On modeling behavior using control theory

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

Date: Sun Dec 19, 1993 10:54 am PST
Subject: On modeling behavior using control theory

[From Bill Powers (931219.0930 MST)]

On modeling behavior using control theory

In discussions of control theory with engineering control theorists on the net, it has become clear that engineers have developed certain ways of viewing the problem of modeling behavior, and have settled on particular ways of representing it. To a great extent, the conception of what has to be explained is dictated by the form in which control diagrams are customarily drawn, and by the mathematical methods already available for analyzing systems with a certain organization. The question of how to apply control theory to behavior thus becomes the question not of how behavior is organized, but how the existing methods of analysis can be used without change, together with their underlying assumptions about what behavior is. The engineers have been arguing, in effect, that the methods they learned in school are completely adequate for the analysis of behavior, and that PCT introduces nothing new. As long as that belief persists we will get nowhere. I think we should focus on the real issues, or give up.

The most common engineering assumption is that behavior is the externally visible effect that an organism has on the physical world around it. The engineering assumption is that there is a "plant" in the world outside the organism, and the objective of organismic control is to bring this plant to a certain static or dynamic state and maintain it there in the presence of disturbances.

The PCT view, in contrast, is that organisms know nothing directly of the properties of the environment. All information they have about the environment is found in the primary sensory signals resulting from the impingement of physical stimuli directly on sensory nerve endings, and in perceptual signals that represent further processing of the primary signals. As a result, all that the organism can control is a representation of the world in the form of neural (or chemical) signals. There is no way for the organism to determine the actual effects of its outputs on the external world, so it has no way of altering its own actions or organization to create systematic effects that are not represented as perceptions.

This immediately rules out "open loop control" as a possible model of behavior. If the organism can't ascertain the effects its output are having, there is no way for it to adjust its outputs to produce any particular effect on the world. All outputs have particular effects on the world, but the organism can control only those effects that appear in its sensory world.

All control is closed-loop control, in a model of organisms. If sensory information is lost, some information dependent on the output signals must be substituted if any form of control is to continue. And when such a substitution takes place, the correspondence between the perceptual signals being controlled and variables in the external world is lost. How much difference that makes depends on the rapidity with which the substituted information departs from the information that would have been received if the sensors were operating normally. ------------------------- In the engineering approaches described on the net, there seems to be a tendency to think of control behavior as one single process, often of great complexity. The PCT approach is to consider behavior as the aggregate of a great many simple control processes, working in parallel and also hierarchically organized. The engineering approach is a matter of preference and custom, undisciplined by the actual organization of the neuromuscular system. The PCT approach is forced on us by the facts we know about the neuromotor systems.

We know of a great many simple control systems in the human organization, which deal with simple scalar variables, their integrals, and their derivatives. These control systems are used in all kinds of behaviors, in controlling all sorts of variables. The 600 to 800 control systems that control the musculature directly are employed by higher systems in every behavior that involves overt action. Their organization remains constant while the behaviors in which they are involved vary over the whole range of human experience, from scratching an itch to multiplying two polynomials together using pencil and paper.

Furthermore, behavior is clearly organized in larger units which also retain their character over a wide variety of behaviors. We learn specific skills like walking and speaking and typing, riding bicycles and driving cars, handwriting letters and numbers, opening doors and windows, putting things down, picking things up, moving things from one place to another, throwing things, pulling things apart, putting them together, and so on and so on into the hundreds or even thousands of elementary control processes. These elementary skills remain the same even though they are used as the means of controlling more abstract variables of all kinds, in all sensory modalities, under all kinds of external conditions.

Representing all these control processes as a single complex expression is simply not possible, and it is probably not possible that they are physically realized as a single complex process. It isn't useful to try to find a general expression that will cover all these varieties of behavior and all these levels of organization. We need to understand the details, because each control system is a general-purpose device that can be called upon in a variety of combinations by higher systems of very different nature. A lower control system provides the means for a higher system to bring one variable of experience to a specified value in a reliable way, without the higher system needing to know how to accomplish that end. The higher system simply generates a signal that is an example of the desired perception, and the lower system takes care of all the necessary details in bringing that perception to the requested state.

The single-system single-level approach is simply too limited to explain all of human behavior, even in principle. It requires that we either understand the entire human system, or understand nothing. Under the hierarchical approach, we can analyze behavior systematically, starting with the least units of organization and progressing to the more complex ones. This is the only practical way to model the entirety of human behavior, and it is also probably the way behavior is actually organized.

To model behavior we must first observe it carefully. I do not see the required degree of care being taken in the engineering approaches. Too much is taken for granted, too much is left out. Common-sense assumptions are used where a model is really needed. Diagrams leave out essential functions and connections. Disturbances are either ignored, or assumed to have convenient characteristics that allow them to be handled in a statistical way instead of in detail. Too many loopholes are left, too much is assumed without demonstration. The approach to PCT is defensive, not open. There is too much blind reliance on mathematics, and not enough attention paid to identifying the constants, variables, and functions with observables. The orientation is theoretical, not experimental.

These are all defects that we have labored, and still labor, to remove from the PCT approach. Until control engineers take them seriously, it will be hard for them to see why their approach is inadequate to the problem.

Best to all, Bill P.