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

This thread discusses what is and what is not research that meets the criteria required by PCT. Towards the end, (after much discussion) some agreement is reached [Comparing views of behavior, Rick Marken (950605.1330) and Stress, Proving PCT, Rick Marken (950605.2145)] on the reasons for difficulties of persuading researchers of the difference between their views and the PCT view.

Beyond the first two posts, note particularly:

Controlled variables vs. side-effects,
[From Bill Powers (950527.0950 MDT)]

PCT observations of behavior
[From Bill Powers (950529.0700 MDT)]

Test; misc
[From Bill Powers (950530.0945 MDT)]

Early PCT research
[From Bill Powers (950602.1505 MDT)]

Early PCT Research
[From Bruce Abbott (950603.1340 EST)]

Old points of view
[From Bill Powers (950603.1550 MDT)]

Something to be Said
[From Bruce Abbott (950604.1510 EST)]

Date: Tue, 9 May 1995 21:23:28 -0700
Subject: PCT Research

[From Rick Marken (950509.2120)] Bruce Abbott (950509.1240 EST)

> What do you DEFINE as PCT research? That is, to qualify as PCT research, what elements must be present? I'd like to be able to examine, say, a piece of published research and using your criteria, classify it as PCT or non-PCT.

Excellent idea!! Here's my take on _real_ PCT research:

PCT research is research aimed at determining the perceptual variables that organisms control and how they control them. For research to be considered PCT research the following elements must be present:

1. Most important: PCT research is characterized by a vision of organisms as _controllers_. PCT research sees the organism as a system that is controlling it's own perceptual inputs. PCT researchers want to find out what the organism is trying to perceive; they want to see the world from the organism's perspective. In PCT research, the term "controlling variables" always refers to what organisms do to variables in the environment, NOT what variables in the environment do to organisms.

2. PCT research is always characterized by a determination of the variable(s) being controlled by the organism. There is always evidence that disturbances to some result of an organism's actions have almost no effect on that result. All analysis and modelling is done after a controlled variable has been identified and while it is being monitored.

3. PCT research is aimed at finding _lack_ of effect of environmental variables on behavioral variables. Behavioral variables (like arm position) are possibly controlled results of action; lack of effect of an environmental variable on a behavioral variable suggests that the behavioral variable is under control.

4. PCT research always looks at controlling on a one-organism-at-a-time basis. Statistical tests are of no use in the study of living control systems; PCT research is based on modelling, not statistics.

5. Published PCT research is always significant (in the normal sense and the statistical sense). PCT research is significant if a controlled variable has been identified and can be monitored. PCT research is not published unless a controlled variable has been identified. PCT research is not of the "manipulate and pray for statistical significance" variety. PCT researchers keep developing their research techniques until controlled variables can be clearly identified and readily monitored. PCT researchers only publish their results when they KNOW what's being controlling and how.

Good luck on finding some good examples of PCT research out there in the conventional literature.

Best Rick

Date: Wed, 10 May 1995 15:29:54 -0500
Subject: Identifying PCT Research

[From Bruce Abbott (950510.1530 EST)]

>Rick Marken (950509.2120) --

Your criteria for PCT research seem clear [they parallel to some extent my "Principles of PCT-Guided Research" posted earlier] but could stand to be reorganized to eliminate redundancy. In the following I quote your statement and then attempt a summary and, in some cases, clarification.

> 1. Most important: PCT research is characterized by a vision of organisms as _controllers_. PCT research sees the organism as a system that is controlling it's own perceptual inputs. PCT researchers want to find out what the organism is trying to perceive; they want to see the world from the organism's perspective. In PCT research, the term "controlling variables" always refers to what organisms do to variables in the environment, NOT what variables in the environment do to organisms.

PCT research acknowledges that organisms act as controllers of their own perceptual inputs.

PCT research is aimed at determining what perceptions the organism is trying to perceive.

PCT research identifies what behavior does to environmental variables, not what environmental variables do to behavior.

> 2. PCT research is always characterized by a determination of the variable(s) being controlled by the organism. There is always evidence that disturbances to some result of an organism's actions have almost no effect on that result. All analysis and modelling is done after a controlled variable has been identified and while it is being monitored.

PCT research is aimed at determining what the organism is trying to perceive, by applying the Test.

After the controlled perception has been identified and monitored, analysis is done and control models constructed which appear to satisfactorily account for performance on the test task.

> 3. PCT research is aimed at finding _lack_ of effect of environmental variables on behavioral variables. Behavioral variables (like arm position) are possibly controlled results of action; lack of effect of an environmental variable on a behavioral variable suggests that the behavioral variable is under control.

PCT research is aimed at identifying controlled perceptions via the Test, i.e., the organism's response to the disturbance of those variables. If the action counteracts the disturbance, the variable on which the disturbance acts is controlled.

> 4. PCT research always looks at controlling on a one-organism-at-a-time basis. Statistical tests are of no use in the study of living control systems; PCT research is based on modelling, not statistics.

PCT research is single-subject research.

In PCT research, models of the control system are constructed; the adequacy of these models is assessed by comparing the model's performance to that of experimental participants, not by statistical hypothesis testing.

> 5. Published PCT research is always significant (in the normal sense and the statistical sense). PCT research is significant if a controlled variable has been identified and can be monitored. PCT research is not published unless a controlled variable has been identified. PCT research is not of the "manipulate and pray for statistical significance" variety. PCT researchers keep developing their research techniques until controlled variables can be clearly identified and readily monitored. PCT researchers only publish their results when they KNOW what's being controlling and how.

PCT research is successful if it identifies and successfully monitors a controlled variable. PCT research is published only if a controlled variable has been experimentally identified and a plausible control model has been developed to account for the behavior of the participants on the experimental task.

Statistical hypothesis testing plays no role in the evaluation of the research findings.

Do those restatements accurately reflect your intended meanings? Is there anything left out? Also, would you consider research on, e.g., the control mechanisms responsible for eyelid positioning and the eyeblink as PCT research? (The above criteria certainly would seem to include such research.)

Regards, Bruce

Date: Wed, 10 May 1995 17:01:26 -0400
Subject: Catching flies

[From John Anderson (950510.1630)]

>Rick Marken (950509.2120)] >>Bruce Abbott (950509.1240 EST)

>> What do you DEFINE as PCT research? That is, to qualify as PCT research, what elements must be present? I'd like to be able to examine, say, a piece of published research and using your criteria, classify it as PCT or non-PCT.

> Excellent idea!! Here's my take on _real_ PCT research:

> [...5 criteria for published PCT research...]

> Good luck on finding some good examples of PCT research out there in the conventional literature.

Bruce and Rick, here's a paper from the 28 April 1995 issue of Science you might consider:

MK McBeath, DM Shaffer, MK Kaiser (1995) "How baseball outfielders determine where to run to catch fly balls" Science 268:569-573.

This paper seems to has a PCT flavor to it, and the topic was mentioned recently (4/10) on CSG-L. Here's the abstract and a brief summary:

'Current theory proposes that baseball outfielders catch fly balls by selecting a running path to achieve optical acceleration cancellation of the ball. Yet people appear to lack the ability to discriminate accelerations accurately. This study supports the idea that outfielders convert the temporal problem to a spatial one by selecting a running path that maintains a linear optical trajectory (LOT) for the ball. The LOT model is a strategy of maintaining "control" over the relative direction of optical ball movement in a manner that is similar to simple predator tracking behavior.'

A previously-proposed model for how fly balls are caught is the optical acceleration cancellation (OAC) model. In it, players run toward or away from the ball along a line in the plane of its trajectory in such a way that the ball rises at a constant rate. This is the model referred to by Bill Powers (950410.0900 MST), responding to Bill Leach (950409.2231 EDT):

> To your example of the ballplayer not chasing a ball that will go well over his position, I can add another twist (that I've mentioned before). An outfielder catching a fly ball headed straight toward him, it is said, moves so as to keep the apparent rate of rise of the ball constant and slow. If this perception is controlled, the ball will arrive within catching range. If you look at the behavior of the ball player, it will seem that the sight of the ball causes the player to run in anticipation of the catch, but in fact something is under control all the way (or at least intermittently as the player casts glances over his shoulder while running toward the fence).

If the fielder runs too far in, the ball rises too fast; if he runs too far out, it rises too slow. In the OAC view, trying to do this on a line out of the trajectory plane of the ball, ie when the ball is not hit directly at you, will be more difficult because then you have to take into account both vertical _and_ horizontal motion parameters. But baseball players agree that catching a fly ball hit directly at you is harder than catching one that's not. Enter the linear optical trajectory (LOT) model.

In the LOT model, the fielder changes his position so that the ball's trajectory traces out a straight line relative to home plate and the background scenery. This requires the player to continuously move directly under the ball, thus guaranteeing that he'll be in a position to catch it.

Quoting the paper:

'The OAC model predicts that fielders select a running path that is straight with constant speed, resulting in a curved optical ball trajectory. The LOT model predicts that fielders select a running path that curves out and has an upside-down-U-shaped speed function, resulting in a linear optical ball trajectory.'

The authors did two kinds of experiments. In the first, they let somewhat-experienced outfielders catch fly balls hit in a number of directions with different initial velocities, and recorded their movements on video tape on a tower above and behind the fielders, to see the shape of the path they ran. In the other, they mounted video cameras on the fielders' shoulders, and taped the ball's path relative to the fielder, to see the shape of the ball's trajectory. To make a long story short, in most of the cases where the balls were caught (31 in each expt), the predictions of the LOT model were supported but those of the OAC model were not. In the few cases where the ball was hit straight at the fielder, the OAC model _was_ supported, but the authors consider this situation as an "accidental view that may require an alternative strategy".

No mention of PCT in the paper, but "control" and "perception" are mentioned several times. Might be worth a look from a PCT point of view.

John

Date: Thu, 11 May 1995 17:19:28 -0500
Subject: Fly chasing

(From Samuel Saunders [950511:1715 EDT])

In view of the recent threads on anticipation and on PCT research, the article "How baseball outfielders where to run to catch fly balls" by McBeath, Shaffer, and Kaiser in 28 April 1995 Science (569-573) may be of interest. The authors propose that the task is performed by varying running speed and direction to control the perception of a linear optical trajectory for the ball. They report 2 experiments, the first filming outfielders chasing fly balls with a fixed camera, the second filming with a shoulder mounted camera. The second experiment appears to be close to the test, and the analysis notes that the presumed variable in fact remains essentially constant, while other controlled variables which have been proposed, such as optical speed, do not remain constant. The one problem I see in the paper from a PCT view is that there is no attempt to model individual runs.

Samuel Spence Saunders,Ph.D.

Date: Fri, 12 May 1995 09:54:53 -0600
Subject: baseball

[From Bill Powers (950512.0900 MDT)]

Samuel Saunders [950511:1715 EDT] --

Seems that several of you out there are starting to think like PCTers.

> "How baseball outfielders where to run to catch fly balls" by McBeath, Shaffer, and Kaiser in 28 April 1995 Science (569-573) may be of interest.

John Anderson (950510.1630) also came up with the same article, and several others noticed it but didn't get around to commenting on it. It's obviously a good PCT experiment!

Best to all, Bill P.

Date: Thu, 18 May 1995 15:10:21 EST
Subject: bee-brains at the ANU

[Avery Andrews 950518]

Today's issue of our campus rag had an article about some very interesting work by a guy called Mandyam Srinivasan, who seems to have been doing some very clever work on perceptual control by bees.

They can be shown to remember how far down a tunnel food has been placed: train them to go a set distance for food, and then if you remove the food, they will stop at the distance they have learned, and mill around looking for it.

How do they know how far to go? Evidently by integrating the visual flow. If, after training them, you make the tunnel wider, they will go further, since walls further away means you have to go further to get the same integral of flow. Make the tunnel narrower, they go a shorter distance, for the same reason. If the tunnel has horizontal stripes, they go in indefinitely, presumably because no flow registers at all.

He also shows that they stay in the middle of the tunnel by equalizing the flow on both sides: if you make one wall of the tunnel move in the direction they are going, they move closer to it.

The article doesn't have citations for publications, but presumably something will be forthcoming eventually.

Avery.Andrews@anu.edu.au
Date: Fri, 19 May 1995 10:24:02 EST

Subject: bee-brain addendum

[Avery Andrews 950519]

And of course Srinivasan sent winds down the tunnels, etc. to show that the bees weren't counting wingbeats, energy expended, time elapsed, etc.

I think he's managed to do what I have only managed to fantasize about, which is produce convincing answers to classic `how do they do that' questions by *implicitly assuming* the truth of the principle that behavior is the control of perception. Since the assumption is implicit, it's hard to challenge, and since the results are so convincing, it tends to incline people to accept it.

Avery.Andrews@anu.edu.au

Date: Fri, 19 May 1995 10:04:08 -0600
Subject: Re: bees

[From Bill Powers (950519.08-- MDT)]

Avery Andrews (950519)--

> And of course Srinivasan sent winds down the tunnels, etc. to show that the bees weren't counting wingbeats, energy expended, time elapsed, etc.

> I think he's managed to do what I have only managed to fantasize about, which is produce convincing answers to classic `how do they do that' questions by *implicitly assuming* the truth of the principle that behavior is the control of perception. Since the assumption is implicit, it's hard to challenge, and since the results are so convincing, it tends to incline people to accept it.

Actually Srinivasan was applying the Test for the Controlled Variable in all the variations he used -- the stripes in the tunnel, moving a wall of the tunnel toward the bee, using the wind, etc. He was testing to see what the bee was maintaining constant by its behavior, applying disturbances that would alter various variables if they were _not_ under control, until he could deduce which variable was being controlled. One of the tests, which involved making parallel stripes down the tunnel so they looked the same no matter how far the bee moved, even mimicked the step in the Test where you check to make sure the perception is what you think it is, by eliminating it. Sure enough, when it was impossible for the bee to visually perceive its progress down the tunnel, the bee lost control of the distance flown.

The only hypothesis that survived was that the bee is controlling the integral of perceived visual flow outward from the center of vision. This was done by expanding the tunnel to change the perception of flow and showing that the bee brought the integral of this perception to the same value as before by flying further.

If you really grasp what the Test is all about, it's nothing but common sense applied to a control process. Srinivasan obviously has a great deal of common sense as well as an intuitive understanding of control theory. How about writing him a nice friendly letter telling him about PCT and inviting him to join us on the net?

Best to all, Bill P.

Date: Fri, 19 May 1995 13:51:35 -0700
Subject: The Test

[From Rick Marken (950519.1350)]

Paul Stokes (950518.1040) --

> What is this 'test' and how is it carried out?

The Test is the basic methodology of PCT; it is how we determine what perceptual variables an organism is controlling. There are different versions of The Test (Avery's post about optic flow control in bees and the Science article about catching baseballs by perceptual acceleration control involve implicit versions of The Test). Here is my view of the "canonical" Test:

1. Identify some behavior of interest (such as "having a conversation").

2. Guess at a controlled variable that might be involved in this behavior; a variable that, if it were controlled, might be responsible for some aspect of the behavior you see. This is the most important and creative step in The Test. For example, you might guess that a person involved in a conversation is controlling for the time between the end of one piece of talk and the beginning of another. So a hypothesized controlled variable is "between talk silent time" (BTST). BTST is a variable -- it can range from 0 to infinity -- so it could be under control. The notion that BTST is controlled is just a hypothesis -- a guess: the person might NOT be controlling BTST.

3. Test the hypothesis by introducing disturbances that would affect the variable if it were not under control. A disturbance to BTST could be introduced by having one of the people in the conversation be a stooge who varies the time until he starts talking after the other person finishes.

4. Monitor the hypothesized controlled variable while it is being disturbed; apply many disturbances and many different kinds of disturbance. If disturbances have the expected effect (the pauses are as long as the stooge makes them) the variable is not under control: go back to step 2 and try again. If the disturbances have somewhat less effect than expected, you are on the track of the actual variable that is under control: go back to step 2 but guess at a variable that is similar to the one already used. If disturbances have NO effect on the hypothesized variable, then stop -- you have found the controlled variable. If, for example, BTST remains the same despite the lags introduced by the stooge (possibly because the subject talks in order to "take up the slack") then BTST is a controlled variable and the average value of BTST is a reasonable estimate of its reference value.

Best Rick

Date: Sat, 20 May 1995 01:48:07 -0400
Subject: Re: bee-brains at the ANU

<[Bill Leach 950519.21:55] >[Avery Andrews 950518]

Avery; that was wonderful! I have heard so many claims for almost mystical navigational ability for bee's that your post was a nice counter point.

-bill

Date: Sat, 20 May 1995 11:31:40 -0500
Subject: More on Bee Brains

[From Bruce Abbott (950520.1130 EST)]

>Avery Andrews 950518 --

The fascinating bee-as-control-system research by Mandyam Srinivasan extends work done much earlier from a similar perspective. Werner Nachtigall (1968) describes some aspects of bee flight as follows:

> In nature it is the insect antenna that is displaced. The negative of its signal is sent to the mechanism regulating wing-beat amplitude, by way of a sense organ in the antenna, then via the nerves to computing centers in the central nervous system, to nerves once again and finally to the flight muscles. If the insect antenna is bent more strongly, the wing-beat

amplitude is made smaller as a result, and thus the undesirable increase in velocity is reduced until the velocity reaches its set value, for example 14 m.p.h.

The bee . . . has particular muscles which can change the position of the antennae slightly with respect to the air stream. Then the air resistance during flight can no longer bend them by the same amount as before but rather, for example, a little less. . . . a new set point for the flight velocity is achieved. Bees and flies have two antennae, one right and one left. It has been shown that each of these antennae regulates only the amplitude of the wing on its own side of the body. This has possibilities for flying in a curved course. If one antenna is cut off, the insect always flies around in circles. It is only when both antennae are operating together that they permit the bee to fly in a straight line and compensate for gusts of wind which might push it a little away from this line. This is critically important for orientation of the flight between hive and food source, which the bee should make as straight as possible.

. . . But there is one more point to be made: the bee does have a second servo system to control its velocity, involving the two large compound eyes. The essential purpose of this system is to hold constant the velocity over the ground; the antennae, on the other hand, regulate the velocity through the air.

. . . But an optical measuring device which fixes points on the ground and computes how quickly they move backwards, can do this. The highly complicated compound eyes are admirably suited to this task. They can recognize the dangerous drifts produced by head, tail, or side winds and compensate for the error in the signals of the antennal control circuit.

Nachtigal, Werner (1968 [English translation 1974]). _Insects in flight: a glimpse behind the scenes in biophysical research. New York: McGraw-Hill, p. 139.

Regards, Bruce

Date: Sun, 21 May 1995 15:36:36 -0400
Subject: Re: More on Bee Brains

<[Bill Leach 950520.17:20] >[From Bruce Abbott (950520.1130)]

Also interesting Bruce as wording reads as though it is that of a person that understands negative feedback control. Thanks to both you and Avery for bring me (at least) "up to speed" on more recent genuine biological research. Very encouraging.

-bill

Date: Tue, 23 May 1995 11:20:41 -0500
Subject: Bee/Fly Flight Motor

[From Bruce Abbott (950523.1115 EST)]

Earlier I mentioned that insect flight depends on an oscillator mechanism to generate the rhythmic motion of the wings. Here is a little information about "flight motors," which I gleaned from Werner Nachtigall's (1974) book, _Insects In Flight_.

There are two basic types of insect flight motor. The first, typified by the locust motor, has the main wing spar extending past the wing joint, with a muscle attached both inboard and outboard of the joint. When the outer muscle contract, it pulls the wing down; the inner muscle then contracts (and the outer relaxes) to pull the wing up again. Synchronization of these contractions is accomplished by a neural oscillator circuit. The locust motor's wing beat frequency is rather low, on the order of 20 to 25 beats per second.

The second type of insect flight motor, typified by the bee or fly motor, does without the neural oscillator and is able to achieve frequencies on the order of 200-300 beats per second. In these insects the top of the thorax is like a stiff lid attached to the box below by means of a stiff membrane hinged at its inner attachment to the lid and outer attachment to the box. Pressure exerted against the sides of the thorax makes the lid "pop" either upward or downward; the center position is unstable. The powerful "indirect" flight muscles pull down on the lid, snapping it into the lower position and jerking against a set of longitudinal muscles. The jerk triggers these muscles into contraction, popping the lid into its upper position and jerking the vertical muscles, triggering them to contract and initiating another cycle.

To continue contracting, the flight muscles must have calcium ions, which enter the muscle cells following neural stimulation. Asynchronous neural stimulation acts as the throttle by governing the rate at which calcium can enter the muscle cells; when the bee or fly wants to stop flying, it suppresses neural activity to the flight muscles, the motor runs out of calcium and, after a few beats, stops. The signal to stop flying comes from sensors on the foot pads. If you glue a stick to a fly's back, the wings will beat until you give the fly something, like a piece of cork, to hang onto.

The "click mechanism" can work only if the lid is under pressure. The bee or fly can alter the pressure on the left and right independently by contracting a muscle anchored to the bottom and side of the thorax. By releasing pressure on one side while maintaining it on the other, the insect can alter the wing-beat amplitude asymmetrically and turn sharply.

As mentioned, the indirect flight muscles are stimulated into contraction, not by neural impulses, but rather by being sharply jerked. So how does the bee or fly start the motor? The answer is that these critters come equipped with a "kick-starter." If you try to swat a fly bear-handed, the looming image of your hand will trigger a volley of neural impulses that travel down special large-diameter (fast) axons to a muscle attached to the base of the middle leg at one end and to the thorax lid at the other. When this muscle contracts it accomplishes two things: (1) it jerks the center leg downward, propelling the animal into the air, and (2) it jerks the flight-motor muscle to kick-start the motor. The whole sequence is completed in about 40-50 milliseconds, and the fly is "outta there" before your hand hits the table.

The next time you try to swat a fly, you might want to think about the marvelous complexity and elegance of the mechanism you are setting out to destroy.

Regards, Bruce

Date: Tue, 23 May 1995 20:41:47 -0700
Subject: S-R, Responsibility, Model-based non-control

[From Rick Marken (950523.2045)] Bruce Abbott (950523.1115)

> The signal to stop flying comes from sensors on the foot pads.

Signalus-response?

Best Rick

Date: Wed, 24 May 1995 10:02:30 -0500
Subject: Buggy Control Systems

[From Bruce Abbott (950524.1000 EST)]

>Rick Marken (950523.2045)

>>Bruce Abbott (950523.1115 EST) --

>> The signal to stop flying comes from sensors on the foot pads.

> Signalus-response?

Not necessarily, but I don't seem to have a fly/bee wiring diagram here in front of me at the moment, so it's a little hard to tell. Bees are able to keep the flight motor running while within the hive in order to provide hive "air conditioning," so for bees at least, there must be some means to override the "stop" signal ordinarily produced by footpad pressure. This would suggest that foot-pad pressure is only one input of a more complex input function. The inhibitory output signal that stops the flight motor probably appears when there is an error in any of several control systems. One might propose that, while in flight, footpad contact with a surface changes the state of (contact AND motor off) to false (reference set to true); hive temperature above reference might alter the reference for this lower-level system to true. On the other hand, footpad input might act as a "kill switch," but one that can be overridden by the temperature control system. Another possibility is that the footpad signal is generated only with a _change_ from no pressure to pressure. With the bee already down, this footpad transient would not occur. Who knows? I don't have the necessary information on which to base a sound model, but it may be out there. Is there a fly/bee expert in the house?

By the way, I'm finding a treasure-trove of control system research in this area. Did you know about it? Next time you're in an academic library, take a look at _The Journal of Experimental Biology_. It would appear that biologists, at least, are having little trouble getting control system studies published, [Example: Kittmann, Rolf (1991). Gain control in the femur-tibia feedback system of the stick insect. _Journal of Experimental Biology_, _157_, 503-522.] Most of these reports do not include a computer simulation of the system under study; nevertheless they often provide valuable information about the structure and function of some of the basic control systems found in insects and other animals. Even at the insect level these systems are impressively complex and subtle. I would encourage anyone who is interested in understanding/speculating about control systems from a basic nuts-and-bolts physiological perspective to take a look at this literature.

Regards, Bruce

Date: Wed, 24 May 1995 10:03:06 -0700
Subject: S-R and Control, PCT Research

[From Rick Marken (950524.1000)] >Bruce Abbott (950524.1000)

Bruce:
> The signal to stop flying comes from sensors on the foot pads.

Me:
> Signalus-response?

Bruce:
> Not necessarily, but I don't seem to have a fly/bee wiring diagram here in front of me at the moment, so it's a little hard to tell.

You won't find evidence against the stimulus-response view in the fly/bee wiring diagram. You have to look at the whole loop. It may be true that what is felt at the foot pad affects flying; but then it is also true (from observation) that flying affects what is felt at the foot pads. So the S-R relationship between foot pad sensor and flying is part of a closed loop; if this is a stable, negative feedback (ie. a control) loop, then the perceptual variable in this loop is under control; what is felt at the foot pads (the perception) is not a "signal": it is a controlled variable.

> By the way, I'm finding a treasure-trove of control system research in this area. Did you know about it?

Only a bit. We heard about some studies of beavers (reported by G. Cziko) that showed that the beaver apparently builds dams to control an auditory perception of "rushing water", keeping that perception at zero.

> Next time you're in an academic library, take a look at _The Journal of Experimental Biology_. It would appear that biologists, at least, are having little trouble getting control system studies published

There are control system studies and then there are PERCEPTUAL control system studies. It's the latter that have been tough to publish. There are a bazillion published studies of manual control in psychology, for example. Many of these are fine but they don't get at the basic PCT question -- what perception(s) is the organism controlling. I suspect that many of the control system studies in biology are equivalent to the "manual control" type studies in psychology. I would bet that studies done to determine the perceptual variables that organisms control are as rare in the biological as they are in psychological literature. My experience is that studies like the bee study described by Avery and the "baseball catching" study in Science are the rare exception rather than the rule-- in biology and in psychology. It may be that there are, in fact, more PCT-like studies in the biological literature. That would be a pleasant surprise, especially of there were a fair number of them.

Best Rick

Date: Wed, 24 May 1995 11:17:44 -0600
Subject: Bees

[From Bill Powers (950524.0915)] >Bruce Abbott (950520.1130)

Good stuff on bees from Nachtigall.

> The bee . . . has particular muscles which can change the position of the antennae slightly with respect to the air stream. Then the air resistance during flight can no longer bend them by the same amount as before but rather, for example, a little less. . . . a new set point for the flight velocity is achieved.

I would guess that this system is actually analogous to the stretch reflex, where the reference signal is converted to a length bias on the spindle and control brings the sensed length to a match. The output of the annulospiral ending is an error signal. The small muscles in the bee's control system would exert a force on the antenna, and flying speed would be varied to apply an equal and opposite force, canceling the bias due to the muscle and bringing the mechanical error to zero.

> Bees and flies have two antennae, one right and one left. It has been shown that each of these antennae regulates only the amplitude of the wing on its own side of the body. This has possibilities for flying in a curved course. If one antenna is cut off, the insect always flies around in circles. It is only when both antennae are operating together that they permit the bee to fly in a straight line and compensate for gusts of wind which might push it a little away from this line. This is critically important for orientation of the flight between hive and food source, which the bee should make as straight as possible.

This isn't quite right: to correct for gusts and direction relative to terrain would require an inertial or visual perceptual system. I seem to recall that bees have a gyro system which is made of small otoliths on the ends of short stiff hairs. The hairs are driven to make the otoliths oscillate, and they tend to retain the same plane of oscillation as the body turns, providing angular acceleration data. I would interpret the differential control of antenna bending simply as a means of controlling the curvature of the flight path relative to the air. In flight, velocity is always relative to still air.

> But there is one more point to be made: the bee does have a second servo system to control its velocity, involving the two large compound eyes. The essential purpose of this system is to hold constant the velocity over the ground; the antennae, on the other hand, regulate the velocity through the air.

This second servo system would operate by varying the reference signal for the flight-curvature and -velocity control systems, and it would maintain the speed and path relative to the visual image of the surroundings. The lower order systems would operate just like those of the "people" in the Crowd program.

> . . . But an optical measuring device which fixes points on the ground and computes how quickly they move backwards, can do this. The highly complicated compound eyes are admirably suited to this task. They can recognize the dangerous drifts produced by head, tail, or side winds and compensate for the error in the signals of the antennal control circuit.

Srinivasan carried this a step further, in determining that the controlled speed variable is visual streaming. The result isn't just "compensation" for "dangerous drifts," but continuous control of speed (and through sensing differential outflow, direction).

Nice observations on fly behavior.

Best to all, Bill P.

Date: Wed, 24 May 1995 18:05:56 -0600
Subject: arm-waving

[From Bill Powers (950524.1730 MDT)]

Bruce Abbott (950524.1000 EST)

I must have missed a post, or maybe it hasn't come yet, but I get the idea:

> Not necessarily, but I don't seem to have a fly/bee wiring diagram here in front of me at the moment, so it's a little hard to tell. Bees are able to keep the flight motor running while within the hive in order to provide hive "air conditioning," so for bees at least, there must be some means to override the "stop" signal ordinarily produced by footpad pressure. This would suggest that foot-pad pressure is only one input of a more complex input function. The inhibitory output signal that stops the flight motor probably appears when there is an error in any of several control systems.

When a bee takes off, the result is to remove foot-pad pressure; when it lands, the result is to produce foot-pad pressure. Anything else you say about it even with a so-called diagram of the nervous system in front of you is just arm-waving. You can tell arm-waving when you see it: as soon as a counter-example shows up, like the bee being able to operate its wings while hanging on with its feet, another glib explanation is tacked onto the argument: now we have to have an override to the supposed "stop" signal, which probably didn't exist in the first place. And then, I suppose, when the bee lets go with its feet and takes off to fly from the hive, the "override" signal that has been countermanding the "stop" signal has to be turned off by an "override repressor" signal -- which has to be turned off by an override derepressor signal to allow the bee to land somewhere else, and so on and so on and so on. This whole approach to explaining behavior "neurologically" is a crock.

I think that all neurophysiologists ought to be required to take courses in circuit design and construction, both analog and digital, before they're allowed to open their mouths about what the nervous system actually does.

I wish we could get off of this kick and get some modeling work done. The world is full of people making wild guesses about how the nervous system works, but it's all verbal and hardly worth paying attention to.

Grumpily, Bill P.
Date: Thu, 25 May 1995 06:59:21 -0600
Subject: bees and PCT

[From Bill Powers (950525.0530 MDT)]

I hope we can keep in mind the main thrust of the PCT Project. The main idea is to reinterpret behavior in terms of PCT, to see whether we become better able to explain behavior. There are many side-alleys that can be explored, many unusual behaviors occurring under rare or extreme circumstances that we could pause to examine and puzzle over. But the main question is and remains, and will remain for a long time, "What difference does it make to see behavior as the control of perception?"

For example, what difference does it make to see the sensory signals from a bee's footpads as inhibiting the mechanisms that cause the wings to oscillate, and to see the same signals as standing for a controlled variable? I think it makes a great deal of difference.

If we see the inhibitory footpad signals as controlled perceptions, we can imagine a reference signal that is set to bring the footpad pressure up to some value, and the means of doing so as changing the drive signals to the flying mechanisms. If the footpad pressure is too low, the resulting error signal reduces the flying efforts, and the flying efforts remain smaller until the footpad pressure is brought to the reference level, the inhibitory signals canceling the reference signal and correcting the error. So settling onto a surface is brought about by specifying a non-zero value for a perception of footpad pressure. The change in flying efforts is a means of increasing footpad pressure. And of course by setting the reference-pressure to zero, the same system will reduce the footpad pressure to zero by increasing the flying efforts until the bee is airborne rather than footborne.

The difference is that in the first interpretation, causality is assigned to the footpad pressure, an environmental stimulus, while in the second case causality is assigned to whatever changes the reference signal for footpad pressure. In the second case, footpad pressure becomes a controlled effect rather than an initiating cause. And in the second case, there can be other systems also contributing to the operations of the wings, even when the error in the footpad-pressure control system is zero. So the bee can happily stand on a surface and buzz its wings to create a draft, without taking off. In that case, incidentally, I presume that the phasing of pitch and roll angles of the wings is changed so that no lift is generated.

Anyway, let's not get hung up on details here, and let's not forget that models are only as good as the experiments which show they are needed.

Best, Bill P.

Date: Thu, 25 May 1995 11:03:32 -0500
Subject: Miss-bee-havior

[From Bruce Abbott (950525.1100 EST)]

>Bill Powers (950524.0915 MDT)]

>>Bruce Abbott (950520.1130 EST) --

> Good stuff on bees from Nachtigall.

The best part is that Nachtigall clearly recognizes that the structures mediating insect flight are organized as perceptual control systems and attempts to interpret them as such. As Srinivasan's bee research illustrates, this approach to insect behavior is still very much alive within biology (although, as the Simmer, Meyrand, and Moulins lobster article shows, not everyone has gotten the message). Much of the research on insect control systems seems to be appearing in German-language journals (which, unfortunately, I am unable to decipher); an exception which I mentioned in an earlier post is the British journal, _The Journal of Experimental Biology_.

>Bill Powers (950524.1730 MDT)

>Bruce Abbott (950524.1000 EST)

> I must have missed a post, or maybe it hasn't come yet, but I get the idea:

>>> Not necessarily, but I don't seem to have a fly/bee wiring diagram here in front of me at the moment, so it's a little hard to tell. Bees are able to keep the flight motor running while within the hive in order to provide hive "air conditioning," so for bees at least, there must be some means to override the "stop" signal ordinarily produced by footpad pressure. This would suggest that foot-pad pressure is only one input of a more complex input function. The inhibitory output signal that stops the flight motor probably appears when there is an error in any of several control systems.

> When a bee takes off, the result is to remove foot-pad pressure; when it lands, the result is to produce foot-pad pressure. Anything else you say about it even with a so-called diagram of the nervous system in front of you is just arm-waving. You can tell arm-waving when you see it: as soon as a counter-example shows up, like the bee being able to operate its wings while hanging on with its feet, another glib explanation is tacked onto the argument: now we have to have an override to the supposed "stop" signal, which probably didn't exist in the first place. And then, I suppose, when the bee lets go with its feet and takes off to fly from the hive, the "override" signal that has been countermanding the "stop" signal has to be turned off by an "override repressor" signal -- which has to be turned off by an override derepressor signal to allow the bee to land somewhere else, and so on and so on and so on. This whole approach to explaining behavior "neurologically" is a crock.

The point I was making was in response to Rick Marken's short question "Signalus-response?" which asks whether I was offering an S-R interpretation of the effect of foot-pad pressure shutting down the bee's flight motor. My point was that without data one could offer all sorts of "glib explanations" (as I then illustrated).

I was also suggesting that whether this particular mechanism operates open-loop or closed-loop is a matter for research to determine (although I would bet that it is closed-loop).

As to whether "this whole approach to explaining behavior 'neurologically' is a crock," I'm afraid that you have me confused. Chapter 3 of B:CP contains a nice discussion of the way neurons might be connected up so as to function as analog computational elements. And there on p. 83 of Chapter 7 is a nice diagram of the "basic first-order control system," the tendon reflex loop. A strong element of the PCT approach to research is its insistence that models should incorporate real (or at least potentially real) neural signals that could in principle be discovered within the nervous (or other) systems of real organisms. Certainly you are not saying that an understanding of the neurological "wiring diagram," together with detailed information about the ways in which individual components operate, is irrelevant to PCT.

If, on the other hand, you are objecting to attempts to understand the operation of a neurological system purely by examining its wiring diagram, I can understand your concern. In a hierarchically-organized perceptual control system, with its intricate, interconnected sets of closed-loop systems, it would be impossible to infer the operation of the system just by tracing neural impulses around the various loops. But what if you are concerned with a simple question: is this particular part of the system operating in open-loop or closed-loop mode? Some actions may not require feedback; it is enough just to issue the order. Now assume that I had the fly wiring diagram spread out before me and wanted to know whether the footpad signal that ordinarily stops the fly's flight motor represents a perceptual variable in a control system or an open-loop shut-down command signal. So I examine the diagram, looking for the sensory feedback loop required if the footpad signal is part of a flight-motor speed control system. The diagram would tell me whether the system is open- or closed-loop. Of course, the diagram does not exist; the point of my bringing it up is simply that if I had such a diagram I could answer Rick's question; without it, getting an answer will require research.

> I think that all neurophysiologists ought to be required to take courses in circuit design and construction, both analog and digital, before they're allowed to open their mouths about what the nervous system actually does.

Experimental psychologists (including psychobiologists), too.

>Bill Powers (950525.0530 MDT)

> For example, what difference does it make to see the sensory signals from a bee's footpads as inhibiting the mechanisms that cause the wings to oscillate, and to see the same signals as standing for a controlled variable? I think it makes a great deal of difference.

I very much like your description of a footpad-pressure control system and the way it might allow the fly to settle gently on a surface. The system is a bit more complicated than that: as the fly approaches a surface on which it intends to land, optical inputs (which indicate approach to the surface) change the reference for leg position from retracted to extended, thus lowering the fly's landing gear. The rate of approach is usually fairly high, so it is likely that the footpad signal increases too rapidly on impact to provide delicate adjustments in wing power under typical conditions (bees may be another story!).

Ever wondered how a fly lands on the ceiling? If the optical signals indicating a looming surface come from the upper portion of the eyes the fly begins to extend its legs while increasing its angle of attack. The forelegs make contact with the ceiling and the forward momentum gets converted to rotational momentum, pivoting the fly's body around the contact point, bringing the second set of legs into contact with the ceiling and shutting down the flight motor. The legs are then adjusted to bring the body level and the hind legs in contact with the ceiling, all accomplished by means of a small set of perceptual control systems. People had long assumed that to land on the ceiling the fly had to either loop or roll to bring its body to the correct position. We can now appreciate that neither view was correct.

There is likely to be quite a bit more known about the control systems of flies, bees, and other insects than I am currently aware of, so perhaps some of the questions we have been speculating about have been explored. Our library is not well-stocked with the journals in which such investigations are being reported. Perhaps some of our Deutsche-speaking friends could give us a hint about what has been reported recently in some of the German-language journals on the control systems of insects.

Auf wiedersehen, Bruce

Date: Thu, 25 May 1995 10:46:43 -0700
Subject: Ain't Miss-bee-havin'

[From Rick Marken (950525.1045)]

Bruce Abbott (950525.1100 EST)--

> The point I was making was in response to Rick Marken's short question "Signalus-response?" which asks whether I was offering an S-R interpretation of the effect of foot-pad pressure shutting down the bee's flight motor. My point was that without data one could offer all sorts of "glib explanations" (as I then illustrated).

I don't think you got my point. You said:

> The signal to stop flying comes from sensors on the foot pads.

This is an S-R description because it takes account of only one half of the actual relationship between variables. If it is true that sensor signals from the foot pad have an effect on the wing movements that produce flying, then it is also true that the wing movements that produce flying have an effect on the sensor signals from the foot pad. There is a closed loop of cause and effect that goes through the environment. The existence of this closed loop could not possibly be determined by looking at just the wiring diagram of the fly. It is a fact that exists because of the nature of fly's relationship to its environment. Even if there is an S-R relationship between sensory and effector neurons in an organism's wiring diagram, the organism is still an input control system (with a fixed and implicit reference for the input) if the effector neurons have a negative feedback effect on the inputs to the sensory neurons via the environment.

> Now assume that I had the fly wiring diagram spread out before me and wanted to know whether the footpad signal that ordinarily stops the fly's flight motor represents a perceptual variable in a control system or an open-loop shut-down command signal. So I examine the diagram, looking for the sensory feedback loop required if the footpad signal is part of a flight-motor speed control system.

You can't see the kind of feedback loops we deal with in PCT by looking just at a wiring diagram. A perceptual control loop goes through the environment so a wiring diagram alone will tell you absolutely nothing about the nature of the system -- whether it is open or closed loop. In order to determine whether a system is a closed loop control system (by inspection rather than by doing The Test) you would have to determine whether the efferent outputs of the wiring diagram have an effect, via the environment, on the afferent inputs to the circuit; so you have to inspect the circuit AND its relationship to the environment.

> The diagram would tell me whether the system is open- or closed-loop.

You might find internal feedback loops in the circuit diagram but you will not be able to tell whether or not the circuit exists in an open or closed loop relationship its the environment by just looking at the circuit.

An interesting example of mistaking a control system for an S-R system based on looking only at the "wiring diagram" occurs in analysis of the "Vehicles" models developed by Braitenberg (sp?). These are simple little software systems that move around by producing output forces in response to sensory inputs. I have seen these Vehicles described as "S-R" systems in journal articles about mobile robots.

The S-R designation for these Vehicles is based on an examination of their internal architecture; they do produce outputs in response to inputs. But their inputs are also, at least in part, a response to the outputs. So these systems live in a closed loop; they have fixed, implicit reference signals set to zero and they are dynamically stabilized by integrations in the output function. These Vehicles are actually control systems that control their input - - which is proven by the fact that they will protect their input (like the perception of target intensity) from disturbance (such as movements of the target).

The people who study these systems seem to have no idea that they are dealing with input control systems -- even though they know every detail of the "wiring diagram" of the system.

Best Rick

Date: Thu, 25 May 1995 12:19:52 -0700
Subject: Flyus goldbergus

[From Rick Marken (950525.1215)]

Well. I was going to let this go but I just can't.

Bruce Abbott (950525.1100 EST) --

> Ever wondered how a fly lands on the ceiling? If the optical signals indicating a looming surface come from the upper portion of the eyes the fly begins to extend its legs while increasing its angle of attack. The forelegs make contact with the ceiling and the forward momentum gets converted to rotational momentum, pivoting the fly's body around the contact point, bringing the second set of legs into contact with the ceiling and shutting down the flight motor. The legs are then adjusted to bring the body level and the hind legs in contact with the ceiling, all accomplished by means of a small set of perceptual control systems.

This is a very strange statement. You have described a sequence of causes and effects:

optical signals from upper portion of eye-->leg extension & increasing angle of attack--> foreleg contact-->rotational momentum-->pivoting around contact--> second set of legs in contact--> shut down flight motor.

Then you say "...all accomplished by means of a small set of perceptual control systems" as though it should be obvious from you description that perceptual control is involved. But you never say what variables the fly is controlling or what means it uses to control those variables.

You describe the mechanism of fly landing as though it were a Rube Goldberg device rather than perceptual control system. Here is the analogous Rube Goldberg device for fly swatting:

optical signals from fly--> picture on TV tube--> child watching TV with fly swatter in lap stands up in horror --> contact between fly swatter and floor --> Mom rotating at the hip to pick it up-->pivoting of fly swatter as Mom picks it up-->swatter hitting ceiling-->shut down of fly's flight motor.

The only hint of a controlled variable in your description of a fly landing on the ceiling is when "legs are then adjusted to bring the body level" suggesting that a perceptual representation of body level is controlled.

There is more to perceptual control theory than just believing in it.

Best Rick

Date: Fri, 26 May 1995 09:34:56 -0500
Subject: Re: Ain't Miss-bee-havin'; Flyus Goldbergus

[From Bruce Abbott (950526.0930 EST)]

>Rick Marken (950525.1045)

>>Bruce Abbott (950525.1100 EST)--

> I don't think you got my point. You said:

>> The signal to stop flying comes from sensors on the foot pads.

> This is an S-R description because it takes account of only one half of the actual relationship between variables. If it is true that sensor signals from the foot pad have an effect on the wing movements that produce flying, then it is also true that the wing movements that produce flying have an effect on the sensor signals from the foot pad. There is a closed loop of cause and effect that goes through the environment.

One could actually argue for a _positive_ feedback relationship here, if one were so disposed. Footpad contact with a surface increases the output of footpad pressure sensors, shutting down the motor, eliminating the lift, placing even more pressure on the footpad sensors. [This is not a serious proposal, just an illustration.] In complex systems like this, nearly every action has perceptual consequences. A closed loop does not necessarily indicate the presence of a control system. What you have to ask is, does the system actually control a perceptual variable.

> The existence of this closed loop could not possibly be determined by looking at just the wiring diagram of the fly.

True enough, but the control systems analyst is not interpreting the diagram in a vacuum and knows what to look for. If it's a control system, somewhere in that diagram there has to be a sensor, pathway for the perceptual signal, comparator, pathway for the error signal, and an effector whose actions feed back through the environment to affect the sensor. As you say:

> In order to determine whether a system is a closed loop control system (by inspection rather than by doing The Test) you would have to determine whether the efferent outputs of the wiring diagram have an effect, via the environment, on the afferent inputs to the circuit; so you have to inspect the circuit AND its relationship to the environment.

No closed loop, no control system. No problem! We ain't miss-bee-havin'.

>Rick Marken (950525.1215)]

> Well. I was going to let this go but I just can't.

Why am I not surprised? (;->

>>Bruce Abbott (950525.1100 EST) --

>> Ever wondered how a fly lands on the ceiling? If the optical signals indicating a looming surface come from the upper portion of the eyes the fly begins to extend its legs while increasing its angle of attack. The forelegs make contact with the ceiling and the forward momentum gets converted to rotational momentum, pivoting the fly's body around the contact point, bringing the second set of legs into contact with the ceiling and shutting down the flight motor. The legs are then adjusted to bring the body level and the hind legs in contact with the ceiling, all accomplished by means of a small set of perceptual control systems.

> This is a very strange statement. You have described a sequence of causes and effects:

> optical signals from upper portion of eye--leg extension & increasing angle of attack-- foreleg contact--rotational momentum--pivoting around contact--second set of legs in contact-- shut down flight motor.

> Then you say "...all accomplished by means of a small set of perceptual control systems" as though it should be obvious from you description that perceptual control is involved. But you never say what variables the fly is controlling or what means it uses to control those variables.

Gosh, Rick, do I have to do _everything_ for you? (;->

> You describe the mechanism of fly landing as though it were a Rube Goldberg device rather than perceptual control system.

I described no mechanism at all, only a sequence of events. I considered providing a (highly speculative) PCT analysis but decided that a brief description of the events would communicate the fly's landing strategy well enough. Based on this description, one could develop a PCT model that would behave as the fly does, but that exercise was left for the reader.

As seen from the outside, most behavior consists of sequences such as I described for the fly. Consider the directions for making coffee: Remove the filter basket and check it for contents. If it is full of old grounds, empty it. If it is empty, place a new filter-cup in the basket. If it contains a new filter-cup, add three measuring spoons-full of fresh ground coffee to the cup. Replace the basket into the coffee maker. . . . and so on. Each step describes the relevant sensory conditions and the behaviors that should occur under each condition.

The sequences are there; what matters is how you explain them: as an S-R chain or as, for example, a program-level perceptual control system. I think I made clear in my brief description of fly landing which view I prefer. If you want to spend the time to develop and describe a PCT model that will account for the observed sequence, be my guest. I thought it was more trouble than it was worth.

Regards, Bruce

Date: Fri, 26 May 1995 09:49:32 -0700

Subject: Fly landings

[From Rick Marken (950526.0900)] >Bruce Abbott (950526.0930)

Bruce:
> The signal to stop flying comes from sensors on the foot pads.

Me:
> If it is true that sensor signals from the foot pad have an effect on the wing movements that produce flying, then it is also true that the wing movements that produce flying have an effect on the sensor signals from the foot pad.

Bruce:
> One could actually argue for a _positive_ feedback relationship here, if one were so disposed.

Could be. It's unlikely, though, because the landing behavior (control of foot-pad pressure?) is stable.

> In complex systems like this, nearly every action has perceptual consequences. A closed loop does not necessarily indicate the presence of a control system. What you have to ask is, does the system actually control a perceptual variable.

That's right! And how do we ask that? All together now...

Re: Flyus goldbergus

> The sequences are there; what matters is how you explain them: as an S-R chain or as, for example, a program-level perceptual control system. I think I made clear in my brief description of fly landing which view I prefer. If you want to spend the time to develop and describe a PCT model that will account for the observed sequence, be my guest. I thought it was more trouble than it was worth.

I think it is worth it and I should have done it myself. I'll try it now. I'm sure I won't get the actual variables right but I hope it illustrates the difference between the PCT and conventional views of behavior.

First, here is your description of fly landing once again:

> Ever wondered how a fly lands on the ceiling? If the optical signals indicating a looming surface come from the upper portion of the eyes the fly begins to extend its legs while increasing its angle of attack. The forelegs make contact with the ceiling and the forward momentum gets converted to rotational momentum, pivoting the fly's body around the contact point, bringing the second set of legs into contact with the ceiling and shutting down the flight motor. The legs are then adjusted to bring the body level and the hind legs in contact with the ceiling, all accomplished by means of a small set of perceptual control systems.

Now here is a version of a PCT description. It is almost certainly wrong because I don't know much about fly behavior and I don't know of any research that has been done to test to determine the variables a fly is actually controlling in this situation.

The appearance of a fly landing on a ceiling is created when the fly, to satisfy unspecified "higher level" references, controls several different perceptual variables: 1) the rate of change of a visual gradient 2) homogeneity of the gradient over the retina 3) pressure on foot pads 4) perceived orientation. The main action of the fly that influences all variables is wing motion; disturbances to variables 1 and 2 include variations in the texture of environmental surfaces. Disturbances to variables 3 and 4 include variations in local velocity of the air. If the fly is near the ceiling, variable 1 and 2 will be brought to their reference levels (zero) by actions (wing movements) that slow the fly's progress toward the ceiling and orient its eyes toward the ceiling; of course, all the fly knows is that the visual gradient is changing more slowly and is starting to become homogeneous over the retina. As this happens there is also a change in the reference for pressure at the foot pads; the actual pressure at the foot pads increases as the visual variables are brought to there references (and the ceiling is encountered), making footpad pressure slightly overshoot the reference. This overshoot is eliminated as the fly's momentum brings the other feet down; muscles are also acting continuously to keep the pressure at all the pads equal (a higher level pressure reference). At the same time the muscles (and wing movements, I presume) are "trimming up" the landing, bringing the otolith based orientation perception to its reference.

When we see a fly land, we are seeing the side-effects of what the fly must do to bring perceptual variables to their (probably continuously varying) reference states. The external view can be misleading; it is not really possible to translate an external description of fly landing (like the one Bruce gave) into a PCT model. This is because it is impossible to tell what variables are actually under control; it is also impossible to see which actions are the ones that are part of the means used to control perceptions; much of what is seen is likely to be an irrelevant side-effect of the actions that are actually involved in control.

For example, in Bruce's description of fly landing, the fly "begins to extend its legs while increasing its angle of attack" in response to "optical signals indicating a looming surface". This maneuver may actually be a side-effect of efforts to get the visual gradient under control; the angle of attack might not actually be under control; it is just varied as part of the means of changing rate of approach (it damps the fly's air speed, maybe); changes in the texture of the surface might affect rate of change in visual gradient in a way that require adjustments that lead to approaches with quite different angles of attack.

The point is that cause-effect descriptions, like those of the fly behavior, may be grist for the start of research -- they suggest possible controlled variables -- but they are not the data to be explained. So it is VERY misleading to say:

> The sequences are there; what matters is how you explain them: as an S-R chain or as, for example, a program-level perceptual control system.

I hope my make-believe PCT explanation shows that this is not correct at all. It implies that existing behavioral data (like that on fly landing) is satisfactory and that all we need to do is to show how PCT can explain it. But this is not the case at all; existing data is of no use to PCT because it doesn't tell us what variables the organism is controlling; we need data like that collected by Srinivasan on bees; that data tells us about at least one of the variables the bee is controlling. Data like the fly landing data Bruce described is simply useless for PCT modelling (though it's fine for flights of fancy).

Apparently one does not have to be a PCTer (as Srinivasan was not) to have the sense to do the right kind of studies of behavior -- studies aimed at finding the variables organisms control. But studies like those done by Srinivasan are (as far as I know) few and far between. One of the goals of PCT science is to make Srinivasan's kind of research the NORM in the behavioral sciences.

As long as people keep looking for PCT explanations of data like that on fly-landing -- data that tells us nearly nothing about the variables that are being controlled -- then the real data we need will not be collected. That's why I get so upset when people (like Bruce) keep saying that PCT has to explain the existing behavioral science data to "prove itself worthy. This is upsetting because 1) most existing data (like the fly landing data) contains only the faintest hints about what variables MIGHT be under control and 2) it just wastes time until the real, systematic PCT data collection begins.

Which reminds me -- how's that PCT rat research coming along Bruce?

Best Rick

Date: Fri, 26 May 1995 12:57:06 -0500
Subject: The Fly as Control System

[From Bruce Abbott (950526.1255 EST)]

Vincent Dethier (quoted below) seems to have the right idea:

> The blowfly (Phormia regina) like all species of life, is a temporary form through which flows energy and matter, the matter becoming, for a while, fly and then passing on. The fly is just another way to reverse entropy on this planet, to defy, apparently, the Second Law of Thermodynamics. It is another way to build orderly complexity in a system characterized by increasing disorder and randomness. . .
>
> Animals are improbable accretions of matter. Real species are every bit as improbable as are the figments of diseased minds or of the imagination of a Hieronymus Bosch or Dali or Escher. The fly is a most improbable beast. . . .
>
> Although the fly and the vertebrate differ in their organization, their goals are the same. Living things are circumscribed cosmoses, designs for utopias where constancy is the goal. The machinery of life operates best under constant conditions. The limits of tolerance are narrow. Existing in a greater world where change is the rule, the minimum constant challenge is to maintain the status quo in the midst of change; the maximum challenge is to exercise autonomous change independently of or even in opposition to change in the wider outside world. These islands of balancing constancy that are organisms, man, the fly in his chitinous box, win their utopias only at cost. Energy is required to maintain temperature stability, to keep water balance, to drive locomotion (which is necessary to exploit new sources of energy when the immediate source is depleted), and to find other flies to make more flies. And energy is required to run the machinery designed to process and utilize energy. The cycle closes upon itself. The primary immediate goal is eating. To this end are so many of the ambitions of men and the activities of flies directed.
>
> Dethier, V. G. (1976). _The hungry fly_. Cambridge, MA: Harvard, Pp. 1-3.

Regards, Bruce

Date: Fri, 26 May 1995 12:12:44 -0600
Subject: Control of consequences

[From Bill Powers (950526.1100)] >Bruce Abbott (950526.0930)

> As seen from the outside, most behavior consists of sequences such as I described for the fly. Consider the directions for making coffee: Remove the filter basket and check it for contents. If it is full of old grounds, empty it. If it is empty, place a new filter-cup in the basket. If it contains a new filter-cup, add three measuring spoons-full of fresh ground coffee to the cup. Replace the basket into the coffee maker. . . . and so on. Each step describes the relevant sensory conditions and the behaviors that should occur under each condition.

The behaviors are sensory conditions, too. In fact these instructions say nothing about the motor outputs that are to be produced. The motor outputs that will actually occur depend on the happenstance state of the environment as each prescribed perceptual result is brought about. The motor acts involved in removing the filter basket depend on whether it's stuck and whether it contains old grounds and the angle of your body relative to the basket and the distance your shoulders happen to be from it. Emptying the basket, if needed, will involve motor acts that depend on what is already in the garbage can, where the garbage can is relative to your body and the basket, whether the grounds are wet or dry, and whether the wet filter tears or stays intact. Placing a new filter into the basket involves motor acts that depend on whether two filters stick together, how far down in the plastic bag they are, exactly where on the filter your fingers happen to grasp it, where the basket is and how it is oriented, and again your bodily configurations relative to the filters and the basket. And so forth.

When we name behaviors we are hardly ever actually naming outputs. We are naming perceptual consequences that we are to bring about by means of whatever output will accomplish them in the environment that exists at a given moment. This is why control systems work as a model of behavior, and open-loop systems don't.

> The sequences are there; what matters is how you explain them: as an S-R chain or as, for example, a program-level perceptual control system.

If you were to watch flies landing on the ceiling over and over, and make quantitative measurements instead of simply classifying similar-looking patterns as if they were identical, you would discover endless variations in the approach and landing pattern. Because of these endless variations in position, velocity, and acceleration, both linear and angular, the motor outputs involved must also be endlessly varying to keep the pattern converging toward the same final result. If the motor patterns were not varying from instance to instance, their consequences would be even less similar. The movements of the fly that you describe are not actions, but consequences of actions. To keep these consequences even moderately similar from instance to instance, the fly must be producing very different motor outputs each time.

The problem lies in naming behaviors and then assuming that all behaviors with the same name must be produced by the same motor actions. The process of naming creates similarities and obscures differences. If a fly "leaps" off a surface ten times, its motor actions might be quite different all ten times, depending on the positions of the legs, the adhesiveness of the surface, the state of the wings, and other processes that were in progress at the time the leap interrupted them. Yet we make all the instances seem identical by calling each of them a "leap." If indeed the leaps are quite similar, we can only conclude that the actions which produced them must have been quite different, because careful observation would show us that the initial conditions were quite different.

Your way of describing the sequence of events is the customary way, the way that behavior is most commonly described. But this level of description is actually a barrier to understanding what is going on. By making differences disappear through classification, you cut out the variations that would actually explain how the outcomes can be _nearly the same_ when the actions are _markedly different_.

Best, Bill P.

Date: Fri, 26 May 1995 12:33:30 -0600
Subject: Re: describing fly behavior

[From Bill Powers (950526.1230 MDT)]

Rick Marken (950526.0900) --

Rick, you brought out a point about modeling the fly's behavior that I had overlooked. Of course Bruce's description is not useful for a PCT analysis! It's a description of appearances from the point of view of a human observer standing outside the fly. The fly itself doesn't perceive any of the things that Bruce described, so there's no way it can be controlling them. The fly doesn't see a side or top or bottom view of itself tracing out a trajectory in three-dimensional space. It doesn't see what angle its body makes to the surface on which it is landing. The fly is controlling a set of fly-perceptions, not human perceptions, and it is doing so by varying what it can vary. We see side-effects of those actions and make sense of them in our own human frame of reference. So we see the problem of landing in our own human way. But the fly has to do it by controlling its own insect sensory impressions, not our human ones.

Best, Bill P.

Date: Fri, 26 May 1995 15:24:39 -0500
Subject: <No subject given>

[From Bruce Abbott (950526.1410)] >Rick Marken (950526.0900)

>> In complex systems like this, nearly every action has perceptual consequences. A closed loop does not necessarily indicate the presence of a control system. What you have to ask is, does the system actually control a perceptual variable.

> That's right! And how do we ask that? All together now...

We seem to have lost the thread here. I was never arguing that we should determine these things from wiring diagrams (although having them would certainly help); I was arguing for doing the research rather than jumping to conclusions. And how do we do the research? All together now...

> Now here is a version of a PCT description.

Plausible but it would have been overkill within the brief aside in which I offered my description. I think my description provided a clear enough picture of the basic strategy the fly employs to land on the ceiling, without recourse to fantasy.

> The point is that cause-effect descriptions, like those of the fly behavior, may be grist for the start of research -- they suggest possible controlled variables -- but they are not the data to be explained.

The fly approaches the ceiling at a steep angle, sticks out its legs, contacts the ceiling with its forelegs, rotates around this contact until its body is horizontal, and stops beating its wings. All this pure descriptive observation and answers, at one very superficial level, a question many have wondered about: how does a fly land on the ceiling? If you want a more detailed explanation in terms of controlled variables, you will have to do the research. Some of this research has been done, which is why I can state some of the sensory inputs involved.

> I hope my make-believe PCT explanation shows that this is not correct at all. It implies that existing behavioral data (like that on fly landing) is satisfactory and that all we need to do is to show how PCT can explain it.

Perhaps it implies that to you, but not to me. I certainly never made that claim. There are explanations and there are explanations. If I want to know why an airliner fell out of the air it may be enough for me to hear that the pilot stalled the aircraft. Someone else may want to understand what variables the pilot was attempting to control and how the interaction of his control systems with each other and with external disturbances led to the stall. Neither explanation is incorrect, they just operate at different levels of detail.

> As long as people keep looking for PCT explanations of data like that on fly-landing -- data that tells us nearly nothing about the variables that are being controlled -- then the real data we need will not be collected.

Skinner said the same thing about cognitive explanations. To paraphrase, "as long as people keep looking for mentalistic explanations of behavior, then the real data we need (on the environmental causes of behavior) will not be collected."

> That's why I get so upset when people (like Bruce) keep saying that PCT has to explain the existing behavioral science data to "prove itself worthy. This is upsetting because 1) most existing data (like the fly landing data) contains only the faintest hints about what variables MIGHT be under control and 2) it just wastes time until the real, systematic PCT data collection begins.

Researchers will have to collect the right kind of data in order to develop and test PCT models. My argument is that the models so developed should then explain (certain) existing behavioral science data. If your PCT model leads you to conclude that a fly can't land on the ceiling, your PCT model is wrong.

Furthermore, (certain) existing behavioral science data can provide strong hints about where to begin when designing a program of PCT research. The fly data, for example, reveal that the fly is sensitive to a number of sensory variables that _appear_ to influence its behavior in certain ways.

Finally, you really should take a look at some of this insect research. In this field, at least, the "real, systematic, PCT data collection" seems to be much more common than you seem to believe.

> Which reminds me -- how's that PCT rat research coming along Bruce?

I'm working on some making and installing some sensors for the operant chamber, thank you. After that I have to get the experiment program written and get the rats used to being handled.

Regards, Bruce

Date: Fri, 26 May 1995 16:59:10 -0500
Subject: A Fly in the Ointment?

[From Bruce Abbott (950526.1700)] >Bill Powers (950526.1100 MDT)

> The behaviors are sensory conditions, too. In fact these instructions say nothing about the motor outputs that are to be produced. . . .

Of course! I think you've gotten the wrong impression.

This whole discussion is beginning to remind me of the following joke. Johnny (who is six) walks up to his mother and says "Mommie, where did I come from?" Mother calls Dad in and together they explain all the details of sexual reproduction to Johnny. Looking very perplexed, Johnny finally blurts out "No! Jimmy came from Cleveland. Where did _I_ come from?"

My description of how a fly lands on the ceiling is more like the answer Johnny was looking for--superficial but adequate for the purpose. I'm not offering an S-R explanation even if it looks that way on the short form.

> Your way of describing the sequence of events is the customary way, the way that behavior is most commonly described. But this level of description is actually a barrier to understanding what is going on.

Not if you understand PCT. I tried to be clear and brief while providing enough information about sensory input that one could imagine how a control system model might be constructed (as Rick was able to do). I tried to avoid the possible S-R connotation by suggesting (without elaboration) that a more detailed account would appeal to a set of control systems. My explanation was neither misleading nor inaccurate, as far as it went. Come on, admit it--you now have the quick answer (Johnny's answer) to how flies land on the ceiling, one you can offer to your friends if the subject ever comes up. It's not wrong, it's just lacking in detail, detail I wanted to avoid in what I intended to be a brief digression.

>Bill Powers (950526.1230 MDT)

> Rick, you brought out a point about modeling the fly's behavior that I had overlooked. Of course Bruce's description is not useful for a PCT analysis! It's a description of appearances from the point of view of a human observer standing outside the fly.

Not exactly--I did hint at some of the perceptual variables involved, which have been deduced from physiological studies, not by guessing what they might be based on external observation. Many of the details I omitted ARE useful to a PCT analysis; I hope my all-too-brief description did not give you the impression that what I had to say is all that is known about the systems involved.

Regards, Bruce

Date: Fri, 26 May 1995 15:22:41 -0700
Subject: Behavioral Fantasies

[From Rick Marken (950526.1520)] Bill Powers (950526.1230 MDT)

> Rick, you brought out a point about modeling the fly's behavior that I had overlooked. Of course Bruce's description is not useful for a PCT analysis! It's a description of appearances from the point of view of a human observer standing outside the fly. The fly itself doesn't perceive any of the things that Bruce described, so there's no way it can be controlling them.

Well, you think that Bruce's description is not useful. I think that Bruce's description is not useful. But it seems that Bruce (950526.1410 EST) thinks that Bruce's description IS useful. In fact, I haven't been able to detect much change in Bruce's position on this matter since he got on the net. Could we be dealing with a controlled variable;-)

Bruce seems to be controlling for the idea that descriptions of the appearance of behavior from the point of view of a human observer (the typical way behavior is described in the behavioral sciences) should be accounted for by PCT. We keep explaining (and demonstrating) why it is impossible to do this. It is because it is impossible to tell what variables are being controlled by simply looking at behavior; descriptions of behavior are, therefore, not particularly useful to PCT.

Perhaps Bruce doesn't buy this because he doesn't know what we mean when we say that the behavior we see is a (possibly irrelevant) side-effect of controlling one's own perceptions. For example, Bruce (950526.1410 EST) says:

> I think my description provided a clear enough picture of the basic strategy the fly employs to land on the ceiling, without recourse to fantasy.

In his brief post on "describing fly behavior", Bill Powers (950526.1230 MDT) explains why Bruce's description of fly behavior provides no picture at all of the fly's "basic strategy". The fly's basic strategy is to control its own perceptions. Bill's post explains why Bruce's description of how a fly lands on the ceiling is more of a fantasy than the PCT version that I invented.

Another way to see the problem with descriptions of behavior is by doing my Mind Reading demo; I think you have a PC version, don't you Bruce? When you move one of the numbers around the screen, notice what is happening to the other numbers (the irrelevant side effects of your controlling). Someone watching your behavior could make up a very complex description of what you are doing in terms of movement of any or all of the numbers on the screen, ie. in terms of irrelevant side effects of your controlling. In fact,the only correct description of your behavior is from YOUR perspective -- in terms of the perception of the location of the number that you are controlling.

Bruce:

> you really should take a look at some of this insect research. In this field, at least, the "real, systematic, PCT data collection" seems to be much more common than you seem to believe.

Why not help out and post some examples. That's what the net's for, no? I have been on the lookout for PCT type research for years. In the last month, there are suddenly two studies (the Science catching article and the Srinivasan bee study) that are models of PCT research. This is QUITE unusual. You claimed that Nachtigall "clearly recognizes that the structures mediating insect flight are organized as perceptual control systems" but then presented evidence (the description of how a fly lands on the ceiling) that he really doesn't. So if there is some "real, systematic, PCT data collection" in the insect research field, please point it out. But the "fly landing" stuff is clearly not it.

I'd love to hear about research, like Srinivasan's, aimed at determining the variables being controlled by organisms. I'm sure there are many others on this net who would like to know about it too, so please feel free to post it.

Best Rick

Date: Fri, 26 May 1995 16:57:41 -0600
Subject: More flies

[From Bill Powers (950526.1605)] Bruce Abbott (950526.1700 EST)

> My description of how a fly lands on the ceiling is more like the answer Johnny was looking for--superficial but adequate for the purpose. I'm not offering an S-R explanation even if it looks that way on the short form.

I guess I see what you mean. It occurs to me that we're using the word "how" differently, in the question "How does the fly land on the ceiling?" One meaning is," Tell me all the things that we can see happening as a fly lands on the ceiling." In the other meaning -- which is the one I automatically assume -- the question is "What is going on inside the fly which would account for what we see happening?"

To have a more detailed description of what happens is useful, in that it gives us more details to account for. But no matter how detailed the description, it doesn't reveal the processes inside the fly that determine what MATTERS about what we see. The pure observation is something like a little kid describing what an orchestra conductor does:

"He lifts up his elbows and waves a stick, and he runs his hand through his hair and pulls his coat down, and he opens his mouth wide with his eyebrows up, and he knocks the music sheets onto the floor, and he points to people, and sometimes he shakes his head real hard. And all the time, there are these guys playing music."

Best, Bill P.

Date: Fri, 26 May 1995 20:02:13 -0500
Subject: Fantasies, Behavioral and Otherwise

[From Bruce Abbott (950526.2000 EST)]

>Rick Marken (950526.1520)

> Well, you [Bill Powers] think that Bruce's description is not useful. I think that Bruce's description is not useful. But it seems that Bruce (950526.1410 EST) thinks that Bruce's description IS useful. In fact, I haven't been able to detect much change in Bruce's position on this matter since he got on the net.

Given the redundancy here, that counts for only one opinion, so the score is one to one. And you're right, I haven't changed my position on this matter. Why would I want to change it from right to wrong? (;->

> Bruce seems to be controlling for the idea that descriptions of the appearance of behavior from the point of view of a human observer (the typical way behavior is described in the behavioral sciences) should be accounted for by PCT. We keep explaining (and demonstrating) why it is impossible to do this. It is because it is impossible to tell what variables are being controlled by simply looking at behavior; descriptions of behavior are, therefore, not particularly useful to PCT.

And I keep explaining why I think they should and can be--and how. I've never heard you try to knock down my argument, or even note it. You seem to prefer to just ignore it and reassert your own position. It's a rather peculiar way to "argue" and, from my position, unconvincing.

How about stating, just for the record, what my argument was? (You can paraphrase if you like.) I'd like some proof that you read it and understood it, even if you disagreed. Then maybe we can discuss whether it has merit.

> In his brief post on "describing fly behavior", Bill Powers (950526.1230 MDT) explains why Bruce's description of fly behavior provides no picture at all of the fly's "basic strategy". The fly's basic strategy is to control its own perceptions. Bill's post explains why Bruce's description of how a fly lands on the ceiling is more of a fantasy than the PCT version that I invented.

I suppose when your kids ask you how to do something you just tell them to control their own perceptions. I'll bet they're pretty perplexed.

I was, of course, describing what the fly does from the observer's point of view, with (as I'm now saying at least for the second time) some additional information about perceptual variables thrown in to help guide speculation about how the control systems involved might operate. People have speculated whether the fly does an inside loop (touching the ceiling at the top of the loop) or a half roll in order to position itself legs up and parallel with the ceiling. As I explained, neither suggestion is correct. The way the fly does it I called a "strategy," but by that I do not mean a conscious (or for that matter unconscious) plan that the fly is following, only which of several alternatives actually occurs.

>>Bruce:

>> you really should take a look at some of this insect research. In this field, at least, the "real, systematic, PCT data collection" seems to be much more common than you seem to believe.

> Why not help out and post some examples. That's what the net's for, no?

You mean in addition to the Nachtigall and Dethier book references I've already posted and my general suggestion to take a look at the _Journal of Experimental Biology_? How specific do you want?

> I have been on the lookout for PCT type research for years.

Try looking somewhere else. The library would be good. (;->

> In the last month, there are suddenly two studies (the Science catching article and the Srinivasan bee study) that are models of PCT research. This is QUITE unusual.

Our library is closed so at the moment I can give you only the small bit I happen to have on my desk. How about:

Kittmann, R. (1991). Gain control in the femur-tibia feedback system of the stick insect. _Journal of Experimental Biology_, _157, 503-522.

Preiss, R., & Gewecke, M. (1991). Compensation of visually simulated wind drift in the swarming flight of the desert locust (schistocerca gregaria). _Journal of Experimental Biology_, _157_, 461-481.

Spork, P. (1994). Adjustment of flight speed of gregarious desert locusts (orthoptera acrididae) flying side by side. _Journal of Insect Behavior_, _7_, 217. [I haven't seen this one yet.]

I could give you another half-dozen or so from the reference sections of the first two articles, but you'll have them when you get these articles. Our library is rather underfunded (each year more journals get cut); even the _Journal of Experimental Biology_ is not longer being subscribed to here, so I'm a bit stuck for the most recent references.

> You claimed that Nachtigall "clearly recognizes that the structures mediating insect flight are organized as perceptual control systems" but then presented evidence (the description of how a fly lands on the ceiling) that he really doesn't. So if there is some "real, systematic, PCT data collection" in the insect research field, please point it out. But the "fly landing" stuff is clearly not it.

Here's Nachtigall's description:

The fly approaches the ceiling obliquely from below, at a steep angle, at a speed of about 25 cm per second. It then flies straight into the ceiling, and shortly before striking it, stretches out all three pairs of legs. The fore pair take up a special attitude, held out stiffly upwards so that they are the first part of the fly's body to make contact. Held in this way they act as shock absorbers and as anchors, adhering at the point of contact by means of their claws and pulvilli (hairy pads). Simultaneously the wings stop beating. Now the fly is clinging firmly to the ceiling with its fore feet, but its body still has a certain forward momentum. Like a flywheel on its shaft the fly rotates about its fore feet and turns its belly upwards, grasping the ceiling with its middle and hind feet--and there it is, sitting upside down on the ceiling, without having had to fly upside down first.

Nachtigall, Werner (1974). _Insects in flight_. New York: McGraw-Hill. p. 119.

This was all determined via high speed photography, is purely descriptive, and is irrelevant to the issue of Nachtigall's understanding of control system theory--and to mine.

Why don't you have a look at this book, Rick, and judge Nachtigall's understanding for yourself? I'll bet they've got it at the UCLA library. Heck, even OUR little library has a copy!

Regards, Bruce

Date: Fri, 26 May 1995 21:05:35 -0700
Subject: Lourdes of the flies

[From Rick Marken (950526.2100)]

(No, I have no idea what the title of this post means, but it is related to the "flies" thread and perhaps Bruce will _miraculously_ convert to my point of view).

Bruce Abbott (950526.2000 EST)]

> And you're right, I haven't changed my position on this matter [the importance of PCT accounting for conventional behavioral science data]. Why would I want to change it from right to wrong? (;-

You wouldn't, of course, unless you enjoyed experiencing error. What I hope you will be able to do (and precious few have) will be to change the reference for what your position is, so that what you currently perceive as the wrong position is perceived as the right one.

You agree that you are controlling for what I said you were controlling for:

> the idea that descriptions of the appearance of behavior from the point of view of a human observer (the typical way behavior is described in the behavioral sciences) should be accounted for by PCT.

You say:

> And I keep explaining why I think they should and can be--and how. I've never heard you try to knock down my argument, or even note it.

> How about stating, just for the record, what my argument was?

I believe that you think PCT should account for descriptions of the appearance of behavior (actually, for the findings of conventional psychology) for at least two reasons: 1) to convince psychologists (including yourself) that PCT is a viable alternative to conventional theories and 2) because a new theory should be able to account for all the data that has been handled by previous theories, and then some (in the way that relativity physics handled all the data handled by Newtonian physics, and then some).

I think we have tried to deal with these points, though perhaps it has not seemed so to you. This whole flurry of posts about "describing fly landing behavior" gets to the nitty gritty of my objection to your argument.

I have tried to answer the first part of your argument by arguing that PCT is not about the same phenomenon that conventional theories of behavior are about. Conventional theories are about "behavior" which is tacitly assumed to be caused results of action; PCT is about controlled results of action -- results that are brought to reference states and protected from the effects of disturbance. Many of the demos and experiments that we have described on the net are aimed at communicating the crucial difference between caused and controlled results. This whole discussion about fly behavior turns on this difference. Conventionally, fly landing behavior is described as a sequence of caused results; PCT requires a description of this behavior in terms of controlled results -- controlled perceptual variables.

The answer to the second part of your argument is based on the same point: the distinction between caused and controlled results of action. PCT cannot be expected to handle data that has been "handled" by previous theories when this data is not about controlled results of action. The analogy is not

PCT : Conventional Psychology :: Einstein : Newton.

It is more like:

PCT : Conventional Psychology :: Newton : Aristotle.

> I suppose when your kids ask you how to do something you just tell them to control their own perceptions. I'll bet they're pretty perplexed.

My kids no longer ask me how to do anything; they just tell me what to do;-). But when they did ask I would unquestionably try to tell them how to do something in terms of the perceptions they should try to control. I have tried to suggest that they do things by controlling MY perceptions of THEM -- but that never went over too well.

> I was, of course, describing what the fly does from the observer's point of view, with (as I'm now saying at least for the second time) some additional information about perceptual variables thrown in to help guide speculation about how the control systems involved might operate.

I'm afraid I did not notice any mention of possible controlled variables in your description of the fly's behavior.

> People have speculated whether the fly does an inside loop (touching the ceiling at the top of the loop) or a half roll in order to position itself

legs up and parallel with the ceiling. As I explained, neither suggestion is correct.

If this is an example of speculating about a controlled variable, then your concept of a controlled variable differs considerably from mine. This is speculation about characteristics of overt behavior; it says nothing about the perceptual variables that the fly might be controlling. The fly cannot perceive itself doing loops, half rolls, touching ceilings, etc. The fact that "neither suggestion is correct" is a foregone conclusion in PCT; if it can't be perceived, it can't be controlled.

Me:
> Why not help out and post some examples [of PCT-type research]. That's what the net's for, no?

Bruce:
> You mean in addition to the Nachtigall and Dethier book references

Of course. Neither Nachtigall nor Dethier did any PCT-type research.

> How specific do you want?

Take a look at Avery's description of Srinivasan's bee research. Avery clearly described just what we need to know: the hypothetical controlled variable (Avery told us it was integrated visual flow) and the methods used to test the hypothesis (Avery explained the different disturbing conditions and how variations in the bee's actions were consistent with the idea that they were controlling visual flow by compensating for these disturbances).

> Why don't you have a look at this book, Rick, and judge Nachtigall's understanding for yourself?

Based on your description of his work and the extended quote from his book, I judge that I can eliminate Nachtigall's book from my list of books that are possibly relevant to PCT research.

Maybe it's best to finish this up with a question related to your first question to me:

Why do you think we (PCTers) don't believe that PCT needs to account for the descriptions of behavior provided by conventional psychology?

Best Rick

Date: Sat, 27 May 1995 09:33:15 -0700
Subject: Rephrased question

[From Rick Marken (950527.0930)]

I can see that the question to Bruce that I posted last night is ambiguous. I asked:

> Why do you think we (PCTers) don't believe that PCT needs to account for the descriptions of behavior provided by conventional psychology?

This could imply that we PCTers actually DO believe that PCT needs to account for the descriptions of behavior provided by conventional psychology and that I am asking Bruce why he thinks we don't believe this.

So let me ask the question again and try to make the intent clearer:

We PCTers believe that PCT should not be required to account for most of the descriptions of behavior provided by conventional psychology, nor should it be required to account for most of the results of conventional psychological research.

Why do you (Bruce or anyone else for that matter) think we believe this?

I would prefer an essay answer. But, if you prefer multiple choice, here are some alternative short answers:

a) We fear that PCT cannot account for these descriptions and results.

b) We have already accounted for some of these descriptions and results and we are tired of doing it.

c) We are rejectionist radicals with a bad case of constipation and nothing but contempt for conventional scientists.

d) We got a bad grade in freshman psychology, sociology or biology and have been bitter about it since then.

e) We are liberals

f) all of the above

g) none of the above

Best Rick

Date: Sat, 27 May 1995 11:49:36 -0600
Subject: Controlled variables vs. side-effects

[From Bill Powers (950527.0950 MDT)]

Just got back from seeing our daughter Barbara off in the start of the Iron Horse bike race, Durango to Silverton. The length is 45 miles, the total climb over two main passes is 5500 feet (the highest pass, Molas, is about 11,000 feet). Last year (her first, at age 35) she did it in 4:20; this year she hopes for under 4:00. The pro winning time last year was 2:10. She should be about halfway right now, starting the four-mile climb to Coal Bank Pass (2500 foot climb to over 10,000 ft). Go Bara!

Rick Marken, Bruce Abbott (continuing) --

==
When you push on a control system, it pushes back.
==

RE: trajectories vs. system organization

In a great deal of modern behavioral research, trajectories of movement are examined in the hope of finding invariants that will reveal secrets of behavior. This approach ties in with system models that compute inverse kinematics and dynamics and use motor programs to produce actions open-loop. These models assume that the path followed by a limb or the whole body is specified in advance in terms of end-positions and derivatives during the transition, so the path that is followed reflects the computations that are going on inside the system.

It is this orientation that explains papers like

Atkeson, C. G. and Hollerback, J.M.(1985); Kinematic features of unrestrained vertical arm movements. The Journal of Neuroscience _5_, #9, 2318-2330.

In the described experiments, subjects move a hand in the vertical plane at various prescribed speeds from a starting point to variously located targets, and the positions are recorded as videos of the positions of illuminated targets fastened to various parts of the arm and hand.

The authors constructed a tangential-velocity vs time profile of the wrist movement for various speeds, directions, and distances of movement. They normalized the profiles to a fixed magnitude, then to a fixed duration, and found that the curves then had very nearly the same shape. Using a "similarity" calculation, they quantified the measures of similarity.

They were then able to compare these normalized tangential velocity profiles across various directions and amounts of movement and show that the treated profiles were very close to the same. They conclude:

> Taken together, shape invariance for path and tangential velocity profile indicates that subjects execute only one form of trajectory between any two targets when not instructed to do otherwise. The only changes in trajectory are simple scaling operations to accommodate different speeds. Furthermore, subjects use the same tangential velocity profile shape to make radically different movements, even when the shapes of the paths are not the same in extrinsic coordinates. Different subjects use the same tangential velocity profile shape.
>
> ... this would be consistent with a simplifying strategy for joint torque formation by separation of gravity torques from dynamic torques and a uniform scaling of the tangential velocity profile ... (p. 2325)
>
> ... if the motor controller has the ability to fashion correct torques for one movement, why does it not use this same ability for all subsequent movements rather than utilize the dynamic scaling properties? Among the possibilities we are considering, the first is a generalized motor tape where only one movement between points must be known if the dynamic components in equation 6 are stored separately....A second possibility is a modification of tabular approaches [ref] where the dimensionality and parameter adjustment problem could be reduced by separate tables for the four components in equation 6. (p. 2326)

This paper was sent to me by Greg Williams as a source of data about actual hand movements, for comparison with the hand movements generated by Little Man v. 2, the version using actual arm dynamics for the external part of the model. The model's hand movements were, as Greg will attest, quite close to those shown in this paper, being slightly curved lines connecting the end-points. Forward and reverse movements followed somewhat different paths, and by adjustment of model parameters this difference, too, could be reproduced.

What is interesting is that the fit between the Little Man and the real data was found without considering tangential velocity profiles or doing any scaling or normalization. In other words, the invariances noted by the authors were simply side-effects of the operation of the control systems of the arm interacting with the dynamics of the physical arm. In the Little Man there is no trajectory planning, no storage of movement parameters, no table-lookup facility, no computation of invariant velocity profiles. The observed behavior is simply a reflection of the organization of the control system and the physical plant.

The path which Atkeson, Hollerbach (and many others at MIT and elsewhere) are treading is a blind alley, because no matter how carefully the observations are made and the invariances are calculated, there will be no hint of the control-system organization, the SIMPLE control-system organization, that (I claim) is actually creating the observed trajectories. No doubt a sufficiently complex trajectory-control model, with just the right tables of coefficients and velocity profiles, would ultimately be able to match the behavior. But this line of investigation, with its underlying assumptions, will never lead to the far simpler and anatomically correct PCT model.

In terms of the current discussion on the net, the observations made by the authors were interesting as checks on the model, but were actually irrelevant to what the control systems were doing. The control systems (the first two levels of the Little Man model) controlled only three kinds of variables that underlay the perceptual signals: angular positions, angular velocities, and angular accelerations. They received no information about wrist position in laboratory space. They contained no provision for computing tangential velocities, or for computing positions of points on the physical arm in space, or for computing space-time invariants. The behavior of the control systems, in other words, took place in a proprioceptive perceptual space that no outside observer could see. In order to translate from this perceptual space into variables that were observable, the computer program generated the resulting arm positions and plotted them in a form suitable for visual inspection. So a side-effect of the actual control process was presented for comparison with a corresponding side-effect of the real control process, as visible to an outside observer.

The approach of Atkeson and Hollerbach appears in many guises. We have already talked about the apparent scaling and normalization of trajectories seen when two hands move rapidly and simultaneously to targets at different distances. In operant conditioning experiments, we have seen how the control of reinforcement by behavior is obscured by the fact that variations in behavior tend to stabilize reinforcement rates, thus making reinforcement rate appear to be the independent variable.

We have also seen a few -- a very few, so far -- studies in which the PCT orientation was used, Srinivasan's being the most recent. What is the difference? I think the difference is in whether the emphasis is on seeing the behavior from the behaving systems's point of view, as best we can imagine it, and seeing it strictly from the human observer's point of view.

From the human observer's point of view, it seems that we must account for the detailed movements and physical interactions that are seen to occur. This leads to trying to find invariances or striking mathematical regularities of some sort in the observed behaviors. It leads to imagining an internal system that is producing explicitly what we are observing; if we observe a trajectory, there must be some generator that is specifically calculating that trajectory.

But from the behaving system's point of view, we can consider only the information that is available to the behaving system; we must look for our explanations there. The trajectories of movement that result from the system's operation are basically side-effects; they are not planned and they are constant only in a constant environment. Furthermore, they are unknown to the behaving system and play no part in the production of behavior. We can deduce from the model of the behaving system what the observable side-effects would be in a given environment, and so can compare those side-effects with our external observations of the behavior. But our explanation of the behavior is not based on those side-effects.

Most important, when we simply describe behavior as a sequence of physical happenings and relationships, we have no way of knowing whether we are describing controlled variables or side-effects. When we see a fly landing on a ceiling, it is perfectly possible that NOT A SINGLE ASPECT OF WHAT WE SEE is perceived and controlled by the fly. When we see the fly extending its legs just prior to landing, the fly may have no perception of the configuration of its legs; to the fly, all that is controlled may be two or three joint-angle signals, not even identified by the fly as representing joint angle. When we see the wings stop flapping, to the fly all that may be controlled is a sensation of vibration. When we see the fly's body making a steep angle with the surface, the fly may simply be experiencing a visual signal indicating, as Rick guessed, a gradient of illumination or texture. Not one of the variables we are observing may ever appear in the ultimate model of the fly's internal organization, just as in the Little Man the actual arm configuration and hand position never appear in the model of the first two (kinesthetic) levels of control. Once we have the right model, we can always compute how its operation will appear to an observer who is focusing on various side-effects of the actions. But the model itself says nothing about those appearances, and makes no use of them.

Best to all, Bill P.

Date: Sat, 27 May 1995 13:23:21 -0500
Subject: Come Fly with Me

[From Bruce Abbott (950527.1320 EST)]

>Rick Marken (950526.2100)

>>Bruce Abbott (950526.2000 EST)]

>> And you're right, I haven't changed my position on this matter [the importance of PCT accounting for conventional behavioral science data]. Why would I want to change it from right to wrong? (;-

> You wouldn't, of course, unless you enjoyed experiencing error. What I hope you will be able to do (and precious few have) will be to change the reference for what your position is, so that what you currently perceive as the wrong position is perceived as the right one.

> You agree that you are controlling for what I said you were controlling for:

>> the idea that descriptions of the appearance of behavior from the point of view of a human observer (the typical way behavior is described in the behavioral sciences) should be accounted for by PCT.

This isn't quite right. What I am attempting to maintain is the perception that my views on various points are correct. I will quite happily change them if you can demonstrate to my satisfaction that they are incorrect; so far, you have failed to do so. However, I am willing, as always, to listen to your arguments and to any evidence for them you may present, and to consider these carefully.

>> How about stating, just for the record, what my argument was?

> I believe that you think PCT should account for descriptions of the appearance of behavior (actually, for the findings of conventional psychology) for at least two reasons: 1) to convince psychologists (including yourself) that PCT is a viable alternative to conventional theories and 2) because a new theory should be able to account for all the data that has been handled by previous theories, and then some (in the way that relativity physics handled all the data handled by Newtonian physics, and then some).

This is a good start, but it needs some work. PCT should be shown to account for _certain_ well-known behavioral phenomena as a demonstration of its explanatory power. The "findings of conventional psychology" is far too big a territory. Many of these findings are irrelevant simply because they are based on group averages. Many studies deal with behavior only as a way to explore aspects of some system, such as perceptual input functions or memory phenomena. These may (or may not) provide information that would be helpful in guiding the construction of a PCT model; there would be no need to model the perceptual control systems involved in generating these data. Moreover, I do not see the need to carry on demonstrations ad infinitum; a certain carefully chosen few should be enough to convince; if they fail to do so, then it is unlikely that further demonstrations would convince either.

The demonstrations I have in mind are actually full-blown PCT research projects, not superficial attempts to show how some PCT model _could_ account for some behavioral phenomenon. Although we did some of the latter in the SD and classical conditioning modeling, these exercises represent only a preliminary exploration of the problem; they also provided an opportunity to learn how to write programs that would collect the data or model the system. The full-blown projects would kill two flies with one swat: elucidate the perceptual variables being controlled in the test situation and provide a cogent theoretical account for the well-established behavioral phenomenon under investigation. This account would also allow one to predict or explain what would be viewed as anomalies under the conventional views.

> I think we have tried to deal with these points, though perhaps it has not seemed so to you. This whole flurry of posts about "describing fly landing behavior" gets to the nitty gritty of my objection to your argument.

> I have tried to answer the first part of your argument by arguing that PCT is not about the same phenomenon that conventional theories of behavior are about. Conventional theories are about "behavior" which is tacitly assumed to be caused results of action; PCT is about controlled results of action -- results that are brought to reference states and protected from the effects of disturbance.

I must have missed something here, because I thought PCT was about behavior, just as conventional theories are about behavior. Bill's book is called "BEHAVIOR: the Control of Perception," is it not? The difference is not in what they are about, but in how they go about explaining it. This difference leads to a different approach to data collection, a different view on what variables should be observed and measured, what manipulations should be performed, and many other differences.

Yet there are points of contact, especially with the experimental analysis of behavior and with ethology (behavioral biology), both of which focus on the behavior of single subjects. The subjects behaving in these studies are autonomous biological control systems; the observed behavioral phenomena are the products of such systems operating in a well-defined environment under known conditions. If PCT has any merit at all it should be able to explain why the subjects of these observations behave as they do--given suitable research. A nice example of this approach is the e. coli research, which identified how the bacterium's input and behavioral output mechanisms, organized into control systems, produce the observed changes in tumble frequency as a function of various types of gradient (heat, nutrient, etc.).

I think where you are getting confused is that you think I am proposing that we should go around developing ad hoc models of known behavioral phenomena without doing the research that would ground these models in appropriate PCT-relevant data. That's not my view.

> The answer to the second part of your argument is based on the same point: the distinction between caused and controlled results of action. PCT cannot be expected to handle data that has been "handled" by previous theories when this data is not about controlled results of action.

Again, I am not proposing that we just try to "plug in" a PCT model and thereby "explain" some set of data. I agree that this would be a waste of effort. I am suggesting that we select some interesting behavioral phenomena for which a conventional explanation has been offered and research those phenomena using standard PCT technology, such as applying "the Test" for controlled variables. The resulting model would be a model of the organism, not of the organism's behavior. This model could then be "run" under the conditions in which the behavioral phenomenon of interest occurs; if at all valid, it should then demonstrate said behavioral phenomenon.

>>Bruce:

>> You mean in addition to the Nachtigall and Dethier book references

> Of course. Neither Nachtigall nor Dethier did any PCT-type research.

Pardon, your closed-minded prejudice is showing. As you have read neither, you could not know that for a fact.

Try reading Nachtigall: 40. How is flight velocity regulated? (Pp. 134-139) It has a familiar ring to it.

>> Why don't you have a look at this book, Rick, and judge Nachtigall's understanding for yourself?

> Based on your description of his work and the extended quote from his book, I judge that I can eliminate Nachtigall's book from my list of books that are possibly relevant to PCT research.

You judge wrong.

Nachtigall's description of how a fly lands on the ceiling explains it from the observer's point of view; it was not INTENDED to provide a control system explanation, which would explain how the fly accomplishes what the observer sees. This has nothing to do with whether Nachtigall does or does not understand control theory (as I said in my previous post to you) and it says nothing about whether his other discussions provide a control systems analysis.

Nachtigall's book is intended to give a lay reader a good, general picture of how insects fly, including basic structures and mechanisms, much as you might describe how an aircraft flies. You would discuss the structural properties of the wings and fuselage, the functions of the various components as they relate to flight, some of the basic physics (aerodynamics), engine characteristics, and so on. Most of this discussion would not be about control theory, not until you began to talk about how the pilot (or the autopilot) manages to fly the plane. But there, you would talk about it. That's the way it comes up in Nachtigall's book.

I note that having rejected the book suggestions out of hand, you made no mention of the research articles I suggested. Does this mean you _will_ look them up, or that these have been rejected out of hand as well?

> Maybe it's best to finish this up with a question related to your first question to me:

> Why do you think we (PCTers) don't believe that PCT needs to account for the descriptions of behavior provided by conventional psychology?

I'm a PCTer, too, but "we" don't subscribe to your view, so it is not true that "PCTers" don't believe this.

To answer your question, my perception is that you have at least two reasons:

1) most conventional research did not collect data that can be used to construct an adequate control system model. A model developed specifically to explain the behavioral data in question, absent the relevant information about controlled variables, would be purely speculative, a mere fantasy, and

2) the better approach would be to start from scratch with studies designed to elucidate the controlled perceptual variables and from those data develop a model of the organism. Eventually such a model should "behave" appropriately (i.e., in agreement with observation, within experimental error) under whatever conditions you choose.

My problem with this is NOT that I advocate just fitting models to extant data (I don't, and thus agree with point 1), but rather, that I see no reason why one should not explore the controlled variables that come into play in some of the situations already extensively studied using conventional approaches, then applying the resulting model to explain the known behavioral phenomena that occur in those situations.

Well, look what just arrived:

>Rick Marken (950527.0930)

> So let me ask the question again and try to make the intent clearer:

> We PCTers believe that PCT should not be required to account for most of the descriptions of behavior provided by conventional psychology, nor should it be required to account for most of the results of conventional psychological research.

> Why do you (Bruce or anyone else for that matter) think we believe this?

Ah, BIG change! The added word MOST changes everything. I AGREE with THAT statement (see above).

> I would prefer an essay answer. But, if you prefer multiple choice, here are some alternative short answers:

I've already provided the requested essay answer. As to the multiple choice question, the best answer is g, but I'd also give some weight to b:

> b) We have already accounted for some of these descriptions and results and we are tired of doing it.

(and of seeing our theory rejected out of hand)

Skinner would say that it's a case of punishment suppressing behavior. I wonder how PCT would account for that? Darn! There I go again, trying to account for conventional findings with PCT. When will I learn? (;->

Regards, Bruce

Date: Sat, 27 May 1995 14:13:01 -0500
Subject: Pushing Back

[From Bruce Abbott (950527.1410 EST)]

>Bill Powers (950526.1605 MDT)

>>Bruce Abbott (950526.1700 EST) --

>> My description of how a fly lands on the ceiling is more like the answer Johnny was looking for--superficial but adequate for the purpose. I'm not offering an S-R explanation even if it looks that way on the short form.

> I guess I see what you mean. It occurs to me that we're using the word "how" differently, in the question "How does the fly land on the ceiling?" One meaning is," Tell me all the things that we can see happening as a fly lands on the ceiling." In the other meaning -- which is the one I automatically assume -- the question is "What is going on inside the fly which would account for what we see happening?"

Tell that to Rick Marken. He doesn't get the distinction.

Johnny, after hearing the Marken account of the fly landing on the ceiling: "No, I mean how does it do it? Does it loop, or roll, or what?"

>Bill Powers (950527.0950 MDT)

Illuminating discussion of the difference between the external description of behavior and the variables used by the organism (and model) to control perception.

I agree that Nachtigall's description of the fly's landing pattern does not tell you what variables the fly is controlling when producing that pattern (although you can make a reasonable guess about some of them based on what you see the fly doing). What I think has gotten lost in the squabble about Nachtigall's description (which was intended to answer "how" only in the sense of what the fly does) is my earlier mention of Nachtigall's control-system analysis of certain aspects of insect flight (on which you commented). Rick Marken has forgotten all about this, to the point that he now asserts that Nachtigall's book is unlikely to have anything to say about control systems. As your use of the fly-landing description in this post again focuses on the external description, I though I'd act like a good control system and "push back" by reminding everyone of Nachtigall's control-system analysis, offered in the same book.

Regards, Bruce

Date: Sat, 27 May 1995 16:01:14 -0700
Subject: A Fly In His Ear

[From Rick Marken (950527.1600)]

Bruce Abbott (950527.1320 EST)--

Me:

> I have tried to answer the first part of your argument by arguing that PCT is not about the same phenomenon that conventional theories of behavior are about. Conventional theories are about "behavior" which is tacitly assumed to be caused results of action; PCT is about controlled results of action -- results that are brought to reference states and protected from the effects of disturbance.

Bruce:

> I must have missed something here, because I thought PCT was about behavior, just as conventional theories are about behavior.

Not quite. An important (the most important, I think) point of PCT is that the word "behavior" is ambiguous. It is typically used to refer to any measurable or classifiable result of an actor's actions. PCT shows that it is important to distinguish between actions, controlled results and uncontrolled results. Without making these distinctions, all behavior is treated the same -- a caused result of action. The conventional, "operational definition" approach to deciding what constitutes a person's behavior treats any result of actions as an instance of "behavior; changes in the distance between my fingers and the ceiling is as good a measure of my behavior (especially if it yields good results in an experiment) as are the letters I type. PCT explains why that point of view is wrong. PCT is about behavior -- but only when behavior is understood to be controlled results of action.

> I am suggesting that we select some interesting behavioral phenomena for which a conventional explanation has been offered and research those phenomena using standard PCT technology, such as applying "the Test" for controlled variables. The resulting model would be a model of the organism, not of the organism's behavior. This model could then be "run" under the conditions in which the behavioral phenomenon of interest occurs; if at all valid, it should then demonstrate said behavioral phenomenon.

I think this is a GREAT idea and I heartily endorse it. Indeed, I thought this what you were going to be doing with those rats.

I guess I have been under the mistaken impression that you have been suggesting that there is something to be learned about control systems from the "interesting behavioral phenomena" that we decide to explain with PCT. Of course, as I now see you know, we can learn nothing at all until we start testing for controlled variables

> Try reading Nachtigall: 40. How is flight velocity regulated? (Pp. 134-139) It has a familiar ring to it.

See, this is the kind of thing that confuses me. Flight velocity is a perception of the observer. How does Nachtigall know that it is controlled (regulated) by the fly? Just because someone uses a control model to explain what they observe doesn't mean that they are using PCT, as I'm sure you are aware. If Nachtigall didn't test for controlled variables, then his research is, of course, of no use to PCT. I saw no evidence in your quotes from Nachtigall that he did any testing for controlled perceptual variables. Perhaps there was some quote I missed that described some controlled variables and how they were detected?

Nachtigall seems to have described some very interesting behavioral phenomena. But, as you know, you can't learn anything about control systems just by looking at "interesting behavioral phenomena". But you keep recommending the Nachtigall book as one I should look at for examples of PCT research. See why I'm confused?

Bill Powers (950526.1605 MDT)--

> It occurs to me that we're using the word "how" differently, in the question "How does the fly land on the ceiling?" One meaning is," Tell me all the things that we can see happening as a fly lands on the ceiling." In the other meaning -- which is the one I automatically assume -- the question is "What is going on inside the fly which would account for what we see happening?"

Bruce Abbott (950527.1410 EST) --

> Tell that to Rick Marken. He doesn't get the distinction.

I understand the distinction and I know that you were talking about "how" the fly lands in the first sense, viz. "describe all the things that we can see happening as a fly lands". I guess I was under the mistaken impression that you were saying that we could tell something about "how" the fly lands, in the second sense -- viz, "what is going on inside the fly which would account for what we see happening" -- by watching "how" it lands in the first sense -- viz. "by looking at all we can see happening as a fly lands".

> I agree that Nachtigall's description of the fly's landing pattern does not tell you what variables the fly is controlling ... I though I'd act like a good control system and "push back" by reminding everyone of Nachtigall's control-system analysis, offered in the same book.

You're confusing me again. You agree that Nachtigall's description of the fly's landing pattern does not tell us what variables the fly is controlling. So of what possible value is Nachtigall's control-system analysis? I could build a control system model of an apple falling from a tree but that would not be evidence that I understood the nature of the controlling done by falling apples (they do none, of course).

In reference to Bill Powers (950527.0950 MDT) you say:

> Illuminating discussion of the difference between the external description of behavior and the variables used by the organism (and model) to control perception.

So you seem to have understood and agreed with what Bill had to say. But then you go on to say:

> I agree that Nachtigall's description of the fly's landing pattern does not tell you what variables the fly is controlling ... I though I'd act like a good control system and "push back" by reminding everyone of Nachtigall's control-system analysis, offered in the same book.

This is the kind of thing that keeps puzzling me, Bruce. Bill's post was about how Atkeson and Hollerback made detailed measures of "how" (in the first sense of "looking at all we can see happening") people move their hand in the vertical plane. In other words, Atkeson and Hollerback did for hand movement what Nachtigall did for fly landing; they gave a detailed description of how it happens (ie. what happens). With respect to this approach, Bill says:

> The path which Atkeson, Hollerbach (and many others at MIT and elsewhere) are treading is a blind alley, because no matter how carefully the observations are made and the invariances are calculated, there will be no hint of the control-system organization, the SIMPLE control-system organization, that (I claim) is actually creating the observed trajectories

I know that you understand what Bill is saying here and why he says it. It is an EXTREMELY important point. So I wonder why you would want to remind everyone of Nachtigall's control-system analysis. Since Nachtigall studies fly landing the way Atkeson and Hollerbach study limb movement, it is not clear how Nachtigall could have gotten hints about the control-system organization that is actually creating the landing behavior when Atkeson and Hollerbach could get no hint of such an organization from their trajectories .

I must not be understanding you correctly when you tell me to look at Nachtigall's book for PCT research. I'm sure what you must mean is that it might be a good idea to go back and look at Nachtigall after we have done the research necessary to develop a reasonable PCT model of the controlling done by a fly. And I agree with you-- that would be an excellent idea. But, of course, unless Nachtigall tested for controlled variables, it makes no more sense to look to Nachtigall's data for hints about the control system organization involved in fly landing than it is to look at Atkeson and Hollerbach's data for hints about the control system organization involved in limb movement.

I'm glad that you understand all this and I'm sorry for any misunderstandings on my part. There must be a fly in my ear.

Best Rick

Date: Sat, 27 May 1995 18:10:08 -0700
Subject: Re: Controller variables vs side-effects

[From Rick Marken (950527.1800)]

Bill Powers (950527.0950 MDT) --

I forgot to mention that this post was SENSATIONAL. I have read it over and over again. I can't recall a clearer or more forthright description of the difference between the PCT and conventional approaches to understanding behavior.

I recommend this one for the "required readings" list (if there is one).

What do think, Dag?

Best Rick

Date: Sat, 27 May 1995 21:31:38 -0600
Subject: Why not read what he says?

[From Bill Powers (950527.2115 MDT)]

Rick Marken (950527.1600) --

> I must not be understanding you correctly when you tell me to look at Nachtigall's book for PCT research. I'm sure what you must mean is that it might be a good idea to go back and look at Nachtigall after we have done the research necessary to develop a reasonable PCT model of the controlling done by a fly. And I agree with you-- that would be an excellent idea. But, of course, unless Nachtigall tested for controlled variables, it makes no more sense to look to Nachtigall's data for hints about the control system organization involved in fly landing than it is to look at Atkeson and Hollerbach's data for hints about the control system organization involved in limb movement.

I wouldn't have known about the Atkeson and Hollerbach paper if I hadn't read it. Wouldn't it be simpler just to read Nachtigall? Maybe he does use a control-system analysis. How can you tell without reading what he says?

Best, Bill P.

Date: Sun, 28 May 1995 10:41:08 -0700
Subject: OK, and why don't you read what He says?

[From Rick Marken (950528.1030)]

Bill Powers (950527.2115 MDT) to me:

> Wouldn't it be simpler just to read Nachtigall? Maybe he does use a control-system analysis. How can you tell without reading what he says?

Ok. Ok. I'll read it. But I wasn't trying to avoid reading it. I was basing my comments about Nachtigall and my judgment about whether or not I should take the time to read his book on Bruce Abbott's description and quotes from the book. I can't read every book that is claimed to be relevant to PCT. Samples from the book help me decide whether or not I should invest a lot of time on it. Bruce was nice enough to post an extended quote from Nachtigall. Here it is, in case you missed it:

> The fly approaches the ceiling obliquely from below, at a steep angle, at a speed of about 25 cm per second. It then flies straight into the ceiling, and shortly before striking it, stretches out all three pairs of legs. The fore pair take up a special attitude, held out stiffly upwards so that they are the first part of the fly's body to make contact. Held in this way they act as shock absorbers and as anchors, adhering at the point of contact by means of their claws and pulvilli (hairy pads). Simultaneously the wings stop beating. Now the fly is clinging firmly to the ceiling with its fore feet, but its body still has a certain forward momentum. Like a flywheel on its shaft the fly rotates about its fore feet and turns its belly upwards, grasping the ceiling with its middle and hind feet--and there it is, sitting upside down on the ceiling, without having had to fly upside down first.

Bruce posted this quote [Bruce Abbott (950526.2000 EST)] in response to my request for an example of "real, systematic, PCT data collection" in the insect behavior literature. Obviously there is no evidence of real, systematic, PCT data collection in the above quote. So, either Bruce was pulling my leg again or he managed to find the one paragraph in Nachtigall's book that was not an example of real, systematic PCT data collection.

I did not ask whether Nachtigall did a "control-system analysis" of the fly's behavior; I have seen plenty of control systems analyses based on data collected by conventional methods and I'm not really interested. I was interested in seeing examples of PCT research. Obviously, the quote from Nachtigall is not it.

Nevertheless, I will get Nachtigall's book and read it, if only to prove that this is not a matter of me being unwilling to "look through the telescope". It looks like Nachtigall writes well (and/or is translated well) so it might be fun. And the paragraph quoted above may be the exception and Nachtigall's book may, indeed, be filled with examples of real, systematic, PCT data collection.

By the way, I found a biology book you all might be interested in. I think it has a lot of stuff in it that anticipates PCT. Here's a quote:

> So God created man in his own image, in the image of God created he him; male and female created he them. And God blessed them and said to them, Be fruitful and multiply, and replenish the earth and subdue it; and have dominion over the fish of the sea and over the fowl of the air, and over every living thing that moveth upon the earth

There it is; PCT. God is obviously the reorganization system which is able to create control systems that allow reproduction, fishing and hunting. The concept of control is clearly articulated in this passage. This book is well-written and readily available:

God (-5068) _The Holy Bible_, Gideons: Any Motel, USA

I think you should read the whole book -- EVERY SINGLE WORD -- before you judge whether or not it is relevant to PCT;-)

Best Rick

Date: Sun, 28 May 1995 15:21:37 -0600
Subject: Re: relevance to PCT

[From Bill Powers (950528.1515 MDT)]

From Rick Marken (950528.1030)--

> Bruce was nice enough to post an extended quote from Nachtigall.

Bruce also said that this was only the descriptive part; that there was also an analysis that could be construed in control-system terms.

I consulted the on-line catalogues and found a copy of Nachtigall in Grand Junction. Mary will put in the interlibrary loan request next trip to town.

> God (-5068) _The Holy Bible_, Gideons:Any Motel, USA

> I think you should read the whole book -- EVERY SINGLE WORD -- before you judge whether or not it is relevant to PCT;-)

I have. Its relevance is spotty.

Best, Bill P.

Date: Sun, 28 May 1995 17:33:17 -0700
Subject: Chopped liver

[From Rick Marken (950528.1730)] >Bill Powers (950528.1515)

Re: Nachtigall quote

> Bruce also said that this was only the descriptive part; that there was also an analysis that could be construed in control- system terms.

I have never expressed an interest in whether Nachtigall provides "an analysis that could be construed in control-system terms". I'm only interested in PCT research.

I had said to Bruce:

> if there is some "real, systematic, PCT data collection" in the insect research field, please point it out.

And Bruce replied with:

> Here's Nachtigall's description:

Followed by the paragraph on fly landing quoted from the book.

I took this to mean that Bruce was describing the Nachtigall work as an example of "real, systematic, PCT data collection".

"L'affaire Nachtigall" started when Bruce said:

> you really should take a look at some of this insect research. In this field, at least, the "real, systematic, PCT data collection" seems to be much more common than you seem to believe.

And I said:

> Why not help out and post some examples.

So I was under the impression that Bruce was describing the Nachtigall studies as an example of "real, systematic, PCT data collection". Maybe I misunderstood. But I am not particularly interested in whether Nachtigall presents "an analysis that could be construed in control- system terms". I am interested in examples of PCT research.

I thought that the insect research literature actually might contain good examples of PCT research; research like the bee studies described by Avery; real PCT-type research where there is a hypothesis about a controlled variable and methods used to test this hypothesis.

I feel like I am getting ragged on for not reading a book after I have found out that it is not what I'm looking for. It is as though the following had happened: I ask if anyone knows of a book about grinding telescope mirrors; someone says "yes" and posts a sample paragraph from the book that describes grinding livers for chopped liver; I don't check out the book because it doesn't seem to be about grinding mirrors at all; I get yelled at for not checking out the book to see whether or not it can help me grind telescope mirrors.

Am I missing something here?

> I consulted the on-line catalogues and found a copy of Nachtigall in Grand Junction.

Linda and I took a lovely walk among the filthy rich; it's not in the Beverly Hills library. I guess I'll have to get it through interlibrary loan and that will take some time.

Best Rick

Date: Mon, 29 May 1995 02:47:08 -0400
Subject: Re: Why not read what he says?

<[Bill Leach 950529.02:40 U.S. Eastern Time Zone]
>[From Bill Powers (950527.2115 MDT)]

I was almost floored by this posting. Rick specifically asked for a posting to provide an example of Nachtigall's work that could be PCT related.

Bruce then posted the interesting but irrelevant (to PCT) example of how a fly lands on a ceiling.

Rick (presumably wondering at this "PCT" example) blasts it rather handily and points out that this example provide little encouragement to read the work.

If I were Rick, I would be more than just a little stunned at your posting on the matter.

-bill

Date: Mon, 29 May 1995 06:51:32 -0600
Subject: Judging a book by a paragraph

[From Bill Powers (950529.0600 MDT)]

Rick Marken (950528.1730) --

> I was under the impression that Bruce was describing the Nachtigall studies as an example of "real, systematic, PCT data collection". Maybe I misunderstood.

Bill Leach (950529.02:40)--

> I was almost floored by this posting. Rick specifically asked for a posting to provide an example of Nachtigall's work that could be PCT related. Bruce then posted the interesting but irrelevant (to PCT) example of how a fly lands on a ceiling.

My impression was that Bruce quoted a paragraph or two from Nachtigall's book in which the landing behavior of a fly was described in detail, and went on to say that elsewhere in the book there were analyses of this behavior (and of control of relative velocity when insects fly in parallel) as control-system behavior. I never assumed that what Bruce quoted was ALL that was in the book, which seems to be what you guys are assuming. For one thing, in the quoted material there was no mention of any experimental manipulations to test hypotheses about what was actually under control. My assumption would be that those subjects would be presented later, after a description of the phenomenon to be explained. That's how I would do it: first describe what we can observe in sufficient detail; then do the experimental tests of hypotheses, to see which observations matter and which don't.

Of course when I read the book I may well be disappointed to find there was no attempt to verify any hypothesis about what the fly was controlling. This wouldn't be the first time I have been disappointed in this way. However, as an author I always hope that people will read more than one paragraph of my writings before deciding what I am writing about. Maybe I have the naive hope that if I stick to that principle in reading other people's works, a kind of sympathetic magic will occur that encourages others to do the same for me.

A quite separate question is whether Bruce understood what we were looking for as an example of "real, systematic, PCT data collection". If he didn't understand, then we should look at the book and explain what is missing, what would be needed to make it such an example. If he did understand, but just didn't happen to cite the experimental details, then we would be jumping to conclusions if we assumed that what he cited represents the totality of his understanding, or the totality of what is in the book.

Rick:

> But I am not particularly interested in whether Nachtigall presents "an analysis that could be construed in control- system terms". I am interested in examples of PCT research.

Well, Srinivasan presented an analysis (cited by Avery Andrews) that could be construed in control-system terms, even though (I gather) he never mentioned controlled variables or the principle of using disturbances to weed out uncontrolled side-effects. It was very sharp of Avery to recognize in Srinivasan's article the experimental manipulations that amounted to the Test for the controlled variable. I think it unlikely that we will find in ANYONE's work, outside PCT, an explicit description of applying disturbances and testing for resistance to them as a method of testing hypotheses about what is being perceived and controlled by the test subject. This is not yet a recognized method. Nevertheless, we might well find other cases in addition to Srinivasan's in which this essential step was taken, simply as a common-sense verification that the variables being considered were actually relevant to the system in question.

I, too, resist reading books just because someone says they say something about control theory, especially when what I know of them seems unrelated to PCT. The problem is that ALL discussions of human behavior have something to do with PCT, but the authors of such discussions don't know that. It's like the old problem that comes up when we mention "purposive behavior." People have noticed that behavior is purposive for a long time, but very little of what they have written on this subject says anything like what we would say. We could read a thousand papers from the past on purposive behavior and never find anything of theoretical interest to us. So after a while, one begins to resist reading such papers on principle, because the payoff is so vanishingly small.

However, when a colleague tells us that a book contains something of relevance to PCT, I think we ought to give it at least a quick read. Then we can either set the colleague straight or thank him.

Best to all, Bill P.

Date: Mon, 29 May 1995 07:47:27 -0600
Subject: Re: PCT observations of behavior

[From Bill Powers (950529.0700 MDT)]

RE: PCT data on behavior

We have a humming-bird feeder hanging on our back deck, with four hummingbirds who vie for the four feeding ports. This feeder consists of a bulb and a red ovoid unit below it with four slits in recessed openings. The assembly hangs on a nylon string about two feet long.

The hummingbirds zoom up to this feeder at high speed, and stop dead in the air to look around before moving up to a feeding port. They change their orientation through 360 degrees, stopping perhaps three or four times in one scan. What actually stops is not their bodies, but their heads. The bodies move through an angle in the vertical plane, but the heads remain level and stationary when sighted against a background object.

When the birds move up to the feeding ports, the whole feeder is usually swinging gently in the wind and twisting slightly around the axis of the supporting string. The head and body of the bird swing back and forth in space, and the orientation of the bird moves in an arc around the feeder as it twists, so a constant distance is maintained between the beak and the feeding port and the beak is pointed directly into the port. I can't see any variation in the relationship of beak to feeding port while this is going on. While the bird is feeding, the angle of the body may change in a vertical plane, but the head remains fixed in relationship to the swinging and rotating feeder. The body appears to rotate vertically about the head.

It takes a while to realize that there is another disturbance beside the swinging and rotation of the feeder. There is usually some variable wind blowing, so the bird is actually not maintaining a position in still air but is flying upwind at varying relative speeds and directions all the time it is stationary relative to the ground or the feeder. There is no visible sign of the effect of the wind or the swinging and twisting of the feeder on the relationship between the bird's head and the feeder port.

From the bird's point of view, it is of course the feeder's distance and orientation that is stabilized. The bird modifies the beat pattern of its wings (a blur to me) to maintain the feeder at a constant distance and in a constant orientation relative to the head. This happens despite disturbances which, without adjustments by the bird, would have an easily visible effect on the relationship of bird to feeder. The controlled variable is probably controlled visually. When the beak is actually in the feeding slot, tactile sensations may also become part of the controlled variable.

As seen through my own perceptual systems, the controlled relationship is not disturbed at all by the movements of the feeder or by the variable wind speed. Artificial aids would be required to measure the residual effects of disturbances, to determine the loop gain and dynamics of the control systems involved. It would take high-speed cameras to sort out just how the figure-eight beat of the wings is modified as a means of varying the relationship between head position and orientation and the position and orientation of the feeder.

In this case, nature provides disturbances and we can easily see that they are resisted.

Best, Bill P.

Date: Mon, 29 May 1995 09:36:25 -0500
Subject: Nachtigall's book and PCT research

[From Bruce Abbott (950529.0935 EST)]

>Rick Marken (950528.1730)

> "L'affaire Nachtigall" started when Bruce said:

\>> you really should take a look at some of this insect research. In this field, at least, the "real, systematic, PCT data collection" seems to be much more common than you seem to believe.

\> And I said:

\>> Why not help out and post some examples.

"L'affaire Nachtigall" started much, much earlier when I provided the following (950520.1130 EST):

\> The fascinating bee-as-control-system research by Mandyam Srinivasan extends work done much earlier from a similar perspective. Werner Nachtigall (1968) describes some aspects of bee flight as follows:

In nature it is the insect antenna that is displaced. The negative of its signal is sent to the mechanism regulating wing-beat amplitude, by way of a sense organ in the antenna, then via the nerves to computing centers in the central nervous system, to nerves once again and finally to the flight muscles. If the insect antenna is bent more strongly, the wing-beat amplitude is made smaller as a result, and thus the undesirable increase in velocity is reduced until the velocity reaches its set value, for example 14 m.p.h.

The bee . . . has particular muscles which can change the position of the antennae slightly with respect to the air stream. Then the air resistance during flight can no longer bend them by the same amount as before but rather, for example, a little less. . . . a new set point for the flight velocity is achieved.

Bees and flies have two antennae, one right and one left. It has been shown that each of these antennae regulates only the amplitude of the wing on its own side of the body. This has possibilities for flying in a curved course. If one antenna is cut off, the insect always flies around in circles. It is only when both antennae are operating together that they permit the bee to fly in a straight line and compensate for gusts of wind which might push it a little away from this line. This is critically important for orientation of the flight between hive and food source, which the bee should make as straight as possible.

. . . But there is one more point to be made: the bee does have a second servo system to control its velocity, involving the two large compound eyes. The essential purpose of this system is to hold constant the velocity over the ground; the antennae, on the other hand, regulate the velocity through the air.

. . . But an optical measuring device which fixes points on the ground and computes how quickly they move backwards, can do this. The highly complicated compound eyes are admirably suited to this task. They can recognize the dangerous drifts produced by head, tail, or side winds and compensate for the error in the signals of the antennal control circuit.

Nachtigall, Werner (1968 [English translation 1974]). _Insects in flight: a glimpse behind the scenes in biophysical research. New York: McGraw-Hill, p. 139.

The quoted material shows Nachtigall clearly appreciates that the systems responsible for the bee's behavior are control systems. If I had stopped there then perhaps this whole silly argument would never have taken place. But nooooooo, I had to mention Nachtigall's description of how a fly lands on the ceiling, offered from the _observer's_ point of view. For some reason, Rick, you seized on the latter as if IT were the evidence I provided for Nachtigall's understanding of control theory rather than the above. You have been pursuing this thesis relentlessly over my repeated protests that Nachtigall's description of fly landing is NOT the relevant quote, that it is in fact only a pure description of observable behavior.

\> I had said to Bruce:

>> if there is some "real, systematic, PCT data collection" in the insect research field, please point it out.

And Bruce replied with:

>> Here's Nachtigall's description:

> Followed by the paragraph on fly landing quoted from the book.

> I took this to mean that Bruce was describing the Nachtigall work as an example of "real, systematic, PCT data collection".

You must have labored mightily to take it that way, because the immediately following paragraph says:

>> This was all determined via high speed photography, is purely descriptive, and is irrelevant to the issue of Nachtigall's understanding of control system theory--and to mine.

> So I was under the impression that Bruce was describing the Nachtigall studies as an example of "real, systematic, PCT data collection". Maybe I misunderstood. But I am not particularly interested in whether Nachtigall presents "an analysis that could be construed in control- system terms". I am interested in examples of PCT research.

> I thought that the insect research literature actually might contain good examples of PCT research; research like the bee studies described by Avery; real PCT-type research where there is a hypothesis about a controlled variable and methods used to test this hypothesis.

The only claim I have made for Nachtigall's book is that it shows that Nachtigall understood what a control system is and that he recognized that insect flight is accomplished by controlling certain perceptual inputs; if you're expecting to find detailed descriptions of investigations designed to identify controlled perceptions, you won't find them there: you will have to look at the sources Nachtigall provides on which he bases his conclusions. What you will get are nice descriptions like the following example.

Nachtigall is asking how to go about photographing a fly in flight:

    Some clever people have hit on the splendid solution of fixing the camera and holding the insect in front of it, attached by the thorax or abdomen to a piece of wood which is held in a stand. As soon as its legs are removed from the ground, the fly begins to beat its wings, and the shutter can be pressed. Very good. We have a beautiful film, but does it tell us anything about how the insect would have moved its wings in free flight? The experimental conditions are unnatural, but does this make a difference?

In the paragraphs following this one, Nachtigall gives the answer: it does make a difference, a very important one. For the fly on the stick,

    The animal no longer has to compensate for its weight because the stand is carrying it. There is no resistance either, since it is not moving through the air. The fly can generate any lift and thrust it wants to.

To get the correct picture,

    . . . the insect must be allowed to balance the forces itself as it does in free flight . . .

Thus, Nachtigall displays a sensitivity to the problems of studying a living system that is controlling its own perceptual variables. You, on the other hand, have been making blanket statements suggesting (based on the fly landing stuff) that Nachtigall does not, and this is what we have been debating about.

I find it curious that you have chosen to emphasize Nachtigall's book as my example of "real, systematic, PCT data collection" and but have neglected to mention the several scientific articles I referenced. Even these do not necessarily represent the best examples (as I noted when I cited them, they were simply a few references I happened to have handy), but they provide a starting point for exploring this literature.

> I feel like I am getting ragged on for not reading a book after I have found out that it is not what I'm looking for. It is as though the following had happened: I ask if anyone knows of a book about grinding telescope mirrors; someone says "yes" and posts a sample paragraph from the book that describes grinding livers for chopped liver; I don't check out the book because it doesn't seem to be about grinding mirrors at all; I get yelled at for not checking out the book to see whether or not it can help me grind telescope mirrors.

> Am I missing something here?

Yes. My suggestion to read Nachtigall was offered because you had asserted based on the fly-landing example that Nachtigall doesn't understand control system theory. I said he does, and to see that, you should read his book. What you are getting "ragged for" is making sweeping assertions based on limited evidence. I never claimed that his book would be filled with examples of "PCT research." Meanwhile, I'm finding that the stakes have been raised: you will now insist that the book contain far more than what I claimed for it. By setting up this straw man you can now obtain the book and then ask where all the PCT research is. Remember, my purpose in mentioning Nachtigall was to point out that the work of insect researchers like Mandyam Srinivasan had precursors already moving in this direction in 1968.

Clearer now? (:->

Regards, Bruce

Date: Mon, 29 May 1995 15:00:52 -0700
Subject: PCT Research

[From Rick Marken (950529.1500)]

Bill Leach (950529.02:40) --

Re the Bill Powers (950527.2115 MDT) "get a book" post.

> I was almost floored by this posting.

Thanks, Bill. Glad to see that "the audience if listening".

Bill Powers (950529.0600 MDT)--

> It was very sharp of Avery to recognize in Srinivasan's article the experimental manipulations that amounted to the Test for the controlled variable. I think it unlikely that we will find in ANYONE's work, outside PCT, an explicit description of applying disturbances and testing for resistance to them as a method of testing hypotheses about what is being perceived and controlled by the test subject.

Of course it was sharp of Avery to notice this; Avery BRILLIANTLY spotted the essence of PCT research in research that was not intentionally based on PCT. That's what this whole thread is about, no? Finding examples of PCT-like research (like the Science study) that was not explicitly based on PCT. Bruce said there were lots of examples of this research in the insect behavior literature -- but he provided none.

Do you really believe that I expected Naftigall to be talking about "disturbances" and "controlled variables"? I think it was clear from the discussion on the net that Srinivasan never used the terminology of PCT or conceived of what he was doing as "The Test". I thought that Bruce was going to do for Naftigall's research what Avery did for Srinivasan's; explain how what Naftigall was doing was equivalent to The Test. He didn't. What are you controlling for, Bill?

Bruce Abbott (950529.0935 EST) --

> The quoted material shows Nachtigall clearly appreciates that the systems responsible for the bee's behavior are control systems.

Manual control theorists appreciate that the systems responsible for driving an automobile are control systems; control systems with "command" (reference) inputs that come from the environment and specify the system's output. Understanding control systems theory and understanding PCT are two different things.

> Rick, you seized on the latter as if IT were the evidence I provided for Nachtigall's understanding of control theory

I seized on IT (the description of fly landing) as a description of Naftigall's PCT like research -- because you said it was. Moreover, most people who understand control theory (like most people who don't) do NOT understand PCT.

> My suggestion to read Nachtigall was offered because you had asserted based on the fly-landing example that Nachtigall doesn't understand control system theory.

No. I asserted that the description of fly-landing was not an example of PCT-like research. Based on the paragraph you quoted, I am willing to bet that there is no trace of PCT-like research in Nachtigall's book. But I would not be at all surprised to find that Nachtigall understands control theory.

> What you are getting "ragged for" is making sweeping assertions based on limited evidence.

Not everyone thinks so. See Bill Leach (950529.02:40).

> I never claimed that his book would be filled with examples of "PCT research."

Are there ANY such examples? Why did you bring up the book in the first place?

> By setting up this straw man you can now obtain the book and then ask where all the PCT research is.

Caught me. I'm always busy setting up straw me. My usual straw man is "reinforcement theory" but I'm trying to get a toe hold in the insect behavior literature;-)

> Clearer now? (:-

Yes. I was working under the assumption that you had joined CSG-L to learn PCT. Now I see that you already know all about it.

While I am busy learning about fly behavior from Nachtigall, I would really appreciate it if you could post a review of my book, "Mind Readings". I am learning from you that I have many misconceptions about PCT: what it is and what it's about. So, before I "strike again" with another book about PCT I would really appreciate it if you could point out some of the more egregious errors in "Mind Readings".

Thanks Rick

Date: Mon, 29 May 1995 19:57:17 -0500
Subject: Buzz off

[From Bruce Abbott (950529.1955 EST)]

>Rick Marken (950529.1500)--

Skipping over the ridiculous, we come to

>> Clearer now? (:-

> Yes. I was working under the assumption that you had joined CSG-L to learn PCT. Now I see that you already know all about it.

I've been learning. How about you? Are you accomplishing what you set out to accomplish on CSG-L? What ARE you trying to accomplish on CSG-L? (Enquiring minds want to know!) It's a bit of a mystery from here.

> While I am busy learning about fly behavior from Nachtigall, I would really appreciate it if you could post a review of my book, "Mind Readings". I am learning from you that I have many misconceptions about PCT: what it is and what it's about. So, before I "strike again" with another book about PCT I would really appreciate it if you could point out some of the more egregious errors in "Mind Readings".

I haven't read it. Why don't you quote a paragraph? I understand that I should be able to judge a whole book from one paragraph. (;->

I don't recall saying anything about your having misconceptions about PCT or making "egregious errors." Do you feel that I have? Do you feel that YOU have?

If you have any curiosity about insect flight, I think you'll enjoy Nachtigall's book despite yourself. It is wonderfully written.

By the way, did you read Bill Powers' nifty little description of hummingbird behavior? On second thought, you wouldn't find it of interest--no application of the TEST and all that, just description from the point of view of the external observer, with a little speculation about possible controlled variables thrown in. Obviously of absolutely no value to a true PCTer like yourself. Forget I mentioned it.

Well, I've got to buzz off.

Regards, Bruce

Date: Mon, 29 May 1995 20:49:26 -0700
Subject: Buzz Bombed

[From Rick Marken (950529.2045)]

Bruce Abbott (950529.1955 EST) --

>What ARE you trying to accomplish on CSG-L?

I'm trying to present what I understand about PCT after having studied and worked on it for over 15 years. I like the CSG-L forum because I can see what people do and do not know about PCT. I try to address questions or misunderstandings about PCT that I think I can answer based on my experience doing PCT research and modelling.

Me:
> I would really appreciate it if you could point out some of the more egregious errors in "Mind Readings".

Bruce:
> I haven't read it. Why don't you quote a paragraph? I understand that I should be able to judge a whole book from one paragraph. (;-

I wasn't asking you to read it for enlightenment. I was asking it as a favor so that you could tell me where I had made mistakes from your point of view. If you don't want to read it, that's fine.

> I don't recall saying anything about your having misconceptions about PCT or making "egregious errors."

I get the impression that you often disagree with me (or consider my statements ridiculous) when I make what I consider to be basic points about PCT. This suggests that your understanding of PCT is quite different from mine. For example, I don't think I have ever heard you agree that "selection by consequences" cannot possibly be a reasonable characterization of the relationship between an organism and its environment. Given these persistent differences, I was interested in how you would evaluate my collection of papers ("Mind Readings") about PCT. I was being only somewhat facetious when I suggested that you would find "egregious errors" in the book; since your perception of PCT seems to be somewhat different than mine, I am sure that some parts of my book will create an error signal in you. I was curious about what would create those errors.

> By the way, did you read Bill Powers' nifty little description of hummingbird behavior? On second thought, you wouldn't find it of interest--no application of the TEST and all that, just description from the point of view of the external observer, with a little speculation about possible controlled variables thrown in.

I did read it but I did see an application of the Test -- and suggestions on how to do more detailed applications of the Test using high speed photography. Here, for example, is an implicit Test for control of perceived distance:

> When the birds move up to the feeding ports, the whole feeder is usually swinging gently in the wind and twisting slightly around the axis of the supporting string. The head and body of the bird swing back and forth in space, and the orientation of the bird moves in an arc around the feeder as it twists, so a constant distance is maintained between the beak and the feeding port

If Naftigall describes how flying insects maintain constant results in the face of naturally occurring disturbances then he has described an application of the Test. I didn't notice such a description in the paragraph you posted but there may, indeed, be such descriptions in other parts of the book.

Best Rick

Date: Tue, 30 May 1995 02:44:26 -0400
Subject: Re: Judging a book by a paragraph

<[Bill Leach 950529.21:34 U.S. Eastern Time Zone]
>[From Bill Powers (950529.0600 MDT)]

I even admit that while I had the same impression as Rick, I too was wrong about Bruce's intentions with regard to the description. I do not believe however, that either Rick or I were too far amiss since even with the disclaimer, the description made no sense from the standpoint of both Bruce's assertions and Rick's requests.

I don't believe that Rick, and I certainly know that I, did not assume that what was quoted was ALL that exists in the book but _please_ both of you guys (Bill and Bruce) how about honoring the quite obvious intent (at least to me) of Rick's original request...

He asked for examples of PCT-like material, which I perceived as being done so that he could judge the desirability of obtaining and reading the work. This is not wholly unreasonable though of course Bruce could have declined altogether for any number of reasons.

> For one thing, in the quoted material there was no mention of any...

I believe that was noticed...

> A quite separate question is whether Bruce understood what we were ...

I did not assume that the quote represented Bruce's ability to identify valid PCT related research data... I was just shocked that the material presented in response to Rick's request was so completely unrelated.

Bruce is certainly not obligated to do "book reviews" for Rick Marken or anyone else for that matter but it still seem to me that it would have been more reasonable to decline to provide an excerpt (with or without an excuse or explanation) than to have provided the one that was given.

OK, I suppose that we have all had about enough of this one and I'm not sure that all the feathers will be smoothed back down at this point regardless of what happens.

Besides I just read Bruce's 950529.0935 and thus probably should not send this at all... I also missed the fact that the Bee description posted by Bruce was a Nachtigall quote.

BTW, it looks like your mail program is working like a Word Processor again. You often have large blank areas in the middle of your paragraphs at about the right point for "page breaks".

-bill

Date: Tue, 30 May 1995 02:56:35 -0400
Subject: Re: Buzz Bombed

Hi Rick;

Hummm, I'm glad your response to Bruce's message of 950529.1955 arrived at the same time as did his message.

I almost considered blasting him twice...

Once for the "saying anything about your having misconceptions about PCT" and once for the "Forget I mentioned it." about Bill's outstanding posting of a pure observation from a PCT perspective.

As it is, you covered things nicely (maybe even more nicely than I would have) and of course I am still a very junior PCTer.

-bill

Date: Tue, 30 May 1995 08:40:46 .
SUBJECT: Fly Boys

{from Joel Judd 950530.0830 CST}

Bruce A. & Rick M. (various):

After looking over five days' postings at once, my hypothesis is that you two are not controlling for understanding at all; rather, this thread is an opportunity to post bad fly pun titles. In fact, I have noticed this trend on the net several times, and question its scientific usefulness.

Rick M. and Bill P. (950529):

> read EVERY SINGLE WORD (of the Bible)

I have, too, and actually find it MOST relevant to PCT. However, since it was never intended to be a treatise on human psychology, maybe it shouldn't consistently be criticized for being a poor one. Or perhaps it should be read with a different purpose in mind...

Joel

Date: Tue, 30 May 1995 11:41:02 -0500
Subject: What Is PCT Research?

[From Bruce Abbott (950530.1140 EST)]

One of the points to come out of the overly-extended discussion between Rick Marken and me about insect research was Rick's definition of what constitutes PCT research. According to Rick, PCT research is aimed toward identifying in a given situation the controlled perceptual variables, by systematically applying disturbances and noting whether these disturbances are resisted (i.e., conducting the Test). I agree that this step is essential, but I see it as only one phase of what I would define as PCT research. In my view, research in which the closed-loop nature of the system under study is explicitly recognized and studied as such would qualify. Such studies might be aimed at identifying the controlled variables OR at elucidating the specific structures involved and their functions within the control system. The focus might be on a quantitative analysis of receptor input functions, motor output functions, comparator functions, and so on. When I encounter work like this in the literature I think, hey, PCT stuff! I'm not sure that Rick (and Bill Powers) would agree, given Rick's narrower definition.

Before we get into another heated debate over whether this or that research is PCT research, I'd like this issue settled. If you prefer to hold to your definition, then I'll simply change nomenclature and refer to the more general approach to understanding control systems as "control systems" research rather than PCT research.

An example of the kind of research I have in mind is Kittmann, R. (1991). Gain control in the femur-tibia feedback system of the stick insect. _Journal of Experimental Biology_, _157_, 503-522. (This is one of the three journal articles I cited in an earlier post as an example of PCT research on insect behavior.) The title is self-explanatory. The reported study follows up on Kittmann's earlier work on the same system; in the current paper he describes quantitative aspects of the system as studied under both closed-loop and open-loop conditions. Kittmann notes the following in his introduction:

> In proprioceptive feedback systems there is a lack of quantitative data concerning . . . changes in the characteristics of the system. The variation in gain -- the ratio between the output and the input of the system -- is particularly important, as it can change the characteristics of the system considerably. Low gains result in ineffective feedback responses, whereas high gains can induce instability, e.g., oscillation of the system. Therefore, to maintain effective feedback, gain must be carefully controlled.

Kittmann's preparation was as follows:

> Extracellular activity of the extensor tibiae motoneurones was recorded from the extensor nerve with 50 micrometer steel wires. Closed-loop experiments were performed under these conditions. For open-loop experiments, the fCO was mechanically stimulated as described by Bassler (1976); a pen motor with a pair of forceps connected to its axis was used to move the chordotonal apodeme, which was cut distally at the FT joint.

Now it seems to me that a study of the gain of the proprioceptive feedback system and its variation under varying experimental conditions is a PCT study. Is it? If not, please explain.

Regards, Bruce

Date: Tue, 30 May 1995 10:40:14 -0600
Subject: Re: Test; misc

[From Bill Powers (950530.0945 MDT)]

Bruce Abbott (950529.1955 EST) --

> By the way, did you read Bill Powers' nifty little description of hummingbird behavior? On second thought, you wouldn't find it of interest--no application of the TEST and all that, just description from the point of view of the external observer, with a little speculation about possible controlled variables thrown in.

I think that my informal presentation may have disguised the Test too well. The controlled variables I had in mind were the distance of the bird's head from the feeder and the orientation of the head relative to a feeding port. The disturbances were the swinging of the feeder, the twisting oscillations of the feeder, and the wind velocity. I observed that the controlled variables (during feeding) remained essentially constant, while the expected effects would be a variation in both the distance and the orientation (expected, that is, under the hypothesis of no control, so that the effects of the disturbances would not be opposed by the bird's actions).

A similar description of the fly's landing would have to include descriptions of disturbances which, unopposed, would have predictable effects on the observed landing pattern. To demonstrate control of any particular variable in that pattern, we would have to show that for each disturbance, the fly altered its behavior in such a way that the controlled variable was immediately restored to its former state, or prevented from changing away from that state.

This is what was missing from the Nachitgall description of the fly's landing pattern. The mere repetition of the pattern is not sufficient to demonstrate control. We must show that disturbances tending to change the pattern are resisted.

Suppose we set up a flow of air parallel to the surface on which the fly is landing. We might find (I'm guessing) that now the fly does not orient its body at a steep angle to the surface, but approaches at a shallower angle when flying upwind. This would tell us that the angle of the body relative to the surface is not a controlled variable, but is part of the variable actions used to control some other variable. Perhaps what we would find is that the direction of approach to the surface rather than the angle of the body is under control. The body angle is simply varied to keep the direction of approach the same. Of course if you always observe without any disturbances present, you can't tell that the body angle is not under control.

A parallel case would be observing how a boater rows across a body of water, a lake. If we observe that the rower points the boat toward the distant dock, we might conclude that the direction of the bow of the boat relative to the dock is a controlled variable. If, however, we do the same test on a river, which disturbs the path of the boat relative to the distant dock, we will find that the bow of the boat is aimed upstream, while the direction of progress of the boat continues to be a straight line toward the dock. So the direction of the bow is ruled out as a controlled variable, and the direction of progress (or some perception related to it) is likely to be under control.

In the Srinivasan article described by Avery Andrews, the Test was actually applied. Various aspects of the environment were deliberately altered, to test the idea that visual outflow was the controlled variable and to rule out other aspects of the perceptual situation as being stabilized against disturbances. That is why this work was considered to be a "good PCT experiment" even if Srinivasan didn't formalize what he was doing in PCT terms. The lack of such disturbances was the reason that the Nachtigall description was not considered a "good PCT description." Nachtigall didn't describe any disturbances. Without disturbances we can't identify controlled variables experimentally.

Bill Leach (950529.21:34) --

> I don't believe that Rick, and I certainly know that I, did not assume that what was quoted was ALL that exists in the book but _please_ both of you guys (Bill and Bruce) how about honoring the quite obvious intent (at least to me) of Rick's original request...

Why run this like a mid-term exam? I've been waiting for someone to explain to Bruce why the Nachtigall description was not a good PCT description, but as you can see, I have finally provided the explanation myself. That would have been an easy way to avoid the squabbles. I can only conclude that there is some other goal involved, like showing who knows most about PCT. What a bore.

Joel Judd (950530.0830 CST) --

Rick M. and Bill P. (950529):

>> read EVERY SINGLE WORD (of the Bible)

> I have, too, and actually find it MOST relevant to PCT. However, since it was never intended to be a treatise on human psychology, maybe it shouldn't consistently be criticized for being a poor one. Or perhaps it should be read with a different purpose in mind...

I think that you and I read the Bible using different premises.

Best to all,

Bill P.

Date: Tue, 30 May 1995 10:26:12 -0700
Subject: Books, Understandings, CSG Goals

[From Rick Marken (950530.1030)]

Bruce Abbott (950529.1955 EST) --

> I haven't read it ["Mind Readings"}. Why don't you quote a paragraph?

Well, I just ordered Nachtigall's book through interlibrary loan. Maybe now you'd be willing to take a look at "Mind Readings". I am rather sorry that you haven't read it since you are certainly in the target audience (scientific psychologists with an interest in PCT). William T. Powers said of "Mind Readings" in the Forward: "This is a book that can show a willing psychologist how to do a new kind of research". Are you willing?

Joel Judd (950530.0830 CST) to Bruce A. & Rick M.

> my hypothesis is that you two are not controlling for understanding at all;

Possibly true. I honestly was interested in seeing descriptions (like Avery's) of PCT-like research in the insect behavior literature. Perhaps my response to Bruce's description of fly-landing behavior should have been more positive -- like "yes, it sounds like Nachtigall's book is all about how insects control". Would that have been more understanding?

Bill Powers (950530.0945 MDT) to Bill Leach (950529.21:34) --

> I've been waiting for someone to explain to Bruce why the Nachtigall description was not a good PCT description, but as you can see, I have finally provided the explanation myself. That would have been an easy way to avoid the squabbles. I can only conclude that there is some other goal involved, like showing who knows most about PCT. What a bore.

Actually, I thought I had explained why Nachtigall's description was not a good PCT description. Not nearly as well as you did, perhaps, but I thought I had mentioned the absence of any evidence of disturbance resistance. Anyway, I have no interest in showing who knows more about PCT. I LIKE it when people understand PCT -- and I like it even better when (like Avery and Tom and you) they understand it better than me. I will try to do a better job of explaining why things (like Nachtigall's description) are "not good PCT" (if they are not).

I thought I had been doing an OK job of explaining (and demonstrating) my points when I disagreed with people about PCT but I'll TRY to do better in the future. I don't know if you have noticed but a lot of squabbling still happens AFTER we explain why something or other is not good PCT. For example, there was quite a lot of squabbling about "control by consequences", after a rather long sequence of computer demonstrations of the impossibility of control by consequences.

As you said, push on a control system, it pushes back. Explanations are still "pushes" if they disturb a controlled variable.

Best Rick

Date: Tue, 30 May 1995 13:03:26 -0700
Subject: Re: What is PCT Research

[From Rick Marken (950530.1300)]

Bruce Abbott (950530.1140 EST) --

> Kittmann's preparation was as follows:

> Extracellular activity of the extensor tibiae motoneurones was recorded from the extensor nerve with 50 micrometer steel wires. Closed-loop experiments were performed under these conditions. For open-loop experiments, the fCO was mechanically stimulated as described by Bassler (1976); a pen motor with a pair of forceps connected to its axis was used to move the chordotonal apodeme, which was cut distally at the FT joint.

> Now it seems to me that a study of the gain of the proprioceptive feedback system and its variation under varying experimental conditions is a PCT study. Is it? If not, please explain.

This is a good question. I would say that such a study is a PCT study as long as one had already identified the significant components of the control loop - - in particular the controlled variable, perceptual signal, reference signal, error signal and output variable. I think it is very possible that Kittmann did identify all the components of the proprioceptive feedback system and that he knew what is most important about this control system (from a PCT perspective), viz. what variable it controls. Once you know the controlled variable, it is not unreasonable to want to know the gain of the control system under various conditions (of disturbance, for example). Assuming that Kittman knew the variable controlled by the proprioceptive feedback system, his studies of the measured gain of this system are consistent with, and useful to, PCT.

I don't fully understand the description of how Kittmann determined the gain of this control loop but I'm willing to assume that he did it correctly.

Best Rick

Date: Tue, 30 May 1995 15:17:44 -0500
Subject: Finis

[Bruce Abbott (950530.1445 EST)]

>Bill Leach 950529.21:34

> OK, I suppose that we have all had about enough of this one and I'm not sure that all the feathers will be smoothed back down at this point regardless of what happens.

>Joel Judd 950530.0830 CST

>Bruce A. & Rick M. (various):

> After looking over five days' postings at once, my hypothesis is that you two are not controlling for understanding at all; rather, this thread is an

opportunity to post bad fly pun titles. In fact, I have noticed this trend on the net several times, and question its scientific usefulness.

>Bill Powers (950530.0945 MDT)

> I can only conclude that there is some other goal involved, like showing who knows most about PCT. What a bore.

One of the well-known little facts of social psychology is that people can be drawn into doing things bit-by-bit that they would never do if they knew at the beginning how far it was going to go. For example, you're driving an old car and the alternator quits. After considering the pros and cons of fixing the alternator versus trading in the car for a new one, you go with the repair. One hundred and fifty dollars later the alternator is working, and then the fuel pump goes out. You've just invested 150 bucks, now it's going to cost $75 more to get back on the road -- and $75 is still cheaper than a new car. So you fork out. Two weeks later the CV joints fail ($350). But you've already committed $225 to keep the car; what's $350 more? You see where I'm going: six months later you've spent over $2000 to keep the old clunker on the road. If you'd been faced with $2000 for repairs versus trade-in at the beginning, you'd have chosen differently.

That's what has happened with the fly discussion. I kept thinking that one more clarification, one more post, would do it. All that has happened is that things have become more and more confused.

I'm sorry I carried on so long and bored everyone to death with it and I apologize. Before I drop it, though, I want to respond one last time in the hopes of clearing up some misunderstandings.

>Bill Leach 950529.02:40

> I was almost floored by this posting. Rick specifically asked for a posting to provide an example of Nachtigall's work that could be PCT related.

> Bruce then posted the interesting but irrelevant (to PCT) example of how a fly lands on a ceiling.

> Rick (presumably wondering at this "PCT" example) blasts it rather handily and points out that this example provide little encouragement to read the work.

> If I were Rick, I would be more than just a little stunned at your posting on the matter.

>Bill Leach 950529.21:34

> I even admit that while I had the same impression as Rick, I too was wrong about Bruce's intentions with regard to the description. I do not believe however, that either Rick or I were too far amiss since even with the disclaimer, the description made no sense from the standpoint of both Bruce's assertions and Rick's requests.

After rereading Rick's request and my reply to it, I can understand the confusion, because it does sound like it was offered as the example Rick requested (however, only if you ignore the second paragraph). Here was the request:

Rick:

>> You claimed that Nachtigall "clearly recognizes that the structures mediating insect flight are organized as perceptual control systems" but then presented evidence (the description of how a fly lands on the ceiling) that he really doesn't. So if there is some "real, systematic, PCT data collection" in the insect research field, please point it out. But the "fly landing" stuff is clearly not it.

I perceived this as a reassertion on Rick's part of his claim that my description fly landing was intended as an example of Nachtigall's understanding of control systems as applied to insect flight. As I had already informed Rick in my prior post that this was not the case, I found it especially irritating that he was using this example, again, to make his case that Nachtigall's research was not relevant to PCT. By posting the direct quote I hoped to show that Nachtigall's description was only description, NOT interpretation (control system or otherwise) and thus in fact provided no evidence one way or the other, and said so explicitly in the paragraph that followed. I was absolutely _astonished_ when the post was again interpreted as an attempt by me to provide a PCT-relevant Nachtigall quote.

RE: hummingbirds

Me:

>> By the way, did you read Bill Powers' nifty little description of hummingbird behavior? On second thought, you wouldn't find it of interest--no application of the TEST and all that, just description from the point of view of the external observer, with a little speculation about possible controlled variables thrown in.

>Rick Marken (950529.2045)

> I did read it but I did see an application of the Test -- and suggestions on how to do more detailed applications of the Test using high speed photography.

>Bill Powers (950530.0945 MDT)

> I think that my informal presentation may have disguised the Test too well.

I knew I was in trouble for this one as soon as I had posted it; the point I was trying to make did not come across. As I'm sure you all remember, this whole affair got started because Rick decided that my description of how flies land on the ceiling looked too much like cause-effect for his taste. I had intended to provide a description much like Bill's of the hummingbird's behavior, but then decided it would take more time that it seemed to merit and settled for offering the pure behavioral description plus some mention of a few of the relevant sensory inputs. To be sure my intentions were not mistaken, I added a vague reference to the fact that the whole sequence was orchestrated by a set of perceptual control systems. This was not accepted.

After much fruitless debate, we end up with Bill's really excellent description of hummingbirds at the feeder, intended to get across to me the difference between a description consistent with a PCT approach and one that is not. The subtle point of my comment to Rick on this description is that it is not really so different from mine: mostly pure description of external behavior, with some suggestions being offered about the variables being controlled. If I had taken the time to develop my description of the variables being controlled by the fly as it lands (including how these variables are protected from disturbance), it wouldn't have sounded so different from Bill's.

I would have done it, too, if I had known what my decision to save on the effort was going to cost in the long run.

Regards, Bruce

Date: Tue, 30 May 1995 15:28:04 .
SUBJECT: Fly on the Wall

{from Joel Judd 950530.1530 CST}

Rick (950530):

My comment about coming to an understanding should have been followed by a "tongue-in-cheek" icon, but I'm getting tired of the winking one (;-)) so I didn't put anything. Joel

Date: Tue, 30 May 1995 20:32:53 -0400
Subject: Re: format; meeting;Test;misc

<[Bill Leach 950530.18:38]
>[From Bill Powers (950530.0945 MDT)]

> skipped lines

none this time

> Why run this like a mid-term exam? I've been waiting for someone to explain to Bruce why the Nachtigall description was not a good PCT description, but as you can see, I have finally provided ...

As far as this fellow is concerned, I see no reason why Bruce could not do a good job of that himself. I don't doubt that Bruce has a "handle" on fundamental PCT in spite of statements made occasionally that are in error from a PCT perspective.

People such as myself really are "junior PCTers" and are prone to easily miss our own failure to consider fundamental PCT concepts when discussing various topics. Even you and Rick "slip" now and then but with a great deal less frequency that the rest of us.

The number of times that I have "zeroed in" on a PCT issue in a discussion (usually almost proud as punch) only to then read a concise and even elegant posting (usually by you) that makes it obvious that I missed the _most_ important concept entirely, is close to depressing.

I am now absolutely convinced that a year or two of a serious attempt to study and understand PCT just _might_ be sufficient to begin to realize just how much one does not understand.

> I can only conclude that there is some other goal involved, like showing who knows most about PCT. What a bore.

Maybe, but Bruce's latest comments upon what might constitute PCT research was not only interesting but also a key to what might have been a part of the problem.

I know that Bruce, like myself, often states things in a manner that is not consistent with PCT. Some of this might be due to misunderstanding but can equally be caused just by not analyzing one's own thought adequately before expressing them.

For me, I find at least now, that I often have to rewrite a paragraph that is supposed to present an idea consistent with PCT several times. Even then, quite often when I read my own message sent out by the listserver, I discover that I still feel that the presentation was poorly done.

I do need to go back over some of my earlier postings though from Dag's disks as I am quite certain that reading those will make me feel much better about the current state of affairs.

-bill

Date: Wed, 31 May 1995 12:19:13 -0500
Subject: Re: What Is PCT Research?

[From Bruce Abbott (950531.1215)]

>Rick Marken (950530.1300) --

>>Bruce Abbott (950530.1140 EST)

>> Now it seems to me that a study of the gain of the proprioceptive feedback system and its variation under varying experimental conditions is a PCT study. Is it? If not, please explain.

> This is a good question. I would say that such a study is a PCT study as long as one had already identified the significant components of the control loop - - in particular the controlled variable, perceptual signal, reference signal, error signal and output variable.

Bill Powers, do you agree?

Regards, Bruce

Date: Wed, 31 May 1995 12:36:17 -0700
Subject: PCT/Non-PCT Descriptions of Behavior

[From Rick Marken (950531.1230)]

Bruce Abbott (950530.1445 EST)]

> After much fruitless debate, we end up with Bill's really excellent description of hummingbirds at the feeder, intended to get across to me the difference between a description consistent with a PCT approach and one that is not. The subtle point of my comment to Rick on this description is that it is not really so different from mine: mostly pure description of external behavior, with some suggestions being offered about the variables being controlled.

The PCT description is, indeed, not very different from yours. But I think the difference is significant. It is not just the suggestions about the variables being controlled that distinguishes a PCT from a non-PCT description of behavior. What I think is most important about the PCT description is that it points to the varying environmental circumstances (disturbances) that should lead to variation in some result of action -- but doesn't. Bill's description of hummingbird behavior included a description of the changing circumstances that exist as the birds feed. For example:

> the whole feeder is usually swinging gently in the wind and twisting slightly around the axis of the supporting string.

What was not explicitly stated was the fact that these changing circumstances should lead to changes in at least one result of the birds' actions -- the distance between the bird and the feeder. One has to know that this is true -- that a constant distance between a floating object (bird) and a swinging object (feeder) is not expected. But it happens:

> a constant distance is maintained between the beak and the feeding port

The description also suggests how a constant distance is maintained, viz. the bird varies its actions appropriately so that this result (constant distance) occurs rather than some other result. The variation in action that produces a constant distance is described as follows:

> The head and body of the bird swing back and forth in space, and the orientation of the bird moves in an arc around the feeder as it twists

Bill never used the term "controlled variable" in his description of hummingbird behavior. Nor did he explicitly suggest that "distance from the feeder" was a controlled variable. What he did was describe the _fact_ that a constant distance is maintained despite the changing position of the feeder (and the changing velocity of the wind). His description suggests that distance is under control; that a constant result is maintained (controlled) because it is being protected from the effects of changing circumstances (disturbance) by the actions of the organism.

In the Naftigall description of fly landing there was no description of constant results being maintained under changing circumstances; there were just descriptions of different results of fly actions: approaching the ceiling at a steep angle at 25 cm/sec, stretching out legs, legs absorbing shock, legs adhering at point of contact, fly clinging to the ceiling, etc. All these are, indeed, results of the fly's actions, just as the distance from the feeder is a result of the bird's actions. But in the Naftigall description there is no indication that the same result is produced under changing circumstances. For example, there is no indication that the angle of approach or speed of approach is the same despite changes in the angle of the ceiling, the velocity of the wind, etc. Including this small detail is the difference between a useless non-PCT description into a useful PCT description of behavior.

> If I had taken the time to develop my description of the variables being controlled by the fly as it lands (including how these variables are protected from disturbance), it wouldn't have sounded so different from Bill's.

This is true as long as you do, indeed, include a description of the disturbances to the proposed controlled variable and (if possible) a description of the actions that protect the variable from disturbance. What is most important in PCT description of behavior is to show that some result of an organisms actions should, but doesn't, vary. The description must include observations of a varying disturbance (like the changing position of the feeder) and of a result of actions (like distance from feeder) that is protected from the effects of this disturbance .

Naftigall's description of fly landing cannot be made into a PCT description just by suggesting that some of the results produced by the fly are controlled variables. For example, it would not have helped if Naftigall had suggested that angle of approach, speed, degree of leg stretch, etc. were controlled variables. This would be mere speculation; no more convincing than saying that any other result of the fly's action is a controlled variable. For example, one result of a fly landing on my ceiling is me swatting it. There is no reason to rule out this result of the fly's actions (getting swatted) as a possible controlled variable based on Naftigall's type of description of fly landing.

A PCT description of behavior, by including descriptions of constant results and the varying circumstances under which they are produced, provides a reasonable basis for guessing that a variable is under control. The actual variable under control is, of course, some perceptual measure, from the organism's perspective, of the result that remains constant despite changing circumstances. The PCT description of hummingbird behavior suggests that the hummingbird is controlling (among other things) a perceptual representation of its distance from the feeder. We can guess that this distance is a controlled variable because the PCT description of behavior included what was needed to make this a reasonable guess -- it included a description of the fact that this distance remained stable under circumstances in which it would not have been expected to do so.

The PCT description of behavior is a description of control.

Best Rick

Date: Wed, 31 May 1995 17:07:23 -0600
Subject: PCT experiments

[From Bill Powers (950531.1540 MDT)]

Rick says:
>> I would say that such a study is a PCT study as long as one had already identified the significant components of the control loop - - in particular the controlled variable, perceptual signal, reference signal, error signal and output variable.

Bruce Abbott:
> Bill Powers, do you agree?

Mostly. Even to talk about such things as the "gain" of a control system in a preying mantis, one must have identified a control system at least in overall terms. It isn't necessary to break down the system into particular components, but in the study in question, at least the output function was identified and tested.

Experiments with perception are useful because they give us hints about the kinds of controlled variables there may be and how they relate to the world we describe in terms of physical measurements. They're not exactly "PCT" experiments because they don't deal with control. But they could become part of PCT experiments and sometimes are. In all experiments where perceptions are tested by the method of adjustment (where the subject adjusts a perception to some criterion appearance), a control process is involved whether the experimenter thinks of it that way or not.

Same with experiments to determine muscle properties. Not strictly PCT, but the results can be very useful in PCT models.

Incidentally, is the name (in the other references) "Naftigall" as Rick gives it, or "Nachtigall" [Nightingale]? I suspect it's "Nachtigall" because when I did a search for a book by an author spelled that way, I got a reference to a book about insect behavior.

Best to all, Bill P.

Date: Wed, 31 May 1995 21:38:00 -0700
Subject: PCT Research, Nachtigall

[From Rick Marken (950531.2130)]

Regarding the Kittmann study referred to by Bruce Abbott (950530.1140 EST)

I said:

> I would say that such a study is a PCT study as long as one had already identified the significant components of the control loop

Bruce asked:

> Bill Powers, do you agree?

and Bill (950531.1540 MDT) said:

> Mostly. Even to talk about such things as the "gain" of a control system in a preying mantis, one must have identified a control system at least in overall terms. It isn't necessary to break down the system into particular components, but in the study in question, at least the output function was identified and tested.

I should add that I don't know whether or not the Kittman study is a "PCT study". If Kittman et al didn't identify the variable controlled by the control system, for example, (an essential part of _identifying_ the control system) then their measures of gain may be meaningless. All I meant to say is that the study of the gain of a control loop is certainly a legitimate topic for PCT research.

Bill Powers (950531.1540 MDT) --

> Incidentally, is the name (in the other references) "Naftigall" as Rick gives it, or "Nachtigall" [Nightingale]? I suspect it's "Nachtigall" because when I did a search for a book by an author spelled that way, I got a reference to a book about insect behavior.

It's Nachtigall. I probably started calling him Naftigall because I was unconsciously confusing him with Nosferatu, the German Dracula. I've been demonizing this poor guy just because his description of flies landing on ceilings is not a PCT description. This is particularly unjust because Nachtigall turns out to be a fine scientist. I just got "Insects in Flight" this afternoon and it is absolutely delightful. I can see why Bruce got upset at my judgment of the quote from the book. There is lots of really great stuff in this book, it is clearly written, beautifully illustrated and one of the translators is a fellow named Roger Abbott (any relation Bruce?).

I haven't read "every word" but it is easy to see that Nachtigall is a careful observer and an excellent researcher. It looks like he has done some very clever analyses of the anatomy, structure and aerodynamics of insects.

I have not been able to find any examples of PCT description or research -- but who cares? Nachtigall did his research in the 1960s, well before anyone but Bill Powers knew that behavior is the control of perception. Nachtigall did correctly describe a control system -- but control engineering had been around for 30 years already. It would have been remarkable if Nachtigall (like Avery's bee person) have stumbled onto the notion of systematically testing for results protected from disturbance.

Even though Nachtigall knew how control systems worked and could imagine the antennae as part of a velocity control system, his perspective on behavior was clearly a product of the behavioral zeitgeist that prevails to this day. Nachtigall can't be faulted for this because he could not possibly have profited from the study of PCT at the time he did his observations of behavior; BCP had just been published when "Insects in Flight" was published. Current psychologists, however, CAN be faulted for their failure to understand that behavior is the control of perception because they've already had 20 years to study PCT.

I'll keep looking for evidence of "PCT-like description" in Nachtigall's book. But I don't think the value of Nachtigall's work to PCT depends on whether or not he made PCT-like observations. Anyone who wants to study the controlling done by insects in flight would be well advised to read Nachtigall to find out something about what insects can perceive and what they can do to influence those perceptions.

Best Rick

Date: Thu, 1 Jun 1995 10:41:55 -0500
Subject: Nachtigall Book Review

[From Bruce Abbott (950601.1040 EST)]

>Rick Marken (950531.2130)

> It's Nachtigall. I probably started calling him Naftigall because I was unconsciously confusing him with Nosferatu, the German Dracula. I've been demonizing this poor guy just because his description of flies landing on ceilings is not a PCT description. This is particularly unjust because Nachtigall turns out to be a fine scientist. I just got "Insects in Flight" this afternoon and it is absolutely delightful. I can see why Bruce got upset at my judgment of the quote from the book. There is lots of really great stuff in this book, it is clearly written, beautifully illustrated and one of the translators is a fellow named Roger Abbott (any relation Bruce?).

I hope so. (:->

> I haven't read "every word" but it is easy to see that Nachtigall is a careful observer and an excellent researcher. It looks like he has done some very clever analyses of the anatomy, structure and aerodynamics of insects.

> I have not been able to find any examples of PCT description or research -- but who cares? Nachtigall did his research in the 1960s, well before anyone but Bill Powers knew that behavior is the control of perception. Nachtigall did correctly describe a control system -- but control engineering had been around for 30 years already. It would have been remarkable if Nachtigall (like Avery's bee person) have stumbled onto the notion of systematically testing for results protected from disturbance.

Thanks, Rick; I feel at least somewhat vindicated now. My initial comment about Nachtigall's work was:

>> The fascinating bee-as-control-system research by Mandyam Srinivasan extends work done much earlier from a similar perspective. Werner Nachtigall (1968) describes some aspects of bee flight as follows: . . .

I then offered an extended quote of Nachtigall's description of the bee's velocity control system. Later, I used this quote as evidence that Nachtigall understood the elements of closed-loop control, to support my thesis that insect researchers have been aware of the concept for a long time; in this context Srinivasan's work can bee seen as an intellectual descendent of Nachtigall's work.

What we have been arguing about, apparently, is whether to classify what Nachtigall did as "PCT research." I thought that Nachtigall's description of flight velocity control was very much in the PCT mold, but it turns out that we have somewhat different criteria as to what is and is not PCT research.

Whether Nachtigall may "have stumbled onto the notion of systematically testing for results protected from disturbance" is difficult to judge from the book (he certainly hints at it when discussing how the velocity control system corrects for wind gusts, etc.); I think one would have to read his _scientific_ publications to answer that question, most of which appeared in German.

At any rate, thanks for the fine review of Nachtigall's book; I wish I could have done as well. Whether we wish to classify Nachtigall as a PCT-style researcher, a "proto-PCT" researcher, or something else, he has indeed written a fine introduction to the elements of insect flight. As you so eloquently said:

> Anyone who wants to study the controlling done by insects in flight would be well advised to read Nachtigall to find out something about what insects can perceive and what they can do to influence those perceptions.

Regards, Bruce

Date: Thu, 1 Jun 1995 11:09:35 -0500
Subject: Open Loop

[From Bruce Abbott (950601.1105 EST)]

I'd like to pick up on a thread that seems to have been left dangling:

>Rick Marken (950527.0930)

> We PCTers believe that PCT should not be required to account for most of the descriptions of behavior provided by conventional psychology, nor should it be required to account for most of the results of conventional psychological research.

> Why do you (Bruce or anyone else for that matter) think we believe this?

>>Bruce Abbott (950529.0935 EST)

>> To answer your question, my perception is that you have at least two reasons: 1) most conventional research did not collect data that can be used to construct an adequate control system model. A model developed specifically to explain the behavioral data in question, absent the relevant information about controlled variables, would be purely speculative, a mere fantasy, and 2) the better approach would be to start from scratch with studies designed to elucidate the controlled perceptual variables and from those data develop a model of the organism. Eventually such a model should "behave" appropriately (i.e., in agreement with observation, within experimental error) under whatever conditions you choose.

>> My problem with this is NOT that I advocate just fitting models to extant data (I don't, and thus agree with point 1), but rather, that I see no

reason why one should not explore the controlled variables that come into play in some of the situations already extensively studied using conventional approaches, then applying the resulting model to explain the known behavioral phenomena that occur in those situations.

> Rick Marken (950527.1600)

. . . . [no response]

This exchange began when I asked you to describe what you took to be my view on the matter. You complied, after which I made some corrective comments. You then asked me to describe my view of your position, and I complied. As it is very difficult to correct error without feedback, I'd very much like to know whether or not my description of your position was accurate. If not, where is it in error?

Regards, Bruce

Date: Thu, 1 Jun 1995 11:05:25 -0700
Subject: Control theory and PCT, Open Loop

[From Rick Marken (950601.1100)]

Bruce Abbott (950601.1040 EST) --

> I used this quote as evidence that Nachtigall understood the elements of closed-loop control, to support my thesis that insect researchers have been aware of the concept for a long time;

> What we have been arguing about, apparently, is whether to classify what Nachtigall did as "PCT research." I thought that Nachtigall's description of flight velocity control was very much in the PCT mold, but it turns out that we have somewhat different criteria as to what is and is not PCT research.

I think one problem is that you believe that an understanding of control theory is equivalent to an understanding of PCT. In fact, understanding control theory and understanding PCT are two different things.

PCT is about the application of control theory to the behavior of living organisms. Control theory was around for years before Bill Powers developed PCT. And there have been many applications of control theory to behavior both before and after the development of PCT. What distinguishes PCT from other applications of control theory to behavior is that PCT is based on the idea that behavior IS control; the idea that organisms produce consistent results under variable circumstances. Non-PCT applications of control theory to behavior assume that behavior is a cause-effect process; control theory is used to explain how inputs cause outputs that produce control. PCT shows that this way of applying control theory to behavior is wrong.

So it's not control theory that distinguishes PCT as a theory of behavior; it is the recognition that behavior IS control that distinguishes PCT. Once you understand that behavior is control, the next step is to understand that the proper mapping of the control model onto behavior puts perception into the loop as the controlled variable.

So it is perfectly possible to understand control theory and not understand that behavior is the control of perception. Evidence of this is provided by nearly everything posted by Hans Blom;-) You can also find it in the psychological literature. Two of my favorites are:

(1) T. B. Sheridan and W. Ferrel(1974) "Man Machine Systems", MIT Press

(2) E. C. Poulton (1974) "Tracking Skill and Manual Control" Academic Press

Check out Figure 9.1 (p. 177) in (1) and Figure 1.1 (p. 5) in (2). In both cases the input to the human is an _error_. The human converts this error into output that reduces the error. So error is viewed as an objective phenomenon ("I know what's wrong when I see it!"). This turns the human into a transfer function, converting objective error into the outputs that reduce error. This, of course, puts the environment (the source of error) in control, just where psychologists always thought it was.

These books are "bristling" (Bill's felicitous expression) with engineering terminology and the differential equations of control theory. But they are definitely not PCT -- not even close. The reason these applications of control theory are not PCT- - the reason they incorrectly map control theory to behavior -- is because they are not based on an understanding of the nature of behavior as control (see the first paper in "Mind Readings" for more detail).

Me:

> We PCTers believe that PCT should not be required to account for most of the descriptions of behavior provided by conventional psychology,

> Why do you (Bruce or anyone else for that matter) think we believe this?

Bruce Abbott (950529.0935 EST) --

> you have at least two reasons: 1) most conventional research did not collect data that can be used to construct an adequate control system model...2) the better approach would be to start from scratch with studies designed to elucidate the controlled perceptual variables

> My problem with this is...I see no reason why one should not explore the controlled variables that come into play in some of the situations already extensively studied

Bruce Abbott (950601.1105 EST) notes that I gave no response to this post and says:

> I'd very much like to know whether or not my description of your position was accurate.

Sorry. The reason I gave no response is because there was no disturbance. Your description of why you think I believe "PCT should not be required to account for most of the descriptions of behavior provided by conventional psychology" was accurate. And what you call your "problem" with that belief was no problem for me. I too "see no reason why one should not explore the controlled variables that come into play in some of the situations already extensively studied". In principle, it's a great idea; in practice, it usually requires access to data that is not available (the kind of data Bill Powers managed to get from Verhave in order to determine the variables controlled by the rats in his shock avoidance study) or it requires just going out and getting data that was not collected.

There are cases where we have been able to guess at reasonable PCT models to account for conventionally obtained data. But my experience has been that there is just not enough known about the conditions under which the conventional data was collected to make it possible to build a PCT model to account for that data. This is true even when the data is fairly noiseless and involves single subjects. What we usually don't have are records of the state of the input variables that might have been under control. That's why it would be fruitless to develop PCT models to account for the data obtained in, say, "stimulus control" and "classical conditioning" experiments.

But there is certainly no law against trying to "explore the controlled variables that come into play in some of the situations already extensively studied." But what that means to me is going out and doing the studies necessary to find out what variables are controlled in the situation already extensively studies.

Best Rick

Date: Thu, 1 Jun 1995 16:40:23 -0500
Subject: PCT Research?

[From Bruce Abbott (950601.1635 EST)]

>Rick Marken (950601.1100)] --

>>Bruce Abbott (950601.1040 EST)

>> I'd very much like to know whether or not my description of your position was accurate.

> Sorry. The reason I gave no response is because there was no disturbance.

So, whenever you give no response, can I assume that you agree with what I've said? Thank's for the clarification: does this mean we'll have to find something else to argue about? (;->

>Bill Leach 950529.21:34

> OK, I suppose that we have all had about enough of this one and I'm not sure that all the feathers will be smoothed back down at this point regardless of what happens.

I don't know about anyone else's feathers, but mine seem to be in good condition. These debates can and sometimes do become rather heated, but that's to be expected when two people are trying to get points across that may challenge each other's conceptions/beliefs; it can be extremely frustrating at times, both when you feel that your views are being misperceived or misrepresented and when you feel that you just aren't getting what the other guy is trying to say. But at times it may be the only way to really break through and come to some understanding (although not always agreement). I don't take it personally, and I hope my statements are not taken personally by others. In my view, it's all just part of the game. It's hot in the kitchen, but that's where all the cooking gets done.

EARLY PCT RESEARCH?

Now that we know what real PCT research looks like, here's a bit of PCT research from the 1970s: you know, the kind that focused on identifying the controlled perceptual variable. Rats were exposed to a schedule of brief footshocks presented at random at an average rate of once per 120 seconds. In one condition these were each immediately preceded by a 5-second warning tone and the houselight illuminating the experimental chamber was on (signaled shock condition). In a second condition the shocks occurred on the same schedule but without warning and the houselight was off (unsignaled shock condition). After being exposed to each schedule alternately for several sessions, the rats were placed in the unsignaled condition. By pressing a lever in the chamber a rat could switch from the unsignaled to the signaled shock condition (indicated by onset of the houselight. Any shocks that occurred in the signaled condition were, as in training, immediately preceded by the signal. Further responses on the lever during the signaled condition had no programmed effect. One minute later the signaled condition terminated (the houselight extinguished) and the rat was automatically placed back into the unsignaled condition. By responding immediately on the lever at this time, the rat could immediately return to the signaled condition for another minute, and so on.

What happened was that the rats pressed the lever quickly and reliably enough to spend 85-95% of session time in the signaled shock condition.

I maintain that this experiment performed the Test for the controlled variable. So long as the signaled schedule remained in effect, the rat did nothing on the lever. However, as soon as the unsignaled schedule replaced the signaled one, the rat immediately approached the lever and pressed it, thus returning itself to the signaled shock condition. The controlled perception was the schedule in effect (as indicated by the state of the houselight), the reference was "signaled schedule in effect," and being automatically switched from the signaled to the unsignaled schedule constituted the disturbance. The experiment showed that the rat would defend against this disturbance by pressing the lever to cancel it. Variables (dependability of stimuli as predictors of shock and safety) were manipulated across blocks of sessions in an effort to identify which specific variables distinguishing the signaled and unsignaled schedules were being controlled. [Badia, P., Harsh, J., Coker, C. C. and Abbott, B. (1976). Choice and the dependability of stimuli that predict shock and safety. _Journal of the Experimental Analysis of Behavior_, _26_, 95-111.]

Apparently I was doing PCT research (testing for the controlled variable) as far back as 1973, when this study began. (;->

Regards, Bruce

Date: Thu, 1 Jun 1995 19:00:38 -0700
Subject: PCT Research?

[From Rick Marken (950601.1900)]

Bruce Abbott (950601.1635 EST) --

>EARLY PCT RESEARCH?

Yes, in the same way that all operant research is early PCT research. In operant research, the subject is able to influence (operate on) some variable that affects the subject himself. In the simple operant conditioning situation, the subject is able to influence food delivery (via bar pressing) that affects how much food the subject gets to eat. In your study the subject is able to influence shock signalling schedule that affects how much shock the subject can avoid.

A disturbance to food delivery in the simple operant conditioning situation is a change in schedule and we know that the subject will compensate for these changes (when the schedule disturbance is not extreme) by changing actions in the way required to maintain food delivery rate constant. The disturbance to shock signalling schedule in your experiment was a change in the shock signalling schedule and we see that the subject did compensate for this change by pressing the bar, restoring the shock signal schedule in which the shock was signalled.

> What happened was that the rats pressed the lever quickly and reliably enough to spend 85-95% of session time in the signaled shock condition.

This is a nice piece of data because it provides at least a rough measure of control. The presumed controlled variable was in a particular state 85-95% of the time; if the rat had done nothing the controlled variable would have been in that state only, what, 50% of the time? So the control loop is definitely keeping shock signalling schedule under control.

> I maintain that this experiment performed the Test for the controlled variable.

Yes. I agree, your experiment definitely involves the Test. I do think you could have spent more time nailing down the controlled variable, though. It seems like there were some other very plausible possibilities, given your description of the study. For example, the subject might have been controlling for having the light on, regardless of the shock signalling schedule. It would also have been nice if you had tried a number of different disturbances to determine that it was, indeed, the signalling schedule that was under control.

> Apparently I was doing PCT research (testing for the controlled variable) as far back as 1973, when this study began. (;-

You were, indeed. And now that you are familiar with PCT I bet you can think of far better ways to find out what the subject is controlling in this rather unpleasant (for the rat) situation.

Best Rick

Date: Fri, 2 Jun 1995 01:26:13 -0400
Subject: Re: Control theory and PCT, Open Loop

<[Bill Leach 950601.23:46]
>[Rick Marken (950601.1100)]

> I think one problem is that you believe that an understanding of control theory is equivalent to an understanding of PCT. In fact, understanding control theory and understanding PCT are two different things.

This IS true but not generally for the reason that you gave. You gave an example of the _error_ input from the environment. This is not valid as an example of Control Theory being applied to a living system. That was rather, an example of someone that does not understand Control Theory misapplying a theory.

Anyone that actually does understand even Engineering Control Theory well enough to understand what is controlled could not possibly put the comparator in the environment.

-bill

Date: Fri, 2 Jun 1995 01:27:00 -0400
Subject: Re: PCT Research?

<[Bill Leach 950602.00:56]
>[From Bruce Abbott (950601.1635 EST)]

I should be in bed (and you probably wish that I was) :-)

> feathers

No I don't think that mine are suffering too badly. I don't take any of this personally and fully agree that "coming to terms" is anything but an easy task. The really important thing here is that everyone honestly make an effort to understand what the other is saying.

Sometimes the "what I was trying to discuss was..." is indeed helpful even though "the other point" may be important too.

> Rats

Did the rats ever press the lever during the signaled condition? If so then what does that mean?

I personally have quite a bit of trouble with this sort of experiment in that I believe that the whole setup is contrived. The rats were essentially placed into a situation that would never exist in their "normal" environment.

Typical behavior for an animal experiencing a shock is, as far as I have observed, to leave the place where the shock occurred and if possible never return.

It seems to me that attempting to learn what the controlled variables are for "typically normal" behavior would be the initial goal of PCT research. Next would likely come studies of frequently observed "abnormal" behavior and then maybe studies of behavioral situations that the subject would not normally ever encounter.

I will accept "flak" for this from Bill P. and Rick on this but it seems to me that asserting that "signaled schedule in effect" is just assigning the observers understanding of the experimental apparatus to the rat.

In other words, I suppose that I am looking for something more basic though I admit that even learning that the rats could learn to do this is more than just a bit of a surprise to me.

You did not mention if there was a way for the rat to avoid the shock if the light signaled its impending arrival.

> Apparently I was doing PCT research (testing for the controlled variable) as far back as 1973, when this study began. (;-

Better minds than I might be able to tell me what is wrong with this but somehow I don't see this as testing for the controlled variable. My difficulty might of course be with trying to think in terms of testing methods for a logical variable.

I easily see (in principle) how one tests for a CEV that is controlled to a fixed reference or even one that varies (if some idea exists as to why it would be varied). In that case all one really has to do is attempt to change the (predicted) CEV and measure what the subject does (recognizing of course that the actual task is a bit more difficult since the subject may change the reference based upon other perceptual input - such as noticing that you are doing something).

BTW, the title sounds anything but PCT! :-)

-bill

Date: Fri, 2 Jun 1995 02:08:27 -0400
Subject: Re: PCT Research?

<[Bill Leach 950602.02:09]
>[From Rick Marken (950601.1900)]

Bruce Abbott (950601.1635 EST) --

> EARLY PCT RESEARCH?

> Yes. I agree, your experiment definitely involves the Test. ...

Ok, so I got my flak even before you saw my message! :-(

-bill

Date: Fri, 2 Jun 1995 02:37:06 -0400
Subject: still at it

<[Bill Leach 950602.02:30]
>[From Bruce Abbott (950601.1635 EST)]

If the pressing of the bar meant that no shock was received then an 85% result would be particularly suspect (I should think). OTOH, if the shock was received anyway then I not sure what it would mean.

> Variables (dependability of stimuli as predictors of shock and safety) were manipulated across blocks of sessions in an effort to identify which specific variables distinguishing the signaled and unsignaled schedules were being controlled.

I am clearly not understanding this statement (if you will forgive a complete ignorance of operant conditioning experimental procedures).

How do you manipulate "dependability of stimuli as predictors of shock and safety"? What are you actually changing and what results confirm or deny the desired state?

I should be in BED!! BYE

-bill

Date: Fri, 2 Jun 1995 07:46:14 -0600
Subject: Re: PCT research

[From Bill Powers (950602.0600 MDT)]

Bruce Abbott (950601.1635 EST)--

> here's a bit of PCT research from the 1970s: you know, the kind that focused on identifying the controlled perceptual variable. .... The experiment showed that the rat would defend against this disturbance by pressing the lever to cancel it. Variables (dependability of stimuli as predictors of shock and safety) were manipulated across blocks of sessions in an effort to identify which specific variables distinguishing the signaled and unsignaled schedules were being controlled.

How did this experiment determine that the rat didn't just prefer the house light to be on? In order to show that the rat preferred the signalled condition, you would have to do the same test with "house light off" indicating the signaled shock condition, and preferably repeat the experiment with other kinds of indicators, too. If the rat always turned on the signalled condition regardless of the kind of behavior required to do that, we might suspect that the rat was perceiving and controlling something related to the signaled condition that was different in the unsignaled condition.

Even after proving that the manipulandum was not itself the primary controlled variable, you would still have a job ahead in proving that the controlled variable was "signaled condition present." That is the human way of perceiving the situation, but my prejudice is that it is a rather abstract perception for a rat to have. I would ask myself, "From the rat's point of view, what is different between what I see as the signaled and the unsignaled conditions, in terms of the experiences the rat is having?" One difference, I would guess, is that in the signaled condition the rat might be able to reduce the experience of shock, or prepare itself in some way to receive the shock, whereas in the unsignaled condition the shock might arrive at any instant and catch the rat unprepared. Whatever it is that the rat is controlling, it has to be something that a rat can perceive. What the human observer can see about the situation is irrelevant.

Probably the best way to gain insight into the rat's experience is to undergo the experiment yourself. While you would be able to characterize the situation as "signaled" and "unsignaled", you could also look for lower-level perceptions that are different even without this characterization. You could ask, "What is better about having a signal indicating that a shock is about to occur?" You might, for example, find that at least some of the time you would be less surprised by the shock if a signal immediately preceded it, so preferring not to be surprised, you would try to find ways to make that signal appear. The ultimate controlled variable would be to keep the experience of the shock as untraumatic as possible, and to do this, the immediate variable that needs to be controlled is the presence or absence of the signal. That is less abstract than a perception of "signaled condition," and I would consider it a more likely prospect for a perceptual variable that the rat could control.

> I maintain that this experiment performed the Test for the controlled variable.

I agree that it did, but it didn't carry it very far. I believe the experimenters could have got a lot closer to what the rat was actually controlling.

> The controlled perception was the schedule in effect (as indicated by the state of the houselight), the reference was "signaled schedule in effect," and being automatically switched from the signaled to the unsignaled schedule constituted the disturbance.

I doubt that what the rat perceived was "the schedule in effect." That is what the _experimenters_ perceived. The experimenters stopped the test when they found a variable that _they_ could perceive as being under control, but that apparent control might merely have been a side-effect of the rat's controlling another, and much simpler, variable. I think that experimenters doing the test should try to distinguish between their own perceptions and those of the test subject, and realize that there can be a considerable difference. Whenever possible they should test perceptual variables of the lowest level they can, to minimize overinterpretation.

For example, in varying "dependability" of a signal, could this not also be interpreted simply as varying the number of times the signal occurred during a session? If the rats controlled for perceiving the signal, would this not make it appear that they are controlling for an abstract condition called "dependability" or "probability?"

Best to all, Bill P.

Date: Fri, 2 Jun 1995 08:49:50 -0700
Subject: Understanding Control Theory, PCT Research

[From Rick Marken (950602.0845)]

Bill Leach (950601.23:46) --

> Anyone that actually does understand even Engineering Control Theory well enough to understand what is controlled could not possibly put the comparator in the environment.

I agree completely. The fact is, however, that there are a number of psychologists who can do a convincing and (for most psychologists) intimidating job of presenting a mathematical analysis of control theory. These people are considered the "experts" in the application of control theory in psychology yet they get what seems to be the simplest aspect of control theory wrong -- the variable controlled by a control system.

So I guess the question is "what constitutes an understanding of control theory?". Apparently there are many aspects to "understanding control theory". One can understand the complex (literally) math while not understanding the basic functional characteristics of a control loop (like control of perception); this seems to characterize the understanding of many of the psychologists who are the experts in control theory. On the other hand, one can understand the basic functional characteristics of control systems while having only a passing familiarity with the complex math; this seems to characterize my own understanding of control theory.

I'm glad that Bill Powers (950602.0600 MDT) agrees with my basic evaluation of Bruce Abbott's (950601.1635 EST) "PCT research from the 1970s". Bruce said:

> I maintain that this experiment performed the Test for the controlled variable.

and Bill said:

> I agree that it did, but it didn't carry it very far.

In response to the same comment I had said:

> Yes. I agree, your experiment definitely involves the Test. I do think you could have spent more time nailing down the controlled variable, though.

I think the fact that this Test was not carried very far (more time was not spent nailing down the controlled variable) is crucial. I would guess that the reason this Test was not carried very far is because the experimenters did not see their goal as identifying a variable that the rat was controlling. It is not clear that the experimenters really performed the first (and most crucial) part of the Test: hypothesizing that a variable was under control. The variable "shock signalling schedule" was not treated as a _possible_ controlled variable (and, as Bill Leach (950602.00:56) points out, an extremely unlikely one since it "is just assigning the observer's understanding of the experimental apparatus to the rat"). It was probably treated as a variable that has a possible effect on behavior (bar pressing) and it did have such an effect. Thus, the experimenters never even considered the many plausible alternative variables that the rat might actually be controlling.

The goal of The Test differs completely from the goal of conventional research. The goal of conventional research is to determine what variables influence the observable behavior of the organisms; the goal of the Test is to see the world from the organisms perspective; to learn what aspects of the organism's own experience it is trying to bring under control.

So, while the research Bruce describes can be seen as having several elements of The Test for controlled variables (mainly, introducing what can be seen as a disturbance to a possible controlled variable) it really doesn't go nearly far enough to achieve the basic goal of the Test -- to determine "beyond a reasonable doubt" the perceptual variables an organism is trying to control.

Best Rick

Date: Fri, 2 Jun 1995 15:01:14 -0500
Subject: Re: Early PCT Research?

[From Bruce Abbott (950602.1455 EST)]

It would seem we agree that the Badia, Harsh, Culbertson, & Abbott (1976) study I described (950601.1635 EST) does indeed apply "the Test" for the controlled variable, but there were some issues raised about the research that I'd like to address here:

>Rick Marken (950601.1900) --

>> What happened was that the rats pressed the lever quickly and reliably enough to spend 85-95% of session time in the signaled shock condition.

> This is a nice piece of data because it provides at least a rough measure of control. The presumed controlled variable was in a particular state 85-95% of the time; if the rat had done nothing the controlled variable would have been in that state only, what, 50% of the time? So the control loop is definitely keeping shock signalling schedule under control.

Actually, if the rat had done nothing it would have been in the unsignaled condition 0% of the time during the testing sessions.

> I do think you could have spent more time nailing down the controlled variable, though. It seems like there were some other very plausible possibilities, given your description of the study.

This research was part of a larger program designed to identify the variable or variables actually being controlled (although these researchers would have SAID that they were attempting to identify the controlling variables, thus putting the cart before the horse). There are many variables that differ in value between the signaled and unsignaled conditions; any of these could be a (or the) controlled variable and they are all confounded in the simple test: state of the houselight, presence/absence of warning, opportunity to prepare for impending shock, identifiable periods of safety during which shock is guaranteed not to occur, and others.

> For example, the subject might have been controlling for having the light on, regardless of the shock signalling schedule.

This had been tested in an earlier study. Rats prefer the signaled schedule whether it is associated with light-on or light-off. Also, if you place the rat on the signaled schedule and lever-pressing produces the unsignaled schedule, the rats learn to avoid pressing the lever, thus remaining on the signaled schedule full time.

> It would also have been nice if you had tried a number of different disturbances to determine that it was, indeed, the signalling schedule that was under control.

I'm not sure what you mean by "different disturbances." We did manipulate a number of variables in an effort to identify which were controlled. For example, in the cited study the "dependability" of the signal as a predictor of shock [i.e., p(shock|signal] was manipulated, as was the "dependability" of signal absence as a predictor of safety [i.e., p(shock|no signal)]. The rats continued to control for the signaled schedule over a wide range of values for p(shock|signal) but not when p(shock|no signal) was degraded. False alarms seem to be less of a problem for the rat than failures to warn.

>Bill Leach 950602.00:56 --

> Did the rats ever press the lever during the signaled condition? If so then what does that mean?

Good question. Yes, they sometimes did, but they will do this even before they have learned the contingency between lever pressing and schedule condition (i.e., during initial training, when lever pressing has no effect on the schedules). The levers are placed in the chamber in a position such that they will occasionally get pressed as a byproduct of exploratory activity--otherwise they would never learn what effect lever-pressing has. This same activity during the signaled schedule can generate lever presses; these were recorded but had no programmed consequences. We assume that these presses are just byproducts of exploration, as opposed to the purposive lever pressing that is observed immediately after the schedule switches to unsignalled. At any rate, they occurred infrequently.

> I personally have quite a bit of trouble with this sort of experiment in that I believe that the whole setup is contrived. The rats were essentially placed into a situation that would never exist in their "normal" environment.

This is a familiar concern; you are worried about what experimental psychologists call the "external validity" of the research: how well the findings can be generalized from the laboratory to the "real" world. An excellent discussion of this issue can be found in

Mook, D. G. (1983). In defense of external invalidity. _American Psychologist_, April, 379-387.

The rat is still the same system and will behave according to the rules imposed by its structure and organization, whether the situation is "natural" or "artificial." When you wish to analyze the performance of a designed control system, you no doubt subject the system to artificial conditions (e.g., step functions, open-loop conditions) that the controller will never encounter in the real world, because the response of the system under such conditions reveals aspects of its performance that are difficult or impossible to see under normal operating conditions. A rat in an operant chamber may not be facing the rich complexity of conditions it usually encounters in the wild, but its behavior will still reflect the basic organization of its system, and the test conditions may reveal aspects of its performance that would difficult or impossible to observe under those more natural conditions. Furthermore, the artificial conditions may be a good enough analog to give results that would generalize to the wider world. How rats respond to signals warning of impending shock may, for example, help us to understand how real-world conditions may contribute to such human conditions as chronic anxiety or the development of peptic ulcers. Once a phenomenon has been identified in the rat studies, one can then do followup studies to determine the generality of the findings empirically.

Artificial conditions may also be used to test specific hypotheses or predictions of theory. The cited study was conducted to compare predictions of the "preparatory response" and "safety" hypotheses.

> Typical behavior for an animal experiencing a shock is, as far as I have observed, to leave the place where the shock occurred and if possible never return.

Yes, but what of the wild rat confronted with the dangerous job of foraging for food? Stay home, stay safe, die of starvation; forage, eat, but possibly get attacked by cats and other predators. Given the choice, would the rat visit locations in which the arrival of predators is always signaled (e.g., by the sounds of rustling leaves or by visual input as the predator crosses open ground) or those in which it could be easily surprised? The artificial situation provided in the operant chamber may provide a reasonable analog.

> It seems to me that attempting to learn what the controlled variables are for "typically normal" behavior would be the initial goal of PCT research. Next would likely come studies of frequently observed "abnormal" behavior and then maybe studies of behavioral situations that the subject would not normally ever encounter.

I agree, but the research program one follows is dictated by the general questions being addressed, which may lead one to focus on particular methods that would seem to provide the best avenue to a clear answer. This particular research was initially spurred by the realization that the rat's choice of the signaled condition was not consistent with extant theory.

> You did not mention if there was a way for the rat to avoid the shock if the light signaled its impending arrival.

Ideally, no. In practice, some rats did stumble on a solution: they learned to roll onto their backs, where the insulation provided by their fir prevented them from receiving further shocks. Rats are excellent little autonomous control systems: with a bit of luck, they will learn how to control variables the experimenter intends to be uncontrollable. Of course, experimenters are good controllers too. If necessary, we shaved the rat's back to remove this source of shock control.

>Bill Powers (950602.0600 MDT) --

> How did this experiment determine that the rat didn't just prefer the house light to be on? In order to show that the rat preferred the signalled condition, you would have to do the same test with "house light off" indicating the signaled shock condition, and preferably repeat the experiment with other kinds of indicators, too. If the rat always turned on the signalled condition regardless of the kind of behavior required to do that, we might suspect that the rat was perceiving and controlling something related to the signaled condition that was different in the unsignaled condition.

I've answered the first question, but I'll add that rats, being nocturnal animals, prefer the houselight off, not on. The suggested test was done. Also, rats choose the signaled schedule whether they have to press a lever, not press a lever, run to the opposite side of a shuttlebox, or stay where they are in order to be in the signaled condition. If pressing the lever does not switch conditions from unsignaled to signaled, the rats stop pressing (i.e., lever pressing extinguishes) even if that action continues to switch the houselight on.

> Even after proving that the manipulandum was not itself the primary controlled variable, you would still have a job ahead in proving that the controlled variable was "signaled condition present." That is the human way of perceiving the situation, but my prejudice is that it is a rather abstract perception for a rat to have. I would ask myself, "From the rat's point of view, what is different between what I see as the signaled and the unsignaled conditions, in terms of the experiences the rat is having?" One difference, I would guess, is that in the signaled condition the rat might be able to reduce the experience of shock, or prepare itself in some way to receive the shock, whereas in the unsignaled condition the shock might arrive at any instant and catch the rat unprepared. Whatever it is that the rat is controlling, it has to be something that a rat can perceive. What the human observer can see about the situation is irrelevant.

A long series of studies was conducted to try to get at that question. For example, in one study the intensity of shock in the signaled condition was systematically varied while holding the intensity of shock in the signaled condition constant. The rats stopped changing to the signaled condition only when the shocks were about three times intense as those in the unsignaled condition. This did not rule out the possibility that the rats were getting prepared for the shock (reducing the aversiveness of the shock experience), but did indicate that if preparation is involved, it is extremely effective. I was able to demonstrate in another study that such preparation, if it occurs, could not involve actions (such as postural adjustments) that minimized shock contact.

I think it would be fair to characterize the entire research program as a search for the actual controlled variable(s) in this rather complex situation, from the rat's point of view and not our own.

> Probably the best way to gain insight into the rat's experience is to undergo the experiment yourself. While you would be able to characterize the situation as "signaled" and "unsignaled", you could also look for lower-level perceptions that are different even without this characterization. You could ask, "What is better about having a signal indicating that a shock is about to occur?" You might, for example, find that at least some of the time you would be less surprised by the shock if a signal immediately preceded it, so preferring not to be surprised, you would try to find ways to make that signal appear. The ultimate controlled variable would be to keep the experience of the shock as untraumatic as possible, and to do this, the immediate variable that needs to be controlled is the presence or absence of the signal. That is less abstract than a perception of "signaled condition," and I would consider it a more likely prospect for a perceptual variable that the rat could control.

Some work WAS done using volunteer college student participants, and they were asked to explain why they thought they behaved as they did. Human participants sometimes behaved differently from the rats owing to their having a deeper control hierarchy than the rats. One student stayed in the unsignaled condition and then explained that "I thought it was worse but I wanted to prove to you that I could take it."

As to the ultimate controlled variable, I have no doubt that the rats are attempting to minimize the aversiveness of the shock experience as you suggest (rats do not appear to develop motives like those of the college student just quoted). This, I think, was taken for granted. The question of interest was _why_ the signaled condition seems less traumatic to the rat than the unsignaled one. After all, the _shock_ schedule in the two conditions is identical. How do the warning signals change the rat's perceptions? How is it that for the mere presence of a signal the rat is willing to take shocks up to three times more intense?

>> I maintain that this experiment performed the Test for the controlled variable.

> I agree that it did, but it didn't carry it very far. I believe the experimenters could have got a lot closer to what the rat was actually controlling.

I hope I've corrected that impression.

> I doubt that what the rat perceived was "the schedule in effect." That is what the _experimenters_ perceived. The experimenters stopped the test when they found a variable that _they_ could perceive as being under control, but that apparent control might merely have been a side-effect of the rat's controlling another, and much simpler, variable.

You should appreciate now that this assessment was premature.

> I think that experimenters doing the test should try to distinguish between their own perceptions and those of the test subject, and realize that there can be a considerable difference. Whenever possible they should test perceptual variables of the lowest level they can, to minimize overinterpretation.

This was done.

> For example, in varying "dependability" of a signal, could this not also be interpreted simply as varying the number of times the signal occurred during a session? If the rats controlled for perceiving the signal, would

> this not make it appear that they are controlling for an abstract condition called "dependability" or "probability?"

Every manipulation introduces confounding with other variables; the only choice one has is which variables to confound. In this case you can manipulate signal dependability by deleting shocks or by adding signals. The first way confounds dependability and shock frequency; the second confounds dependability and signal frequency. Both ways were investigated, and they yielded the same result.

On a somewhat different topic, my master's thesis examined whether rats preferred to have physical control over shock (being able to prevent it or to terminate it once it occurred) in the same way that preference for signaled over unsignaled shock schedules was investigated. Yes, this is control in the PCT sense of the term (shock is a perception). The answer was that they were not interested in changing from a condition in which they lacked control over shock to one in which they had such control, so long as the shocks in the two conditions were identical (i.e., same frequency, intensity, and duration). In fact they were indifferent to the two conditions. The procedure equated the actual shock experience in the two conditions; apparently just the perception of _having_ control is not itself a controlled variable, at least for the rat.

Regards, Bruce

Date: Fri, 2 Jun 1995 16:26:00 -0600
Subject: Re: Early PCT research

[From Bill Powers (950602.1505 MDT)]

Bruce Abbott (950602.1455 EST) --

Your expanded description of the series of experiments answers all the objections I raised. Considering the thorough way in which you tested for variables that the rat acted to control, the main question left in my mind is "what kept you from tumbling to PCT?" Was it simply the custom in that field of interpreting variables controlled by the rat's actions as variables that controlled the rat's actions? If that were all there were to it, I should think that the transition to PCT would be relatively easy for behaviorists. As you look back on your own frame of mind during those experiments, can you find any illumination on this subject?

> I think it would be fair to characterize the entire research program as a search for the actual controlled variable(s) in this rather complex situation, from the rat's point of view and not our own.

Yes, I concede that. Not "concede" -- agree willingly.

> Every manipulation introduces confounding with other variables; the only choice one has is which variables to confound. In this case you can manipulate signal dependability by deleting shocks or by adding signals. The first way confounds dependability and shock frequency; the second confounds dependability and signal frequency. Both ways were investigated, and they yielded the same result.

This isn't really a case of confounding, because "dependability" is an abstraction that is a direct function of shock frequency and signal frequency. If you vary the relationship between shock frequency and signal frequency, you necessarily vary dependability, or pr(shock|signal). You also vary the ratio of total shocks to total signals, the ratio of shock rate to signal rate, and the number of unsignalled shocks.

All of these interpretations of the actual events involve perceptions of various levels. The question is, which of them, if any, is reasonable to attribute to a rat? The simplest interpretation, I would think, since the rats showed a strong preference for signaled versus unsignaled shocks, would be that the rats simply acted against the occurrence of unsignaled shocks. Even that is probably too abstract; what they acted to control toward zero, I would guess, was whatever it was about the experience that was worsened when the shocks were not preceded by a signal. If this were the controlled variable (number of shocks not preceded by a signal, with a reference level of zero), would it not explain all that you observed?

> On a somewhat different topic, my master's thesis examined whether rats preferred to have physical control over shock (being able to prevent it or to terminate it once it occurred) in the same way that preference for signaled over unsignaled shock schedules was investigated. Yes, this is control in the PCT sense of the term (shock is a perception). The answer was that they were not interested in changing from a condition in which they lacked control over shock to one in which they had such control, so long as the shocks in the two conditions were identical (i.e., same frequency, intensity, and duration). In fact they were indifferent to the two conditions.

This is another (well-conceived) example of this issue of overinterpretation. The idea of having control is just that, an idea. It is an abstraction formed over many experiences of many kinds, a highly cognitive sort of thing. To a human being, it is perfectly natural to think in terms of having or not having control, regardless of what is being controlled. What you showed is reason to doubt that a rat perceives the situation in the same way -- that the rat says "Oops, I've lost control," or "Goody, now I have control again." Having control was not a goal for its own sake; only actually controlling the effects of shocks mattered, and as you describe the experiment, that control was minimal because the shock would already have started. It's hard to guess what would happen to the rat's control systems in the time immediately following the onset of a shock.

I trust you misspoke when you included "being able to prevent it" among the conditions to which the rats were indifferent. If a rat could press a bar before the shock occurred and thus prevent it, it would learn quickly to do so -- my rat paper was a study of just such a situation, an experiment done by Verhave, a "Sidman avoidance schedule." But that would not raise the issue of a preference for an abstract capacity to control; it merely would verify that rats will control shocks if they can. It wouldn't tell us what they think about being able to do so, if anything.

I seem to recall other studies in which rats given a choice between producing food by pressing a bar or eating freely-available food will spend a respectable amount of time pressing the bar. Now that might suggest a preference for being in control, but of course that idea would have to be tested as thoroughly as you tested for controlled variables in your experiment.

The issue that I am calling overinterpretation has shown up before on the net, in the form of using too high a level of perception to interpret what is really a low-level process. When you take the point of view of a high-level system, all processes seem to partake of the kind of perception typical of that level. If you're working mainly from the category level, you can see all lower-level processes as involving categories. Even a sensation-signal looks like a category of intensities. An event is a category of transitions and configurations.

But explaining the operation of lower systems in terms of categories is simply a mistake of "overinterpretation." It is forcing a high-level point of view onto a process that uses a simpler mode of perception.

Recognizing this problem can make a profound change in the way we see the behavior of animals and even other people. When I commented some months ago that behaviorists seemed to take too cognitive an approach to animal behavior, this is what I was trying to talk about. It is all too easy to impose one's OWN cognitions on the world of perception, to see goals and purposes of a kind that do not actually exist except in our own perceptions. Even when we're looking at other people, what WE see them doing is not necessarily -- or perhaps even not often -- what THEY see themselves doing. This is called "anthropomorphizing," which is not necessarily a sin, but which needs to be considered carefully especially when the entity about which we anthropomorphize is not human.

Way, way back, Korzybski admonished us that we should develop a "consciousness of abstracting," an awareness that our generalizations are subjective and impose structure on the world that is not necessarily there. The map, he said, is an abstraction; it is not the territory where things actually exist and happen. In

PCT, we say this a little differently: "It's all perception." One of the primary insights of PCT, which Clark McPhail wrote about eloquently in a paper on Dag's PCT Text disc in the file "Epiphany.s", is that ALL that we can experience (including our actions) consists of our own perceptions, not of a direct connection to the objective world. To internalize this idea is to experience a scientific awakening, one that arouses us from the dream of objectivity.

The Test is a direct expression of that basic realization. No longer can we look at someone else's behavior and simply assume that what we see going on is what matters in the ecology of the other. We must try to guess at what the other organism is perceiving and in what states it wants its perceptions to be. Nothing is obvious any more; every guess must be tested. We may see that a rat prefers certainty to uncertainty -- but those are our own perceptions, and we may find that there are far simpler interpretations that will work just as well, or better, for the rat.

Is this, perhaps, the key to the transition from your kind of behaviorism to PCT? When you realize that you control only your own perceptions, it is hardly any step at all to seeing that this must be true of all other organisms as well. You don't need to interact with other people or with animals very long, from this new viewpoint, before realizing that you have no direct way to know what they are perceiving, and thus what they are controlling, and thus what they are doing.

Best, Bill P.

Date: Fri, 2 Jun 1995 23:05:19 -0400
Subject: Re: Early PCT Research?

<[Bill Leach 950602.18:33]
>[From Bruce Abbott (950602.1455 EST)]

> Actually, if the rat had done nothing it would have been in the unsignaled condition 0% of the t me during the testing sessions.

Huh?

I want to express some additional critical comments on this experiment but I believe that the experiment was carried out with even greater care than I can infer directly from your last posting (which is itself considerable). The comments will obviously be my opinion and of course may in part be due to additional aspects of the actual project with which I am not familiar.

> ... or variables actually being controlled (although these researchers would have SAID that they were attempting to identify the controlling variables, thus putting the cart before the horse).

I really do believe that this is a matter of a GREAT deal more than that the researchers "would have SAID...".

The real subject view vs. experimenter view is the most critical aspect of designing and analyzing experiments. The problem with "controlling variables" is NOT that the wrong CEV is found it is deeper than that. Indeed, the researcher may well HAVE identified the correct CEV (even if he thinks it is a controlling variable).

The first problem is that the approach of a research that is looking for something in the environment that controls observed (by the researcher) behavior of the subject is that what the researcher sees is incidental to what the subject is doing.

The second problem is that the researcher "projects" onto the subject his own perception as though the subject also perceives the same thing (as indeed you did with the "control of schedule" statement).

The third problem is that the focus really IS on the environment an not on the subject. Thus, it is easier to miss a basic CEV because the observer's perception is possibly at a higher or even lower level than the subject's.

As I see it, the recognition that one can not control the behavior of a control system and thus must always focus on the subjects perceptions is the single most important view point for a PCT researcher. And I mean by "can not control" that while it is often possible to create a situation where YOU will observe behavior on the part of the subject that appears to be "controlled" in the manner that you wish to perceive, the actual reasons why what you observe happening may have little to do with what you think "is causing" what you see.

> Good question. Yes, they sometimes did, but they will do this ...

Was there any data collected then to distinguish between what appeared to be intentional actuations vs. accidental?

Was their any noted change in that data as a function of different experimental configurations?

The assumption may well be incorrect though if they occurred infrequently then it might well not matter.

> The rat is still the same system and will behave according to the rules imposed by its structure and organization, whether the situation is "natural" or "artificial." When you wish to analyze the performance of a designed control system, you no doubt subject the system to artificial conditions (e.g., step functions, open-loop conditions) that the controller will never encounter in the real world, because the response of the system under such conditions reveals aspects of its performance that are difficult or impossible to see under normal operating conditions. A rat in an operant ...

I will give you that this is fundamentally correct as long as one recognizes that the experimenter IS introducing disturbances to CEVs with which the experimenter is completely unaware and that the errors in the perceptual systems associated with these disturbances. And also recognizes that "extreme" disturbance or disturbance that produces an error in a controlled perception such that the error can not be "controlled away" (or more accurately, reduced to acceptable (to the subject of course, control limits) will result in reorganization or other control actions (reference changes) that will result in changes that are unique to the specific subject.

Additionally, many of the variables that are manipulated (intentionally or otherwise) may be changing perceptions for several CEVs. I believe that all of these sorts of things are at least marginally easier to identify only if the researcher continually reminds himself that it is not his perceptions that count.

> Once a phenomenon has been identified in the rat studies, one can then do followup studies to determine the generality of the findings empirically.

I can imagine for various reasons that this step would literally be forced without consulting the researchers (or ignoring their recommendations if such represent a disturbance to whoever wants to "take the next step")... Unfortunately linking studies together that were made with as fundamental an error as thinking that a Stimulus "triggers" behavior has a very low probability of producing better results or more useful knowledge.

> Artificial conditions may also be used to test specific hypotheses or predictions of theory. The cited study was conducted to compare predictions of the "preparatory response" and "safety" hypotheses.

Naturally the problem here is that if the underlying philosophy for the hypothesis is wrong then an experiment "designed" to allow only two options is going to produce results subject to misinterpretation. Possibly the MOST significant observation made was the one that you mentioned about the rats that would roll over on their backs to apparently avoid the shock altogether!

> Yes, but what of the wild rat confronted with the dangerous job of foraging for food? Stay home, stay safe, die of starvation; forage, eat, but possibly get attacked by cats and other predators. Given the choice, would

the rat visit locations in which the arrival of predators is always signaled (e.g., by the sounds of rustling leaves or by visual input as the predator crosses open ground) or those in which it could be easily surprised? The artificial situation provided in the operant chamber may provide a reasonable analog.

I will provisionally accept this as long as you accept that the comparison is probably stretching things a bit. I will also give you the fact that attempting experiments with the experimenter having some control over the limits to the options for the subject is necessary in animal research as long as you will "give me" that when you begin establishing such limits you really do not know what perceptions you are disturbing for the subject, how you are disturbing them (from the subject's view point of course) and that the answers to these questions _could_ be crucial to the results obtained.

> conditions is identical. How do the warning signals change the rat's perceptions? How is it that for the mere presence of a signal the rat is willing to take shocks up to three times more intense?

Rats don't like surprises and will control to reduce the probability for having one?

Bill P.

>> under control, but that apparent control might merely have been a side-effect of the rat's controlling another, and much simpler, variable.

> You should appreciate now that this assessment was premature.

I don't think so. I do think that the experiment was well run to the point of obtaining what is probably a correct evaluation of what the rats were observed to be doing but I still don't think that a rat intentionally control such an abstract concept as "the schedule in effect". My earlier "off the cuff" comment about "control to reduce probability for having one?" is not much better.

What might be better would be if I said "That the rats have a reference level for surprises that is set rather low (subjective at this point) and the signaled event represents a smaller disturbance (or no disturbance) TO THIS perception than does the unsignalled event. The rats should reorganize if unsignalled events are experienced that represent a sufficient disturbance to the controlled perception to result in an error. When a means, sequence of controlled perceptions or program is discovered that brings this CEV under control then the rat will use the new found method (unless other error(s) exist in higher priority CEV(s) or reorganization "disassembles" the learned method.

Even still there is in addition always the question of CEVs of which the experimenter has no knowledge and their possible effect(s) on the results.

It is beginning to occur to me as to why Rick has always had (for the year and a half that I have known him and presumably for quite a bit longer than that) had such a strong aversion to attempting to review past work using a PCT analysis.

Though it does sound to me as though in your experiments there seemed to be a high level of recognition that the rats themselves seemed to "have something to say" about what went on even to the point of thinking about and testing for "what the rat might be doing", the initial design, conducting and analysis for an experiment without a rigorous PCT methodology will fall short at every point largely because many of the "right questions" will not be asked before and during design.

I suspect that at some point, experiments of the sort that you described would ultimately have to result in an understanding of control theory as applied to living systems. The question would be how long?

-bill

Date: Fri, 2 Jun 1995 23:05:56 -0400
Subject: Re: Understanding Control Theory, PCT Research

<[Bill Leach 950602.21:14]
>[From Rick Marken (950602.0845)]

> ... systems while having only a passing familiarity with the complex math; this seems to characterize my own understanding of control theory.

This is also one of the reasons for some of the "heated" discussions between you and Martin. Martin has several times presented a point that was absolutely a true statement in control theory but was irrelevant to the type and design specifics of the control systems that we study.

I almost feel that this is similarly true for a portion of Hans presentation concerning "modeled control" system application to living systems. Hans has several times discussed the outstanding ability of model based control to reduce the effects of white (or pink) noise, achieve significantly higher accuracy for certain specialized control situations.

It seems to me that an argument for "Fuzzy Logic" would stand a much better chance of being won! Living systems explicitly do not exhibit outstanding control accuracy when compared to all but the shoddiest of engineered control systems.

Thus hyping a system because of its "optimal" control capability seems to me to be presenting a case for why such a system might not exist.

Some of the "adaptive" ideas could be valid in the sense that we do have experimental evidence that some sort of tuning must occur.

We have pretty good evidence that perceptions that are directly perceived "can be controlled" IF they are rather high level perceptions. The evidence for controlling low level perceptions upon a loss of input is completely contrary to model operations -- that is at low levels it seems pretty conclusive that loss of perception means loss of control. When low level control is maintained anyway then the evidence is pretty strong that what happens is that other perceptions "replace" the desired one and "control" of the original perception is an artifact of the physics of the situation.

Rather than go back and edit the message I just wrote in reply to Bruce's message I add this thought here:

Operant conditioning might be at least headed in the right direction. The major difficulty in any behavioral research with animals especially is that perceptions of what is observed by the researcher never have any significance to the subject. What is observed really is absolutely incidental to the subject.

Researchers doing experiments such as the one that Bruce described are faced with performing the sort of analysis that we often carry on here when discussing PCT questions with respect to unmodeled hypothesis. The levels of the hierarchy that are involved, the multitude of unknown CEVs make objective testable hypothesis concerning the postulated CEV very difficult to formulate.

It may well be that for the foreseeable future that the only really good PCT data will come from 1) Continued PCT modeling of the nature currently used 2) Relatively "simple" experimental observations such as the insect studies and 3) Complex studies with human subjects under PCT conditions.

The higher animals being without communications methods may well be too complex to study until more knowledge is obtained about the nature of HPCT.

-bill

Date: Fri, 2 Jun 1995 23:06:58 -0400
Subject: Re: PCT research

<[Bill Leach 950602.22:41]
>[From Bill Powers (950602.0600 MDT)]

>In engineering parlance, the "output" is the controlled variable. ...

Yes, I know but for whatever reasons, I have personally always tried to distinguished between the two.

To me when thinking of a power supply with voltage regulator in terms of control theory calling the actual supply voltage the "output" is "OK". That is, doing so doesn't generally get you into any trouble.

The same attitude exists for the "output" of a communications radio transmitter.

When the output causes an action in the environment that involves what to my mind is a significant transform then I view the output as a really separate and distinct "thing" from what is being controlled. I suppose that my choice of distinctions is probably based upon personal experience with the various control systems that I have worked with and the fact that most of the time there were no "experts" available to call upon.

As an example, when I was called upon to assist with a rudder control problem on a submarine that I served aboard (while at sea), the idea that what is controlled is supposed to be actual physical position of the rudder is absolutely essential to any troubleshooting effort. If one fails to appreciate that the control system will be trying only to control ITS' OWN perception of the controlled variable then one will immediately be in trouble and success in "fixing" the problem will be mostly a matter of luck (or using a term you may also be familiar with - easter-egging).

I would guess that in all the control system troubleshooting work that I have done that the actual problems were almost even split between "effector" problems and "perception not tracking CEV". Of course I would not have worded it at all that way in the past though I did often use something like "Your controller 'thought' everything was fine; input was messed up."

I'm disappointed -- sniff

You did not comment upon my "Bill will like this one" remark.

I did go to bed... finally and will try again tonight :-)

-bill

Date: Fri, 2 Jun 1995 23:53:10 -0400
Subject: Re: Early PCT research

<[Bill Leach 950602.23:18]
>[From Bill Powers (950602.1505 MDT)]

... And this message indicates to me that I was possibly being too critical of Bruce's posting.

Once again fearful of my wording raising the issue of "alerting phenomenon" but;

Many people express a clear desire to be "shocked" or "surprised" and I really mean "shocked" even by the term surprised. These are people that are almost "hooked" on horror films or other "thrill seeking" activities. They often describe their "adrenalin rush". They appear to control for being in situations where their system experiences a disturbance to a controlled variable of sufficient magnitude to cause a reference change for metabolic rate (or whatever is properly termed the result of a control system setting the reference for energy output sharply higher).

In trying to analyze my own attitudes about this; In general I do not like horror movies and I specifically do not like "fearful surprises". Anything that results in my feeling emotionally similar to say a situation where I am certain that I am not going to be able to avoid a vehicle collision is decidedly "unpleasant" to me.

This attitude of mine makes it rather easy for me to project the idea that a rat in Bruce's experiments, for example, could actually have a reference for a low level of surprises (shocks but emotional as opposed to electrical).

Without further data and assuming that the rats that were shaved then behaved similarly to the others I would also conclude that rats have a low level reference for the effects perceived from electrical shock also (naturally I can project both my own feelings concerning such experiences, many, as well as informal observation of the behavior of some other animals and humans with respect to electrical shock).

A serious problem in designing experiment with animals that I believe must exist is that CEVs have relative priority and the experimenter is facing an almost insurmountable problem in that he not only do not know the priorities but he does not even know the CEVs. As long as one does not know what the CEVs are for a particular complex control system then it become extremely difficult to ensure that one is not disturbing perceptions either accidentally or through the intended disturbance(s) that are unknown but at a high priority.

I agree that the experiment that Bruce described did indicate that a specific CEV was being disturbed in the subjects and that while the label was probably incorrect what was labeled was likely really there. I am impressed with several of the aspect of their attempts to be sure that what they postulated was really what was being controlled.

For example when Bruce clearly stated that the experiment was also set up so that pressing the bar would turn on the light but cause the test apparatus to switch to "unscheduled" mode that the rats (and I am assuming that I correctly understood that we are talking about rats that previously learned that pressing the bar would shift to the "scheduled" mode of operation) quickly changed their behavior such that they did not press the bar (presumably after some learning experience).

Bruce's comment about the nocturnal "nature" of rats was one that I already accepted.

The discussion with of our own perceptions presented in view of the present discussion was particularly worthwhile from my viewpoint. Sometimes it helps to be actually trying to think in PCT terms about a problem to really begin to appreciate what otherwise might still be seen as a wonderful presentation.

-bill

Date: Sat, 3 Jun 1995 07:13:11 -0600
Subject: Re: Bill P. will love this; preventing control

[From Bill Powers (950503.0630 MDT)]

Bill Leach (950502.2241) --

> I'm disappointed -- sniff

> You did not comment upon my "Bill will like this one" remark.

[And Bill P. will truly _love_ this next one]

> Compensation The basic idea of feedback is intuitive and simple. From the perspective of a human operator attempting any control action, whether that of positioning a lamp on a table, steering an automobile, or any of the innumerable actions we take continually and instinctively, our action is

> almost invariably tempered by our continuing observation of any discrepancy between intent and status thus far. [The next one however is in serious error unless the author meant "results" when he said "output"] This is negative feedback: The control action is a function of the difference between the desired output and the actual output.

I commented on the final sentence, but not the first part. You are right that I love the first part in that it recognizes the role of control in behavior as a whole. But I hate words like "tempered" , especially in a phrase like "almost invariably tempered". This allows room for actions that would occur anyway, but are only "tempered" by the error. And of course this implies that we observe error signals rather than "status." Words like "tempered," "modified," "modulated," "influenced," and others like them are used mainly when the author knows there must be some sort of effect but doesn't have any idea what it is.

Most of your comments on "Early PCT research" are what I would have said. Your hands-on experiences with real control systems go a long way toward explaining our uncanny agreements on most details of control systems.

One aspect of Bruce's experiment that I didn't comment on (to avoid spoiling the occasion by nit-picking): I wonder why, when a rat found it could prevent the shock (when warned) by flipping over on its back, the experimenters simply ignored this control action and shaved the animal's back. This is the sort of reaction one expects when the experimenter is determined to apply a stimulus, and finds that the animal manages to avoid being stimulated. Overwhelming physical force is then applied to make sure that unwanted control system loses control. Then the experimenter proceeds to "take data", having thrown out a good part of it.

There are other means of avoiding shocks, such as balancing on one foot or two feet for a moment on a single grid wire or standing with the feet on different grid wires that have the same polarity, but I presume that rats were not observed to do this -- that is, were observed not to do this (there's a big difference). I understand that in some shock-administration systems the polarity of the floor's grid wires is rapidly scrambled to keep animals from controlling the shock.

It seems to me that in behavioral experiments there is a lot of thwarting of the animal's control systems, so the experimenter can get the effect he or she wanted. It's hard to get a good picture of control behavior when you start by opening the loop. All this does is prevent you from noticing the most important control processes.

Best, Bill P.

Date: Sat, 3 Jun 1995 12:17:03 -0400
Subject: Re: Bill P. will love this; preventing control

<[Bill Leach 950603.10:10 U.S. Eastern Time Zone]
>[From Bill Powers (950503.0630 MDT)]

In retrospect I see that I was "taken in" just because the author used a human example (which itself implies that the human is recognized to be controlling) in an otherwise rather dry work.

Being critical about it now, it is almost embarrassing that I was so enamored by this text.

In the first place he used the phrase "attempting any control action" which implies the possibility of "uncontrolled actions". There is also however the "or any of the innumerable actions we take continually and instinctively" which seems to contradict the first phrase by implying a recognition that control is not optional.

The "our action is almost invariably tempered by" would have tempted me to send this guy a copy of B:CP. It is almost like this guy might have just about failed a psych class for asserting that the behavior that he has observed appeared to be control system action!

I think that I realize now why his "observation of any discrepancy between intent and status" was not recognized for its error. For some reason, the "positioning a lamp on a table" reminded me of a specific instance. The instance involved "control of a perception not perceivable while actually touching the lamp thus the overall control action involved "stepping back and evaluating the appearance, estimating the position change that might "make things look right", stepping back to the table itself and controlling the new position related perception, then iterate.

Of course I should have recognized that I should not have missed what he was saying - basically that the comparator is in the environment (though I doubt that is what was intended).

This "observation of the discrepancy" thing is a real problem in living control system discussions.

In my own example of moving a lamp, "observation of the discrepancy", I believe, is exactly what was going on. Or at least to describe it that way is not a gross error. I would control a perception of the lamp's position while the real goal was not perceived and therefore uncontrolled. Stepping back and observing allowed me to again perceive the perception whose control was my ultimate goal. If an error existed, I was cognizant of the existence of an error (but probably not all of the detail concerning the nature of the error). In some complex fashion the error was a reference for a program level (?) perception that resulted in a new reference for my perception of the lamp's position.

If this description is at all reasonable for what might have actually been going on in my "lamp moving episode" then it is easy (at least for me) to see how one could write something like the analogy we are discussing. His description _is_ wrong as far as describing a closed loop negative feedback control operation.

It is also probably wrong for describing a complex "iterative" control action because the term "continuing observation" implies continuous observation which is precisely what can not exist is one is "observing the discrepancy ..." of a "controlled" perception.

You raise a good point for me with:

> It seems to me that in behavioral experiments there is a lot of thwarting of the animal's control systems, so the experimenter can get the effect he or she wanted. It's hard to get a good picture of control behavior when you start by opening the loop. All this does is prevent you from noticing the most important control processes.

In fairness to the designer and implementer of an experiment this matter is very difficult to resolve completely (impossible?). My objection to Bruce about the contrived experiment was the result of my uneasiness about the issue (but was also a wholly unacceptable position for most detailed research).

The causality principle in combination with "valid" testing of hypothesis is the foundational principle of modern science. The application of the process to perfection necessitates determining all of the variables, their values and their relationships to each other. To the extent that this is _not_ done the experimental "evidence" is degraded. We believe that achieving perfection in this regard is completely out of the question so we should recognize that all experimental evidence is limited in its' certainty as are all conclusions based upon such evidence.

Since the goal nevertheless is to reduce the uncertainty, two of the possible ways are to limit the number of variables possible and to control others... thus the "controlled" environment experiment.

In the physical sciences this method has had stunning results. Many of the "failures" in controlled experiments have been the most valuable learning tools... when the researcher recognized that the failure WAS a significant challenge to current understanding.

Of course many of the failures have also been due to failure on the part of the experimenter to identify variables and therefore fail to measure or control relevant variables.

In behavioral studies, the sheer number of variables potentially involved or relevant is bound to be staggering. In general the attempt to simplify the experimental conditions with respect to the study of a complex control system will always consist of a balance between having too much data with too many uncontrolled (or at least not understood) variables and unintentionally overwhelming the control system being studied.

In ANY experiment an assumption or hypothesis concerning the direction of causality is "fatal" until that particular error is noticed and corrected. This particular error results in the experimenter asking the wrong questions and designing the experiment around manipulation of the wrong variables.

The experiment then is based upon observer control of an "effect" while at the same time trying to "track" the causal effects upon the CAUSE. I can not think of a real example of this sort of thing in the physical sciences field history (at least for a simple "A causes B" relationship). I don't doubt for a moment that with sufficient research and effort that I could find specific examples for more complex cases.

In the behavioral sciences the situation is almost infinitely more complex. It is precisely the fact that living beings are complex closed loop negative feedback control systems with INTERNALLY set references that makes it possible for an experimenter to "prove" that an "effector" exists in the environment and "works to produce its' effect".

As long as the experimenter has sufficient control over the environment that the subject encounters it is possible to obtain almost any "conditioning" behavioral response desired. This is particularly true if observed behavior that is not in accordance with the "theory" is ignored (such as avoiding shock altogether).

While it might sound like I am trying to "pick" on Bruce and/or his work in this field, I really am not. A rat is "behaving" all the time and from what I have seen of the critters, except when sleeping they are so active as to almost make the observer tired just watching.

The accidental level presses is a good example of the problem, I think. Behavioral results that are not intended by the subject are difficult to evaluate. In his experiment they, at least, were looking for and testing for "intended results" on the part of the subject. The problem then is that there are thousands of intended results mixed in with at least potentially, many thousands of observable unintended results. How and where do you "draw the lines of distinction"?

I think that this problem is true for real PCT research and the significant difference is that the viewpoint demanded by PCT brings the problem to the fore.

-bill
===================
Date: Sat, 3 Jun 1995 11:56:06 -0600
Subject: Re: behavioral experiments

[From Bill Powers (950603.1100 MDT)]

Bill Leach (950603.10:10)--

> This "observation of the discrepancy" thing is a real problem in living control system discussions.

Yes, it keeps coming back. The way to straighten it out is to remember that the relationship you're controlling can be in many states, not just the first one that comes to mind.

Suppose you're trying to put the lamp in the center of the table. So you look at it and see it's not centered, move it, look again, move it again, until it's centered relative to the table as near as you can see. So that looks as if the center of the table is the goal, and the error is the difference between the actual position and the center position.

To show that this is wrong, just change the reference condition. Adjust the lamp so it is one foot to the right of the center of the table. Now it becomes obvious that you're controlling a perceived relationship of the lamp to something else, and that you can set the reference relationship any way you want. Zero error might correspond to the lamp's being one foot to the right of the center of the table. Now the comparator is back where it belongs, in your head. You're comparing a perceived relationship to a desired relationship.

> The causality principle in combination with "valid" testing of hypothesis is the foundational principle of modern science. The application of the process to perfection necessitates determining all of the variables, their values and their relationships to each other. To the extent that this is _not_ done the experimental "evidence" is degraded. We believe that achieving perfection in this regard is completely out of the question so we should recognize that all experimental evidence is limited in its' certainty as are all conclusions based upon such evidence.

When we deal from a belief in open-loop causality, we have to consider controlling experimental conditions to an extreme degree. For example, if we put a rat on a T-maze and want to measure its running speed, we would have to make sure the runway was absolutely level. If it tilted left or right the rat would fall off it; if we tilted it uphill, the rat would (theoretically) slow down and perhaps even stop unless we boosted the stimulus to give it a bigger push.

In reality, of course, the rat will behave the same way over quite a range of tilts. But suppose you're determined to see the effect of tilts on running speed. You'll have to do something to prevent the rat's speed-control system from working, like deafferenting it. Then you'll be able to write your paper on the effect of tilt on running speed, a phenomenon you created yourself.

Shaving the rat's backs removed the possibility of the rat's behavior controlling the receipt of shocks. This enabled the experimenters to do a study in which shocks were treated as an independent variable, being unaffected by any behavior of the rat (or so it seemed). With that problem solved, the experimenter could then study the effect of a warning signal prior to the shock, and the effect of its presence or absence on the rat's pressing a bar that enabled the warning signal.

However, if enabling the warning signal really had no effect on the rat's experience of the shock, then why would the rat turn on the warning signal? Because it looked nice? Obviously (to a PCTer), the rat may have suffered less error when the warning signal was present than when it wasn't present. The experimenters may not have noticed how this happened, but maybe it did. If the experimenters _had_ found some way in which the rats suffered less error from the shock with the signal present, they would have had to take steps to remove it, because they were trying to keep the shock trauma constant, in order to study some other effect. That's what they did when they saw the rats flipping over onto their backs after the signal.

If in fact the experimenters had managed to keep the shock trauma exactly the same whether the warning signal occurred or not and regardless of the rats' behavior, they might not have had any effect left to observe.

With regard to the extreme complexity of behavior, this actually works in a direction that helps with the Test. If behavior is subject to so many unpredictable influences, then the chances that some particular consequence of behavior would prove to be unchanging in the presence of everything that conspires to vary it are pretty small. Finding a controlled variable that is held constant should be a lot easier than verifying that a response variable, which is always changing anyway, is changing in a way that correlates with some particular stimulus variable.

Best, Bill P.

Date: Sat, 3 Jun 1995 13:42:50 -0500
Subject: Re: Early PCT Research

[From Bruce Abbott (950603.1340 EST)]

>Bill Powers (950602.1505 MDT) --

> Your expanded description of the series of experiments answers all the objections I raised. Considering the thorough way in which you tested for variables that the rat acted to control, the main question left in my mind is "what kept you from tumbling to PCT?" Was it simply the custom in that field of interpreting variables controlled by the rat's actions as variables that controlled the rat's actions? If that were all there were to it, I should think that the transition to PCT would be relatively easy for behaviorists. As you look back on your own frame of mind during those experiments, can you find any illumination on this subject?

This is a good question and a difficult one. I think the best answer I can give is that work such as this is done in a context: a particular set of questions being asked from a particular point of view. Within that context what you are doing seems to make sense; it fits in with previous research and builds on it by providing new pieces to the puzzle that is being assembled. To use Kuhn's descriptive, you are doing "normal science," the ordinary pick-and-shovel work within an accepted paradigm.

In this case the context was traditional learning theory dating back to the work of Thorndike, Pavlov, Hull, and Skinner and supported by nearly a century of conditioning studies. The work we were doing in Pete Badia's lab in the mid-1970s was an extension of a line of investigation begun by Pete's mentor Charlie Perkins, who had predicted on theoretical grounds that subjects would prefer a shock schedule in which shocks were signaled over an identical schedule in which shocks were not. We were evaluating Perkin's hypothesis (preparation) against an alternative we favored (safety signal) and a few others, while trying to generate clear, systematic data (functional relationships) whose value would stand (we hoped) independent of any particular theory.

The fact was that rats given a choice between the two otherwise identical schedules chose the signaled one (as Charlie had predicted) and that with the schedule parameters we were using this preference was powerful. In the context of traditional reinforcement theory, this meant that there was something about the signaled schedule, _relative to the equivalent unsignaled schedule_, that was acting as a powerful reinforcer for lever pressing in our operant conditioning paradigm. From the theoretical point of view this was extremely interesting because computation of the so-called reinforcing values of various elements of the two schedules suggested that, under traditional assumptions, the rats should be indifferent between the two options; thus the observed preference seemed an anomaly. In this context we viewed our studies as a search for the source of (differential) reinforcement. What was motivating lever pressing was the greater overall reinforcing value (equivalently the lesser punishing value) of the signaled schedule relative to the unsignaled one. What was the source of this difference?

It wasn't until I was nearly finished with my graduate training that I encountered B:CP; by that time I had become very unhappy with the traditional view and was already beginning to explore the notion that a control-systems approach might provide the alternative I was looking for (I was reading Ashby, Simon, and others). Unfortunately I was also in the middle of preparing for qualifying exams and running three different studies, including my dissertation; I found B:CP very appealing but did not really have an opportunity to sit back and think about the broader implications. It seemed to me at the time that the two views (traditional reinforcement theory and PCT) were dealing with the same observables, but I needed to step back and analyze how the concepts of PCT and of reinforcement theory mapped onto one another. That is, I wanted to understand how control theory would explain those findings for which the traditional view offered what others were accepting as a satisfactory account.

Perhaps another difficulty that got in my way was that I did not really understand the distinction between functional models (models of the behavior of a system) and mechanistic models (models of the system itself). Without this appreciation one can be impressed with models (such as the matching law) that seem to explain observed functional relationships but which turn out to be mostly exercises in curve fitting. There were plenty of such models to admire in the late 70s.

There is a sense in which viewing a set of results such as those obtained in the signaled versus unsignaled shock studies is like looking at one of those "reversible figures" such as the Necker Cube. Such figures can be interpreted in alternative ways that are each self-consistent but mutually contradictory. You can see the Necker cube as if viewing it from above and to the left or as if viewing it from below and to the right. Similarly, you can view our experiments as methods to identify the source of reinforcement for the observed lever-pressing behavior or as methods to identify the perceptual variable being controlled by the rat via its lever-pressing actions. If you've already been trained to recognize the cube-seen-from-above in your data, it may require considerable effort and coaching before you can see things the other way. Even if you can see both cubes, both may appear to be perfectly reasonable representations of the sense-data, at least until you have made an effort to evaluate the deeper implications of each view. In the late 1970s I was beginning to see that other cube, but had not yet done the work required to appreciate its full implications for behavior analysis.

At this point, with what I think is a much better grasp of PCT and its implications, I can flip back and forth between PCT and traditional reinforcement-based descriptions and recognize where they are actually dealing with the same phenomena and in some cases, using essentially the same methods. In those cases it is indeed only a small step from the traditional view to PCT, but for those used to seeing the cube in its traditional orientation, getting the cube to reverse can be a difficult--and disorienting-- perceptual task.

> I trust you misspoke when you included "being able to prevent it" among the conditions to which the rats were indifferent. If a rat could press a bar before the shock occurred and thus prevent it, it would learn quickly to do so -- my rat paper was a study of just such a situation, an experiment done by Verhave, a "Sidman avoidance schedule." But that would not raise the issue of a preference for an abstract capacity to control; it merely would verify that rats will control shocks if they can. It wouldn't tell us what they think about being able to do so, if anything.

No, I did not misspeak. In the avoidance experiment the rat performed on a Sidman shock-avoidance schedule on one session and the actual temporal pattern of shock delivery was recorded (rats on this schedule occasionally make mistakes and receive shocks); this pattern was "played back" on the next session, in which the rat had no control over shock delivery. As with escapable versus inescapable shock schedules, the rats failed to resist the disturbance when the apparatus switched them from avoidable to unavoidable shock schedules or vice versa. The key here is that during training the rats had learned that the shock frequency was the same whether they controlled shock delivery let the apparatus determine when shocks would be delivered.

>Bill Leach 950602.18:33

>>[Bruce Abbott (950602.1455 EST)]

>> Actually, if the rat had done nothing it would have been in the unsignaled condition 0% of the time during the testing sessions.

> Huh?

Oops! Sorry, I meant to say "in the _signaled_ condition 0% of the time."

As to your rather extensive comments in this post, I'm a little pressed for time at the moment, so I'll just respond to the suggestion you raise that doing PCT research on animals is especially difficult because the researcher has no way of knowing what variables the animal is actually perceiving and controlling as opposed to those the experimenter perceives in the situation. This problem has long been recognized by researchers of animal behavior, but it is not an insurmountable one as one can always devise tests specifically designed to determine what these variables are, as was done, for example, in pigeon navigation to determine what variables the pigeon monitors and controls when homing. On the other side of the coin, it may be easier to map out the control systems of animals that lack the higher levels of the control hierarchy found in humans. I never found a rat that chose, say, higher over lower shock intensity because it thought I was trying to assess its ability to "take" the higher shock level and wanted to prove to me that it could.

As to the rat rolling over on its back to avoid receiving shock, the fact that it did so was no reason to throw out the data so long as I could then prevent that behavior and thus collect valid preference data on choice between the two schedules. If you want to know whether Johnny would choose apple over cherry pie, allowing him to choose cake spoils the test and fails to give you the answer you seek, even though it might be of paramount interest in another context. So while Johnny is making his choice, you try to keep the cake out of sight.

Regards, Bruce

Date: Sat, 3 Jun 1995 16:44:08 -0600
Subject: Old points of view

[From Bill Powers (950603.1550 MDT)]

>Bruce Abbott (950603.1340 EST) --

That was a beautiful answer to my question about the transition from traditional behaviorism to PCT. Your analogy to the Necker cube is particularly powerful. It's been said often that PCT is a "valid perspective" on behavior, but that there is something to be said for other approaches, too, such as S-R. This gives the impression that if you use the PCT perspective, you see one set of interesting phenomena, while from another perspective you might see other phenomena that the PCT view misses.

But your discussion makes a strong case that this is not the relationship between the old behaviorist view and PCT. The points of view, while each seems to make sense when it is dominant, are not supplementary or complementary, but _mutually exclusive_. When you are "in" one of the points of view, the other simply appears wrong, backward. When you switch viewpoints, you also switch which view seems right and which wrong. It's as though you're using the same mental machinery to support either point of view, and it can be used only for one of them at a time.

I've tried to learn to switch from the PCT view to other views when a conflict comes up between them. I don't do very well at it, because my supply of lore is far greater in the control-theory field than in any other field, so I can't really get into that place where another view seems to be right and natural. To get into another theorist's shoes, you have to do more than just pretend to accept different conclusions; you have to see how the other guy's reasoning ties together a whole web of observations and ideas in a way that invites belief. I think your experiences with PCT might support this concept; as your experience with phenomena seen from the PCT standpoint grew, it became more possible for you to find a comfortable position in either camp. This, I think, tends to remove the familiarity factor and makes it possible to look for other criteria by which to compare the usefulness of the viewpoints.

> No, I did not misspeak. In the avoidance experiment the rat performed on a Sidman shock-avoidance schedule on one session and the actual temporal pattern of shock delivery was recorded (rats on this schedule occasionally make mistakes and receive shocks); this pattern was "played back" on the next session, in which the rat had no control over shock delivery. As with escapable versus inescapable shock schedules, the rats failed to resist the disturbance when the apparatus switched them from avoidable to unavoidable shock schedules or vice versa. The key here is that during training the rats had learned that the shock frequency was the same whether they controlled shock delivery or let the apparatus determine when shocks would be delivered.

Amazing how your concepts have paralleled ours. Many years ago we did a very similar experiment in which a person did a tracking run while the computer recorded the cursor position; then, under some pretext, the person was asked to "repeat" the run, while the recorded cursor positions were played back. Some people got completely though the second run without realizing that the handle no longer had any effect on the cursor. Most others, when they discovered that the loop was broken, would simply stop tracking, realizing that there was no point to it.

Because the same disturbance was applied in the second run as in the first, we could compare how well the person "tracked" by generating the cursor movements that would have resulted from the handle movements in the second pass. The errors were very large. In some cases, the phantom cursor would go clear off the screen. Interestingly, it was possible for experienced controllers to deliberately believe that they were controlling and go through a whole run, producing results much like those of people who were fooled.

This "deliberate belief" seemed to have the effect of preventing higher-level systems from turning off the tracking control system. As a result, we could look (I think) at the open-loop behavior of the control system, something that is normally very hard to do because when control is lost, people generally just stop trying.

Interesting that apparently the rats did not seem to maintain a believe that they were controlling -- maybe they didn't have one in the first place. Comparative PCT is going to be an interesting field because we can do the same nonverbal experiments with people and different animals. One day, when we call someone a bird-brain, we may know what we are talking about.

Thanks again for the great post. I hope this stuff will take its place in your published writings before long.

Best to all, Bill P.

Date: Sat, 3 Jun 1995 19:32:52 -0700
Subject: Points of View, Early PCT Research

[From Rick Marken (950503.1930)]

Bill Powers (950603.1550 MDT) to Bruce Abbott (950603.1340 EST) --

> Your analogy to the Necker cube is particularly powerful... The [old behaviorist view and PCT] points of view, while each seems to make sense when it is dominant, are not supplementary or complementary, but _mutually exclusive_. When you are "in" one of the points of view, the other simply appears wrong, backward.

I have a feeling that Bill might think that there is reason to suspect that one point of view IS right and the other wrong.

I like Bruce's analogy to the Necker cube, too, but I would suggest that an even more appropriate analogy may be the three dimensional version of this illusion, where an actual wire cube is viewed monocularly. The cube alternates between two perspectives, as in the two-dimensional case, but now one perspective is, in fact, correct.

> Amazing how your concepts have paralleled ours. Many years ago we did a very similar experiment in which a person did a tracking run while the computer recorded the cursor position; then, under some pretext, the person was asked to "repeat" the run, while the recorded cursor positions were played back.

This study is described on pp. 67-75 of "Mind Readings" (Please read it, Bruce. Pretty please. It's just $18.00. Cheap). In fact, it was exactly like the shock-avoidance study Bruce described. There was no pretext for the repeat run; like the rats, the subjects were suddenly dealing with a variable over which they actually no longer had control. Only one subject I tested noticed the change (I set it up so that there was no "hitch" in cursor movement when the replay began). The results were just as you described them -- the open-loop control actions would have produced little or no control if they had actually had an effect on the cursor.

This discussion of Bruce Abbott's early PCT research reminded me of other examples such research. In particular, I remember seeing a film where an infant controls the focus of a picture by sucking on a nipple. The goal of the study was (as I recall) to see whether infants (like two or three months old, perhaps) could perceive focus; the assumption (I presume) was that if they could, then "in focus" would be a reinforcer. In the film, the infant does, indeed, keep a picture in focus by sucking (Maggie Simpson eat you heart out).

It was obvious that the infant was controlling a perception; there was resistance to disturbance since the picture would fall out of focus when the sucking eased up. Nevertheless, as I recall, this result was interpreted in the context of reinforcement theory; "in focus" reinforces sucking, resulting in high sucking rate. If this was the interpretation, they could have quickly disabused themselves of this notion by disturbing the picture _into_ focus occasionally and watching what happened to the rate of sucking under those circumstances.

Best Rick

Date: Sun, 4 Jun 1995 12:08:41 -0400
Subject: Re: Early PCT Research

<[Bill Leach 950603.23:43 U.S. Eastern Time Zone]
>[From Bruce Abbott (950603.1340 EST)]

> ... _signaled_ condition

Thanks, that sure corrected an error!

I think that I am gaining an appreciation for at least some of what you have been doing.

Some of the information that you seek is not itself relevant to PCT (ie: does Johnny prefer apple pie over cherry pie when indeed he prefers cake to either). The difficulty here is that to do the test you may need to create a situation that Johnny would never permit if you did not overwhelm his control systems. That is, given free control he might never be willing to eat either apple or cherry pie.

This is the danger that I think that you always face in such experiments (to some extent unavoidably). If you run your test and he chooses Apple pie you could well have an item of "knowledge" that might actually be useless (unless you are planning to incarcerate Johnny, offer only apple or cherry pie and want to know which to order)!

Of course not all experimental result data will necessarily be useless but just about anytime that you have to prevent the subject from doing what the subject wants to do you are in danger of making a generalization that will really only be valid when applied to the conditions that existed in the experiment.

I DO NOT think that your rat experiment was of this sort however. While this will be conjecture on my part, I think that the results of you experiment support the hypothesis.

Sometime back Rick made the statement that skeptical curiosity had to be taught (I think I have that right). I retorted that it did not but rather WE do suppress it in children. My position was that it was not so important to teach children to use or how to use skeptical curiosity but rather to teach adults to recognize that such a characteristic in children should be valued and not suppressed by the adult.

My basis for this starts with the idea that newborn infant humans begin live by activating every muscle that they have. Soon they are observed moving their own limbs (erratically at first) within their field of vision. Later they are observed moving just about anything that the can grasp and move into their field of vision.

Their entire life seems composed of eating, sleeping and exploring. This remains their primary activity set for several years. Some activities will become regular as they discover things that they learn to enjoy, but the exploring will continue to death (if they are fortunate).

The exploring "sees" its' first curtailment when some large control system physically prevents exploration (overwhelms the smaller control system). Further curtailment occurs when a consistent link is discovered between certain exploratory activities and an error associated with the reference value of pain caused by overwhelming control actions by one of these large control systems.

These probably do not suppress the exploratory activity much but the next disturbance probably does major damage. At the verbal level of exploration an "normal healthy child" can ask more questions about more subjects that any other living being. Unpleasant experiences in this exploration, including the often irrational explanations provided by the large control system(s) likely begin to affect the systems concepts in a negative way with respect to seeking understanding of what is observed.

I posit that this continual barrage of "social conventions" ("good little boys and girls don't ask questions like that" or whatever) results in the formation of systems concepts that suppress the otherwise close to intrinsic desire for understanding of observations.

I am sure that exploration of our environment must be closely related to some intrinsic reference(s). Virtually all animals that I have had experience in observing will explore extensively when introduced to a new environment. The exploration activity for individual animals varies greatly but this could well be due to "bad experiences" in previous exploration attempts.

Most people (though certainly not all) will admit to some discomfort in "unfamiliar" surroundings.

Certainly a "case" can be made for this characteristic existing because of the consequences of "natural selection" but regardless of why it exists I really do believe that at least all humans are "born with it" and it takes a pretty dedicated, persistent and overwhelming force to suppress it.

-bill

Date: Sun, 4 Jun 1995 15:10:46 -0500
Subject: Something to be Said

[From Bruce Abbott (950604.1510 EST)]

>Bill Powers (950603.1550 MDT)

> That was a beautiful answer to my question about the transition from traditional behaviorism to PCT. Your analogy to the Necker cube is particularly powerful. It's been said often that PCT is a "valid perspective" on behavior, but that there is something to be said for other approaches, too, such as S-R. This gives the impression that if you use the PCT perspective, you see one set of interesting phenomena, while from another perspective you might see other phenomena that the PCT view misses.

Thanks, Bill. But the critics are right. There is indeed "something to be said" for other approaches: they're _wrong_.

> To get into another theorist's shoes, you have to do more than just pretend to accept different conclusions; you have to see how the other guy's reasoning ties together a whole web of observations and ideas in a way that invites belief. I think your experiences with PCT might support this

concept; as your experience with phenomena seen from the PCT standpoint grew, it became more possible for you to find a comfortable position in either camp. This, I think, tends to remove the familiarity factor and makes it possible to look for other criteria by which to compare the usefulness of the viewpoints.

The problem is that those who find the accepted paradigm compelling may be unwilling to expend the (perhaps considerable) effort required to truly see things from the other view's perspective. The current view seems to make logical sense, appears to be supported by common sense and by considerable data, and provides a framework for new research. You know what is known and what questions need to be addressed next. You've invested quite a bit of time and energy to thoroughly understand the current view and your experience seems to show that it works. You have little reason to expect that some other view will do better and every reason to want to continue doing things the way you've always done them. So when some other proposal comes along, you may give it a quick once-over, but you probably will not invest the effort needed to really understand it and to work out its implications. It's just human nature.

Re: shock controllability study

> Interesting that apparently the rats did not seem to maintain a believe that they were controlling -- maybe they didn't have one in the first place.

Actually I took pains to make sure they would NOT have such a belief. When the rats entered the condition in which they could not control the shock, the lever they had used for this purpose in the controllable-shock condition was retracted into the chamber wall. I didn't want people arguing that the rats were indifferent between the controllable- and uncontrollable-shock conditions because they "believed they still had control in the uncontrollable condition."

> Comparative PCT is going to be an interesting field because we can do the same nonverbal experiments with people and different animals. One day, when we call someone a bird-brain, we may know what we are talking about.

It may even be a compliment. Anyone who has read about pigeon navigation or looked at the ability if these birds to identify representatives of complex categories in a picture cannot help but be impressed with the kinds of things a "bird brain" can do.

>Rick Marken (950503.1930) --

> I like Bruce's analogy to the Necker cube, too, but I would suggest that an even more appropriate analogy may be the three dimensional version of this illusion, where an actual wire cube is viewed monocularly. The cube alternates between two perspectives, as in the two-dimensional case, but now one perspective is, in fact, correct.

Your point is well taken, but the analogy with the 3-D Necker cube breaks down when you consider that the "correct" orientation is what people tend to see first. In the many times I've done this demonstration it invariably takes a bit of looking before they are able to perceive the "incorrect" orientation so that the cube reverses its apparent direction of rotation. (Apparently there are still subtle monocular cues present that help to guide interpretation). With reference to reinforcement theory and PCT it is the _incorrect_ orientation people seem to grasp most easily.

> This study is described on pp. 67-75 of "Mind Readings" (Please read it, Bruce. Pretty please. It's just $18.00. Cheap). In fact, it was exactly like the shock-avoidance study Bruce described. There was no pretext for the repeat run; like the rats, the subjects were suddenly dealing with a variable over which they actually no longer had control. Only one subject I tested noticed the change (I set it up so that there was no "hitch" in cursor movement when the replay began). The results were just as you described them -- the open-loop control actions would have produced little or no control if they had actually had an effect on the cursor.

It's an interesting study but I'm afraid not quite parallel to mine. In my study the rats knew when they did and did not have control over shock, but did not care to control which situation they were in.

> This discussion of Bruce Abbott's early PCT research reminded me of other examples such research. In particular, I remember seeing a film where an infant controls the focus of a picture by sucking on a nipple. The goal of the study was (as I recall) to see whether infants (like two or three months old, perhaps) could perceive focus; the assumption (I presume) was that if they could, then "in focus" would be a reinforcer. In the film, the infant does, indeed, keep a picture in focus by sucking (Maggie Simpson eat you heart out).

A similar study I recall reading about allowed infants to control the motion of an overhead mobile by sucking on a pacifier. When sucking produced the motion, the amount of sucking increased enormously; when the contingency was broken the sucking returned to baseline levels. These results, too, were given a traditional reinforcement interpretation.

> It was obvious that the infant was controlling a perception; there was resistance to disturbance since the picture would fall out of focus when the sucking eased up. Nevertheless, as I recall, this result was interpreted in the context of reinforcement theory; "in focus" reinforces sucking, resulting in high sucking rate. If this was the interpretation, they could have quickly disabused themselves of this notion by disturbing the picture _into_ focus occasionally and watching what happened to the rate of sucking under those circumstances.

It's not that easy. If the "reinforcer" is having the picture in focus, and the picture is already in focus, then the response does not _produce_ the reinforcer and thus should not be maintained. Rather than being "disabused of this notion" by the results of your suggested demonstration, reinforcement theorists would have found support in them.

> This study is described on pp. 67-75 of "Mind Readings" (Please read it, Bruce. Pretty please. It's just $18.00. Cheap).

Oh, all right, if it'll make you happy. [What some people will do to sell a book! (;->]

Regards, Bruce

Date: Sun, 4 Jun 1995 14:50:55 -0700
Subject: Control, Reinforcement

[From Rick Marken (950604.1450)]

Bruce Abbott (950604.1510 EST) --

> When the rats entered the condition in which they could not control the shock, the lever they had used for this purpose in the controllable-shock condition was retracted into the chamber wall.

Then the study is quite different than the one I mentioned from "Mind Readings" (though you should read the book anyway). Apparently, the rats in your study had a choice between working to keep the shock rate at some level versus not working to keep the shock rate at the same level. I'm not surprised that they were indifferent to the two situations. Indeed, I would have expected a preference for the second (no control) condition.

> I didn't want people arguing that the rats were indifferent between the controllable- and uncontrollable-shock conditions because they "believed they still had control in the uncontrollable condition."

Why do you think the rats were indifferent? As I said above, one might actually expect them to prefer the "no control" situation. The rat is getting the same perception (number and distribution of shocks) in two situations; in one situation the rat has to work to get this perception, in the other it doesn't.

If I were a lazy, Democratic welfare rat, I might prefer the second way of getting the perception. But, of course, I might be a free enterprise Republican rat and hope that I can get a better perception (fewer shocks) when I have to work for it. I presume that you took steps to control for "hope for a better perceptual result" as carefully as you controlled for "belief in having control" in this experiment;-)

> the analogy with the 3-D Necker cube breaks down when you consider that the "correct" orientation is what people tend to see first.

We could fix that with a peep hole and a blank background. What's important about the 3-D Necker cube version of the analogy is that there is a right and wrong interpretation. Testing can show which one is right. The same is true of the behaviorist and PCT perspectives on behavior; we can test to see which is right.

Me:

> If the "reinforcer" is having the picture in focus, and the picture is already in focus, then the response does not _produce_ the reinforcer and thus should not be maintained. Rather than being "disabused of this notion" by the results of your suggested demonstration, reinforcement theorists would have found support in them.

I think you're right. I've found it impossible to disabuse a reinforcement theorist of the notion that reinforcement strengthens behavior. Demonstrations (like E. coli) don't work; modelling doesn't work.

Do you think there is any way to convince a reinforcement theorist that there is no such thing as reinforcement?

Best Rick

Date: Mon, 5 Jun 1995 12:13:52 -0500
Subject: Controlling Rats; Reinforcement

[From Bruce Abbott (950605.1210 EST)]

>Rick Marken (950604.1450)]

>>Bruce Abbott (950604.1510 EST) --

>> When the rats entered the condition in which they could not control the shock, the lever they had used for this purpose in the controllable-shock condition was retracted into the chamber wall.

> Then the study is quite different than the one I mentioned from "Mind Readings" (though you should read the book anyway). Apparently, the rats in your study had a choice between working to keep the shock rate at some level versus not working to keep the shock rate at the same level. I'm not surprised that they were indifferent to the two situations. Indeed, I would have expected a preference for the second (no control) condition.

In the absence of other information, yes, so would I. But what if just HAVING control is itself a controlled perception? Maintaining that perception might have been "worth" a little extra effort. My study was designed to find out.

So why did I think rats might _prefer_ having control? Earlier research had shown that, in rats at least, shock that was controllable was less aversive and had milder physiological effects than otherwise equivalent shock that was uncontrollable. Given this, a logical inference would be that rats given the choice between controllable and uncontrollable shock would prefer the former. My study showed quite conclusively that they don't, at least under the conditions I tested. There remains, of course, the possibility that they might prefer control under other circumstances.

> Why do you think the rats were indifferent? As I said above, one might actually expect them to prefer the "no control" situation. The rat is

getting the same perception (number and distribution of shocks) in two situations; in one situation the rat has to work to get this perception, in the other it doesn't.

I suspect that any preference that might have existed for the "no control" situation was weak and fell below the limits of experimental error.

There is an apparently common-sensical idea that people who are experiencing stress in their lives should be given more control over the sorts of things that cause the stress, and that this control (even if illusory) will diminish the impact of those events. My research (and a little thought) suggests that this idea is not necessarily true. Having the perception of control will probably reduce the stressfulness of the situation if your experience suggests that this control will generally lead to an objectively better outcome than if you lacked control. For example, you may feel more relaxed when YOU drive the car than when your teenage son, who has had two weeks of driving experience, is behind the wheel. However, having the perception of control will probably _increase_ the stressfulness of the situation if your experience suggests that this control will generally lead to an objectively _worse_ outcome. For example, you might feel much more comfortable allowing the pilot to land the 737 you are flying in than taking the wheel yourself. In my study the rats were able to compare the objective outcome under controllable- and uncontrollable-shock conditions and determine that they were the same. Given that determination, there was no basis for choosing one condition over the other and the rats remained indifferent (within the limits of experimental sensitivity).

>> If the "reinforcer" is having the picture in focus, and the picture is already in focus, then the response does not _produce_ the reinforcer and thus should not be maintained. Rather than being "disabused of this notion" by the results of your suggested demonstration, reinforcement theorists would have found support in them.

> I think you're right. I've found it impossible to disabuse a reinforcement theorist of the notion that reinforcement strengthens behavior. Demonstrations (like E. coli) don't work; modelling doesn't work.

> Do you think there is any way to convince a reinforcement theorist that there is no such thing as reinforcement?

Yes, but I think it's going to take a whole series of demonstrations to drive the point home. I've already driven the stake through the heart of more than one theoretical vampire only to find, much to my surprise, that I hadn't killed it after all, so I'm not surprised that you've had the same result. Winning this argument is going to be more like piling blocks on the lighter of the pan balance until it finally tips. The other guys are keeping their hands pressed firmly down on the other pan, and you're going to have to overcome that force and not just the weight of the theory itself.

One difficulty is that there _is_ such a thing as reinforcement--by definition. So when you tell a reinforcement theorist that there "ain't no such thing," you're bound to be met with a bit of skepticism. I'll try to clarify this point with an example. Let's say you meet some people who believe that lightning bolts are weapons thrown by Thor, the god of thunder. You try to explain that their belief is wrong, because there are no such things as thunder bolts. Thor's followers might be forgiven for dismissing you out of hand--after all, lightning bolts are empirical facts! Of course, what you MEAN is that lightning bolts, DEFINED AS WEAPONS THROWN BY THOR, do not exist, but that subtle distinction gets missed.

If I food-deprive a rat for a few hours and then give the rat the opportunity to earn a bit of food by pressing a lever, the rat learns to press the lever. The frequency of lever-pressing will increase over time and will be maintained at some higher level so long as the rat remains hungry and the contingency between lever-pressing and pellet delivery remains in effect. The observed increase in responding is _by definition_ reinforcement. Thus the argument is not whether there is or is not a phenomenon which has been labeled "reinforcement," but how this phenomenon is to be explained.

Reinforcement theory holds that certain sensory consequences of responding cause changes to occur in the nervous system that make it more likely that the response will occur again. PCT holds that responding increases because error between a controlled perception and its reference level generates output that tends to reduce the error; given the environmental feedback function, reduction of error requires an increase in (lever-pressing) output. Reinforcement theory makes reinforcement the central explanatory principle for behavior change; PCT makes it a side-effect of control. The argument is not about the objective phenomenon of reinforcement but its theoretical significance.

Regards, Bruce

Date: Mon, 5 Jun 1995 13:40:03 -0700
Subject: Comparing views of behavior

[From Rick Marken (950605.1330)]

Bruce Abbott (950604.1510 EST) --

> The current [non-PCT] view [of behavior] seems to make logical sense, appears to be supported by common sense and by considerable data, and provides a framework for new research... So when some other proposal comes along, you may give it a quick once-over, but you probably will not invest the effort needed to really understand it and to work out its implications.

I'm wondering whether the current view of behavior is anything more than just that -- a view. I say this because I think it is true that most psychologists would, indeed, agree that the current view is supported by considerable data (and, apparently, rejected by none).

I think of the "current view" in the behavioral sciences as any version of the cause-effect model of behavior: S-R, selection by consequences, output generation. If behavior is, indeed, the control of perception, then it is rather surprising that there is not more data that rejects the current view. Although the behavior of a control system can look like S-R, selection by consequences or output generation, the behavior of a control system is VERY different than an S-R, selection by consequences or output generation _model_ of behavior.

The fact that the results of virtually all the studies that have been done in the behavioral sciences are seen as being consistent with the current view suggests, to me, that people are not really comparing these results to the behavior of a working model that embodies the current view; they are only looking for qualitative match between what the see and what they think they should see. So, the fact that there is SOME relationship between an independent and a dependent variable, for example, is enough to prove that the basic assumptions of the cause-effect model are correct; stimuli DO cause variations in behavior.

I think this qualitative approach to evaluating theories is one reason why we (PCTers) are often accused of developing "straw man" theories when we try to test working versions of S-R, selection by consequences and output generation models against our experimental results. The models we develop based on our understanding of the "current view" always fail - - rather dramatically in most cases -- and I don't think conventional behavioral scientists are used to this; they certainly don't expect to see their "current view" of behavior just flat out rejected by the data. They are used to seeing data that is ALWAYS consistent with their view -- data that can be explained (verbally) and accounted for (by curve fitting).

So I'm not sure that any amount of data could convince conventional psychologists that their current view is wrong. Any result produced in a PCT experiment could be "explained" by the conventional view. If the data clearly reject the current view, conventional psychologists can always say that our model is a "straw man". This is what happened to Tom Bourbon and Bill Powers in their paper demonstrating the inability of S-R and output generation models to account for the simplest tracking behavior. It is what happened to me in a paper demonstrating the inability of a selection by consequences model to account for goal directed behavior when the consequences of actions are random.

I guess what I'm saying is that most conventional behavioral science is actually pre-scientific and, thus, not really vulnerable to disproof. I suppose I would change my mind about this if some conventional behavioral scientist would propose a working model (based on the current view) of some behavior that we could test against PCT. Does anyone know of such a model?

Best Rick

Date: Mon, 5 Jun 1995 21:52:06 -0700
Subject: Stress, Proving PCT

[From Rick Marken (950605.2145)]

Bruce Abbott (950605.1210 EST) --

> There is an apparently common-sensical idea that people who are experiencing stress in their lives should be given more control over the sorts of things that cause the stress, and that this control (even if illusory) will diminish the impact of those events.

According to PCT, stress is caused by error -- not things or perceptions of those things. So any lack of control IS a cause of stress (that's why Ed Ford's book about PCT is called, quite appropriately, "Freedom From Stress"; we are free from stress when we are in control). You can't really "give" people control, though you might be able to show people how they can improve their ability to control (again, that's what Ed's book is about; ways to improve your own ability to control your perceptions).

If control is illusory (as it was in my closed-open loop experiment) then it is not control -- so it won't help reduce stress (error). A case could be made for the notion that most of the control we seem to have over other people is illusory; that's probably why much of our stress is associated with perceptions that involve other people; it's just not possible to reliably control those perceptions (as Ed's book makes clear; if you won't read my book, you could at least read Ed's;-)).

> My research (and a little thought) suggests that this idea is not necessarily true.

I'm not sure that your research really addressed the causes of stress. My interpretation of your results is that the rats were in stress in BOTH conditions: control of shock and no control of shock. I think that any level of shock greater than 0 (the rat's probable reference level for shock) is stressful for a rat -- it creates error. Thus, your methodology ensured that the rats experienced the same amount of stress in the control of shock and no control of shock conditions.

I take the rats' failure to reliably select either control or no-control of shock as evidence of reorganization. If the control of shock condition allowed the rat to actually control the shock, keeping it at the reference level (0) then I suspect that the rat would have selected the matching no control of shock condition every time; that way the rat gets the reference level of shock (0) with no work. The observed "indifference" refers (I think) to the fact that the rat selects the control of shock condition 1/2 the time and the no control condition the other 1/2 . This suggests that the rat is reorganizing -- randomly selecting behavioral strategies because the existing strategies are not working; the rat is getting shocked -- it is in stress; it is reorganizing because its current control organization is not working (no thanks to the experimenter, I might add;-)

Me:

> Do you think there is any way to convince a reinforcement theorist that there is no such thing as reinforcement?

Bruce:

> Yes, but I think it's going to take a whole series of demonstrations to drive the point home.

I basically agree with you; there is no one experimentum cruxis that demanded a Copernican view of the universe, for example. But (per my previous post) at least there was a working model of the solar system (Ptolemy's) that could be compared to observation. I don't believe that the equivalent of a Ptolemaic model (of behavior in general; there are some good models of specific phenomena, like Mach Bands) exists in psychology.

Best Rick

[This thread is long enough. The file ReinforcementTheory.pdf contains continuing discussions which focus on Reinforcement Theory. The file ShockingExperiments.pdf contains continuing discussion of rat experiments].