Controlled versus controlling; a tutorial

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

Date: Tue Feb 21, 1995 5:05 pm PST
Subject: Re: controlled versus controlling; a tutorial

[From Bill Powers (950221.0850 MST)]

Bruce Buchanan (950220.20:30 EST)--

> My understanding from B:CP has been that a controlled variable is the measure or goal, a perception that behavior is organized to maintain, by correcting for differences between perception and that quantity. So I understand that the variable is a perception to be controlled by behavior.

This is essentially correct, but the small remaining differences make a difference. Here is a somewhat extended but not exhaustive tutorial, written for you but also for a more general audience.

I try to maintain a consistent way of speaking about what we can observe and what we conjecture as part of a model. A controlled _variable_ or _quantity_ is an observable quantity; a controlled _perception_ is not observable from outside another organism. Likewise, the reference _level_ of the controlled variable or quantity is observable; the reference _signal_ which, in the model, determines the reference level, is not observable from outside another organism. So we can observe variables or quantities, but we must infer signals. This distinction reminds us of the difference between what an experimenter has to work with and what a theoretician has to work with.

Experimentally, we detect control in the following way. We see some variable in the environment that is being acted upon by an independent influence in the environment. From examination of the environment and from physical, logical or other principles, we can predict how much effect the independent influence should be having on the physical variable. We _observe_ that the actual effect is at least an order of magnitude smaller than the predicted effect.

The only physical explanation for the failure of the independent influence to alter the affected variable in exactly the predicted way is that there is a second influence acting on the same variable in opposition to the effects of the first influence.

The opposing influence must, of course, be traceable to the action of the system we think may be a control system. If we apply a known force directly to the variable we are investigating, we find that the organism exerts an opposing and very nearly equal force, with the net result that the variable scarcely changes.

Having shown that the action of the organism cancels the effects of external influences on the putatively controlled variable, we must now prove that the organism must sense the variable in order to produce the cancellation. Sometimes this is easy and sometimes not. If control involves vision and we can communicate with the organism, we can simply ask the organism to close its eyes (or blindfold it) and repeat the experiment. If the actions are no longer systematically opposed to the effects of independent influences on the controlled variable, we have completed the proof that the organism is (when able to see) controlling the variable. If, on the other hand, the action of the organism _continues_ to oppose the effects of external influences on the variable, we have misidentified the controlled variable and must look further for an explanation.

When control has been verified, we have also found the _reference level_ of the controlled variable. Formally, it is that state of the controlled variable at which the action of the system does not tend to change the variable. More dynamically, it is the state toward which the actions always tend to restore the variable when external influences tend to make the variable change.

This is the basic experimental paradigm of PCT, the Test for the Controlled Variable. As you can see, it does not make use of any model, nor does it contain any guesses about the internal organization of the organism. It was constructed that way deliberately, to satisfy the normal scientific requirement that a phenomenon be reproducible by anyone; that it be public. If you can do the Test, you can check to see whether any variable is under control and prove to anyone's satisfaction that a given organism is controlling it (or not). You can do this even if you can't explain how the control is brought about.

In effect, you are defining what is meant by the term "control" in PCT, then seeing if the relation of the organism to the environment fits that definition. Other usages of the term "control" which do not conform to this definition refer to other phenomena, not the phenomenon with which PCT deals. This is important to remember when you hear others talking about "control." They may or may not be talking about the phenomenon with which PCT is concerned.

The Test, as just described, suggests that control is static because the candidate controlled variable is maintained constant. In other circumstances, however, with somewhat more elaborate procedures, it is possible to show that a controlled variable is being caused to vary in some pattern that is essentially independent of external influences on it. Now we find that the behavior of the controlled variable follows a pattern which conforms neither to the pattern of actions alone nor to the pattern of external influences alone. At all times, of course, the vector sum of the organism's actions and the independent influences physically explains the behavior of the controlled variable; what remains unexplained is why the vector sum is as it is. Once again, the proof of control is completed by showing that when sensing the state of the controlled variable is prevented, control is lost.

We can show by extended observation that (a) the action is systematically opposed to the independent influences when measured in units of effect on the controlled variable, and (b) that the correlation of changes in the controlled variable with either changes in the action or changes in the independent influence approaches zero.

This latter result is highly non-intuitive, yet it is true of all control systems in the PCT meaning of control. It is therefore critical to understand this result in order to understand PCT. I won't elaborate further here.

One precaution: a PCT control system does not control "things" but _variables_. If you want to investigate control of a "thing" like a car, you must first break it down into variable attributes such as velocity, direction, position, cleanliness, color, price, and other aspects of the car which are at least conceivably variable under independent influences and the influences of actions by an organism. A person does not control "a car" but _something about_ the car.

Everything up to here has employed only observable phenomena, public phenomena. Since no restriction has been placed on what we choose to call a "variable," an "independent influence," and an "action," the Test can be used to demonstrate the presence of control under any circumstances and at any level of abstraction -- if it exists. If control does not exist, the Test will show that it does not. So the hypothesis of control is falsifiable.

This brings us to the model of control, the heart of PCT. PCT is a theory which attempts to answer the question of how an organism must be internally organized in order to create the phenomenon of control that we can see from outside it. The Test tells us if control is occurring. The model is supposed to explain how it can occur.

The basic PCT model is a simple translation of the conditions we see externally into a system that can reproduce those conditions. Because sensing the controlled variable is essential to control, the control model contains a perceptual function which represents the state of the controlled variable as the state of a perceptual signal.

We observe that the action depends on the relation of the actual state of the controlled variable to some specific reference state, so we need to specify the reference state inside the model. This is done by supplying a reference _signal_ of the same physical nature as the perceptual signal. These two signals are combined in a comparator, which by convention subtracts the perceptual signal from the reference signal and emits an error signal representing the sign and amount of the difference.

The error signal enters an output actuator which produces the action that we observe.

All that remains to make the model workable is to make sure the signs of the effects in each internal function are correct to produce the observed result, and that any dynamical modifications required for stability are met inside the system. There are ample resources for doing this; the entire field of control system engineering.

Generally, in PCT we are not concerned with dynamic stability in control systems for the simple reason that all normal organismic control systems are, in fact, dynamically stable. There are good reasons for trying to guess how the living control system achieves dynamic stability, but when we are looking at the overall organization of the organism, we are more concerned with understanding steady-state relationships.

A "steady-state" relationship does not imply that variables are not changing. For example, it is possible for a living control system to control the rate of rotation of a shaft, by cranking it. Here, the "steady state" condition is a constant rate of turning of the shaft, implying a continuing cycle of ever-changing cranking movements. We consider this a steady state because when the actual rotation rate is constant at the right reference level, the perceptual signal representing rate of rotation will match a constant reference signal. The internal dynamics of the system that make it stable will show up if there are brief disturbances, like increases and decreases in the load which induce an error into the control system and call for a change in the cranking efforts. In a relatively unstable system, the cranking rate and shaft rotation would speed up, slow down, speed up, slow down, and so forth in a series of diminishing undershoots and overshoots _of cranking velocity_. In a stable system, the control process would immediately adjust to the disturbance and maintain a steady rotation rate thereafter.

Another example is the way an aggressive policeman follows your car when waving you over to one side for a lecture or a ticket. You see the unnerving sight of the police car ten feet behind you, moving with you at 75 miles per hour. You hesitate to hit the brake lest he run into you, but you needn't worry: the policeman is controlling the _relationship_ between his car and yours; as you slow down and pull to one side, the police driver maintains the same (reference) distance from your car until you both have stopped. The steady state is defined by a constant reference signal (inside the policeman) specifying the intended distance between the cars. The dynamics of control enter only when you hit the brake; the policeman, being a very stable control system, simply slows down with you, not oscillating between nearer and farther (which would indicate some degree of dynamic instability).

By looking at steady-state relationships we can begin to understand the organization of behavior as a collection of control processes. It is clear that there are many control processes going on simultaneously in independent control systems, as when you control your arm and leg positions at the same time while swimming. It is also clear that there are different _levels_ of control, because something is coordinating the control systems concerned with arm and leg positions.

A start at sorting out the multiple parallel and hierarchical control processes in the human system is contained in an extension of basic PCT called HPCT -- Hierarchical Perceptual Control Theory.

The basic relationships among levels of control are seen in the lowest level of control, which is concerned with control of muscle efforts and joint angles. The neuromuscular organization of the systems involved here clearly conforms to the basic PCT model: sensing, comparison, and error-driven actuation of outputs, the muscles. These systems have been treated as "reflexes," but they are really control systems. What is interesting about these control systems is how they are used by higher systems.

The traditional neurological picture of motor behavior is that commands formed high in the brain are fanned out to the muscles, where they enter a "final common pathway" and cause the muscles to contract. But when we look at the details of the reflex systems, we find that muscle contraction is almost independent of the signals reaching the spinal motor neurons from higher systems. There is strong and immediate feedback from sensors that detect tension in tendons and changes in length of the muscles. The downcoming signals do not cause muscle tensions directly; they are reference signals that specify the amount of _sensed_ tension and muscle length that is to be produced by _varying_ the contraction of the muscles. Unexpected disturbances will cause the amount of muscle contraction to begin changing within less than 10 milliseconds, with no change in the command -- actually, reference -- signal from above. The control system maintains the _sensed effects_ of the muscle tensions in a match with the reference signals issued by higher systems. The muscle tensions automatically vary as required to keep these low-order perceptual signals at the specified levels (constant or changing).

We see the general principle: higher systems act not by sending commands to act, but by adjusting reference signals for lower-level systems. The lower-level control systems continually vary their own outputs in whatever way is required to keep the upcoming perceptual signal matching the given reference signal. If disturbances occur, the lower-level systems will simply resist them, and the higher system will not experience any significant disturbance.

HPCT is a proposed 11-level structure of control systems which are related in this way, higher systems acting to control their own perceptual signals by means of varying reference signals for sets of lower-level systems. In doing so, of course, they produce effects on the outside world, which are sensed and ultimately become the controlled inputs of higher systems. All control loops are closed via the environment, which means that all levels of control can be experimentally investigated.

Other kinds of control links from one level to lower levels are possible; one very likely kind is control through adjustment of the sensitivity of lower systems, or other parameters. The primary emphasis in HPCT so far, however, has been on control via adjustment of reference signals.

With the "standard" concept of control via reference signals, we can see that many behaviors are brought about by varying reference signals. If there is a control system that can make sensed arm position match any specific reference signal, then arm position can be varied in any desired pattern by varying the reference signal in that pattern. The natural dynamics of the arm do not come into play; they are absorbed into the details of lower control systems. So an arm can be made to move in a wide variety of patterns quite independently of its natural dynamics; the movements reflect not the physical dynamics but the pattern of variation of the reference signal(s).

Finally, a little bit of discussion of translating between common language and PCT. Translating between PCT and other scientific theories of behavior is actually much more difficult and often impossible, because other theories propose a very different internal organization of the brain and nervous system, usually incompatible with PCT and HPCT. Common language, however, is largely congruent with the ideas in PCT, so the translation is often direct.

The simplest translation, and the only one we will spend any time on, is between commonsense terms like goal, purpose, intention, desire, and expectation into the one PCT term, reference signal. All of the common terms refer to somehow specifying or establishing a conception of a state of affairs which does not yet exist, yet which somehow guide actions toward making them into realities. This is precisely the function of a reference signal entering a control system.

A reference signal, when given a specific state by higher systems, represents the way a perceptual signal will be when the external world is in some particular state. The action of the associated control system is always based on the difference between the reference signal and the perceptual signal. The perceptual signal indicates the ACTUAL state of the world at all times (of course, as perceived), while the reference signal specifies a _particular_ state that is to be brought about and in general does not yet exist. As the difference drives actions, the world outside is altered and many perceptual signals change; when the control system is properly organized, the perceptual signal becomes more and more like the reference signal until the error is as small as possible. At that point, some aspect of the outside world will have been made to come to a particular state which we see as the reference level of some controlled variable. It will come to this state regardless of external influences which would determine its state if the controlling organism were not acting.

The main difference between the PCT conception of a goal or purpose and the commonsense idea is that it relates directly to perceptions and only very indirectly to states of the outside world. What we intend is not to act, but to perceive. Only the perceptual consequences of acts are specified by reference signals; the means of achieving the goal-state of a perception will depend on what disturbances, what independent external influences, are also acting on the controlled variable. The reference signal specifies a perceptual outcome, not a motor output. This is why my first book was called "Behavior: the control of perception."

There are many more interlocking aspects of PCT and HPCT which bear on other problems of explaining human behavior, such as the way higher perceptions are functions of lower ones. The reader who wants more details is advised to read the above book as well as other materials that are reviewed every month on CSG-L.

Best to all, Bill P.