994 18:14:20 CST
Subject: Re: Appealing to Complexity Worshippers

Tom Bourbon [940805.1801]

>[Paul George 940805 16:00]

>[Rick Marken (940508.0945)], also Avery Andrews (940805.1347)
. . .

> However, if propagation of truth or knowledge is your goal, rather than a cloistered purity of thought and purpose, you need to communicate with the unenlightened.

I'll let that one pass. It would be too easy for me to slip into thinking of it as an insult, and I don't believe that is what you intended.

> IMHO you need to first communicate that the theory can encompass the complexities that many psychologists and commoners ;-) observe in human behavior, and then focus them towards the basics. Then they can be amazed at how much so called complex behavior simple hierarchical control can model. But first you have to catch the mule's attention.

Paul, why don't *you* take on that role? Seriously. Make a good-faith attempt to publish some genuine material on PCT applied to the complex issues of the day in psychology. Don't submit the watered down and distorted versions so many others have published. You would do a great service for "the cause."

Actually, I believe the only way to get through to the academics in the behavioral sciences is for us to demonstrate that we can help someone answer a very simple, but very important question, or solve a real problem, then tell the academics that if they want to do the same thing, here is how it is done -- PCT. Nothing else will work. We are trying that approach now. (Bill P., I'll fill you in on latest developments here at the med school a little later -- maybe Monday.)

> If you make it easy for them to classify you as monomaniacal simpletons, then they will. It is much easier to dismiss you than to consider changing their way of thinking or studying a new approach. Why should they waste time investigating something that "obviously can't" scale up or map to reality?

But, Paul, even when we have tried to "dumb down" the theory, they have rejected it. There is more to the problem than has yet met your eye.

> Being "despised and rejected of men" may produce a kind of group pride or feeling of belongingness, but is not very effective for propagating your mnemes. The history of science deals with the mnemes that spread widely enough and survived long enough to make an impression. I have no doubt that university (and monastery) archives are littered with seminal dissertations that never again saw the light of day.

I love it when people tell us we *elected* to live in this relation to "real" scientists! ;-))

Later, Tom

Date: Sat, 6 Aug 1994 02:26:42 -0400
Subject: Re: B:CP Reactions - Cut 2

<[Bill Leach 940805.23:50)] >[Paul George 940804 1600]

Paul, since no one else seemed to mention this...

An additional reason for what you see (or don't as the case may be) is that PCT is THEORY and HPCT is HYPOTHESIS.

In a very real sense, there is no argument about the "correct" theory for behavior because there IS only one THEORY and that is PCT.

For a physicist use a label of hypothesis for most of the proposals in the rest of behavioral sciences is a serious affront to the term (indeed even the word "science" is "insulted" in such application).

Thus, the theory of PCT has support of experimental evidence and is in the unique position (for this field) of no been challenged by any experimental evidence.

HPCT OTOH, has a "great deal" going for it but does not have the same sort of experimental evidence. It is a hypothesis in the same sense as such would be found in other "hard" sciences. Careful consideration of the experimental data suggests that the hypothesis is sound but does not suggest that this can be the only answer.

You must also remember that this crowd is rather "steeped" in the traditions of "pure sciences" and is not particularly "big" on trying to do an "end run" on the "community". Little doubt (in my mind) that a major reason for this reluctance is the "massive" abuse of scientific principles that are seen daily in the popular media.

It is already "bad enough" as far as what can happen to a scientific treatise in "peer review" without foregoing that step and just "handing" it to the media. A media whose commitment truth an objectivity might be considered to be less than always honorable <cough, gag> (watering down that last sentence was tough for me to do).

Having said that however, I think that Rick emphasized what is the most significant reason; Until one begins to grasp the significance of what closed loop control means at the simplest level and begins to realize that just the phenomenon of control itself teaches that much of what is viewed as "complex behavior" is rather no more than the physical consequences of the action of a control system.

A difficulty that I see is that decomposition of higher level references could conceivable proceed in many different ways. People could (and do) argue about how these different structure details may vary. However, the evidence for the phenomenon of control is overwhelming and models that can be constructed of only a few loops (or even one) provide remarkable fidelity.

There is not the slightest doubt in anyone's mind that the number of control loops in each of the models is vastly smaller in number than in any of the human subjects. However, the issue is that the models are control systems and work with such high fidelity that the models can reasonably be used for predictive functions.

I guess what I am trying to say is that PCT is much like electronics itself. We do not need to explain exactly how a specific electron works its way through semi-conductor lattice to understand the operation of a transistor nor is such knowledge required to build a computer.

However, the ability to build such a computer is greatly improved when the phenomenon of amplification (in the transistor) is understood.

Thus in PCT, we don't even know how many "wires" there are much less how they all route. We don't know if there are a bunch of "math co-processors" "up there" or not. We don't know what sort of "special function processors" might exist if any. What is rather well established though, is that once a controlled variable is identified, it is possible to demonstrate reliably that living beings control perception.

It is the phenomenon of control that is the "sore spot" not the details.

-bill

Date: Sun, 7 Aug 1994 11:26:48 -0700
Subject: What's to explain?

[From Rick Marken (940807.1130)] Paul George (940805 16:00)

> Of course a major reason for 'slow' progress may have been that affordable computers with sufficient power only became available in the last 5-10 years.

Or it could be because, over the last 20 years only three people have been doing PCT research, and doing it in their spare time to boot, while other people were leaning over their shoulder saying "yeah, but why don't you guys study something really important, like the kind of things those AI and complex systems people are studying?" Sheeez.

Tom Bourbon (940805.1719) and (940805.1801)--

I agree with everything you say, Tom, only more so.

I think one of the things that Paul George doesn't quite appreciate yet is that PCT is not an alternative explanation of the "facts" that psychology has already "discovered". According to PCT, nearly all the facts of psychology are hogwash; there is almost nothing in psychology (including most cognitive science, AI, etc) for PCT to explain. When we are asked "how would PCT explain such and such" it usually turns out that "such and such" is just words used to describe an (often non- replicable) statistically significant result (averaged over many people) that seemed "psychologically significant" for a brief, trendy time. PCT doesn't explain random phenomena; it explains control.

Best Rick

Date: Tue, 9 Aug 1994 17:34:28 EDT
Subject: Re: B:CP evaluation, part 2

Paul George 940809 1730]

[From Bill Powers (940805.1150 MDT)] and others

Thanks for the response.

> Well, sheesh, how much can we talk about at once? We talk about the things we have done the most work with. Actually, if I had responded to your initial inquiries with a complete summary of everything that is in the PCT model, you would have thought I was some kind of nut.

I think there is a middle ground. If you had posted the final PCT model and a description of the hierarchical layers _I_ would have been much happier. The intro stuff (monthly post) is good but IMHO doesn't identify the existence of a lot of PCT concepts, much less the HPCT hypotheses (ack Bill Leach 940805.23:50).

If Dag had posted the counter-response to his response on comp.infosystems you would see how just asserting the wonders of PCT is received [e.g. "In a very real sense, there is no argument about the "correct" theory for behavior because there IS only one THEORY and that is PCT."- Bill Leach 940805.23:50].

It was to the effect of "you can convince yourself you can control everything if you like" (I trashed the article without thinking). The evangelical approach comes across as "Cast aside your illusions and open yourself to the truth of PCT. Through faith and study will come true enlightenment". (I exaggerate a mite ;-). Facts are less important than perception (speaking of preaching to the choir)

> I'd love to have an FAQ. How much do they cost, and can I get a used one cheaper? By the way, what is an FAQ?

Assuming you are serious, it is a Frequently Asked Questions list. Customarily a newsgroup has one which is posted periodically and is available for downloading from a FTP site. Similar to the monthly posting.

> ...It might turn out to be a really effective document. "How I finally got the idea of PCT," or something. In fact, your suggested outline sounds pretty good to me.

I might be able to in the future, but I will have to get a new copy of B:CP to avoid misstatement. I had to send mine back to the Cincinnati library last Friday. However, if the Book exists in softcopy (word processors were rare in the old days;-), the extraction would be faster for you. I would be happy to edit/embellish.

>Tom Bourbon [940805.1801]

> Paul, why don't *you* take on that role? Seriously. Make a good-faith attempt to publish some genuine material on PCT applied to the complex issues of the day in psychology. Don't submit the watered down and distorted versions so many others have published. You would do a great service for "the cause."

First of all because a paper from a BSBA in Management Information Systems would be unlikely to be well received by a PhD review committee of a psychology journal. However, If I get into Case Western's Organizational Development Masters Program next year, I would like to use PCT in some of my projects and perhaps publish. Any suggestions?

> But, Paul, even when we have tried to "dumb down" the theory, they have rejected it. There is more to the problem than has yet met your eye.

Granted, but I don't think 'dumbing down' the theory is the key. The idea is to state that structures are hypothesized to address their concerns, and that they are at least as well justified as their own. Then you can direct them to the details of control theory and 'the test'. Actually, selling to Psychologists may be harder than to others due to vested interests. The internet however allows you to reach a larger audience who can benefit (say students). Censorship is far more difficult. But.... you have to make the materials easily accessible.

> I love it when people tell us we *elected* to live in this relation to "real" scientists! ;-))

Nope, but after being knocked around for a number of years you tend to 'control for rejection'. Note Bill P's comment to a new poster last week (?) saying 'we've pretty much given up trying to publish outside our own little journal' (paraphrased).

>Tom Bourbon [940805.1719]

> Like Rick Marken (940508.0945), I think you understate the case. I would place Powers's B:CP, and his theory of behavior in the same league as Newton's "Principia,"...

I tend towards understatement. However, the ideas are perhaps more revolutionary to Psychologists than to others such as engineers. Ask me again after I have re-read the book a couple of times. One pass probably isn't sufficient to extract all it contains.

> Sorry, but have we read the same book? And which net are you reading from? What you say about the levels we focus on is a pretty serious distortion of the truth. You have repeated that assertion many times, even though it is patently false. Why do you do that? I don't understand. Can you clue me in?

That Is my perception based upon what I see discussed, and my understanding of the levels. There may have been other discussions at other times, but the net has an extremely limited memory. Consider that you as a group may assume a large body of common knowledge or understanding that may not be shared by all on the list. Also consider that you may not have successfully communicated your ideas. From discussions of the tracking program, it appears only to have 3 hierarchical loops. IMHO they all appeared to be focused on just level 2 & 3(1 being simulated). However, I haven't seen the code or detailed papers.

[From Rick Marken (940807.1130)]

> I think one of the things that Paul George doesn't quite appreciate yet is that PCT is not an alternative explanation of the "facts" that psychology has already "discovered". According to PCT, nearly all the facts of psychology are hogwash; there is almost nothing in psychology (including most cognitive science, AI, etc) for PCT to explain.

I appreciate it, but _saying_ it initially to a body of psychologist (or most 'outsiders') guarantees that they will dismiss anything else you have to say out of hand.It makes you _sound_ like crackpots. How did you all respond when I just suggested that the canonical model _possibly_ needed elaboration (In fact it had such an elaboration in B:CP)? When you thought I didn't take you seriously?. I'm talking marketing or salesmanship here. As I said above, the facts or truth don't matter, it is perception and assumption. They must think there is value in PCT before they will evaluate it. And learning depends upon recognizing that there is something to be learned.

Later Paul George

Date: Wed, 10 Aug 1994 09:15:23 -0700
Subject: PCT Research

[From Rick Marken (940810.0915)]

Tom Bourbon said:

> I love it when people tell us we *elected* to live in this relation to "real" scientists! ;-))

Paul George (940809 1730) replies:

> Nope, but after being knocked around for a number of years you tend to 'control for rejection'.

This implies that we would treat acceptance of PCT as a disturbance. Why not try "The Test" and see if it is. But you have to be careful about how you define the controlled variable. If someone says "I love PCT and I accept everything you say about it; I think there can be no doubt that PCT is an important approach to understanding how people control their behavior", would you consider that a disturbance or a non-disturbance to a person controlling for perceiving "rejection of PCT"?

Tom Bourbon again:

> I would place Powers's B:CP, and his theory of behavior in the same league as Newton's "Principia,"...

Paul:

> the ideas are perhaps more revolutionary to Psychologists than to others such as engineers.

Unfortunately, this is not the case. Engineers reject PCT as vigorously as do psychologists. Does it really seem to you that the notion of "control of perception" is taken for granted in control (or any other kind of) engineering? B:CP explains in detail why "control of perception" is the central fact of purposeful behavior - - ie. the behavior of all living systems (and some non-living ones). The power of this fact is extraordinary; it explains how purposes can be carried out, why behavior appears to be S-R, selected by consequences or planned output, why organisms get into intra- and interpersonal conflict, etc. No one (NO ONE) before William T. Powers noticed this fact about control or understood its implications. That's why B:CP ranks with the Principia -- in my mind, anyway.

Me:

> I think one of the things that Paul George doesn't quite appreciate yet is that PCT is not an alternative explanation of the "facts" that psychology has already "discovered". According to PCT, nearly all the facts of psychology are hogwash; there is almost nothing in psychology (including most cognitive science, AI, etc) for PCT to explain.

Paul replies:

> I appreciate it, but _saying_ it initially to a body of psychologist(or most 'outsiders') guarantees that they will dismiss anything else you have to say out of hand.

I agree -- and we never do that. We begin (or should begin) by trying to convince psychologists that behavior IS control -- and showing why this is the case. Nevertheless, we still get questions about how PCT explains "such and such" or other non-control phenomenon. And, indeed, when we answer these questions, we are "dismissed out of hand" -- because, of course, it was the favorite phenomenon of the person who asked. Any suggestions about how to deal with this problem?

Best Rick

Date: Wed, 10 Aug 1994 11:10:31 -0600
Subject: Hype

[From Bill Powers (940810.0930 MDT)] Paul George (940809.1730)

You make some good points, again. Obviously we all could have done better in communicating PCT, had we known what kind of response we would get from the mainstream. I was pretty naive: I thought that just describing the model as clearly as I could would be enough!

I strongly agree with you about evangelical overstatements of the accomplishments of PCT. There are many areas in which we have no actual data about the worth of PCT as a useful model of behavior. We're working on that, but I truly wish we could all remember to limit our claims to what we can back up with demonstrations and data. There's nothing wrong with extrapolations, as long as they're labelled as such. But touting PCT as the One True Faith is a certain turnoff for any intelligent person. I'm especially disturbed by claims that PCT is a "proven theory" in areas where it is no such thing. If asked, I will disavow such claims.

Best to all,

Bill P.

Date: Wed, 10 Aug 1994 14:58:10 EDT
Subject: Re: PCT Research

Paul George 940810 15:00 >[Rick Marken (940810.0915)]

Me

>> Nope, but after being knocked around for a number of years you tend to 'control for rejection'.

> This implies that we would treat acceptance of PCT as a disturbance. Why not try "The Test" and see if it is.

Nailed by sloppy phrasing again. I intended 'minimization of rejection'. More accurately tending to reduce the pain or discomfort brought about by rejection as people tend to desire approval and recognition. You tend to expect rejection in certain quarters. Eventually you stop banging your head against the wall. Consider the image of an adult lion being tied up with a string that restrained him in cub-hood.

> Unfortunately, this is not the case. Engineers reject PCT as vigorously as do psychologists. Does it really seem to you that the notion of "control of perception" is taken for granted in control (or any other kind of) engineering?

Well, to me it _is_ (on the other hand I am not renowned for conventional thought processes). Your universe is what you can perceive as compared to what you would like it to be (dodging for the moment whether these are intrinsic values or remembered perceptions). Any control you attempt can only be detected through perception. People control for illusion all the time ('image', the Vedic concept of 'maya'). An engineer or scientist relies upon measurements and assumptions (or associations) about how various measurements relate. In software engineering we are dealing almost entirely with intellectual artifacts using symbolic techniques.We are dealing with perceptions and assumptions about the world (a model) or another set of symbols. The rules are what ever we agree (or appear to agree) they are. It seems obvious to me that you try to control results, not actions.

PCT did reveal to me a number of twists I hadn't previously considered and an amazingly elegant and powerful mechanism to describe the behavior. Your experiments demonstrated that a lot can be done by simple feedback and that a lot of 'computation' can be eliminated. To that extent PCT _is_ truly a breakthrough.

> Nevertheless, we still get questions about how PCT explains "such and such" or other non-control phenomenon. And, indeed, when we answer these questions, we are "dismissed out of hand" -- because, of course, it was the favorite phenomenon of the person who asked. Any suggestions about how to deal with this problem?

A tough one. In some cases you may be able to rephrase the phenomenon as a control problem as Bill P (?) did with the operant conditioning data for the rats. In other case you can wave your hands ( cries arise: Heretic! Heretic! :-) and indicate that it would be handled in the higher levels or reorganization in a mechanism not yet studied. You could even be real sneaky and suggest that their model might be incorporated within PCT and encourage them to try to formulate the model. The approach of 'we haven't addressed it yet, but show me how your theory/concept can be justified in fact or experiment'(a.k.a. 'we're no worse than you') can also work, albeit not winning friends. (of course you have never tried anything like this ;-})

In reality however, it is almost a no-win situation. The person in question (particularly a department head or other professor emeritus) would reject _any_ theory not his own, as his identity or reputation is based upon it. But, if you can get them to acknowledge that PCT as a possible interpretation, though inferior to theirs, you may be able to convince the next generation.

There is a theory that ideas in science don't change, the proponents of the old ones just die off eventually and no longer sit on review boards, thesis committees or chair departments. "When a respected senior scientist says something can be done, he is almost certainly right. When he says it cannot, he is almost always wrong"- R.A. Heinlien

Back to the bushes, Paul

Date: Wed, 10 Aug 1994 16:55:50 CST
Subject: Re: B:CP evaluation, part 2

Tom Bourbon [940810.1655] >Paul George 940809 1730]

>>Tom Bourbon [940805.1801]

>> Paul, why don't *you* take on that role? Seriously. Make a good-faith attempt to publish some genuine material on PCT applied to the complex issues of the day in psychology. Don't submit the watered down and distorted versions so many others have published. You would do a great service for "the cause."

> First of all because a paper from a BSBA in Management Information Systems would be unlikely to be well received by a PhD review committee of a psychology journal.

Do you *really* think they would treat you any worse than they do the rest of us? ;-)) . . .

>> But, Paul, even when we have tried to "dumb down" the theory, they have rejected it. There is more to the problem than has yet met your eye.

> Granted, but I don't think 'dumbing down' the theory is the key. The idea is to state that structures are hypothesized to address their concerns, and that they are at least as well justified as their own.

Oops. That bit about "at least as well justified as their own" would bring down a nuclear barrage. Believe me, I've lived through a few. I *never* try to tell editors and reviewers that I am as good as they are -- no any more. And any attempt I have made at saying, "I'm trying to address your concerns; really, I am," has met a similar fate -- "how can he (me) presume that continuous data on a computer game could be of any relevance/interest to us?" and things like that. Even saying something like, "Gee, folks, I know this isn't what you are used to seeing, but I'm offering it in the hope that *someone* might find it useful," brings down the hammer -- or is it the ax? Either way, it hurts.

> Then you can direct them to the details of control theory and 'the test'.

I (we) rarely get that far. You see, we have this real problem: we are talking about a phenomenon -- control -- that isn't even recognized by most behavioral scientists. Somewhere up front, we must say so, and try to show them what we are talking about. Just about there is where all hell breaks loose, long before we can get to theory and models and esoteric things like that. :-(

> Actually, selling to Psychologists may be harder than to others due to vested interests. The internet however allows you to reach a larger audience who can benefit (say students). Censorship is far more difficult. But.... you have to make the materials easily accessible.

Hey, here we are, for all the world to see. How much more "accessible" can we get?

>> I love it when people tell us we *elected* to live in this relation to "real" scientists! ;-))

> Nope, but after being knocked around for a number of years you tend to 'control for rejection'. Note Bill P's comment to a new poster last week (?) saying 'we've pretty much given up trying to publish outside our own little journal' (paraphrased).

Hmm. Like Rick, I wonder what it would be like to have a reference for non-zero perceptions of rejection. How much of it do you think each of us prefers? (It's obvious Rick has a reference for much more than I do.) How would you test for that controlled perception? Accepting a few of our articles should send us into deep bouts of depression and reorganization, shouldn't it? I wish someone would test that possibility, just so I could find out if it is true.

>>Tom Bourbon [940805.1719]

>> Like Rick Marken (940508.0945), I think you understate the case. I would place Powers's B:CP, and his theory of behavior in the same league as Newton's "Principia,"...

> I tend towards understatement. However, the ideas are perhaps more revolutionary to Psychologists than to others such as engineers. Ask me again after I have re-read the book a couple of times. One pass probably isn't sufficient to extract all it contains.

OK. Let me know when you finish it again and I'll ask again.

>> Sorry, but have we read the same book? And which net are you reading from? What you say about the levels we focus on is a pretty serious distortion of the truth. You have repeated that assertion many times, even though it is patently false. Why do you do that? I don't understand. Can you clue me in?

> That Is my perception based upon what I see discussed, and my understanding of the levels. There may have been other discussions at other times, but the net has an extremely limited memory.

The net has *no* memory, and some subscribers approach that condition. ;-))

> Consider that you as a group may assume a large body of common knowledge or understanding that may not be shared by all on the list. Also consider that you may not have successfully communicated your ideas. From discussions of the tracking program, it appears only to have 3 hierarchical loops.

In my programs I model the person with only one loop, but it represents control by a system with about a sixth-level reference signal. Don't be mislead by the number of loops in the diagrams in our papers.

> IMHO they all appeared to be focused on just level 2 & 3 (1 being simulated). However, I haven't seen the code or detailed papers.

Ah, then may I recommend a few papers? Time to run, Tom

Date: Thu, 11 Aug 1994 09:51:38 -0600
Subject: Re: acceptance of PCT

[From Bill Powers (940811.0615)] Paul George (940810.1500)

I'm following your conversation about PCT research and acceptance at the same time I'm working on a paper that compares the inverse- kinematic/dynamic model of arm control with a PCT model. It seems to me that our problem goes much deeper than that of merely gaining acceptance.

There is something seriously wrong with a science that could accept the idea that the brain achieves organized behavior by computing the joint torques that will produce it. To do these computations, the brain requires information about the masses and moments of inertia of the arm segments, the properties of muscle contraction both static and dynamic, the properties of nonlinear muscle springs, the variations in mechanical advantage as the joint angles change, the physics of dynamic movements, the trigonometry of spatial relationships, the location of targets in visual space, and the initial state of the arm in terms of positions and velocities relative to a possibly moving target. The brain is required to do computations involving signs and cosines, multiplications and additions and divisions and subtractions -- and to do all this in real time with so much precision that after a double time integration, the final result is the kind of pointing accuracy we observe in the real system. The body is assumed to be as stable as a rock, the muscles to be immune to fatigue, and the environment to be free of unpredictable disturbances.

There is an air of dreamlike unreality about this model. It is assumed that anything a mathematician can do with pencil and paper and symbols, the brain can do with neural currents (without symbols). It is assumed that all the knowledge that external observers have obtained in 300 years of studying the physics and geometry of the arm and environment is available in real time to the neural computers, even those of a monkey or a mouse. Just how this information becomes available is not even considered.

So the question naturally comes to mind, "Why should PCT researchers be interested in acceptance of PCT by people who could bring themselves to believe in such models?" Just what would such people be accepting? Another abstract mathematical scheme with no more justification than any others they have believed in? Another form of magic? Another kind of prediction that is true some of the time, for some systems, under some conditions? Is there the slightest chance that they would grasp PCT and use it as a real theory? Or would they just use it as another "perspective" on nature, to be adopted or not at one's convenience?

PCT does not explain all behavior of all organisms, just as physics does not explain all behavior of all matter. But the failures of explanation are of a kind different from those found in psychology. They aren't statistical; they are total. There are phenomena that we simply don't understand, and we know we don't understand them. There is no point pretending that we do understand, until we actually do. What psychology is missing is a concept of this point of understanding where you are simply backed into a corner, and no matter what alternatives you may think of you are continually forced back to the same view. Everything else is ruled out. If there is a better explanation, and one knows there will always be a better one some day, it is simply not available now. The trail, for now, ends here. When we reach such a point of understanding there is no choice but to proceed as if it is true.

The kind of understanding that comes out of psychological theories is the same kind one could obtain with an understandingness pill, like the sense of godlike comprehension that some people seem to get from cocaine. It is a decision to believe rather than to look further.

Psychological theorizing does not stop when there is no other place left to go; it stops when an already weak sense of skepticism about one's own ideas is completely suspended. So we have dozens of competing "microtheories," all existing at once and all accepted as part of psychology, and hardly a one destined to last more than five years.

Is this the field in which we aspire to gain acceptance? Best, Bill P.

Date: Thu, 18 Aug 1994 15:23:35 EDT
Subject: Re: Ayn Rand and PCT

[Paul George 940818 15:25]

For those who are interested in testing whether PCT would be accepted by the followers of Objectivisim there is a newsgroup alt.objectivism where current advocates hang out. I don't have access to it, but there are occasional cross posts to sci.econ. They may also be found on various of the libertarian forums.

Date: Fri, 19 Aug 1994 08:06:57 -0700
Subject: Accepting vs Learning PCT

[From Rick Marken (940819.0800)] Paul George (940818 15:25)

> For those who are interested in testing whether PCT would be accepted by the followers of Objectivisim there is a newsgroup alt.objectivism where current advocates hang out.

I am no longer interested in whether or not anyone in particular would accept PCT; clearly they would. What I am interested in is whether or not anyone would _learn_ PCT, by which I mean learning how to identify controlled variables (The Test) and how to produce models that control these variables in the same way that the organism controls them.

In order to learn PCT one must have at least a basic grounding in mathematics, programming, physical science, neurology, physiology, and control engineering. I think it also helps to be somewhat familiar with conventional psychological science, if only to be able to explain what PCT is _not_.

Right now I estimate that there is probably a 4 order of magnitude difference between the number of people who have accepted PCT (something on the order of 10,000; I include in this group Reality Therapists as well as research psychologists like Carver, Scheier and Lord) and those who have learned it in the sense I describe above - - that is, done the research and modelling (something on the order of 7). Because a large proportion of those who have accepted PCT have not learned much about what they have accepted, there is a good chance that the PCT they accept is not the PCT reflected in the control phenomena and models studied by those who are doing the PCT research and modelling.

All this will change when the Center for the Study of Living Control Systems becomes a reality. One of the main goals of the Center will be to teach willing members of the approximately 9,993 who already accept PCT how to study and model what PCT is actually about -- the purposeful behavior of living control systems.

Best Rick

Date: Fri, 19 Aug 1994 17:45:20 CST
Subject: Re: Accepting vs Learning PCT

Tom Bourbon [940819.1741] >[Rick Marken (940819.0800)]

>>Paul George (940818 15:25) --

>> For those who are interested in testing whether PCT would be accepted by the followers of Objectivisim there is a newsgroup alt.objectivism where current advocates hang out.

> I am no longer interested in whether or not anyone in particular would accept PCT; clearly they would. What I am interested in is whether or not anyone would _learn_ PCT, by which I mean learning how to identify controlled variables (The Test) and how to produce models that control these variables in the same way that the organism controls them.

Yes. The problem is not merely one of finding people who accept PCT. As Rick said later in his post, if we count people who have accepted one or another version of control theory as a way of thinking about and talking about behavior, they probably number in the low tens of thousands, at least. The problem is finding people who accept control theory who are also willing to go the next part of the journey and learn how to do PCT science -- testing to identify variables that people control, modeling control in living systems, and applying the theory _rigorously_, wherever they can. I can probably count the number who have gone that far on my personal (limited) assortment of fingers and toes. . . .

> All this will change when the Center for the Study of Living Control Systems becomes a reality. One of the main goals of the Center will be to teach willing members of the approximately 9,993 who already accept PCT how to study and model what PCT is actually about -- the purposeful behavior of living control systems.

Surely among that many people are a few who will drop everything else and do the necessary work. Nothing less will do, not because PCT is a religion that requires all neophytes pass through a monastic experience at The Center prior to full initiation, but because PCT is a demanding science that cannot be mastered by taking short cuts. It is _not_ about slogans, catchy one-line phrases, and clever ways to manipulate other people. It is _not_ about how to gain a new "perspective" or "framework" from which to reinterpret (and thereby cling to) every old idea you already had about behavior. It _is_ about frequently staring your own ignorance, prior conceptions, and lack of skills squarely in the face, taking a deep breath, then getting to work, doing whatever _is_ necessary and throwing away whatever is _not_ necessary.

I have joined Rick in giving up on the idea that we should spend any more of our time trying to get more people "interested" in PCT. Another radical is born.

Later, Tom

_