

Paper presented By William T. Powers
at the CSG conference in Durango, July 1993.

This paper evolved in the spring of 1993 as a collaborative effort by CSG netters to identify misunderstandings and "myths" by "devils advocates" among people in academe who have misinterpreted the fundamentals of Control Theory and rejected manuscripts submitted for publication in scientific and psychological journals. What follows is version 4 in this effort, crafted by Bill Powers. The supporting collection of misleading quotations in the literature is large. See file DevilsBibliography.pdf.

=====

THE DISPUTE OVER CONTROL THEORY
William T. Powers and
The Editorial Board of the Control Systems Group

Feedback control has been seen as a central concept in the behavioral sciences for five decades. But its actual nature has been widely misunderstood, and for this reason its potential and significance have been seriously underestimated, especially since the mid-to-late 1970s.

This article is intended to set the record straight on the content and claims of control theory in general, and PCT or perceptual control theory in particular, in the context of the behavioral sciences. While the particular adaptation of control theory to be described here has attracted many supporters over its 40-year history, it still represents a minority position. A great many readers of refereed journals know it primarily through the way it is represented by critics -- or in some cases by would-be supporters -- who have grasped its basic approach to understanding behavior less than completely.

Perceptual control theory stands in a peculiar relationship both to traditional disciplines in the behavioral sciences and to branches of engineering which have explored applications of engineering control theory to the behavior of organisms. Control theory itself, under any label, offers a way of analyzing behavior that has been unknown in the behavioral sciences during most of their history. Perceptual control theory involves some rearrangements and reinterpretations of the engineering models, specifically designed to facilitate the application of control theory to living systems. The result is that behavioral scientists and engineers alike find objections to PCT, but of totally different kinds. Objections from behavioral scientists focus on departures from traditional ways of interpreting behavior, while those from engineers focus on the unfamiliar ways of representing control systems that are necessary to match the model to an organism. We will try here to thread our way between the shoals on the one side and the rocks on the other, and show why at least some students of behavior consider PCT to be as valid an approach as any others that have been offered.

THE ORGANIZATION OF A CONTROL SYSTEM

The first thing we must do is to correct several wrong impressions created long ago by the too-literal adoption by psychologists of an engineering diagram of a typical control system. Fig. 1 is taken from (_____), but has appeared in many other publications going all the way back to the cybernetician Norbert Wiener (1948, p. 132, Fig. 5) and the engineering psychologist _____.

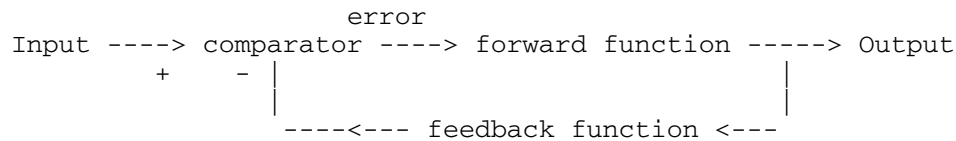


FIGURE 1

When a behavioral scientist sees the terms input and output, the natural translation is into *sensory* input and *behavioral* output. The above diagram thus seems to suggest that a control system is basically a stimulus-organism-response device, with a feedback loop added. Overall, it converts inputs into outputs in the normal cause-effect way.

However, in the engineering diagram the input is *not* a sensory input. It is a *reference* input, the means by which the user of a control system can set the desired value of the output. Feedback action brings the output, or a sensory representation of it, to a match with the setting of the reference input. The actual sensors are not even shown in this diagram; they are the means by which the output of the system is sensed and converted into an internal signal, the one entering the feedback function.

This diagram is also misleading as to the meaning of "output." In an artificial control system, the output is the variable that the customer wants to be controlled. But the state of this so-called output is not, in general, an immediate consequence of the effector action of the system. In a household temperature control system, the controlled variable is not the heat energy being output by the furnace into the ducts, but the temperature of a room somewhere far from the furnace. In the exposure control system of a camera, the controlled variable is not the opening of the iris diaphragm or the shaft speed of the actuating motor, but the amount of light falling on the photocell and the film. The physical output of a control system's effector is seldom identical to the quantity that is under control. Instead, the effector output is linked to the controlled variable through some physical process which may be quite indirect and involve changes from one kind of physical variable (iris area, energy output) to another (light intensity, temperature).

Therefore what the engineer calls the output of a control system is only in special cases the same as the physical output of the system's effector. The engineer simply moves the definition of the output to the position of the controlled variable, and measures it in units appropriate to the controlled variable. This, however, is not at all apparent from Fig. 1.

Any non-engineer looking at Fig. 1 would assume that the output represents the physical, visible behavior of the system -- in an organism, the patterns of motor output that have effects in the environment. The assumption would be that these physical outputs are under control. Working under that assumption, and the previous one that the "input" on the left is a sensory input, a behavioral scientist might well wonder why so much is made of the feedback connection. At most, such a connection can only modify the stimulus-response law governing the effects of stimulus inputs on behavioral outputs.

Looking at just the output part of Fig. 1, let us add some details, shown in Fig. 2.

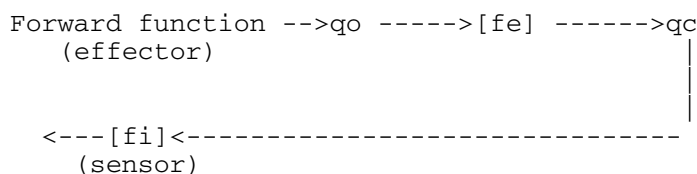


FIGURE 2

The forward function is, or contains, the effector of the control system. The immediate output of the effector is the state of some physical quantity such as a force or a torque. This is labeled q_o , for output quantity.

The output quantity is linked to the controlled variable q_c through some environmental path in which various physical laws come into play. This path is indicated as an environmental function f_e .

Finally, the controlled quantity q_c affects a sensor, or more generally an input function f_i (which can include a sensor and signal-processing computations). This is the feedback function of Fig. 1. The signal emitted by the input function connects to one input of the comparator of Fig. 1.

Fig. 2 is identical with the output part of Fig. 1 except for some details that remove ambiguities. Instead of showing the feedback path as starting vaguely in the vicinity of the "output" line, a specific physical variable is named as the controlled quantity and the feedback path represents its effects on a sensor. The "output" line is broken down into an effector output, q_o , and a place to put representations of physical laws, f_e , that connect that output to the controlled quantity.

Functionally, Fig. 2 indicates the same control system that Fig. 1 indicates. An engineer looking at the two diagrams would see the second as simply making explicit some details he or she normally takes for granted. But to many behavioral scientists, Fig. 2 will bring in some new considerations. We can make those new considerations even more explicit by completing and rearranging Fig. 2, as shown in Fig. 3.

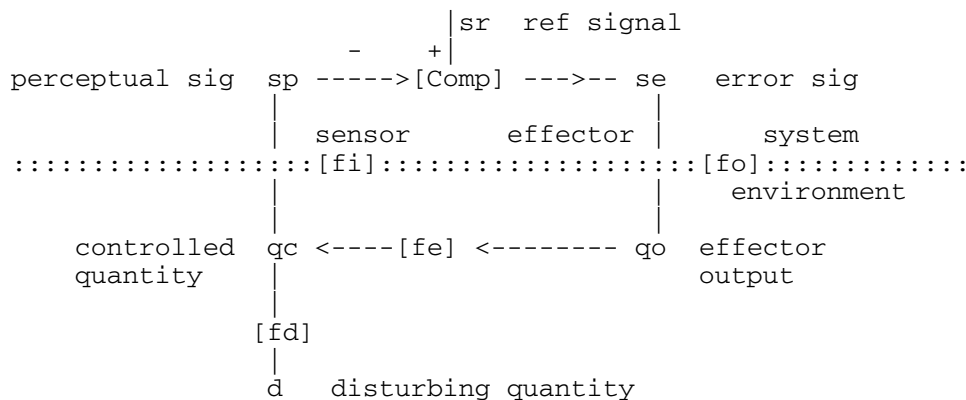


FIGURE 3

This figure 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, one would have to draw a vertical line through the forward and feedback functions, with the environment on the right and the active control system on the left.

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 effectors lie in 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 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. The "user" of this system who specifies the desired level of the controlled variable by adjusting the reference signal is not an external person, but a higher system internal to the organism.

If a behavioral scientist were looking for a place in these diagrams to introduce a stimulus, where would it be placed? Working from Fig. 1, the appropriate place would seem to be at the location labeled "input." But as seen in Fig. 3, that location can't be affected from outside the system, at least not directly. The stimulus input obviously must enter where the sensors are in Fig. 3. However, in Fig. 3 the sensors are already detecting the state of the controlled quantity, which in turn is already being influenced by the output quantity generated by the effector. The most the stimulus can do is add a contribution to the state of the sensor, which, from a point of view inside the system, is the same as influencing the controlled quantity. So the stimulus input becomes a disturbance of the controlled quantity. The external cause of the stimulus may or may not actually be sensed. As drawn in Fig. 3, it is not sensed: only the net result of the disturbance and effects of the system's effector output is sensed. This is the most common case.

The controlled quantity corresponds to a proximal stimulus, a physical variable in the immediate vicinity of the sensor. The disturbance is like a distal stimulus, a change in the environment that affects the system only through its influence on proximal stimuli. But these remote causes are not the only influences on the proximal stimulus: the physical action of the system itself also affects the proximal stimulus, the controlled quantity. This is the situation that holds for any control system. The controlled quantity, directly sensed by the system, is affected equally by remote events and by the actions of the control system itself. And the actions of the control system clearly depend, at least in part, on the state of the controlled quantity.

This is the famous "closed loop." In the relationship between the active system and its environment, there is a little vortex of causality without a beginning or an end. External influences from the environment do not affect just the sensory input to the system; they affect the causal loop. The reference signal coming from above does not just produce output; it enters the same causal loop in a different place, its effects adding to those from the input function.

It should be clear by now that a control system, properly diagrammed, is a new kind of organization as far as the behavioral sciences are concerned. The unit of organization is not simply a link connecting higher systems to behavioral outputs, as assumed in the cognitive sciences; it is not simply a link connecting sensory inputs to behavioral outputs, as assumed in behavioristic or other empirical input-output approaches. It is not any normal kind of causal process in which one can pick a starting point and follow a chain of events through to an ending point. The closed loop of causation gets in the way of any conventional kind of attempt to trace cause and effect.

Control theory is the body of mathematical methods that permits us to analyze and understand the behavior of closed-loop organizations like those in Fig. 3. It is *not* the theory that "organisms are servomechanisms." Any time that the relationships in Fig. 3 are found, anywhere, control theory is necessary to understand what is happening and to predict the behavior of the system-environment relationship.

Perceptual control theory, PCT, is the version of engineering control theory based specifically on the organization of Fig. 3. In that figure, it can be seen that a control system controls its own input, not its output. And in doing this, it makes a perceptual representation of that input -- a perception, for short -- match an internally-given reference signal that specifies the desired state of that perception.

This interpretation of behavior is not like any conventional one. Once understood, it seems to match the phenomena of behavior in an effortless way. Before the match can be seen, however, certain phenomena must be recognized. As is true for all theories, phenomena are shaped by theories as much as theories are shaped by phenomena.

THE PHENOMENON OF CONTROL

Traditional scientific approaches to understanding behavior have recognized two kinds of behavior which we will call stimulus-driven and command-driven. Stimulus-driven behavior we can define as changes in behavior that follow from changes in the surrounding environment. Command-driven behavior is a more mentalistic concept, in which at least certain kinds of complex behavior are commanded by internal processes not strictly or immediately dependent on the environment. Control theory, as we shall see, offers a third choice.

According to the concepts of stimulus-driven behavior, the actions of organisms are to be explained in terms of visible occurrences in their environments. Under a reflexological theory, stimuli are literally stimulations of sensory organs, with the resulting neural impulses being routed to the muscles to make them contract; complex behavior results from complex sets of stimuli that excite many reflexive responses at the same time. This was the initial concept that psychology inherited from biology and neurology.

Early behaviorism took a somewhat more global view, in which the basic reflexive picture was assumed, but the approach was more abstract. It was not practical to trace every effective stimulus to the specific stimulations of nerve-endings that it produced, so stimuli came to be defined as stimulus objects or events observed at some distance from the organism. Given the underlying reflexology, it could be assumed that if responses were observed when distal stimulus objects were manipulated, the appropriate proximal stimulus effects must have occurred too. At the response end, it was also impractical to get immersed in details. Rather than describing behavior in terms of individual muscle tensions, the practical approach required looking at larger consequences of those muscle tensions, the typical patterns they produced. If some such pattern were observed, it could be assumed to have followed from muscle tensions in some regular way. This more global concept of stimulus and response has had a long life.

Even this more global approach was too restrictive for many behavioral scientists. Human beings did not just react to simple objects and events by producing simple responses; they could react to quite abstract aspects of the environment such as complex situations, social influences, and events far in the past. Such environmental situations could create not only conditioned responses to current events, but longer-term ways of responding which could show up as habits, attitudes, preferences, complexes, conflicts, traits, tendencies, preferences, and biases. Such effects were seen not in specific acts but as overall patterns of action. This kind of concept has also had a long life.

All these versions of stimulus-driven behavior account for the bulk of what has been and still is being published about human nature. While proponents of the various versions may contest vigorously and sometimes bitterly with one another, they are all in agreement that the scientific way to study behavior is to trace out the ways in which the environment shapes and directs it. They disagree mainly about what aspects of the environment and of behavior are meaningful.

The phenomenon addressed by all these versions of stimulus-driven behavior is the observation that behavior is often closely and conditionally related to events in the environment, as effects are related to causes in the worlds of physics and chemistry.

The command-driven approach has less scientifically-reputable roots. In this area, the interest has always been in cognitive processes, for instance: consciousness, language, emotion, intelligence, logic, insight, and goal-seeking or purposiveness. While many proponents of this view have paid lip-service to the more "scientific" ideas across the aisle, the emphasis has not been on environmental causation but on the person as a conscious, interpreting, knowing, active agent. The phenomenon addressed in this approach

is the fact that many behaviors seem uncaused; even if one calls them "responses," there is no obviously corresponding "stimulus" to account for them. The appearance is that there is organized activity in the brain that can command behaviors independently of immediately antecedent inputs from the environment.

The computer revolution rescued this branch of theory from being orphaned from science by offering a link between phenomena of cognition and consciousness and the workings of computer models of brain functions. A way was found to handle "mental" phenomena in physical terms. Many psychologists have chosen not to take advantage of this potential link, regarding the psychological world as permanently separated from the physical one, but cognitive science as a whole now has at least the potential of working in the same universe as the rest of science.

Whatever the views of a particular cognitive scientist, there is one assumption held in common if the brain is thought to be involved at all: that complex brain activities are translated into action through commands originating high in the brain and being elaborated, step by step, until they become (primarily) commands for the tensing of muscles. Hence the classification "command-driven." Those patterned outputs produce, through ordinary physical laws, the global patterns we recognize as behaviors. On the output side, at least, the concept of command-driven behavior is in agreement with the concept of stimulus-driven behavior. Where they disagree most directly is in accounting for the immediate causes of behavior.

This somewhat superficial summary of two main branches of behavioral theory is not intended to give deep insights into either of them, but to set the stage upon which control theory appeared in the 1930s. The phenomenon of behavior that the new control engineers of the 1930s chose to investigate and model does not fit comfortably into either of these branches, although both branches have often laid claim to it. It is the phenomenon we see when a person controls something.

What the new control engineers saw people doing can be seen through an example. A person operates some piece of equipment, or just that person's own arms and legs, in a way that affects some variable aspect of the environment. This aspect is also affected by other influences, so the outcome is affected more or less equally by the person and by independent forces in the environment. The reason for the person's action is that the person wants to, or has been ordered to, maintain that aspect of the environment in some particular condition. "I want you to keep an eye on that gauge," the boiler attendant is told, "and adjust the burner to keep that steam pressure nailed at 300 pounds per square inch." Before control theory, there was no theory of behavior that could correctly explain how carrying out such a task is possible. This, of course, did not prevent explanations from being offered.

Explaining the phenomenon

From the standpoint of command-driven behavioral theory, the attendant's behavior is explained easily: the person adopts the goal of maintaining the gauge at 300 pounds per square inch, constructs a plan for varying the output action on the burner control, and carries it out. There seems to be little to explain except why the person is motivated to act in this way. How the task is actually carried out is a matter for physiologists to explain.

The answer is almost as easy from the stimulus-driven point of view. The person is conditioned to react to fluctuations of the gauge by moving the hand on the lever or knob that changes the burner heat output. Or: The person is reinforced by receiving a paycheck for responding to the discriminative stimulus indicated by the supervisor's verbal order in a way that keeps the gauge at a steady reading, that steady reading eventually becoming a secondary reinforcer. Or: The person is influenced by social pressure to do a good job and please his superiors. Or perhaps: the person is a member of an oppressed class forced by the capitalistic system to bow to the orders of others and

engage in this demeaning task. There are many ways to explain behavior by pointing to something going on in the environment.

None of these answers would have helped the early control engineers, because they were trying to build a device that could do what the person was observed to be doing. The engineers didn't need to be told that the person was actually doing the task. They didn't want to know about antecedent causes or motivations or social influences. They just wanted to understand how anyone could do what that person was observed to be doing, whatever the reason for doing it and whatever the person's or society's attitude toward doing it. As it turned out, nobody had ever before figured out how a person could actually do such a task, although some, like William James and John Dewey, had guessed roughly what the right answer would be.

There were three puzzles to be solved (as we can now see the problem).

First, the causes of disturbances of the variable to be controlled were invisible. The steam pressure shown on the gauge would fall if somewhere in another room or building someone turned on a machine that used steam pressure, or if the energy content of the fuel dropped, or if the line voltage dropped and slowed the furnace's blower, or for any number of other mysterious and unknown reasons. The invisibility of the causes of disturbances made no difference in the attendant's ability to keep the steam pressure at the requested level.

Second, the steam pressure was indicated on the gauge, but because different steam pressures might be required at different times, there was no indication of the right steam pressure. All the gauge did was report the actual pressure; its needle did not also indicate the pressure to be maintained. The only input corresponding to the right pressure was in the initial instructions, which occurred only once. From then on, no input corresponded to the right steam pressure. And there was no input to specify what actions would result in this end.

Third, the steam pressure, although it was certainly sensed by the attendant, was also in part caused by the attendant's actions. A chain of causation could be traced from the gauge reading, into the attendant's eyes, through some hypothetical connections in the attendant's brain (which the engineers had to simulate), through the attendant's muscles, to the burner control, and back to the pressure reading on the gauge. The chain formed a closed loop.

This third part of the puzzle was the critical part. How could such a closed loop of causation be analyzed? The gauge reading, in order to be called a stimulus, had to be independent of behavior. To be called a response, it had to be dependent on the motor outputs of the attendant. It could not, according to traditional thought, be both at the same time.

The stimulus-driven approach necessarily would have to find a way to separate stimulus from response. In similar situations, this was usually done by separating them in time. It would be assumed that first a fluctuation of the gauge reading occurs, which leads to a series of events that results in a movement of the attendant's hand and a change in the heat output of the burner. That change alters the steam pressure in a way that, one hopes, is opposed to the original fluctuation. Then the cycle can begin again. This at least sounds like a plausible analysis.

The command-driven explanation would have an even harder time with this closed loop of causation. In order to formulate a command that would oppose the fluctuation in the gauge reading, a cognitive system would have to know about all potential causes of disturbances. But none of the multiple causes of pressure fluctuations can be sensed. There is no one command that can be sent to the muscles that will result in opposing an unpredictable fluctuation of the gauge reading. A top-down command-driven system can't handle this situation at all.

The control engineers used neither type of analysis. Instead, using techniques well-known in their underlying disciplines, they first characterized each

subprocess with a descriptive input-output equation, and then *they solved all the equations as a simultaneous set*. Sequential cause and effect never entered the picture. Neither plans nor commands were involved. The rule was simply that each variable could have only one value at a time. No matter how one variable depended on others, whether delays or integrative lags were involved, whether there was amplification in one part of the loop and losses in another, whether static or dynamic relationships were involved, all the variables in the system had to satisfy all the equations that pertained to them at a single instant -- and at every instant.

That approach, in a nutshell, is control theory.

Comparing the three models

Control theory was expressed by engineers in diagrams like Fig. 1, which, topologically transformed, becomes Fig. 3 in PCT. In Fig. 3 two pathways can be identified, as indicated in Figs. 4a and 4b.

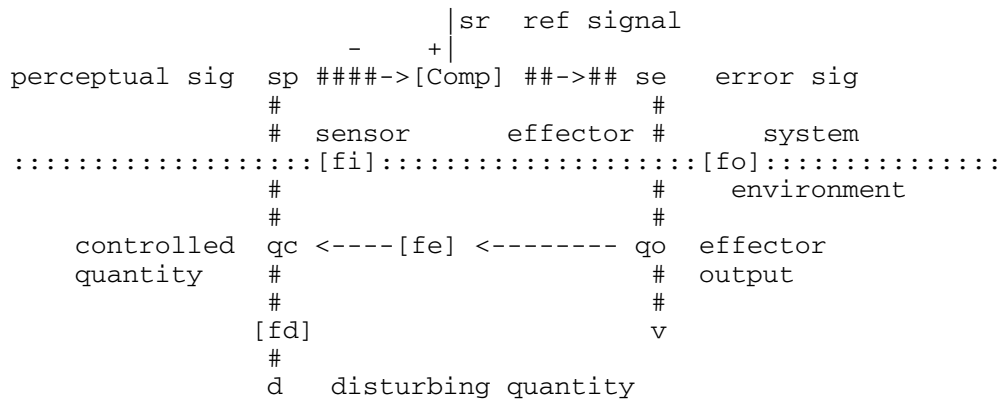


FIGURE 4a

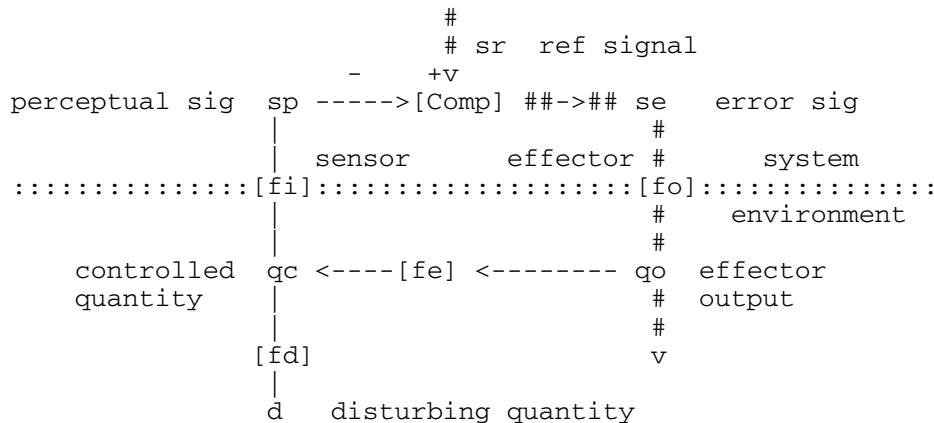


FIGURE 4b

In Fig. 4a, we see a shaded path (#####) starting with the disturbing quantity, going through the controlled quantity at the input, running through the perceptual function into the comparator, coming out of the comparator and going through the output function, and ending in the output quantity. This is the pathway equivalent to those envisioned under the various versions of stimulus-driven behavior. The disturbance plays the role of a distal stimulus, the controlled quantity the role of a proximal stimulus, and the output quantity the role of the motor behavior (simple or complex) that results. The part of the pathway inside the system (above the line) corresponds to well-known neural pathways.

One link in the environment is not on this pathway, the one from the output quantity, through f_e , to the controlled quantity. Recognition of this pathway is an afterthought in command-driven theories if it appears at all. It appeared in the middle years as the concept of response chaining, or as reinforcing consequences of actions. The function f_e would correspond to a "contingency of reinforcement" in Skinnerian theory. It was not, however, considered an integral part of behavior: the shaded path was the primary path and all other effects were secondary or consequential.

Fig. 4b shows a shaded path (#####) starting with the reference signal descending from above, then passing through the comparator and the output function to the output quantity. This is the path envisioned in command-driven theories. Somewhere above this diagram lie the cognitive systems that formulate and generate commands that follow the shaded path. A second path is added originating somewhere in the environment and rising toward the cognitive systems; this path represents informational inputs that form the basis for cognitions. That second path would not influence the downgoing path, so it is not shown following any existing path in Fig. 4b. If we actually drew a box at the top and labeled it "cognition," the result would look suspiciously like Fig. 4a. In this diagram, too, the feedback path through f_e is an afterthought, indicating effects of one commanded action on the input situation leading to the next command.

Both stimulus-driven and command-driven theories treat the closed loop of Fig. 3 either by ignoring it or by trying to split it into separately operating parts. Control theory encompasses both paths, and makes the external feedback connection a concurrent part of the closed loop.

Under the control-theoretic analysis, this diagram is seen as one single entity, the closed loop, with one independent input in the form of a reference signal coming from above and another in the form of environmental disturbances impinging on the loop from outside. The signal from above does not act on the outputs of the behaving system, but on the closed loop as a whole. The disturbances, likewise, act on the whole closed loop, not on inputs to the system.

Each of the two traditional concepts recognizes one of the two independent variables and ignores the other. Cognitive theories recognize that command signals are generated independently by higher systems, but fail to recognize that the output is subject to disturbance from outside the system. Hence cognitive theories assume that the command signal produces a corresponding outcome in a disturbance-free environment. Stimulus-driven theories recognize independent variables in the environment that act on the system, but fail to recognize that the resulting behavioral outputs may also be affected by signals generated independently inside the behaving system. Thus stimulus-driven models assume that the output is determined solely by the stimulus input. Also, they do not consider the possibility of independent disturbances acting directly on the behavioral outcome; disturbances are identified as inputs to the sensory interface, and the outputs are the sole determinants of observable behavior.

Neither approach reveals the external feedback path as creating effects concurrent with both independent variables, effects that greatly modify their assumed effects.

We can now see that the PCT model acts as a synthesis of the concepts of command-driven and stimulus-driven behavior, showing how each one corresponds to one aspect of a control system. The synthesis, as often happens, shows both what is right and what is wrong with the older ideas, and brings out new considerations never covered by either one. The situation is similar to that of the blind men and the elephant (Marken, 1992).

HPCT

The diagram of Fig. 3 represents a unit of behavioral organization; it is not the whole behaving system. This unit can be duplicated and organized into levels of control, a hierarchical control model referred to as hierarchical perceptual control theory, or HPCT.

Some idea of the construction of a hierarchical model can be seen by looking (in a simplified way) at the organization of motor behavior. The lowest level of control in a human being consists of spinal reflexes involved in muscle force generation. The input functions correspond to sensors embedded in muscles and tendons, which monitor both muscle length and forces applied across joints. The signals from these sensors travel to the spinal cord where they synapse with motor cells, the length signals with a positive sign and the force signals with a negative sign. The motor cells also receive signals from higher in the nervous system; these correspond to reference signals, and the spinal motor cells play the role of comparators. The output of the spinal cells is the error signal, which enters an output function made of the contractile elements of the muscles. The environmental feedback path is composed of the series spring components of the muscle and the mechanical laws that convert contractions into forces on the tendons. Those forces tend to swing the limb segments about the joints, and this alters the muscle lengths, producing physical effects that alter the stimulation of the muscle length receptors.

Muscle receptors are also stimulated by effects of minute muscles in the length-detecting spindle cells. Thus the length signals are really error signals: they represent the difference between actual muscle length and the length of the spindle muscles, set by gamma efferent signals from higher systems. The muscle spindle is actually a combined mechanical input function and neuromechanical comparator. We can thus discern a two-level control system, shown in Fig. 5.

[Insert Fig. 5 about here]

[Footnote: The actual arrangement is somewhat different from the one shown in Fig. 5, but is functionally equivalent to it. In the real system, there are independent alpha reference signals entering the force control system; these become effective mainly when the limb is physically prevented from moving and the controlled variable becomes applied force. With the limb free to move, the alpha and gamma reference signals have equivalent functions.

The lowest level, shown as Loop 1, controls sensed force, which is equivalent to torque about a joint. The sensed force is compared with the reference force by the spinal neuron, and the difference operates the muscle. This makes the actual force relatively independent of changes in the response of the muscle to driving signals, and also independent of other sources of force variation such as inertial effects from limb motions. To some extent this control system makes the generated force independent of joint angle and limb motion.

The second level, Loop 2, detects a mechanical effect on muscle length of the forces controlled at the first level. This effect is produced by any limb motions that the forces create. The mechanical comparator compares that effect with a reference effect specified by a reference signal, and the resulting error signal enters the spinal motor neuron as a *reference signal for the force-control system*. Disturbances can enter this second level as mechanical disturbances of the limb and as loads carried by the limb. The control action renders the sensed muscle length, and hence joint angle, relatively independent of such disturbances.

Because of the way we have drawn the boundary between the behaving neural system and its environment, both control loops pass through the environment. The force-control system is the shortest and simplest loop, the controlled quantity being the physical force generated in a tendon. The muscle-length control system is a larger loop in which an effect of the controlled force

operates through laws of mechanics and physical dynamics to alter the controlled quantity of the higher system, the sensed muscle length (which corresponds roughly to joint angle). This sensed muscle length is compared with a reference muscle length; the resulting error signal enters the second-level output function, *which consists of the entire first-level control system*. The second-level loop is also closed through the environment, now through a slightly less direct path that brings in more global properties of the physical world. A single second-level control system may alter reference signals in several first-level control systems, particularly in those employing opposing muscles. All those first-level systems comprise the second-level output function.

Now, skipping a few levels, consider how the muscle-length control system is used in visual-motor coordination, a task like reaching out a finger to touch a target. The controlled quantity is a spatial relationship between finger and target, sensed visually. The perceptual signal indicates the finger-target relative distances in three dimensions. If the goal is to touch the target, the reference signals for these three perceptions will all be set to zero (setting them to a nonzero magnitude would indicate that some non-zero target-finger distance is to be brought about and maintained). After the required comparison stage, the signals representing the three errors are routed to all the relevant sets of two-level control systems as shown in Fig. 6, with appropriate signs, altering the reference signals for muscle length and thus, through the action of the force-control systems, altering the configuration of the limb segments and the position of the finger relative to the target. The output of the third-level system thus acts through a complex output function consisting of two levels of kinesthetic control (with many systems operating in parallel) to bring the visual perception closer to the reference state in the three specified dimensions.

So the visual control systems, too, involve a closed loop passing through the environment. We can easily extend this layering. For perceived finger position we could substitute the perceived position of the tip of a pencil held in a hand. Varying the position reference signals could result in moving the tip of the pencil to make loops and lines; the varying reference signals could be the outputs of higher control systems concerned with forming letters. The reference signals specifying letters to be perceived could be the outputs of systems controlling word-perceptions, sentence-perceptions, perceptions of grammar and syntax, and so on as far as one can find reasonable need for more layers of control. And in every instance, the control loop would pass through the environment, where evidence of its operation can be seen. Disturbances of various kinds would be resisted by systems at the appropriate level, through automatic variation of reference signals for lower level systems.

This is an essential aspect of the HPCT approach. Any specific HPCT model is falsifiable, because every layer of the model controls specific variables in the publicly-observable environment, and each type of actually-observed controlled variable must show resistance to disturbance and hierarchical relations to more global variables that match the same phenomena as predicted by the model.

The ideal of a working multilayer HPCT model that matches many levels of real behavioral organization is still only a distant prospect. But the experimental methodology is clear, and rough fits of this model to behavior are easy to find. A model of the above-described operation of visual-motor pointing behavior has actually been simulated on a computer, including realistic muscle properties and physical arm dynamics, and it does indeed behave quite realistically (Powers, 1992).

The basic organization of the HPCT model has higher systems acting strictly by varying reference signals for lower systems. It is possible for other inter-level modes of control to exist; for example, control through variation of parameters. Also, provision has to be made for reorganizing processes that account for the acquisition of new control systems and modification of old ones. All these subjects are somewhat peripheral to HPCT; they are considered

as subjects for future work, when application of the basic model shows the need for organizations outside its scope.

Building a hierarchical model that works is clearly a huge undertaking. However, the model of Fig. 3 can be applied to behaviors at arbitrary levels of organization, without the need to specify the lower-level systems employed for output or the sources of reference signals. Tasks can be set up in which reference signals remain reasonably constant, and parameters of a working model can be adjusted for a fit to real behavior. This sort of approach has been used, in an exploratory fashion, in many experiments. The predictivity of simulations constructed in this way is extremely high; correlations of modeled to actual behavior higher than 0.99 are common. While the behaviors involved are quite simple, the facts discovered in this way are of very high quality, their probability of truth being high enough -- millions to one in favor -- for use in deductive arguments. For examples, see (Refs)

PCT experimenters, to keep their spirits up, sometimes like to compare themselves with Galileo rolling little balls down inclined planes. The results may not be obviously earth-shaking, but they are highly reproducible and highly predictive of natural phenomena. There is no telling what sort of science might arise from accumulating such simple but utterly reliable facts.

MISAPPREHENSIONS AND MISSTATEMENTS

The introduction of control theory to the behavioral sciences has been gradual and spotty. In part this has been the result of its proponents' only gradually coming to understand its full meaning with respect to living systems. But a very important component of the problem has been the difficulty that conventional behavioral scientists have had in grasping the fundamentals of control theory and seeing how control theory differs from more established interpretations of behavior.

As indicated in Figs. 4a and 4b, there is enough overlap between the concept of a control system and more conventional concepts to provide tempting opportunities for assimilating the PCT model into older frameworks. Compounding the difficulty has been a surprising tendency for scientists who are normally careful to know what they are talking about to leap to intuitive conclusions about the properties and capabilities of control systems, without first having become personally acquainted with the existing state of the art. In many cases there is a strong suggestion of defensiveness in the misinterpretations, as normal sequential or causal analysis is used to show that control systems either can't work or else work according to conventional principles. Neither assertion is true. The controversy over control theory in general and PCT in particular has involved factors not strictly of a scientific nature.

We will deal now with some of the major misinterpretations and misstatements about control theory that have appeared in the refereed literature. In citing specific examples, we do not mean to blame specific authors for the errors; most of these errors originated long ago and have been propagated by hearsay, attaining the force of myths elevated to the status of facts. If any criticism is warranted, it is for promulgating statements with an authoritative air without having verified personally that they are justified. Most of the mistakes we will cite are common and understandable; most beginning students of control theory go through the same process of trying to make the principles of control fit into the causal world with which they are familiar. But beginning students of control theory do not publish their guesses.

Some of the points to be made below will concern the basics of control theory. Errors at this level will simply be corrected because they indicate a misunderstanding of the basic idea. Other will concern misrepresentations of PCT as it applies to living systems. Here there is no assertion that PCT is necessarily the best or only version of control theory or even the best or only explanation of behavior. The point here is that if arguments against PCT

are to be published, they should deal with what PCT actually says and not with a misrepresentation of it.

[HERE ENDS THE PAPER AS OF JULY 24, 1993. THE REMAINING SECTION WILL BE CONSTRUCTED FROM EXAMPLES IN THE LITERATURE CONTRIBUTED BY CSG MEMBERS AND PARTICIPANTS IN CSGNET.]

=====

See file DEVILS.BIB for the posts where such examples in the literature were discussed. This topic was also discussed as the first item on the agenda during the 1993 CSG conference in Durango.