

An Introduction to Perceptual Control Theory

Standing at the Crossroads

Paper distributed at the Control Systems Group Meeting,
August 15-19, 1990, at Indiana, Pennsylvania.

By William T. Powers

[*Comments by Dag Forssell, January 1, 1997]

I'd like to try today to give you the sense that psychology is standing at a crossroads—and not only psychology, but all the sciences of life. We are about to experience the advent of something for which many people have searched, an organizing scheme that pulls together all the disparate schools of thought, specializations, movements, and evanescent fads that make up various fragmented branches of the life sciences.

The organizing scheme is called “Perceptual* Control Theory.” This theory explains a phenomenon, as theories are supposed to do. The phenomenon in question is called control. Everyone has heard this word, and most people have occasion to use it from time to time, but in science it has become part of the metalanguage rather than designating a subject of study. A scientist does a control experiment, or demonstrates how manipulation of stimuli and rewards can control an animal's movements, or advocates a proper diet to control cholesterol level or competes for control of a department. This word is used as part of a background of ordinary language, but it has not been part of the technical language of the life sciences.

[* The word Perceptual was added at a later conference to distinguish Bill Powers's creation from competing, non-functional interpretations by other authors. This term is technically more precise, since all control systems actually control their perceptions, not their outputs.]

The reason is quite simple: nobody in or out of science understood the process of control until about the beginning of World War 2. By understanding the process, I mean being able to define it, characterize it,

measure its parameters, predict how it will proceed, and recognize it in a real system. This doesn't mean that control was impossible to accomplish before World War 2: after all, most people accomplish digestion without understanding any biochemistry. But control is as natural a process as digestion, and like digestion can be understood in a scientific way only by studying it and learning how it works.

World War 2 started only about 50 years ago. Perhaps you can see why this fact implies some problems with studying control as a natural process. If control is

a natural process, it was occurring in 1840, 1740, 1640, and so on back to the primordial ooze. In 1940, the sciences of life were already something like 300 years old (and their prehistory was far older than that). If nobody understood control until 1940, it's clear that these sciences went through a major part of their development without taking it into account.

The next question is obvious:

how did they explain the phenomena that arise from processes of control?

Many of the puzzles and controversies that occupied early researchers could have been resolved if scientists had realized that they were dealing with control processes. Purpose could have been studied scientifically instead of merely theologically. We can see now that all these early researchers, not recognizing a control process when they saw one, were drastically misled by some side-effects of control. The principal side-effect that deceived them resulted from the way control systems act in the presence of disturbances of the variables they control. When a disturbance occurs, a control system acts automatically to oppose the

We are about to experience the advent of something for which many people have searched, an organizing scheme that pulls together all the disparate schools of thought, specializations, movements, and evanescent fads that make up various fragmented branches of the life sciences.

incipient change in the controlled variable. But if this opposition is not recognized (it's not always obvious), the observer will inevitably be led to see the cause of the disturbance as a stimulus and the action opposing its effects as a response to the stimulus. Furthermore, this opposition results in stabilizing some aspect of the environment or organism-environment relationship. That stabilization conceals the role of the stabilized variable in behavior; the better the control, the lower will be the correlation between the controlled variable and the actions that stabilize it. The variable under control is the one that is actually being sensed, but the logic of control makes it seem that the disturbance is the sensory stimulus.

Donald T. Campbell [Late Professor of Psychology, Lehigh University] has proposed a "fish-scale" metaphor of scientific progress. Each worker constructs just one small scale that overlaps those already laid down by others. Eventually the whole fish will be covered completely. But what if the fish is a red herring? Then all these patient workers will devote their lives to covering the wrong fish. The converse of the fish-scale metaphor is that a person who is concentrating on fitting one little scale to others already laid down is bound to have a very localized view of the problem. Seeking to extend the accomplishments of others, a single worker can make what seems to be progress—but it is unlikely that a single worker will discover that something is wrong with the whole design. The result can easily be the diligent application of fish-scales to a giraffe.

I submit that something like this has happened in the life sciences. A fundamental misconception of the nature of behavior, natural but nevertheless horrendous, has pointed the life sciences down the wrong trail. Nearly all life scientists, particularly those who try to achieve objectivity and uniform methodology, have interpreted behavior as if it were caused by events outside an organism acting on a mechanism that merely responds. This hypothesis has become so ingrained that it is considered to be a basic philosophical principle of science. To explain behavior, one varies independent variables and records the ensuing actions; to analyze the data, one assumes a causal link from independent to dependent variable and calculates a correlation or computes a transfer function. This leads in turn to models of behaving systems in which inputs are transformed by hypothetical processes into motor outputs; those models lead to explorations of inner processes (as in neurology and biochemistry) predicated on the assumption that one is looking for

links in an input-output chain. One assumption leads to the next until a whole structure has been built up, one that governs our thinking at every level of analysis from the genetic to the cognitive.

Perceptual control theory, by showing us an alternative way of understanding this entire structure, therefore threatens the integrity of practically every bit of knowledge about behavior that has ever been set down on paper.

This is, of course, a message of the type that leads to a high mortality among messengers. That is why you are listening to a person with no reputation to lose and no fame to protect, instead of a Nobel Prize winner. In an utterly predictable way, scientists have for the last 50 years gone to great lengths to avoid learning control theory or else to assimilate it into the existing picture of behavior. Failing that, they have simply declared it irrelevant to their own fields, with the result that the authoritative literature of perceptual control theory is almost completely insulated from the mainstream. It appears in publications like proceedings of the Institute of Electrical Engineers division on Man, Machines, and Cybernetics, or in human factors and manual control publications, or in Xeroxed papers passed from hand to hand. There is a scattered literature on perceptual control theory in the life sciences, but nothing on this subject gets past the referees into a standard journal without first having its teeth pulled.

Despite all the defenses, the concepts of perceptual control theory are spreading. When our descendants look back on the latter half of the 20th Century, they will probably be amazed at the speed with which perceptual control theory became accepted: 50 years in the course of a science is nothing. We control theorists have nothing to complain about. Our greatest successes have come not through pounding at locked doors, but through continuing to explore the meaning of this new approach and learning how to apply it in many different disciplines. If we do our job correctly, acceptance will take care of itself. That job is not something one can toss off overnight, nor can it be done by just a handful of people. We are coming to a time of rigorous re-evaluation of all that is known or presumed to be known about the nature of organisms. The more people that are involved in this enormous project, the sooner it will be accomplished. That is why we are all so glad to welcome our guests at this session: after the party, you will be invited to help do the dishes.

There has been progress in understanding how organisms work, the wrong model notwithstanding. Biochemical reactions are not going to change because of perceptual control theory. Muscles and nerves will continue to operate as they are known to operate. Even at more abstract levels of analysis, many phenomena will continue to be accepted as valid observations; for example, phenomena of perception, of memory, of cognition. If competently observed, these phenomena will still be part of the legacy of earlier workers. When we pull the stopper on the old theories, we must keep a strainer over the drain and let only the bath water out.

Part of the task of reconstructing the sciences of life consists of separating valid observations of components from invalid conjectures about how they work together. Consider biochemistry as an example. Biochemistry is an odd mixture of solid research and wild leaps of undisciplined imagination. The research reveals chemical processes taking place in the microstructure of the body. The wild leaps propose that the chemical reactions somehow directly produce the behavioral effects with which they are associated. It's as though a specialist in solid-state physics were to propose that electrons flowing through wires and transistors are responsible for the music that comes out of a radio. While it's true that a shortage of electrons will make the music faint, and that without the electrons you wouldn't get any music, the physicist would be laughed out of town for suggesting that electrons cause music, or that you could fix a weak radio just by putting some more electrons into it. You can't understand the role of the electrons without grasping the principles of organization that make the radio different from a radio kit.

In the same way, if shortages or excesses of chemicals like enzymes and neurotransmitters are found to be associated with functional and behavioral disorders, all we then know is that these substances play some role in the operation of the whole system that creates organized behavior. If there's a shortage of some chemical substance, then some other system has reduced its production of that substance, and some other system still has decreased its effect on the driving system, and so on in chains and causal loops. Nothing in a system as complex as the human body happens in isolation. If biochemistry is to have anything to say about the organism at any higher level, biochemists are going to have to study whole systems, not isolated reactions. We need a functional theory to supplement the microscopic laws of chemistry.

There are workers in biochemistry who are investigating feedback control processes. One significant process involves an allosteric enzyme that is converted into an active form by the effect of one substance, and into an inactive form by the effect of another. When these two substances have the same concentration, the transition from active to inactive is balanced; the slightest imbalance of the substances causes a highly amplified offset toward the active or the inactive form. In one example, the active form catalyzes a main reaction, and the product of that reaction in turn enhances the substance that converts the enzyme to the inactive form—a closed-loop relationship. The feedback is negative, because the active form of enzyme promotes effects that lead to a strong shift toward the inactive form. This little system very actively and accurately forces the concentration of the product of the main reaction to match the concentration of another substance, the one that biases the enzyme toward the active form. This allows one chemical system to control the effects that another one is having on the chemical environment.

A person without some training in recognizing control processes might easily miss the fact that one chemical concentration is accurately controlling the product of a different reaction not directly related to the controlling substance. The effect of this control system is to create a relationship among concentrations that is imposed by organization, not simply by chemical laws. This is the kind of observation that a reductionist is likely to overlook; reductionism generally means failing to see the forest for the trees. Even the workers who described this control system mislabeled what it is doing—they concluded that this system controls the outflow of the product, when in fact it controls the concentration and makes it dependent on a different and chemically-unrelated substance.

To shift through several gears, consider the lines of research that began with Rosenblatt's perceptron. This device was conceived as a behavioral system that could be trained to react to patterns contained in its input information. First this idea was shown, by something of a hatchet job, to be impractical, and then it was shown to be practical again if several levels of training could occur within it (I haven't seen any apologies to Frank Rosenblatt, who died without vindication). In all its incarnations, however, the perceptron has been thought of as a system that learns to "respond correctly" to a stimulus pattern.

From the standpoint of perceptual control theory, however, organisms do not respond to stimuli but control input variables. So does that invalidate all that has been learned about perceptrons? Not at all. Perceptual control-theoretic models desperately need something like a perceptron to explain how abstract variables can be perceived. In a perceptual control model, however, the perceptron is only one component: it provides a signal that represents an aspect of some external state of affairs. It's easy to show that behavior can't be explained simply by converting such a signal into an output action. But behavior can be based on the difference between the perceptron's output signal and a reference signal that specifies the state of the perception that is to be brought about. The control-system model shows where the functions that are modeled as perceptrons belong in a model of the whole system.

Shifting gears again: some theorists are trying to model motor behavior in terms of "motor programs" and "coordinative structures." In these models, command signals are presumed to be computed such that when applied to elastic muscles they produce the movements of a real limb. These models contain some impressive mathematics, taking into account the linkages of the limb and the dynamics of movement of the limb masses. But perceptual control theory says that behavior is not produced by computing output; it is produced by comparing inputs with desired inputs, and using the difference to drive output. No complicated "motor program" computer is needed. Does this mean that the mathematical analysis by the motor program people is spurious and ought to be discarded?

Again, not at all. At some point in elaborating the perceptual control model, we must show how the driving signals actuate muscles to cause the movements we actually see. This entails solving all the physical equations for muscle and limb dynamics, just as the motor programmers have done. If they did their arithmetic right, it will still be right when we substitute the perceptual control-system model for the central-computer model. Both models have to produce the same driving signals. The only thing that will change is that perceptual control theory will show how the required driving signals arise naturally from perception and comparison against reference signals, instead of being computed blindly from scratch.

Finally, shifting to overdrive, what do we do about Artificial Intelligence? We take advantage of whatever it really has to offer, modifying it only where we know it fails to explain enough. One place where it fails to explain enough is in the way it deals with action. Basically, it doesn't deal with action. It starts its analysis with perception of abstract variables in the form of symbols, constructs models that imitate human symbol-handling processes as well as possible, and finishes by generating more strings of symbols that describe actions to be taken. It says nothing useful about how a description of an action, in symbols, gets turned into just those muscle tensions that will in fact produce an action that fits the description. When devices are built that are run by symbol-processing computers, the critical transformations that make action out of symbols are simply put into the device by its builders. Many of those critical parts turn out to be servomechanisms—perceptual control systems.

The assimilation of perceptual control theory into the life sciences will require a lot of this kind of reanalysis. Some old ideas will have to go, some will stay. This job is best done by people who are already competent in existing fields. Of course these also have to be people who can see that there is room for improvement along lines other than the standard ones.

In the current membership of the Control Systems Group we have representatives of at least a dozen disciplines of the life sciences, and a few persons representing some unlikely occupations such as piano teaching and law. When these people meet, there is little difficulty in communicating because all of them have a basic understanding of perceptual control theory. But communication isn't the only factor that makes these meetings valuable. The most important lesson comes from seeing how perceptual control theory applies in someone else's field.

The biggest problem with introducing perceptual control theory to scientists in conventional disciplines is that each scientist tends to think only of the scientific problems that are defined in that one field. The problem in question may involve behavior, but behavior is generally taken on faith to work the way some other specialist says it works. In fact most scientists tend to dismiss details involving other fields, assuming (often quite wrongly) that somebody else understands them well enough. We therefore find some very detailed biochemistry or neurology or

personality-testing, all done competently, being used to explain behavioral phenomena that are very poorly analyzed and in many cases don't actually occur. The sociobiologist concludes that behavior patterns are inherited, not knowing that only the consequences of motor outputs, not the outputs themselves, repeat. What does a geneticist really know about the actions through which a bird catches a bug? You can inherit the perceptual control systems that are capable of catching bugs, but you can't inherit acts that happen to take you where a particular bug is going next. The combination of narrow expertise in one field and naive conceptions in every other field leads to facile explanations that are right only at one point.

Specialists must see the need for a model of behavior that applies in all disciplines, even those in which the specialist is not competent. Once the Artificial Intelligence researcher understands exactly why organized behavior cannot be produced by computing outputs, he or she will modify the AI model so it will work correctly with more detailed systems actually capable of organized behavior. Important effects of learning how perceptual control theory applies in other fields will occur at the boundaries between disciplines—exactly where we need to work if we are ever to have a unified science of life. At Control Systems Group meetings, specialists from many fields hear other specialists talking about the way perceptual control theory has made them rethink the problems in a different field. Because of the common understanding, this inevitably reveals one's own hasty assumptions, and encourages still more rethinking.

One last remark about the CSG. The CSG does not represent any one scientific discipline. It has no agenda of its own beyond encouraging the application of perceptual control theory within existing disciplines—no agenda, that is, except perhaps lowering the barriers between disciplines. The psychologists in the group are still psychologists, the sociologists are still sociologists, the therapists are still therapists, the engineers still engineers. This is not a political movement nor an alternative to established science. It is simply a vehicle for promoting interaction among people interested in using or learning more about perceptual control theory in any specialty whatsoever. When all the branches of the life sciences have assimilated and begun using perceptual control theory, the CSG, its work accomplished, will have no further reason to exist.

In this presentation I have talked around perceptual control theory, alluding to some of its conclusions without attempting to justify or explain them. Learning perceptual control theory can't be done by listening to a half-hour's talk. I hope that some of you will find the promise of a unifying principle for the life sciences appealing enough to go further into this subject.