Posts about engineering, terminology, and PCT

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

Date: Sun Jun 20, 1993 4:51 pm PST
Subject: Re: emergence, etc.

[From Bill Powers (930620.1600 MDT)] Hans Blom (930620) --

> The slogan "it's all perception" is much too static. It has a connotation of having to act in an a priori circumprescribed way given a set of perceptions and a set of top-level (very slowly changing; innate?) reference levels.

I agree. Perceptions are learned, too. There is a reason for the slogan, however. It's to remind us that each of us sits inside one of these gadgets, and that ALL we know either of the world or of our own actions and inner being consists of perceptions. When we produce a particular carefully crafted action, it is a perception of that action that we know about and control. When we see an effect of the action on the world, it is a perception of the effect that we observe and adjust. When we feel joy or anger at the result, that, too, is a perception. There is nothing else to experience. Our actual outputs, the signals moving in the outward direction, are not part of experience at all. This is of no concern to an engineer, who looks at his control systems strictly from outside. But in psychology it is the key to understanding what control theory means for behavior.

> Blaming what you see for what you do is usually considered a defense mechanism or, if more forceful, for criminality.

True. This is one of the things that PCT is trying to change in psychology: the idea that perception causes behavior. PCT says that the person selects some experience as a goal, and acts to make present-time perception conform to it. Behavior controls perception, not the other way around.

> Optimal control theory applied to humans says that you can fine-tune your actions as well, that you have control over how you control, i.e. that you have SELF-control.

Again, I agree. Of course you can't fine-tune your actions directly unless you can perceive them; otherwise, you can only reorganize until the perceived result of the actions is what you want, without any direct knowledge of the actions themselves. And it remains true that even when you specifically adjust your actions, it is a perception of the actions you must adjust; the action itself is output, and not sensible.

> Yes, see how far this reaches when you consider humans. We do not only operate on the outside world but on the inside world as well.

The distinction between inside and outside is a perceptual classification; both, as far as the brain (or PCT) is concerned, are inside (or both are outside, it makes no difference). Everything the brain can deal with exists in one space, the space we call the experienced world. This world is derived completely from signals generated by sensory receptors; there is no other way to get information about an external world. The nature of that world has to be inferred by the brain from the behavior of the signals and how they respond to attempts to affect them.

>We can tune our responses finer and finer, and reach ever
>higher qualities of response and perception.

I'm not sure how you mean this, but it sounds like one of the concepts we're trying to destroy. "Response" is the conceptual opposite of "control." It implies a blind reaction to an input, and carries overtones of jab-and-jump psychology. If we can "tune responses" we must be sensing something that depends directly on them; all we can actually tune is the sensory consequence, for a pure response (of your own) is not itself experiencable. Whether you intended this or not, this way of speaking about what is learned encourages the old idea that perceptual inputs cause motor outputs -- the very idea that allows people to blame what they see for what they do. This is one of the many basic conventional concepts that stand arrayed against PCT.

>Control over control is self-control, perceiving your own
>perceptions is self-perception, consciousness.

That sounds nice, but I don't believe it. If you diagram a system that senses the stability of a control system and adjusts parameters to control stability, you do not have a system controlling itself: you have a system controlling something about a different system. If you perceive your own perceptions, one subsystem is perceiving the perceptions originating in another subsystem, and most likely interpreting them in a different way. The moment you say "I am thinking," you have denied the statement: the system that is aware of the thinking is not thinking, it is making a statement about a system that is thinking. The "I" of which you speak is never the "I" that speaks.

The only way to make sense of self-reflexive ideas is to treat a person as if that person were solid, like a potato: only the whole person perceives and acts. Only in that way can one say that the referring self is the self referred to. That view is contrary to the modeling approach, in which we try to understand the whole in terms of interactions among its subsystems.

I said:

>> In other words, I could ask the question, "How is competent adult human behavior organized when its organization is not being changed?" This is what HPCT is about.

And you said:

> Whereas my focus is more on how human behavior can become even more competent, i.e. more on learning (and evolution as a kind of learning).

I think that my goals have to be reached before yours can be reached (at which point yours would be mine, too). Before you can study how to make the human being more competent, you have to have a way to measure its competence. Psychology has fallen down on that job; nothing it says about behavior can be taken as a clear fact, because its factual statements are riddled with important exceptions and counterexamples. We need a highly predictive and accurately descriptive model of how behavior works when it is not changing. Only then can we measure change in any reliable way, and know whether our attempts to improve competence are having any effect, good or bad.

> One can, at our level of discourse, see a feedback amplifier or some such device in two very different ways. The first is as a device that transforms an input (voltage, current or power) into an output (voltage, current or power). The feedback is not really relevant here. The (voltage, current or power) gain may be any value; both gains and losses can be similarly realized. The second way is to see the device as a power modulator, where the input modulates the transfer of power from its power supply to its output. The feedback is not relevant here either.

Neither of these concepts is especially relevant to PCT. The first is the usual idea in which the "input" (meaning, really, the reference input) is confused with sensory inputs, so it appears that an input from the environment is causing an output into the environment. In living control systems the reference input does not come from the environment, but from higher systems.

The second applies primarily to the output function. It's not often necessary to draw the power supply of a control system; the behavior of the system is quite insensitive to changes in the power supply, as you know.

> Then, internal in the device, we see the LOOP GAIN, which has nothing to do with power at all, but is the gain in the loop that the SIGNAL travels. A decent control system has a loop gain much greater than one,

although a loop gain of less than one is not unthinkable. In the latter case we might not want to consider the system a CONTROL system.

Loop gain, in PCT, is not "internal to the device." The relevant closed-loop path passes through the environment. I suspect that you haven't yet understood just how the PCT diagram differs from the standard engineering one. Actually, I'm working on this subject right now for the joint paper, but I suppose it won't dilute the writing too much to show you three diagrams that will be in the paper, and preview the discussion. It's important to understand exactly what we have done with the standard diagram.

Fig. 1 is what the behavioral sciences have taken from engineering, for the most part:

```
            < ------------"the device" -------->

                       error
 Input ----> comparator ----> forward function -----> Output
         +    - |                                |
FIG 1           |                                |
                 --<----- feedback function <---
```

If you draw a circle around everything under "the device" above, that is everything but the Input and the Output, you have the first case you describe above, as well as your third statement in which feedback is "internal to the device."

Now let's just add a few details without altering the overall appearance: I'll have to use some abbreviations to fit it all in.

```
        < -- "device"----><-- environment ------------- >

sr --->[Comp]- se -->[fo]--> qo -->[fe] --> qc <--- [fd]<-- qd.
    +   - |                                  |
FIG 2     |                                  |
        sp <------------[fi]<---------------
```

Here sr = reference signal, comp = comparator, se = error signal, [fo] = forward or output function, qo = output quantity (the immediate effector output), [fe] = environmental function (which transforms the effector output into an effect on the controlled quantity), qc = controlled quantity (the physical quantity actually under control), [fi] = input function (which includes the sensory receptors and any computations that immediately follow), and finally, sp = perceptual signal, the internal analog of the controlled quantity. [fd] and qd provide for representing independent disturbances and their influence on the controlled quantity.

I hope you agree that the organization of this model is identical to that of Fig. 1, except for the explicit inclusion of a possible disturbance and insertion of some stages implicit in Fig. 1. I have relabelled the "input" as the reference signal, which does no violence to engineering custom, and the "output" as the controlled quantity, which is also an acceptable alternative in engineering parlance.

However, I have expanded the details at the system's output a bit. I have distinguished between the immediate effector output and the controlled quantity, and introduced an environmental function expressing the dependence of the controlled quantity on the effector output. An example would be a control system that controls a shaft's angular position. The controlled quantity qc is the angle at the end of the shaft where the load or workpiece is; the output quantity qo is the torque output of the driving motor. The intervening [fe] expresses the way torque is converted into shaft position, given the way load resistance depends on angular position (which could be assigned to the disturbing branch).

This separation is always important in detailed control-system design, but especially so when the effector is coupled to the controlled quantity loosely or through complex intervening processes. Then we clearly would expect the effector output to be changing far more than the controlled quantity is changing. Even when we're just talking about output torque and controlled shaft position, the system may have to vary the output torque radically, even changing direction, in order to maintain a specified shaft position, as twisting disturbances are applied one way and the other to the controlled quantity at the end of the shaft. I know you know all this; I'm just making the description complete.

Now consider Fig. 3:

```
                             |sr  ref signal
                         -  +|
                  sp ------[Comp] --->-- se
                  |                       |
                  |  sensor     actuator  |      system
    :::::::::::::[fi]::::::::::::::::::::[fo]:::::::::::::::
                  |                       |     environment
     controlled  qc <------[fe] --------- qo
       quantity   |                  actuator output
                  |
FIG 3            [fd]
                  |
                 qd   disturbing quantity
```

This is organized exactly like Fig. 2. It is simply rearranged. It is actually just like Fig 1., with details added. The plane of separation between system and environment, however, is not the one suggested by the first diagram. To locate it in the first diagram, you would have to draw a line like this:

```
                                 [SYSTEM]    | [ENV]
                                             |
 Input ----> comparator ----> forward function -----> Output
                 |                           |      |
                 |                           |      |
                  ---------- feedback function <---
FIG 4                                        |
                                             |
```

This distinction means little in engineering, but in PCT it is essential for getting the correspondences between the engineering diagram and the physical organism right. In Fig. 3, the horizontal line separates the nervous system of the organism from all that is not nervous system. Sensors and actuators lie on the boundary. Notice that in Fig. 3, there is no chance of mistaking the reference input for a sensory input. The reference signal comes from higher up, inside the behaving system. The sensory inputs are strictly associated with the feedback path through the environment. In living control systems, unlike artificial ones, the reference signals are not accessible from outside the behaving system.

In those control systems that have been traced out in human beings and animals, Fig. 3 corresponds closely to the sensors, intervening cells, and output paths. The reference input corresponds physically and functionally with what are traditionally called "command" signals, signals which carry outputs from systems higher in the brain. Those command signals have been thought of traditionally as carrying commands to the muscles, causing them to contract (the feedback paths are ignored even though they are mentioned in a sort of puzzled way). The control-system diagram, with parameters filled in to make it fit real behavior, shows that the so-called command signal is really a reference signal. Its primary effect is to specify the level to which the perceptual signal will be brought. The actual outputs could be in any state, depending on what disturbances happen to be acting on the controlled variable. The closed loop system varies the output in any way required to make the controlled quantity, and thus the perceptual signal, match the reference signal. It does this without any instructions from the reference signal.

Fig. 3 was drawn as it is with full knowledge of the engineering diagram of Fig 1, for a specific purpose. Almost without exception, behavioral scientists have interpreted the "input" of Fig. 1 to mean "sensory input." When that is done, Fig. 1 becomes nothing but a stimulus-response diagram with an internal feedback loop having no obvious function. Wiener said it "reduced the dependence of the output on the load." This has been taken as cybernetic justification for the old model in which sensory inputs cause behavioral outputs.

The reason for emphasizing the distinction between the actual effector output and the controlled quantity (usually absorbed into a single equation in engineering) is to show the difference between the physical action of the system and the sometimes remote outcome of that action which is actually under control. When we see the controlled variable separated from the effector output, we can much more easily understand that the visible behavior of an organism is really just its actuator output, while the focus of the control action is an effect of that actuator output -- a joint effect, because disturbances act on the controlled quantity, too. Thus, with this diagram, we can point out the specific difference between what we see an organism doing and the controlled outcome of those variable actions.

You might think that this rearrangement would be easy to explain to real control engineers, but that has not always proven to be the case. Control engineers get just as set in their ways as psychologists. Long experience only seems to make matters worse. One old control engineer with whom we went around and around for six months on the net ended by saying that he saw what we meant, but he just couldn't get used to talking about a controlled variable as associated with input. So he bade us farewell, wishing us luck in a gentlemanly way. Of course a much younger one, encountered in a different venue, thought this was terrific, and adopted the PCT model for teaching control theory to graduate students. I guess there are a couple of control engineers on this net who have seen the light. I don't know whether you have or not; it's hard to tell from what you say.

All this volume of output was necessary to explain why I object to your statement

> Then, internal in the device, we see the LOOP GAIN, which has nothing to do with power at all, but is the gain in the loop that the SIGNAL travels.

The loop gain in Fig. 3 is the product of the partial derivatives of [fi], [Comp], [fo], and [fe]. That is the gain that determines how tight the control will be. It specifically must include the path through the external environmental feedback function. The control loop in PCT is NEVER "internal to the system." It ALWAYS passes through the environment, no matter what level of control is involved. This is what makes PCT models testable. There may in fact be closed loops totally above the line in Fig. 3, but in behavioral experiments we can do nothing with them. Their effects will simply be absorbed into the basic model of the control system.

And I think that's quite enough of my Sunday and your time to spend on one post.

Best, Bill P.

Date: Mon Jun 21, 1993 9:52 am PST
Subject: information, stress

[From Rick Marken (930621.1000)] Bill Powers (930620.1600)

> I suspect that you [Hans] haven't yet understood just how the PCT diagram differs from the standard engineering one.

This was a wonderful discussion, Bill.

It shows, once again, that PCT is trying to just make a simple point; that is, behavior is the process if controlling INPUT perceptual variables. That is a simple point, but it is basic. If one doesn't fully understand and accept this simple fact about the organization of living systems, there is no chance that their more complex analyses of behavioral phenomena can be worth much because they are based on the wrong premise. I think that many scientists in all fields these days are happy to skip the fundamentals in order to get on with the real interesting, complex modelling. I think you alluded to this, Bill, in you discussion of the information theory material that you had read. It is this apparently irresistible urge to get on with the complex stuff and skip lightly over (or just ignore) the fundamental assumptions is what leads to 1) a lack of interest in PCT 2) Karolyan PCT and 3) trendy science.

Best Rick

Date: Fri Sep 24, 1993 1:05 am PST
Subject: PCT and manual control

[From Bill Powers (930923.1930 MDT)] Chris Wickens (9309xx)

I'm gradually learning from experience: before we get into a big discussion about PCT vs. the manual control literature, we're going to have to get our terminology straightened out. We use terms differently in PCT from the way they're used in standard control engineering. When we say "input" we don't mean the reference signal, but the variable being sensed by the system at the input boundary between system and environment; when we say "output" we don't mean the external variable that's under control, but the immediate output of the system's effector(s) acting on the local environment. The reason I suspect that we have a terminology difference is this:

> I don't think any of us in the serious manual control literature believe that error (as defined explicitly by the difference between the state of the output..its position and derivative -- and the input -- its position and derivatives) MUST automatically be the signal that drives control (with an opposite sign, to reduce the error).

If you mean "controlled variable" by output and "reference signal" by input, then that paragraph makes sense. But if you mean "effector output" and "sensory input" respectively, you're not even describing a control system as we understand and model it. We clearly have to get this sorted out.

One thing we do in PCT that's not explicitly done in standard control engineering is to clearly delineate the input and output boundaries between the active system and its environment, and specifically identify the sensors and effectors involved, plus their relationships to environmental variables. This is important in modeling organisms in order to distinguish what is a property of the organism from what is a property of its environment. The environment part of a control loop can be defined in a way common to all organisms, but the organism part may be different with every different organism. In particular, a given condition of the same environment may constitute an error for one control system, and no error or the opposite error for another one.

Hans Blom has suggested that in talking with control engineers we use the topology of their customary diagrams instead of our usual "canonical" PCT diagram. So here's what I'm talking about in the standard engineering form:

```
                               ||
 ref      ----  e   -----      ||
 signal->|comp|--->|outpt|     ||       ----------
 [INPUT]  ----     |funct|--->---qo -->|f.b. funct|->
                    ------     ||       ----------   |
    percep|tual    ORGAN-      ||                    |
       sig|nal      ISM           ENVIRONMENT        |
          |         -----      ||                    |
           ---<---|input|<-------qi--------<--------
                  |funct|      ||  (controlled var)
                   -----       ||   |  [OUTPUT]
                               ||    \
                                      ---<-- disturbance
```

The vertical double line is the organism-environment boundary, the organism being on the left. The output function includes the effector, which turns the error signal (e) into an output quantity that is the physical effect on the immediate environment. The state of the output quantity acts, through properties of the environment represented as the environmental "feedback function", on the variable that is to be controlled, here labelled qi. The controlled variable is also subject, in general, to independent disturbances, so its state is the sum of two (or more) influences, one being the behaving system's own action. Note that what we call qi, the input quantity or controlled variable, is called in standard engineering terminology the _output_. Also, in the standard engineering diagrams, the output is shown coming directly out of the control system without explicit mention of the effector or intervening functions in the environment, and the input path generally contains no explicit representation of the sensor.

In PCT, the input function senses the state of the input quantity and converts it to a signal representing its state. This is the perceptual signal, which enters the comparator. The other input to the comparator is a reference signal (called _input_ in standard engineering diagrams).

The perceptual signal represents the controlling system's only knowledge of the controlled quantity. What the system controls, therefore, is the state of the perceptual signal, not necessarily the state of the external (observable to others) controlled variable qi. If the sensor calibration drifts, the perceptual signal will still be maintained in a match with the reference signal, while the visible controlled quantity's value changes. The variable most reliably controlled by this system is the perceptual signal. Thus the name of my first book: _Behavior: the control of perception_.

Notice where the error signal is in this system: inside the organism. There is no error in the environment. The input function does not detect any error conditions: it simply reports the state of the controlled variable in the form of a perceptual signal. Because the reference signal could be set to any value, there is no "natural" error condition in the environment. What states of the input quantity constitute an error depends entirely on the setting of the reference signal inside the organism.

This is a "compensatory" configuration. In a "pursuit" configuration the situation inside the organism is exactly the same. However, the perception is now derived from two environmental variables, one that is under direct feedback control, and a second that is independent. The input function senses the difference between the controlled environmental variable and the uncontrolled one, reporting the difference between them (in visual tracking, the distance between the target and the controlled cursor). The perceptual signal now stands for the magnitude of this difference. There is still no error in the environment; the perceived difference is simply whatever it is. It is the reference signal that specifies the desired difference, which may or may not be zero. A person can just as easily keep the cursor one inch to the right of the target as on it, by setting the reference signal to a nonzero value corresponding to one inch of separation.

So in PCT there is never an error in the environment, under any conditions. The variable under control is defined by the nature of the input function. The reference signal determines what state of the resulting perception is the zero-error state. In the hierarchical PCT (HPCT) model, the reference signal is the output of a higher-level control system which acts to control its own more abstract variable by varying the reference signal of the lower-order system(s). There are many levels in the HPCT model, related in this way.

You can see that the organization of the PCT model is identical to that of normal control systems. The main difference is that we don't define control in terms of objective consequences outside the organism, but in terms of perceptions and their associated reference signals. Reference signals are generated entirely inside the organism; the only inputs from the outside world are sensory inputs representing the states of environmental variables which are defined by the nature of the input function involved.

Our "test for intentions" is really a general test to determine what external variables the organism appears to be controlling. Presumably, these are represented by perceptual signals. The Test requires applying disturbances to the supposed controlled variable, and seeing whether the system's output action varies so as to have an equal and opposite effect on it. When such a relationship is found, we presume that there is a perception inside the system corresponding to the observable variable being controlled relative to an internal reference signal. We can infer the setting of the reference signal by finding the condition of the controlled variable toward which the system's actions always urge it in the presence of disturbances. The intention of the system is the setting of its reference signal. The qualitative aspect of the intention is defined by the nature of the controlled variable found with the Test. The quantitative aspect is the particular state in which the system appears to be stabilizing that variable. The intentions or goals of the system relate to perceptions, not objective aspects of the environment. To intend is to intend to perceive.

This very brief summary doesn't get into hierarchical control, or into "reorganization" which is the primary learning mechanism of this model. But perhaps you can see enough of the PCT approach in this to make some comparisons with the manual control approach. I'd be very interested in seeing them.

Best,   Bill P.

Date: Fri Dec 10, 1993 10:20 am PST
Subject: What PCT says about control (For Osmo)

[From Bill Powers (931210.0715 MST)]   Osmo Eerola (931209)

My earliest direct experiences with the properties of negative feedback systems (not counting my U.S. Navy experience as a technician during World War II) began with a Philbrick analog computer, in about 1954. Here I became acquainted with operational amplifiers, which are high-gain (x100000 and up) DC amplifiers with a bandwidth of perhaps 20 kilohertz and a frequency characteristic that falls off in amplitude by 3 db per octave. By connecting various passive components in series with the negative input and in the path from the output to the negative input, one can create functions ranging from simple proportionality to first and higher integrals. By using diodes, one can make an absolute-value function; by using an amplifier open-loop, one can create a step function; by using a capacitor in series with the input, one can make a derivative-taker. Adding a feedback path to the positive input can create a flip-flop with hysteresis. Multiplication was even possible using special circuit elements

```
                              -FEEDBACK ELEMENT--<-
                             |                     |
                             |    |\               |
                             |    | \              |
                             |  - | OP \           |
IN-->INPUT SERIES ELEMENT|---x-I--|      -------------> OUT
                                + | AMP/
                        GROUND -->|  /    x and I: see text
                                  |/
```

FIG 1: An operational amplifier in
    an analog computing setup.

My insight into "control of input" arose when I read, and verified, that an operational amplifier in one of these computing configurations acts to keep its negative input terminal at a voltage matching that of its positive input terminal. When you understand that, it's easy to see what functions are created by various combinations of passive series input elements and feedback elements. The current generated by a voltage applied to the input element can be computed by realizing that the other end of the input element is effectively tied to a fixed voltage; that current must run through the feedback element to the output, which immediately shows the way the output depends on the input current, and hence on the input voltage.

It took me a little while to make the connection between operational amplifiers and the control systems I was simulating by using them. The engineering control-system diagrams I was using were not organized the way the analog computing diagram was organized, and there was an entirely different emphasis on various aspects of the whole system. Finally, however, I realized that the controlled variable of the control system did NOT correspond to the output voltage of the operational amplifier, but to the voltage at the negative input terminal. I had been deceived by the fact that engineers talked about controlling "outputs," stabilizing them against disturbances, while I could see that the "output" of the operational amplifier was not stabilized against disturbances. The disturbances were clearly the input voltages applied to the input series element, and what was stabilized against them was the voltage at the negative input terminal, not the output voltage.

The reference signal was the voltage at the positive input terminal; if that voltage varied, the voltage at the negative input terminal would follow it accurately. The output of the operational amplifier corresponded to something like the torque in an output effector motor, while the feedback element corresponded to the physical laws determining the effects of that torque on the angular position of a shaft. Under that analogy, the shaft position would be the controlled variable at point x, and the sensor that feeds the shaft position back to the control system's input would be in the position labelled I above. Disturbances would be things like loads, and the series input element would express the physical effects of loads on the same shaft position.

Now consider one of the elementary physiological control systems in a human being, the so-called "tendon reflex." It is organized like this:

```
                            | alpha command signal
                            |
                 -       +  V
      ------>-----  Spinal Motor Neuron----->-
     |                                        |
     |sensory signal                          |
     |                                        |
     |  force                                 |
  TENDON <------(passive spring)---<-MUSCLE CONTRACTION
 RECEPTOR<--(physical effects)-
                               |
                          disturbances
```

Fig. 2: Spinal control system

Clearly, this diagram can be converted into a diagram of the form of Fig. 1 just by rotating and moving things around a little without changing any connections. The spinal motor neuron and the contractile part of the muscle combine to make up an operational amplifier with a high-powered output; the negative input is the sensory signal input to the spinal motor neuron, and the positive input is the alpha command signal. The feedback pathway consists of the physical effects of the muscle contraction on the tendon receptor, and the sensory signal completes the path to the negative input.

The only missing component is the series input element, and that is clearly supplied by external disturbances that tend to alter the force applied by the muscle to the tendon. While the disturbances are not applied directly to the negative input, they are connected though a physical path to the negative input and that path can be represented as in Fig. 1. The relationship between disturbances and muscle contractions is clearly that between the operational amplifier's output and the input to the series computing element. The form of the relationship is determined by the forms of the functions in the feedback path and in series with the effects of the disturbance. If the gain in the neuromotor "op amp" is high enough, the characteristics of the spinal neuron and muscle are of minor importance: the passive external components determine the nature of the observed relationship.

The physical output of this system is a contraction in a muscle, a shortening of contractile fibers. But that is not the controlled variable. What is controlled is the force created by the muscle, and that force is what the tendon receptor senses. The input to the sensor is controlled. When disturbances appear, the output of the system, the degree of muscle contraction, changes equally and oppositely, maintaining the sensory signal almost unchanged and matching the alpha command signal. If the alpha command signal changes in magnitude, the muscle contraction will vary as required to keep the input signal tracking the alpha command signal, and as much more as required to offset the effects of any disturbance.

Clearly, the alpha command signal is not a command signal, but a reference signal. The spinal motor neuron is a comparator, and the muscle is an output function.

Compare the above two diagrams with a standard engineering diagram of a control system, found in any elementary textbook.

```
            +       error  ("controller")
INPUT ------- Comp ------>FORWARD ELEMENT --------->OUTPUT
               |-                           |
                ---<-----FEEDBACK ELEMENT<----
```

Fig 3. Standard engineering diagram
of a control system

How would you match Fig. 3 to Fig. 2? The feedback element clearly includes the tendon receptor, which is affected by contractions of the muscle that produce force through the spring characteristics of the passive part of the muscle. Those passive spring components would have to be part of the forward element as the diagram is shown; there is no provision for showing the difference between the physical action of the effector and the resulting creation of an output effect, the force. The feedback element would be located at the junction labeled "I" in Fig. 1. The tendon receptor's signal is carried to the negative terminal of the spinal motor neuron. The error signal is the output of the spinal neuron, so the muscle corresponds to the "controller" or forward element. The "output" of the muscle is contraction producing a force acting on the tendon receptor, closing that part of the loop. What, then, corresponds to the line labelled INPUT? It is the alpha command signal, which is the reference signal for this control system.

Where, in the spinal control system, does the alpha command signal or reference signal come from? It does not come from any sensory input, but from higher in the nervous system, in some cases from the brain stem, in others from the cerebellum, and in still others directly from the primary motor area of the cerebral cortex. It is not an input from the environment, but the pathway through which higher centers in the brain operate the muscles. In other words, higher centers in the brain produce actions by varying reference signals that enter the comparators of spinal control systems that control sensed force. The reference signal specifies a particular amount of force, and the feedback action of the spinal control system alters the contraction in the muscle until the sensed force matches the requested force. If the requested force varies, the control action makes the sensed force track it.

Clearly, the label "output" in the standard diagram is misleading. The physical output of the muscle that is immediately caused by the error signal is a contraction, a shortening of contractile elements. That output stretches a passive spring element, which results in production of a force. The force is the variable that is sensed and controlled, and that accelerates the limb, but that force is most closely connected to the input, not the output, of the control system. Independent physical effects can alter the force just as much as the muscle contraction can; even changing a joint angle alters the force by stretching or relaxing the spring element. The only aspect of muscle function that corresponds reliably to the driving error signal is the length of the contractile elements in the muscle: that is the true output of the system. And that output is obviously not controlled; any disturbance of the force can cause it to change. The force is a consequence of applying that contraction to the physical world in parallel with other physical processes that also affect the force. The force is a controlled variable, but it is not a measure of the output of the control system.

In our PCT diagram, we carefully separate the elements of the control process so that things which are critically different are not lumped together and represented as a single function, as shown in the "canonical" diagram below.

```
                               |reference signal
                               |
perceptual signal---->---COMP------>----- error signal
                  |                         |
            input funct            output funct
                  |                         |
input quantity  i<-----feedback funct-<--o output quantity
                  |
                   ----<--disturbance function <---disturbance
```

Fig. 4: The standard PCT diagram of a control system

Note that the output quantity, a measure of the immediate physical effect of the output transducer, is separated from the input quantity by a feedback function. This feedback function contains the physical relationship that exists between the output quantity and its effects on other variables. In a shaft angle control system, the output quantity would be the motor torque which varies directly with driving current, and the feedback function would express the things like shaft twist and moment of inertia that are involved in translating motor torque into an effect on shaft position at the end of the shaft. The disturbance would be some other physical variable, such as a varying load or an externally applied torque, while the disturbance function expresses the effects of the disturbance on the controlled variable, the shaft angle.

The controlled variable is treated as an _input_ quantity, because it is directly sensed by the input function (which includes a sensor and any signal-processing functions). It is the input quantity that is controlled. When disturbances vary, the output quantity varies in an equal and opposite way, keeping the input quantity from changing (nearly, if control is good). The overall effect is to keep the perceptual signal, which is what we call the output of the sensor, in a continuing match with the reference signal.

Note that this discussion is not about advanced control theory and such matters as achieving stability under various conditions. It is about how we think about control systems, how we set up our ideas about their organization before we begin to analyze them. You will notice that each element of the PCT diagram in Fig. 4 has an exact counterpart in the diagram of a spinal control system, Fig. 2. In contrast, the standard engineering diagram in Fig. 3 is very hard to fit to Fig. 2, and in addition it leaves false impressions about the nature of the "input" and the "output." From the engineering diagram, one could easily get the impression that reference signals are inputs from the environment, and that muscle contractions are controlled outputs: neither of those interpretations is correct.

My contention is that the engineering diagram of a control system has been extremely misleading to scientists trying to apply control theory to living systems. When an engineer designs an artificial system for a specific purpose, any diagram will suffice if its elements contain all the necessary transformations. But in trying to analyze an already-existing living control system, understanding is impeded if the diagram is not organized in detail like the real system. I am sure you can see that the PCT diagram is exactly like the engineering diagram, except that elements and relationships that are laid out specifically in the PCT diagram are lumped together in the engineering diagram. All engineers understand that the "input" is a reference signal -- but even such engineers are misled when they think about living control systems into believing that this "input" is a sensory input from the environment. In the living control systems about which we know the specifics, like the spinal reflexes, iris reflex, and others, reference inputs do NOT come from the environment, but from systems higher in the brain's organization.

I should mention that Fig. 2 is a simplification; there are actually two layers of control loops in the motor reflexes, the higher level being the stretch reflex which senses and controls muscle length. The error signal from that system enters the spinal motor neuron, adjusting the reference signal for the force (acceleration) control system. The gamma efferents are the reference signals for this second layer of control, which controls muscle length and thus, approximately, joint angle. Gamma and alpha efferents are often co-activated.

While the engineering diagram is sufficient for engineering purposes, in my readings about engineering designs I have not been impressed by the orderliness of the initial approach to design. As far as I can see, there are no systematic design principles taught to students; at least none appear in the textbooks I have seen. The basic approach seems to be to define what it is about the plant that one wants to control, and to find inputs and forward transforms that will produce approximately that effect, with feedback being used more or less as a way of trimming the performance and almost as an afterthought counteracting disturbances. The design phase is brief and sketchy; the student is plunged immediately into complex mathematics without further consideration.

I think that the principles made clear in the PCT diagram could improve engineering design. The basic principle is clear in this diagram: the primary consideration is to provide a sensory signal that represents the variable in the plant that one wishes to control, so that by comparing the signal against a reference signal, one can generate an error signal indicating the degree to which control has not been achieved. Then the design of the forward part of the system can be filled in, to provide the gain and output power necessary for good control and the filtering necessary for stability.

With this kind of design, the reference signal represents directly the desired state of the plant variable, and the feedback signal directly represents the actual state of that variable. Oddly enough, in many designs I have seen (including modern examples), this simple concept seems not to be used. The potential meanings of these signals are lost in the shorthand mathematical representations of the system, and the simple logic of control as seen under PCT is buried in the mathematics. I suspect that in many cases this has led to awkward and unnecessarily complex designs.

But that's really not my problem, and engineers may know of difficulties that preclude this simple and orderly approach. All I know is that the PCT approach makes good sense out of behavioral organization, and the models we have built to represent real behaviors on that basis work very well.

Best, Bill P.

Date: Tue Dec 28, 1993 9:17 am PST
Subject: PCT and CT

[From Bill Powers (931228.0750)] Osmo Eerola (931228.0845 GMT)

(In replying to Martin Taylor):

> Cannot see why "multiple input" systems could not control "single output".

... and ...

> ... it is very well possible that multiple inputs can control a few outputs - the output O(t) is a function of several inputs I1(t), I2(t),...,In(t). There can even be several sensed (perceived) variables of a process and all they control the output in combination in a closed loop manner.

The most important parts of your statement are "is a function of" and "in combination." A function or combination of a set of inputs has just one value at a time. Thus what seem to be multiple degrees of freedom when seen from outside the system become a single degree of freedom inside the system. This is the basis for the hierarchical perception aspect of HPCT. When multiple inputs are reduced to a single representative signal, it is no longer possible for each input to have a unique effect on the output. Instead, only a function of all the inputs has a relationship to the output.

Because of this reduction in degrees of freedom, it is possible for the same set of inputs to be controlled in several different and independent respects. The simplest example is a system with two input variables. One perceptual function can compute the sum of the two variables, and a second one the difference. Two independent control systems can be set up, using two output actuators, such that one controls the sensed sum while the other controls the sensed difference. The reference signals for each system can be adjusted independently to cause the sum and the difference to be maintained at the specified levels, independently of each other. One control system's error signal affects both external variables in the same direction, and the other affects them in opposite directions. So both actuators are used by both control systems, but the effect is that of independently controlling two degrees of freedom of the environment. This is easy to simulate, by the way.

Control of this kind is handled in standard control theory, but in a way that obscures the principle involved (at least for me). When I see matrix control equations, I don't see any of these basic relationships. Maybe other people do. I've never seen them discussed in a CT text.

From me:

\>> I think of controlling as something done by the whole closed loop, not by any one part of it.

From you:

\> Sure. But one can observe the closed loop functioning by measuring it from different parts:

That's true of the engineer looking at an artificial control system. But now consider your own control systems (the ones inside which you live). Now your only access to the loop is in the perceptual signals, the feedback signals. You can't experience your outputs, or the mechanisms in the external part of the loop. Everything you know (even about the forces you generate with your muscles) has to come to you through sensors, so all you ever experience are the feedback signals in various control systems.

When you look at another person behaving, you see the immediate and indirect environmental effects of the muscle forces, but now you have no access to the feedback signals or the other parts of the control systems inside the other person. You can't see directly what aspect of the environment the other person is perceiving, so you can't see immediately what the other person is controlling.

Building up a control-system model of human behavior requires putting these two views together into a model that explains them both (never forgetting that while you are watching another person behaving, you are still observing only your own perceptual signals). I haven't seen this discussed in CT textbooks, either.

\> You have applied Ct to human behavior and call it PCT. The main "difference" between CT and PCT is that you control perceptions, not outputs (but actually the outputs are affected, too).

Yes, the outputs are affected, too, as well as the external variables that the outputs act upon. Let's try to reach agreement on what "output" means in PCT. The basic output of a higher organism consists of muscle contractions. Muscle contractions are a physical effect of neural signals that depends ONLY on the neural signals (disturbances capable of altering the muscle contractions independently of the neural signals' effect would hardly ever occur, and then only under very unusual circumstances). In general, we use the word output to mean the immediate effect of an actuator device on the part of the physical world that it directly acts upon.

The parallel in the world of electromechanical system would be found in a motor used as an output device. The electronic signal entering the motor is like the neural signal in the living organism. This electronic signal (operating through a power output stage) causes current to flow through the windings of the motor, with the result that a torque is exerted on the armature. It would be very unusual to find any external agency that could alter this torque, so the torque (the static torque at least) is a direct measure of the immediate effect of the electronic system on its environment. We define that as the output of the control system.

This lets us distinguish between the physical outputs generated by a control system and the effects those outputs have on the rest of the physical world. In the human system, the muscle contraction has the immediate effect of stretching the elastic part of the muscle between its attachments. This is a purely passive, mechanical, physical effect involving the series spring constant of the muscle. It is the stretching of this "spring" that produces a force tending to pull the attachments of the muscle together. The shortest feedback loop in the human body involves sensors embedded in the tendons that fasten the muscles to the bones. These tendon signals do not measure the output of the control system (the shortening of the contractile part of the muscle), but a physical effect of that output (the stretching of the series spring element, and hence the force applied to the tendon). This is a negative feedback loop, with the comparator residing in the spinal cord as a spinal motor neuron. The error signal is the output of this neuron, and it enters the muscle to cause the contractile elements to shorten. The feedback from the tendon receptor inhibits (proportionally) the output of the spinal motor neuron, making the feedback negative

We can apply this same principle to the artificial system with a motor as an output device. The output of the electronic control system is the torque applied to the armature of the motor. The immediate result is to create a twisting moment on the shaft that connects the motor to the load. The shaft "winds up" slightly, and this slight torsion produces a torque at the other end of the shaft, determined by the torsional spring constant of the shaft. That torque tends to accelerate the load.

In the human system, the actual force applied to a tendon when the contractile fibers shorten depends not only on the contraction, but on the angle at the joint. If the joint angle changes, the series spring element in the muscle will lengthen or shorten even with the contractile part in a constant state of contraction. The joint angle, in turn, depends on external loads and inertial forces (if the angle is changing). So the net force sensed in the tendon is affected not only by the muscle contraction, but by external disturbances of various kinds.

This applies in the artificial control system as well. The angular position of the load can be disturbed by external forces and inertial effects, independently of the motor's torque. The acceleration of the load is thus not completely determined by the torque applied to the armature of the motor and the springiness of the shaft; it is also affected, just as much, by externally applied torques and by the moment of inertia of the load. The acceleration, velocity, and position of the load are affected BOTH by the output of the motor (the torque applied to the armature) AND external torques and loads.

Therefore measurements of the state of the load are not the same as measurements of the output of the motor, with output defined as we do in PCT. I trust that you agree with this analysis so far. Any actual engineering design has to take these factors into account. These details tend to get lost in conceptual diagrams, but in building any real control system they must be considered just as laid out here.

The tendon "reflex" is a force or torque control system. When you press a finger against a tabletop, this force control system is used to make the applied force vary in a precise way, with the tabletop preventing any angular accelerations or velocities from developing. The reason that this is a force control system is that force is _sensed_ by the system. The muscle contraction can have different effects under different physical circumstances; it is the fact that tendon receptors sense force, and not velocity or position, that makes this a force control system.

In the artificial system, what is controlled about the load depends on what is sensed. If a strain gauge is used to detect the torque applied by the load to something else, something that is prevented from spinning freely, then we will have a torque control system. If we mount an angular accelerometer on the load, the system will be an angular acceleration control system. If a tachometer is mounted on the load, sensing angular velocity, we will have an angular velocity control system. If we mount a potentiometer on the load, sensing angular position, we will have an angular position control system. If the load is a pinion gear engaging a rack, and we use a linear position sensor on the rack, the system will be a linear position control system. If we let the load be an iris diaphragm and we mount a photocell behind the iris, we will have a light-intensity control system.

What is controlled depends on which physical effect of a control system's output is sensed. The sensor, not the immediate output of the system, is the determining factor. In the above paragraph, we used exactly the same output, a torque applied to the armature of a motor, to control six different physical variables. What made the difference was not the kind of output used, but the kind of sensor used, and where among all the indirect effects of the output torque we chose to place the sensor.

By making a distinction between the output of a control system and physical effects of that output, we separate the controlled variable from the output. When disturbances are applied to the controlled variable, the output can now change in a way that has an equal and opposite effect on the controlled variable. If we shine a light on the iris diaphragm, disturbing the illumination of the photocell behind it, the motor will rotate and close the diaphragm until the light intensity is at its former value. If we increase and decrease the illumination slowly and smoothly, the motor will rotate in one direction and the other, and the light intensity at the photocell will not change appreciably (of course it will change a little bit). So the output of the control system is changing, but the _controlled effect_ of the output is not changing.

Because the sensor determines which physical effect of the system's output will be controlled, the controlled variable is defined primarily by the sensor and the signal emitted by the sensor. This is what "control of perception" means in the context of standard control theory.

I have not seen any discussion of this kind in a textbook of control engineering. When it comes down to a specific design, then of course all the details I have mentioned must be considered and the appropriate calculations must be done. But these details get left out of conceptual discussions. In your diagram showing G(s) and H(s) and S and P, none of the details was considered. The output of G(s) was simply labelled S, the controlled variable, with no provision for a disturbance that might also affect S. When such a disturbance is drawn into the diagram, it becomes necessary to separate the output of G(s) from the state of S, because now the disturbance can alter S independently of effects from the output of G(s). And when you do that, calling S the "output" becomes less justifiable -- S is not a measure of the actual output of the controller any more.

There is much more to say on this subject, particularly about controlled variables that exist only as functions of multiple external variables, but this is enough for now.

The control of perception concept is important in the analysis of human behavior, because the physical outputs of a human being have innumerable effects on the local environment. Most of these effects are not under control. The only way to find out which effects are controlled is to find out which are being represented in the perceptions of the person, and which of those perceptions is being compared with a reference signal to produce an error signal that drives the outputs. This is not like the engineering problem, where you start by knowing what is to be controlled and design the rest of the system around that. When we apply control theory to living systems, we do not see what is being controlled. All we see are the physical outputs and their multiple physical effects. That is not enough to tell us what the person is doing -- that is, which effects of the observable actions are intentional and under control, and which are accidental side- effects.

Sorry to go on for so long, but it is important to me to convince you that PCT is not JUST standard control theory. It is a detailed and orderly way of applying CT that is not taught in any control-system courses I have ever seen. Whether it is necessary for engineering students I can't say, but this careful approach is very necessary if we are to understand human control systems. Human behavior is so complex and multidimensional that one can easily get lost in trying to apply control theory. Without a detailed and systematic approach, important details can easily be overlooked, with the final result that the beautiful overall picture doesn't come into view.

Best, Bill P.

Date: Thu Feb 24, 1994 8:58 pm PST
Subject: Re: Feedback & Making the Grade

<[Bill Leach 940224.22:36)] >[Richard Thurman (940224.1000)]

Hey Richard, nice striking up a conversation with you.

I date back to about 1957/58 in my introduction to the concept of feedback. The first real "grounding" in the term was Frederick E. Terman (sp?) and that was the use of feedback in electronic amplifiers for stability and achieving specific characteristics. Also, the electronic oscillator.

Next came "control systems" where "actuator" drive was modified using feedback to improve response and avoid "overshoot".

In all of the above, the feedback is essential a signal that is almost completely dependent upon the "internals" of the system under discussion. To some extent, particularly in the mechanical actuator case, the "resistance" provided by the "external" environment has a decidedly direct effect on the feedback. But conceptually, the important point with feedback as I am referring to it is that the feedback is a function of output and not input.

I am afraid that these uses of the term have probably all but permanently taken a "set" in my mind. I think that my mind "hears" the term feedback and automatically does a test to see if the term is being applied in the above fashion or not and if not then completely ignores the idea that what is being discussed is actually "feedback".

> This statements lends me to believe that PCT is using the term feedback differently than in classical control theory.

Could well be but what worries me about that is that there could very well be feedback loops within the human physiology that DO directly correspond with this more classical usage... then what will we call them?

Whoa there! Don't align me with the Systems Theory camp yet! ...

After reading some of the postings today, I'm not so sure that has to be the "unforgivable sin". :-)

-bill

Date: Fri Feb 25, 1994 7:08 am PST
Subject: More thoughts on the term

<[Bill Leach 940225.08:13 EST(EDT)] >NET

I my previous comments on "feedback" I finally got to the point of saying essentially that feedback is a "do it to yourself" operation.

I "hear" support for that position in that comments by others are made to the effect that "one that is providing feedback to another is really only able to provide a disturbance variable".

I believe that the term as used most of the time here has been with regard to very high level behavior. That is, in general someone tries to provide "feedback" to re-enforce a particular behavior (or maybe feedback to discourage a particular behavior).

This sort of use of the term feedback is very common. If you are considering two or more people to be a SINGLE system, then such use of the term is probably correct at some level(s) of abstraction.

As was amply pointed out here (as far as I can tell), determining whether something is even feedback much less "positive" or "negative" can well nigh be impossible for the observer. The concept becomes even harder to get "a handle on" when you consider that the observer has to consider the references (unknown in the absolute sense for sure) of all of the interacting parties. To me, that relegates such use of the term feedback to the realm of "Classical Psychology".

I don't have as much trouble looking at the idea that a change in the disturbance signal that is a direct result of an action by the subject being considered to be feedback. To me then, IF the subject is trying to control perception to reference where from the subject's view the other party's disturbance IS the perception being controlled then again, feedback as an applicable term is "ok" with me.

It is my belief however that such conditions are seldom known or considered when people throw out the term feedback. Even in such use, I still have a feeling that the term is being abused and looses its power to describe a valid control phenomenon.

If we do want to use the term feedback (within PCT) in such a high level way on interactive control systems, I will attempt to adjust my thinking.

-bill

Date: Fri Feb 25, 1994 12:50 pm PST
Subject: feedback

[From Rick Marken (940225.0900)] Bill Leach (940223.18:51 EST) --

> I don't personally believe that the term "feedback" has any place in PCT unless there is an actual situation where it is determined that a portion of the output signal is returned to the input to enhance control.

I see that you come to PCT via engineering control theory. I hate to say it but this does not bode (no pun intended) well; control engineers (on this list, anyway) have been quite reluctant to abandon some of their well-entrenched ways of looking a control systems -- making it impossible for them to learn PCT. The control theory of PCT is identical to the control theory of engineering; what is different is the mapping of control functions to system functions, a process that often makes explicit some functions that engineers take for granted (and, thus, ignore). Bill Powers described the difference between the engineering and PCT mapping of control theory to actual systems in a wonderful post several months ago. Perhaps someone can find it -- it may have already appeared in a closed loop.

I will just say here that the difference between the engineering and PCT mapping of control theory to systems is most profound in the part of the control loop where (as you say) "a portion of the output signal is returned to the input". In engineering diagrams (the maps of the system used by engineers) this "output takeoff" is, indeed, just a line connecting the output back to the comparator. This line looks like a wire (and in some control systems it may be) but in a living control system this line is a neuron carrying a perceptual signal. The place where this line connects to the "output" variable is a LOT more complex than the engineering drawing lets on because this is the place where the "output" variable is sensed. The sensor is the perceptual function in PCT and in living organisms it is the afferent nervous system -- which provides us with our perceptions of the world. The "output" variable in the engineering drawing is really an environmental variable that is sensed by the organism. In PCT we call this "output" variable a controlled (environmental) variable. The actual "outputs" of an organism (according to PCT) are the physical actions that affect the state of the "output" (controlled) variable; in engineering drawings, the PCT output variable is called the "plant output".

In PCT, the relationship between action and result (between plant and "output takeoff" in the engineering model) is the feedback function; it is the set of physical laws that determines how actions affect the controlled perceptual consequences of those actions. (The PCT model makes it explicit that the perceptual variable is controlled; this, of course, is also true in the engineering model -- it's the same model -- but the engineer would say that it is the "output" -- really the "output takeoff" -- that is controlled). The relationship between action and perception (in a living control systems) is called feedback because the effect of these actions are fed back onto the actions themselves (because they occur in a loop). In PCT, "feedback" is a FUNCTIONAL DEPENDENCE of variables on themselves -- it is not a THING that can be GIVEN. Feedback is the central concept of PCT: because of the functional dependence of actions on themselves ("feedback") we are dealing with systems that have completely difference behavioral characteristics than systems that do not have it. When the sense of the feedback is negative (actions have effects that reduce their own tendency to occur) the behavior of a feedback system is PURPOSEFUL -- it maintains perceptual variables at internally specified values.

Best Rick

Date: Fri Feb 25, 1994 9:55 pm PST
Subject: Electronics metaphors and PCT

[From Bill Powers (940225.1030 MST)] Bill Leach (940224 etc.)

Bill, I'm glad to know you're an old electroniker. So have I been since the early 1950s (actually since 1946 as a Navy "electronic technician's mate" if you can accept that term without jumping to conclusions). We ought to be able to work out this problem about feedback with no difficulty.

The basic problem is that you're using electronics metaphors that need to be expanded and adjusted to fit PCT.

One metaphor is "feeding back a portion of the output to the input." If you consider this idea carefully, you will realize that what is fed back is not a "portion of the output," but a signal that is proportional to the output. The fed-back signal doesn't have to be in the same form as the output.

For example, the output could be a voltage, and the input could be a current going into the base of a transistor. You can't feed back a portion of the output voltage to achieve feedback, because the transistor is current-sensitive, not voltage-sensitive. So what do you do? You insert a device into the feedback path that converts the output _voltage_ into a proportional input _current_: a resistor. The fed-back current now subtracts (for negative feedback) from the signal input current at the same base junction.

Now go slightly further afield. Suppose we want to make a super hi-fi audio system. Rather than feeding a voltage or current back from the output transformer secondary to get the feedback we want, let's mount a high-quality capacitor microphone right in front of the speaker, very close to the diaphragm (less than half the highest wavelength). This microphone is used to pick up the air pressure caused by the speaker and convert it to a voltage instantaneously proportional to air pressure. We use that voltage as the feedback signal to an earlier stage of the power amplifier (using a resistor to convert it to current if necessary). By this means we make the feedback signal proportional to the actual momentary air pressure, which is really what we want to control (who cares what the diaphragm movements are, if the air pressure fluctuates exactly as the input signal does?). The gain of the system makes the feedback signal match the input signal, which now means that the air pressure fluctuations match the electrical fluctuations in the input signal. Perfect hole-in-the-wall hi-fi.

The metaphor of taking a portion of the output to feed back is now converted to something more realistic: we are using a feedback path that creates a signal proportional to the output effect we want to control. By using the output to make that signal track the input, we thereby make the output effect also track the input.

One more step is needed: to get away from the idea that it is the output of the system that is controlled. Suppose that instead of making music, we use the same amplifier to position a load (in the early days, I actually used a commercial 100-watt audio power amplifier with response to DC in servomechanisms -- very handy). Now the output of the electronic system is a torque applied to a motor armature.

We don't want to control the torque, but something that is affected by that torque. The armature turns a shaft, the shaft turns a gear train, the gear train turns a screw, and the screw moves the load. So the position of the load is a rather remote effect of generating an output torque.

What we need now is to turn the position of the load into an electrical signal. We do this by mounting a linear potentiometer on the load, long enough to span all desired positions. We apply a fixed voltage across the pot. The load moves the wiper, and the wiper then picks off a voltage proportional to position. We can now feed back this position voltage, which is really a perceptual signal in PCT terms, to the input of the amplifier. If there is high enough gain, the fed-back voltage will be made to track the input voltage. Since that fed-back voltage is determined by the load position, the result is to make the load position track the input voltage, too. We now have a system in which a physical output (a torque) is mechanically converted into a different variable in different physical units (a position), and in which the feedback effects are converted again into still another physical form (a voltage analogous to position).

One last set of adjustments to the electronic metaphors needs to be made. We have to change a bunch of labels around to meet the needs of modeling the behavior of an organism.

First off, we actually have two "inputs" to the control system. One is the input signal that is used to specify the desired load position (or that is the electronic signal input to the audio amplifier). The other is the electronic signal generated by the position-measuring potentiometer (or the signal representing the state of the output in the audio amplifier). One (or both) of these signals has to be renamed so we know what is meant when we refer to the "input."

In PCT we rename the input that is driving the whole system; we call it the "reference signal," which is compatible with normal engineering usage. The other signal, the one that represents the state of the controlled variable (load position or air pressure), we call the "perceptual signal" because it is an internal representation of the external variable that is to be controlled. The device that actually generates the perceptual signal as a function of the state of the controlled variable (the potentiometer) is called the "input function" (surprise!). Servo engineers often call it the "sensor."

On the output side, we call the output of the control system the actual physical effect produced by the electronic (or neural) signals at the electronics-environment boundary: the torque in the servo motor, or the force created by the voice-coil. We call this the "output quantity" because it is a measurable physical quantity that is the first possible measure of the macroscopic physical effects of the system on its environment.

There is one last label. The output quantity (torque or force) is connected to the controlled quantity (load position or air pressure) through physical linkages of some sort (gears or laws of air compression and wave propagation). These linkages make up what we call in PCT the _feedback function_.

So to sum up for the load position servomechanism and add a few remaining terms:

A reference signal enters the system at one input, along with a feedback or perceptual signal. The difference or error signal (computed by a comparator or differential amplifier) is highly amplified to create a macroscopic physical output effect, a torque, which we call the output quantity. This torque feeds back through mechanical linkages to affect the state of a controlled quantity located at the _input_ to the control system: the position of a load which is sensed by a potentiometer, the input function. The resulting input or perceptual signal which represents the state of the controlled variable connects back to the place where the reference signal came in, at the comparator, and we have completed the loop. If the amplifications are large enough, and dynamic stability is achieved (not usually difficult), the result will be that varying the reference signal will result in varying the load position so the perceptual signal representing load position almost exactly tracks the reference signal.

Kind of makes you want to go out and build one, doesn't it? Building one is exactly as simple as it sounds, nearly.

Now we need to orient this diagram in a standard way and decide where to put the system boundaries.

First, tilt the system 90 degrees so the reference signal comes in from above instead of (as usually) from the left. Draw the boundary between the control system and its environment as a (now horizontal) line that passes through the output actuator (motor) and the input sensor (potentiometer). The only input from the environment is the sensory input representing the state of the controlled quantity. The only output to the environment is the output of the actuator, the output quantity. In the environment, draw a path from the output quantity to the controlled quantity, which is placed just below the sensor. This path is the environmental feedback connection through which the output quantity affects the controlled variable at the input.

The reference signal comes from higher up inside the overall system in which we find this control system. It comes from inside the organism, not from its environment. This is a break with engineering tradition, because when engineers build a servomechanism, they make its reference input accessible to adjustment by external agencies in the environment -- users. In organisms there is no external user, and reference signals are generated by higher organizations inside the organism. The only inputs from environment to organisms are _sensory_ inputs, which simply report the current state of the controlled variable.

I just don't feel like going through the tedium of making an ASCII diagram of all this. You'll find the standard diagram in nearly all PCT publications, in some form or other. If the above is too confusing I'll byte the bullet.

Incidentally, you are quite right to reject the implication that a disturbance is "feedback." A disturbance is an independent influence applied to the controlled variable (directly or through effects anywhere else in the loop). It would be friction in the screw moving the load, or someone pushing on the load. That is simply a disturbance, which is automatically countered by variations in the control system's output quantity.

The term "feedback" itself is almost always misused if you want to be a purist. Feedback is a property of an entire loop; the effect of a variable in the loop _on itself_. When I call the output-to-input connection the "feedback function" I am speaking loosely; it is really the external connection that makes feedback possible. If we understand the organization of a control loop clearly, we can be a little loose sometimes, although it's probably a bad habit.

Best, Bill P.

Date: Sat Feb 26, 1994 1:06 am PST
Subject: Re: Electronics metaphors and PCT

<[Bill Leach 940225.21:23)] >[Bill Powers (940225.1030 MST)]

"Electronic Technicians Mate" should have been the forerunner to the "Electronics Technician", yes? Hummm, 1946... probably not involved in Radio though as I think that Radiomen serviced and operated their equipment.

Bill, thanks! Your lucid explanation did a great deal for me, I think.

I really don't think that I ever did have a problem with the reference signal. Even in "regular control theory" I have always "sorta" considered changes to set-points as a change in "purpose". The idea that the "set-point" "setter" is wholly internal doesn't bother me a bit.

Just to make sure that I do have this straight though:

The entire perception signal is considered feedback, yes? A disturbance then, if perceived becomes a _part_ of the feedback, yes?

Thus, in the skater example: If we assume that the coach's words are a significant factor in the perception of the subject, then calling what the coach is doing "feedback" is NOT really out of line.

Now then calling it "positive" or "negative" feedback would be wrong as far as PCT is concerned BECAUSE the in the sense that such modifiers are typically used the paradigm is not PCT. If I think that I am "giving positive feedback" to support a behavior, that is ok but from a PCT standpoint, the "feedback" still results in a perception and that perception is still a part of a negative feedback control loop?

-bill

Date: Sat Feb 26, 1994 1:06 am PST

<[Bill Leach 940225.21:12 EST(EDT)] >[Rick Marken (940225.0900)]

I'm glad that I waited to read Bill's post before replying to yours.

I am quite happy to say that I believe that I fully understood Bill's posting. The functional mapping is no problem to me BECAUSE PCT really does map the functions and not just play games with the terms!

I accept that many times, even here, when we discuss human behavior we will likely get a little "loose" with the use of our terms (particularly something like "feedback"). I now think that such occurrences will not bother me too much since as I see it there is both the rigorous testing of the theory (where terms must be used in their most precise and strictest fashion if results are to be meaningful) and there is the looser, more subjective conjecture sort of discussion of the possible implications of PCT. In these "typical real life" type situations we can look for and talk about PCT type behavior but really only get a "feels right" sort of "proof". Of course such discussions may "prove" to be quite valuable for some troubled individual if a PCT approach is taken in the solution to their problem(s).

-bill

Date: Sat Feb 26, 1994 9:20 am PST
Subject: Functional mapping; basics of PCT

[From Bill Powers (940226.0800 MST)] Bill Leach (940225.2112)

> I am quite happy to say that I believe that I fully understood Bill's posting. The functional mapping is no problem to me BECAUSE PCT really does map the functions and not just play games with the terms!

Why haven't I ever thought of putting it that way? That's exactly what we try to do! We try to map the functions (and signals and physical variables) in a simple control system model to functions, signals, and physical variables in real behavior. Naturally, this requires some guessing about the functions and signals inside the real person, but that's the modeling game. From various indirect evidence, including personal and behavioral evidence but also taking into account what physiological and neurological information is available (that we can make any sense of), we try to deduce functions and signals that actually exist in the central nervous system. Of course we can directly observe and measure the physical variables and functions outside the system in an experiment, so we don't have to guess about that part of the model.

What we're always after is the _minimum_ model of a person's insides that would account for what we observe. Our models have to be taken as representative or equivalent models, not as literal maps: they accomplish the same overall result, but not necessarily in the same detailed way that the nervous system does it. However, if we propose an input function (for example) that receives position information and generates a neural signal indicating rotation rate, as in controlling the angular velocity of a rotating line on a screen, we are saying that there IS some input function in the real nervous system capable of computing rotation rate given a series of position signals, even if we don't know how it works. The model contains a number of black boxes labeled according to the input-output relationship they have to establish, but we don't try to guess at the inner structure of the black boxes. As you know there is an infinite number of ways to construct a computing network that will create a given input-output function. The only way to find out what the actual wiring is is to open the black box and look. So far, none of us is licensed to do that.

As a linguist on the net in Australia, Avery Andrews, said, what we're doing is reverse-engineering behavior.

One of the most useful pieces of evidence is perception. A major postulate is that when we consciously experience the world or our bodies, we are being aware of the content of the perceptual signals emitted by input functions at various levels in a hierarchy of perceptions in our brains. In other words, if I see a cubical paperweight on the table in front of me, I am being directly aware of perceptual signals representing all the attributes of what I see: the cube-ness, the table, the "on- ness", the "in-front-ness", the "paperweight-ness", and so on -- a large number of perceptual signals each standing for some function of the basic inputs to my visual receptors. Most of these perceptual signals can be controlled by acting on the environment.

If this is really what these experiences signify, then it should be possible to deduce some relationships in the hierarchy of perceptions by examining all these perceptual signals to see how they depend on each other. Some perceptions should be functions of others: the dependent perceptions should change when the ones they depend on change, and the dependent perceptions should be nonexistent when the ones they depend on don't exist. Controlling a dependent perception should require, absolutely, varying the perceptions on which it depends. If a perception is under control, disturbing one of the perceptions on which it depends should result in behavior that either opposes that change directly, or results in changing one of the other perceptions so that the dependent perception remains (nearly) the same.

It is possible to find classes of perception that show these relationships -- that's what my proposed 11 levels of perception and control are about. For example, if you examine "configurations" (level 3, roughly objects), and try to find perceptions on which they depend but which are not just smaller configurations, you come up with shading, color, edge, curvature, and so on -- the class of perceptions I call "sensations." And if you look at any one sensation, you find that it depends only on the amounts, magnitudes, of stimulation of given kinds -- the level I call "intensities," the bottom level. Only the bottom level is directly affected by stimulation from the environment. And going upward, you can ask what kind of perception depends on the existence of configurations (looking for the least possible step), and you come up with, I think, "transitions," which includes motion. Then, seemingly, continuing upward, you get events, relationships, categories, sequence or ordering, programs (rule-following), principles, and system concepts. Each level depends on the levels below. Similar levels can be found in all sensory modalities.

All this is very much subject to revisions, deletions, and additions, and of course the whole thing needs experimental testing. The nice thing about PCT is that you can actually test for control of perceptions, and thus find out if the perception really exists (from the experimenter's point of view). If you want to know whether a person can control the state of a logical proposition, for example, (the 9th level as currently numbered) you make the truth-value of the proposition depend on some lower- level perceptions, ask the subject to keep the proposition true (or false), and disturb the lower-level perceptions to see if the person's action opposes the effect of the disturbance either directly or through altering other perceptions (other variables in the proposition). If the person can control the state of the logical proposition, then obviously the person must be perceiving it. Ultimately, this sort of testing, done for perceptions at all levels, should enable us to sort out what the levels really are and how they are really related, if in fact they exist at all. The hierarchical control model is set up so that all loops, at every level, are closed through the environment. So you can always test for control at any level through a behavioral experiment.

Before that sort of major project can be undertaken, PCT has to become airborne enough so that we don't have to spend all our time just defending our existence and overcoming the massive opposition to expending funds on this silly, narrow, old-fashioned, elementary, simple-minded, unscholarly, unintellectual, naive engineering approach to understanding behavior. This is not a small project and it can't be done by four or five people scattered around the country doing experiments in their spare time on themselves and their families with home computers and no money, while making a living doing something else. As the nucleus of a scientific revolution, PCT has developed far enough to suggest numerous lines of research. The main problem now is for competent understanding of it to spread far enough that this sort of research becomes respectable and fundable.

As funding is a zero-sum game, the implication is that other lines of research will have to be perceived as less promising, so funding can be shifted appropriately. That, of course, is the BIG POLITICAL PROBLEM. Nobody following another line of investigation wants to be the one whose pet theory is deemed less promising than PCT. For the threatened, there are two solutions to that problem: (1) fight like hell to show that PCT is (a) wrong, (b) trivial, or (c) just a subset of what you are already doing, or (2) adopt PCT and join the revolution. So far even on this net and more so elsewhere there is a lot more of (1) than (2). One way to judge the power of PCT is to look at the energy expended in defending against it (that's just good control theory).

I think that any disinterested party who has been following csg-l for two or three years must have noticed how many _different_ ideas are ranged against PCT. Practically every approach to human behavior, it seems, is disturbed by the statements that PCT generates. Objections to PCT arise from every quarter, even quarters that staunchly disagree with each other. Practically every objector also claims that the basic concepts of PCT were anticipated in his field or are contained in it. This is a very interesting circumstance, because it says that PCT is relevant enough to many different fields to have some of its conclusions accepted by them, and also that it differs from them enough to cause considerable pain. One might conclude that PCT agrees with what is right in these fields, and disagrees with what is wrong in them. But of course a single proponent of a single field of investigation sees the situation only from one point of view, claiming the agreeable parts of PCT and rejecting the parts that seem flawed. The fact that a person in a different field might accept and reject different parts of it isn't evident unless you're paying attention to ALL the objections, as we who sit here in the middle necessarily do.

Well, Bill L., you got back more than you asked for, but I have a hunch that you're about to join us, so maybe all this is relevant.

Best, Bill P.

Date: Sat Feb 26, 1994 10:24 am PST
Subject: Feedback

[From Bill Powers (940226.1000 MST)] Bill Leach (940225.2123)

In reply to your second post (the one to me):

> The entire perception signal is considered feedback, yes? A disturbance then, if perceived becomes a _part_ of the feedback, yes?

What we call "feedback" is pretty indefinite. Probably the best way to think of this is to divide the control loop into two parts, the forward part and the feedback part. The forward part is the path from the sensors, through the organism, to the output effectors. The feedback part is the path from the effectors, through the environment, to the sensors. By thinking in terms of path segments instead of signals, you can make the division clearer.

We divide the external feedback path into several parts: the output quantity, which is the immediate physical effect of the effectors on the environment, the "feedback function" which converts that output effect into an effect on the controlled quantity, and the controlled quantity itself which is directly sensed by the input sensors. The input sensors and the output effectors sort of straddle the system-environment boundary.

A disturbance comes into the external part of the loop from some independent source. It can affect any part of the feedback path, but we always express it as an equivalent disturbance applied directly to the controlled variable. Multiple disturbances, even if they affect different parts of the feedback path, are similarly expressed, as a single equivalent disturbance applied to the controlled variable.

The perceptual signal, which is the output of the sensor and which indicates only the state of the controlled variable, is affected by disturbances that change the controlled variable, but it is also affected by the output of the system via the feedback function, which also can change the controlled variable. So you can't say that the perceptual signal is a perception of the disturbance. It is a perception of the state of the controlled variable, period. There's no way to tell how much of the perceptual signal is due to the output and how much to the disturbance, especially since the output opposes the disturbance. One of the neatest counterintuitive facts about control is that the state of the controlled quantity (and hence the perceptual signal in the model) shows a very LOW correlation with both the output and the disturbance, while the output shows a very HIGH negative correlation with the disturbance. Of course the state of the controlled variable shows a very high positive correlation with changes in the reference signal, in the model.

Note that the disturbance itself is NOT sensed.

> Thus, in the skater example: If we assume that the coach's words are a significant factor in the perception of the subject, then calling what the coach is doing "feedback" is NOT really out of line.

If, as I proposed before, the coach is just calling out angles of bend and the skater knows that the reference angle is 90 degrees, then the coach is acting as a feedback function, converting the output of the skater (the angle of bend) into a string of verbal numbers, 86, 88, 91, 95, 93, 90, 90, etc. that the skater is perceiving and controlling by varying his angle of bend. The coach is then part of the skater's feedback loop.

If the coach calls out "86, 88, 101, no I mean 91, 95 ...", the coach has now injected a disturbance into the external part of the skater's control loop. The spurious "101" was not generated by the skater's varying the bend angle, but by a mistake that the coach made. This mistake caused a sudden change in the controlled variable which was independent of the skater's actual angle of bend. The critical thing about a disturbance is that it tends to alter the controlled variable in a way that is independent of the control system's own effect on the same variable.

When you understand the basic organization of a behavioral control system, the words you use to describe it are only mildly important. A listener who also understands can correctly interpret even loose usages. If you say the skater is depending on "feedback from the coach," you know that this means the coach is responding to what the skater does in a regular way, so the cause of the coach's feedback is, reliably, the skater's action. But if the coach spontaneously offers advice like "use your head, use your head," this provides no regular or understandable perception to the skater of the skater's own actions, and is just a disturbance.

> Now then calling it "positive" or "negative" feedback would be wrong as far as PCT is concerned BECAUSE the in the sense that such modifiers are typically used the paradigm is not PCT.

Remember that PCT is based on real control theory, not on verbal conventions. This means that "positive feedback" already has a meaning, the meaning defined long before the words were taken up by laymen to mean "encouragement." Positive feedback results from a wrong sign in the control loop so that a small error makes the error larger instead of (negative feedback) smaller. To compute the sign of feedback, you multiply together all the signs associated with all the functions (including the comparator) encountered in one complete trip around the loop, starting anywhere. This number must be negative if the feedback is to be negative and therefore error-correcting. Positive and negative feedback are properties of the WHOLE LOOP, not just the external feedback path.

The lay usage of the terms came from association of "positive" with "good" and "negative" with "bad." In control theory, exactly the opposite connotation exists: for control, positive is bad and negative is good.

You are feeling depressed and, hoping to hear something that will make you feel better, you tell me "I'm really feeling terrible today." I can create a positive feedback situation by saying something like "Why are you always whining to other people about your problems?" If I respond to everything you say with words calculated to make you feel worse, your actions based on your feelings of error will simply make the error larger. That is positive feedback. It has nothing to do with what I say, but only with the _relationship_ between what you feel now and what you feel next after one trip all the way around the loop.

You can't tell whether a response to someone's words amounts to positive or negative feedback without knowing how the person is hooked up inside. If you're talking to a masochist, the above response to the above complaint will be highly satisfying, being exactly right for enabling that person to accomplish the goal of feeling rotten -- you're completing a negative feedback loop.

So the strict usage of positive and negative feedback turns out to be correct, and the lay usage is wrong, for the lay usage assumes that encouraging words will necessarily help the other person in some way. If you really understand what positive and negative feedback are, you'll wait to find out what kind of response the other person is trying to get, instead of assuming that everyone with a complaint wants to be bucked up. The problem with picking up jargon and free-associating on it is that you are likely to misunderstand the situation because of your misinterpretation, and thus behave inappropriately.

Just remember: if a person wants encouragement, and says "I'd like to be a better person," and you say "I think you are much better today than yesterday," that is acting to create a NEGATIVE feedback loop.

Best, Bill P.

Date: Mon Feb 28, 1994 5:49 pm PST
Subject: Re: Functional mapping; basics of PCT

<[Bill Leach 940228.19:14)] >[Bill Powers (940226.0800 MST)]

Bill;

I have been "getting more than I bargained for" since I first arrived here. :-) And I don't mind it a bit.

As is usual for about anything that I have seen you post, there is plenty there and a great deal to think about. Thank you again.

-bill

Date: Mon Feb 28, 1994 8:12 pm PST
Subject: Re: Feedback

<[Bill Leach 940228.21:32)] > Tom Bourbon [940228.1533]

Thank you Tom. What Bill said need no clarification to answer my direct question about feedback. However, your posting does clarify any question I might have had about other variables (unasked).

-bill