Perceptual Control Theory at 40

William T. Powers
in: Journal of Living Control Systems,
Winter 1994 Volume 4 Number 1
Edited by Greg Williams

As this issue of *Closed Loop* is the first one carrying the subtitle *Journal of Living Control Systems,* readers encountering our approach to this subject for the first time might need an overview of perceptual control theory (PCT) to get started. So this paper will be Yet Another Introduction to PCT. I will slant it, however, toward those coming into to this subject from the physical sciences; the relationship of PCT to physical approaches has been discussed at some length lately on CSGnet.

Rather than just reviewing the history or the principles of PCT, I'll try to develop an argument that leads from conventional views of behavior to the new view that PCT gives us, emphasizing in the end the odd role that organisms, seen through the eyes of PCT, play in a world otherwise dominated by physical laws. The point will be to show that control theory provides us with the germ of a radically new understanding, a break with all traditional theories of behavior—and many new ones as well. The future progress of PCT depends on understanding just how different a view of behavior we get by understanding the logic of control, the logic of a controlling organism's relationship to its environment

The Etiology of Perceptual Control Theory

All living systems are sensitive to their environments; all act on their environments. This is ancient knowledge. The puzzle presented to the behavioral scientist is only how that sensitivity becomes converted into action. What are the rules, if any?

The most obvious and straightforward scientific approach to this question was realized long ago. In the physical sciences, if you want to know the properties of an assemblage of matter, you apply experimental forces and other influences to the object and observe what it does as a consequence.

In the worlds of physics and chemistry, this is a relatively easy task. Objects tend to be simple and have few properties; they are normally homogeneous or made of simple repeating units. It is not hard to make sure that experimental effects on them are the only effects of any importance. All similar objects made of the same materials behave in essentially the same way, and they will continue to do so no matter how many times an experiment is repeated in fact, measurements of properties can be almost indefinitely refined by repeating them. A physical or chemical experiment can be clearly described and can be replicated by anyone who wishes to check the results.

Several issues of *Closed Loop* have been recreated and are / will be posted at www.PCTresources.com

The reasoning about the meaning of an experiment can be communicated in clear and formal language, and even the reasoning process itself can be made public by being expressed in mathematical terms that anyone can learn. The history of a material object is entirely expressed in its present condition; the path by which it got into that condition is irrelevant, and only the current environment is of any importance in determining what will happen in the future to that piece of matter.

These confidence-building thoughts about the physical-science approach were, of course, tried out on organisms. The results were anything but confidence-building. A behavioral scientist reading the preceding paragraph might well experience mounting despair and envy of the physicist. While it is true that organisms are made of matter and must therefore obey all of the laws of physics and chemistry, it is not true that they are homogeneous or made of simple repeating units. They are, in fact, immensely more complex internally than the objects studied by physicists and chemists. They are too sensitive to their environments for any scientist to be sure of having control of everything important that happens to them. Not only are they sensitive, but they adjust themselves internally to external circumstances. It is not possible to perform the same experiment over and over on an organism to refine measurements of its properties—just imagine giving the same physics test over and over to refine a determination of a student's state of knowledge of physics, or giving a weight-lifting test to an athlete, day after day, to refine measurements of the athlete's strength.

The initial attempts to apply the methods of physical science to organisms were moderately successful at answering questions about perception. When the same methods were extended to the study of behavior, the results were not so encouraging—in comparison with expectations, they could only be called failures. Organisms were so subject to unpredictable influences, it seemed, that extraordinary precautions had to be taken to eliminate unwanted and unpredicted behaviors. This seemed at first to be a technical problem, to be overcome by greater attention to controlling the environment during experiments. As the years went by, however, it became apparent that no amount of attention to detail was enough. Not even the simplest phenomenon, such as a blink of an eye in response to a puff of air, could be made to occur with complete reliability. More complex behaviors simply went all over the map. The dream of creating "Newton's laws of behavior" was apparently unattainable.

This did not cause a loss of faith in the methods of physics. Most behavioral scientists continued to assume that behavior was created by environmental influences. This assumption led to an attempt to find suggestions of regularity in behavior through statistical means, and then to a conclusion that this was the only possible means of exploring behavior, because behavior is inherently variable. The basic concept was retained: what organisms do is caused by what is done to them by the surrounding environment. But the requirements for formal language, public means of reasoning, ability of anyone to reproduce results, and refinement of measurements by continued experimentation were mostly impracticable. Despite the failure of the physical-science approach, the assumption was that the failure of organisms to behave as predictably as planets was due to technical difficulties, not errors in basic principles. The alternative conclusion, that something was wrong with applying the physical principle of cause and effect to the behavior of organisms, was simply not considered.

This alternative eventually came into play by a roundabout path.

Control theory was invented by engineers of the 1930s trying to build devices that would behave like human beings carrying out a specific kind of task: a control task. Even though the engineers did not realize it (many still do not realize it), the concept of control introduces a new principle, one that denies the basic idea that organisms do what the environment makes them do. While cause and effect still work in control theory as anywhere else, the organization of a control system creates apparent cause-effect dependencies that are different from the actual ones. Part of understanding control processes in organisms is the understanding that conventional cause-effect interpretations can be more misleading than informative.

Organisms are sensitive to their environments, and they act on their environments. The old assumption was that the sensing was the primary process, with the acting following from it. But that is an arbitrary assumption. It is just as plausible to assume that the acting is the primary process, and that the sensing, at least in certain critical regards, follows from the acting. It is even more plausible to say that sensing and acting are processes that go on simultaneously, in continuing streams that can't be clearly separated into cause and effect. This is basically what the inventors of control theory discovered: a type of system in which behavior affects the inputs on which behavior appears to depend. This is the type of system they had to use to imitate the human behavior called controlling.

This discovery led eventually to cybernetics, which endorsed this concept of closed causation without exploring more than its general philosophical implications. Years had to pass before more detailed implications came to light. Still more years had to pass before the basic concepts of control theory could be boiled down to a systematic model of control behavior now called PCT that could replace the old systematic cause-effect model based on the approach of physics.

The basic difference between the physical approach and that of control theory is that the physical approach deals with properties of energy and matter, while control theory deals with the properties of *particular organizations* of energy and matter. George Herbert Mead pointed out early in this century that physics doesn't deal with forms, with the entities into which we divide the world of experience. The physicist explores what is the *same* between a horse-cart and an ox-cart. The systems approach is concerned with what is unique to each vehicle, with differences in behavior brought about not by the differing physical or chemical composition of different objects, but by the differing organization of forms made of the same materials differently arranged.

It stands to reason, therefore, that physical laws will have a different significance when seen in the context of an organized system. We can admit that they make the behavior of the system possible, without also admitting that they *explain* the behavior of the system. Physics and chemistry can explain how it is that a neural signal liberates energy that causes a muscle fiber to contract, and how it is that this contraction leads to accelerations, velocities, and positions of limb segments connected to a joint spanned by the muscle. But they can't explain how it is that this signal arises under just these circumstances to reach that particular muscle. Physics and chemistry can't even be applied until the organization is specified. It is at the level of organizational understanding that control theory confronts older conceptions of the organization of behavior in living systems.

The Phenomenon of Control

I am going to avoid semantic arguments about what "control" really is. I will use the term in a particular sense; if others interpret it in a different sense, they will have difficulty following this exposition. I use it in this sense: A system is said to control a variable if it acts on that variable, in the presence of other unpredictable influences of comparable size on the same variable, so as to maintain the variable in an arbitrary state. The "arbitrary state" might mean a state of constancy, or any arbitrary pattern of change. The critical aspect of this definition is that physical influences that normally account entirely for the state of the variable are no longer effective, while the action of the control system causes the variable to behave independently of those other physical influences. When that is true, the variable is called a controlled variable.

The first important fact about control to notice is that the controlled variable is being acted upon by many forces, only one of which is attributable to the control system. The driver of a car can apply a lateral force to the front end of a car by turning the steering wheel. But there are many other influences that create forces acting laterally on the car at the same time: crosswinds, bumps in the road, tilts in the roadbed, unevenly inflated tires, and asymmetries in the aerodynamics of the car's shape, to mention a few.

If we observe, as we commonly do, that the path of the car does not follow strictly from the sum of all of the external forces acting on the car's mass, we can only conclude that it is the driver's contribution that makes the difference. If we see the car moving in a straight line, we can only conclude that the sum of all forces, including the one that the driver can alter, is zero.

So, if any of the external influences is seen to vary, but the path of the car does not vary as Newton's laws and engineering principles would predict, we have to deduce that the driver must be producing a varying force that just cancels the sum of the external forces. Indeed, we can observe the driver continually making adjustments of the steering wheel angle, while the car continues in a straight, or very nearly straight, line.

Likewise, if we observe the car moving along a smooth curve, but we see that the sum of all extraneous forces would tend to make it move along some other path, we can deduce that the varying forces created by the driver add just enough more force in just the right way to produce the curved path. If we see the car moving along a straight expressway, then turning off to take an exit ramp, then making other turns until it ends up parked in a parking lot, we can be quite sure that normal external forces would not have made the car follow just that path (an easily tested assumption). We can be sure that the varying forces created by the driver's motor actions on the steering wheel must have been exactly those necessary to add to the natural forces to create this overall result.

To anyone accustomed to normal physical or engineering analyses of the motions of objects, there must be a jarring note in this account What is generally done is to observe all of the independent contributing forces and the initial conditions, and then to deduce through physical laws what the resulting motion must be. The driver's steering forces and the external forces due to winds, road tilts, and so forth simply occur as they occur, and the car's path is the outcome.

But here we are speaking as if one of the determining forces, the varying force being generated by the driver, is being adjusted so as to create a *preselected* outcome. Instead of the outcome varying randomly as the unrelated applied forces make it vary, the outcome conforms to some predetermined pattern. One of the causal forces which adds to the other forces continually changes in just the way needed to maintain that pattern. We would appear to be saying, and we are in fact saying, that the outcome we observe is being produced *on purpose.*

The vast majority of behavioral scientists have always rejected this interpretation. When the concepts of PCT were first being developed, this resistance was massive and almost universal (it is considerably less today). To say that outcomes are produced intentionally has seemed to most scientists to call for a reversal of cause and effect, or for giving the future an effect on the present. Many have argued that if all of the causal influences are known, the outcome must be whatever it is, and to call it "intentional" adds no explanatory power. Clearly, the outcome is an effect of converging causes, not a cause of the converging causes. Even if behavior does seem to entail intended outcomes, a scientist must stick to normal cause and effect, and find some other explanation.

There have been centuries of attempts to find some other explanation. But prior to the advent of control theory, all other explanations, we now know, were spurious. Even now there are many who strongly resist admitting that outcomes are indeed intended, and that organisms are the loci of these intentions. This resistance is misplaced, because now we can explain exactly how it is that an outcome can be controlled.

How Control Works

Once again: All living systems are sensitive to their environments; all act on their environments. So far we have talked only about actions and other physical influences on the environment. To see how control works, we must now talk about how organisms sense their environments.

Sensing is a process by which an external variable comes to be represented as a neural (or chemical) signal inside an organism. This looks like normal physical causation, but it is not like most causal processes. There is *amplification* involved. Metabolic processes in an organism maintain the sensing nerve-endings in hair-trigger states of readiness to fire. Only a tiny added stimulus is needed to cause a neural impulse to be generated, and metabolic processes instantly restore the sensor to the brink of firing again. So a small continuing stimulus causes the sensory nerve ending to fire again and again, at a frequency that corresponds to the amount of stimulation. The signals that leave the nerve ending involve the expenditure of many times the energy that causes the sensory ending to fire, nearly all of the energy being supplied from stores within the organism itself.

These neural signals can be further amplified, and eventually they can be routed to effectors such as muscles that provide a final amplification up to levels that can have significant effects on physical processes in the environment. The result is that organisms can create physical forces of large magnitudes which are produced without any significant reverse effect on the physical variables being sensed. This creates a novel relationship between the organism's output forces and other physical processes.

I remember inventing my first perpetual motion machine, at the age of perhaps 12. I had read that a certain kind of motor could be used either as a motor or as a generator. So I thought of putting fan blades on two of these motors and using one to blow air onto the other, the idea being that the generator would supply the current needed to run the motor while the motor supplied the wind that would run the generator. It took a few more years of education to realize that one has to think of physical processes quantitatively, not just qualitatively. It makes a difference *how much* air can be blown, and *how much* current can be generated, and *how fast* the driving fan can be spun by the available current. High school physics was enough to show me the embarrassing truth: that in physical systems, there are balances that are maintained: balances of forces, of momenta, and of energies. The world studied by physicists is rigorously constrained by these balances, these conservation laws. You can't get any more out of a physical system than goes into it. This is how I and most other people learned to think about physical processes.

This is also true of organisms, of course. No more energy comes out than goes in. But the energy that goes in is of a different form from the energy that comes out it is the chemical energy in food and air, obtained independently of the physical processes involved in behavior, and stored for future conversion into actions. So when an organism, a person, comes across some natural physical process in its environment, it is in a position to throw a monkey-wrench into the machinery by spending some of its store of energy.

Let's switch examples now. Suppose a person sees a fat child and a thin child sitting on opposite ends of a teeter-totter. The end with the fat child on it is, of course, on the ground, and the thin child is high in the air. The upward force of the ground on the fat child's side, plus the upward force from the thin child pressing down on the other end, just equal the fat child's weight. The physical system is in equilibrium.

Now the person places a hand on the thin child's end of the teeter-totter and pushes down, spending a bit of metabolic energy from the last few days' meals and several thousand breaths of air. The thin child descends and the fat child rises. If the amount of downward push follows a certain law, the teeter-totter will end up horizontal and stationary again.

What is the required law? If the force applied is large when the fat child is low and small when the fat child is high, with a continuous transition between the two states, there will be one state in the middle where the force is just right to bring the teeter-totter to the horizontal with all forces in equilibrium But what could make the force applied by this helpful person follow that law?

Suppose we tried to mechanize this effect. When the fat child's end goes down, a cable pulls a weight at the center of the teeter-totter toward the thin child's end, and vice versa. The history of perpetual motion machines is full of such clever devices. All such devices, however intricate and devious their designs, fail because you can't get more out of a physical system than went into it.

But the person helping balance the teeter-totter is exerting a force of just the right amount without any linkage from the teeter-totter that produces that force. The only link from the teeter-totter to the person is through the person's visual sense, which registers the angle of the teeter-totter as feeble neural signals inside the person's brain. This requires only intercepting some of the light reflected from the physical apparatus and the children, a process that supplies only an infinitesimal amount of energy to the person and exerts no measurable force at all, either way.

The neural signals that now represent the angle of the teeter-totter are further amplified, and they finally enter muscles where the greatest (by far) amplification of all occurs, producing a force that acts downward on the teeter-totter. This force is greatest when the fat child is accelerating upward, smallest when accelerating downward. Stored energy is used by the person in applying the force to the moving teeter-totter. That's vital; none of this could work without the independent source of energy that comes from the eggs and roast beef and peanut butter sandwiches that the person has been eating.

What happens in the end is that the neural signals representing the angle of the teeter-totter come to some particular state representing the horizontal position, and the force applied to the teeter-totter is just the difference in weight of the two children. The physical system is now being maintained in a state far from equilibrium, but if you include the helpful person in the physical system, everything is in equilibrium again: forces, momenta, and energy inputs and outputs.

The factor that determines where this equilibrium will occur is now in the person, not the external physical system. There is some particular condition of the sensory signals that corresponds to the observed equilibrium. If the sensory signals indicate a deviation from this condition, the force will either increase or decrease in the direction that tends to restore the equilibrium. The rule is simple: if the angle slopes downward toward the fat child, increase the force; if upward, decrease it. This rule, which is applied inside the brain of the person, is what determines the equilibrium point.

There's one more factor to consider. The person balancing the teeter-totter might decide to maintain the board at some angle other than horizontal. This amounts to redefining the condition of equilibrium. In a control system model, this is done by providing an adjustable reference signal against which the signal representing angle can be compared. This occurs inside the person's brain. The final amplification of signals that drives the muscles is applied to the *difference* between the reference and sensory signals, so the opposition to even small deviations from equilibrium can be very strong.

With the addition of the variable reference signal, the person can now cause the teeter-totter to behave in any arbitrary way at all, as long as the available muscle forces are large enough and the person doesn't exhaust the stores of metabolic energy. As the reference signal varies, the teeter-totter's angle varies in exact correspondence. It can be made to vary regularly or irregularly, quickly or slowly, with or without a child sitting on either end—or not at all, even though the children climb on and off the board. The angle of the teeter-totter is now completely determined by a reference signal inside the person's brain, and the normal physics of the teeter-totter is totally overridden. The person is inserting extra force, extra momentum, and extra energy—whatever is required to make the desired behavior appear.

This same analysis could have been applied to the driver of the car. The lateral position of the car is represented in the driver's brain as some sort of neural signal. Another neural signal, a reference signal, specifies the lateral position that is to be maintained, and amplification of the difference between the two signals produces muscle forces that act on the car to make its lateral position, as sensed, match the specified position. Varying the reference signal will then cause the lateral position of the car to change in a parallel way, independently of other forces acting on the car. The normal physics of car motion is overridden; external forces lose their determining effects.

Organisms in Control

In the world of physics, there are physical objects linked to each other by properties of the environment and physical laws that cause the behavior of one object to depend on the behavior of other objects. Even in the most complex of physical systems, there is a kind of natural bookkeeping that accounts for all of the interactions. The sum of all forces acting on and inside the system, counting both actions and reactions, is zero. The sum of all changes in energy content, including energy inputs from outside and energy outputs to the outside, is zero. All momenta add up to zero, or at least a constant.

If we want to make one variable in a physical system depend on another one, the normal approach is to establish a physical link. This link connects forces from one object to another object, which involves transfers of energy and momentum and sometimes flows of matter. The new link participates in the balances of the system; it can generate no new energy, and it can create no unbalanced forces. The affected object is in physical equilibrium with the affecting object. If A is pushing on B through the new linkage, then B is pushing back on A with exactly the same force.

An organism is, of course, a physical system subject to all of the same laws and balances. But the organism can create linkages among objects in its environment which, at first glance, seem to violate physical principles.

First, the organism can move about in its environment and dispose itself to create forces on many different objects in many different ways. This means it is in a position to affect objects that are not normally affected by such actions.

Second, the organism can orient its sensors to create internal signals representing many aspects of physical objects around it. The visual sense is particularly potent in this regard: simply by looking in different directions, the organism can create internal signals that stand for the states of objects in many different ways: their position, velocity, size, color, relation to other objects, shape, and so forth. It can do this without affecting those objects in any measurable way.

As a result, an organism can position its muscles and limbs, and its sensory apparatus, in ways that create arbitrary linkages between the objects it can sense and the objects to which it can apply forces. Furthermore, because of the high amplification that takes place inside the organism, this linkage can be made one-way—that is, one object can be made to affect another object without being affected by the reverse path through the same link. There is a violation of the normal energy balance in the physical system, because any normal physical link would require energies, forces, and so on to remain in balance.

The unbalances are made up by the organism from its internal energy stores, and from the way it braces itself against the world as it exerts forces. If we consider the physical environment and the organism as a single system, there is, of course, no violation of any physical principles. The point, however, is that the physical environment linked to an organism can no longer be treated as if no organism were present.

Consider the car and driver again. With no driver in the car, but with the car rolling along the road, physical influences on the car can be calculated according to normal physical principles. From the speed and direction of the wind and the aerodynamic properties of the car, the wind force acting on the car can be calculated. Similarly, forces arising from tilts and bumps and soft tires can be calculated. All of these forces can be added up, and their effects on the car can be computed. From these forces and the properties of the car and road, the motion of the car can be computed with, in principle, as much exactness as we please.

But now put a driver in the car. Suddenly, the path of the car ceases to follow from the sum of all external forces and the properties of the car and the road. Instead, we find that a new physical linkage has been created. Now when the wind blows and the road tilts, the result is a movement of the steering wheel which *prevents* the car from obeying the physical laws that previously applied.

Even more important, we find that the physical linkage that has been created is *not* between the steering wheel and the wind or the tilt of the road, but between the steering wheel and the lateral position of the car. What the driver is sensing is the *outcome* of all of the applied forces (which now include the effects of turning the steering wheel). The driver watches the visual appearance of the hood of the car against the road ahead and acts to maintain that visual appearance in a specified state (either constant or changing in a specified way). The only thing that gives the car's lateral position an effect on the path of the car is the fact that the driver is sensing that lateral position, internally specifying an intended state for that perception, and producing steering forces based on the difference between what is actually sensed and what is intended to be sensed.

From outside the driver, this critical perceptual linkage is invisible, undetectable in terms of any changes in the physical world. Nothing in the world changes measurably because of being sensed. Nothing in the physical outside world indicates the driver's internal reference signal that specifies the intended state of the perception. As far as any external measurements are concerned, the force that turns the steering wheel has no observable external physical cause. It is an arbitrary force generated for no physically observable reason.

The strangest thing about this force is that after it is added to all of the other independent forces that are applied to the car at the same time, the result is an outcome that is repeatable with great accuracy for long periods of time, even if the external forces change and even if there are changes in the properties of the car and the road. When all of the external forces change, the outcome does not change; instead, the remaining force applied to the car changes in just the way that keeps the outcome the same. The cause changes in order that the effect be preserved.

An organism can attend to any perceivable aspect of the environment. If the forces that the organism can generate are comparable to the external forces that exist at the same time, that aspect of the environment can be made to conform to the organism's intention for it, and to cease behaving as the natural forces on it would otherwise dictate. The actions of the controlling organism supersede the physical laws that normally govern that part of the environment, in the respect that the organism is controlling.

Conclusions

Organisms are physical systems, and they exist in a physical world. But the laws of physics do not explain their behavior or its effects on the physical world. Organisms force the world around them into highly improbable forms, states of motion, and organization, and they act in a way that keeps normal physical forces from having their normal effects. It is organization, not physics, that explains how they do this.

To understand human behavior in these new terms is to seek a kind of explanation completely different from what behavioral scientists, modeling their approach after physics, have sought. This is what PCT is about, and where its promise for the future lies.