PCT in 11 Steps

By William T. Powers

1 Behavior as Control

Control is a process of acting on the world we perceive to make it the way we want it to be, and to keep it that way. Examples of control: standing upright; walking; steering a car; scrambling eggs; scratching an itch; knitting socks; singing a tune. Extruding a pseudopod to absorb a nanospeck of food (all organisms control, not only human beings).

The smallest organisms control by biochemical means, bigger ones by means of a nervous system. Whole organisms control; the larger ones have brains that control; most have organs that control; if they are composed of many cells, their cells control; the DNA which directs their forms and functions controls; even some molecules, certain enzymes, control by acting on the DNA to repair it when it's damaged. Control is the most basic principle of life and can be seen at every level of organization once you know what to look for.

In this series[1] we will examine the process of control to see how it works, how it explains the behavior of organisms, how we can recognize it when we see it, and how understanding it can change our theories. In the first 11 mini-chapters we will see how PCT, Perceptual Control Theory, grows out of and replaces its main theoretical predecessors.

We will start by seeing how the mainstream of behavioral science found itself in channels that led to confusions and impossibilities, and how engineers who had no interest in psychology at all managed to discover the one basic principle that could have saved the sciences of life from a 300-year search down one blind alley after another. The problem is not that the life sciences got everything wrong; it's just that they got the most important things wrong: what behavior is, how behavior works, and what behavior accomplishes.

1 Bill Powers wrote this compact series of 11 brief statements to serve as an outline for a proposed TV program. The program did not come to pass, but this is an excellent summary of PCT.

2 Behavioral Science I

Before PCT, there was behavioral science. The "behavioral" part indicates that if we're behaviorists, we're interested in what we can see organisms doing, not in what we might guess goes on inside their minds, or brains, or other insides. Others have tried guessing, but without much success.

When a person accidentally moves a bare foot too close to a fire, an observer can see the foot pull away from it. In Descartes' *Treatise on Man* (1631) he says "If the fire A is close to the foot B, the small parts of this fire, which, as you know, move very quickly, have the force to move the part of the skin of the foot that they touch, and by this means pull the small thread C, [running up the back to the brain] ... simultaneously opening the entrance of the pore d, e, where this small thread ends... the entrance of the pore or small passage d, e, being thus opened, the animal spirits in the concavity F enter the thread and are carried by it to the muscles that are used to withdraw the foot from the fire."

This sounds like an attempt to understand responses to stimuli, but 380 years later we can understand it as a description of a negative feedback control system, which we will get to before long.

If the observer happens to be the organism with the overheated foot, one more effect can be observed: it hurts. This leads to noticing that the foot is generally moved according to whether the sensed warmth is too little, too much, or just enough. The fire affects the sensed temperature of the foot in one direction; the response affects the same sensed temperature in the opposite direction. This turns out to be an exceptionally important observation. It's a pity that nobody could have analyzed it in 1631, but Newton's calculus then lay 73 years in the future. A differential equation would have explained this circle of causation that baffled philosophers of science until, 400 years later, control system engineering appeared.

3 Behavioral Science II

Just as PCT began to get organized, a new branch of behavioral science appeared: cognitive science. The emphasis moved from externally visible variables to those experienced by each individual. Now it was permissible to explore processes inside the brain and try to analyze them, but the phenomena to be explained scientifically were still basically the way stimuli cause responses. Theoretically, stimuli from the environment were now analyzed by cognitive processes in the brain, which then would formulate plans for generating responses appropriate to the stimuli.

The main task for the brain was now to figure out what commands should be sent to the muscles to generate appropriate results, given all the information coming into the brain from outside. This required the brain to have knowledge of neural and physiological processes as well as physical processes in the external world, and entailed rapid computation of the "inverse kinematics and dynamics" of body and environment ("kinematics" = properties of linkages, " dynamics" = movements of masses). Once this plan of action was turned into the set of necessary commands, it could be executed to produce the actions and their anticipated results.

There is something wrong with this picture. Rabbie Burns observed that the best-laid plans of mice and men gang aft agley, which is true not because we are bad at analyzing and planning but because plans of *action* are always close to their expiration dates. A planned action such as turning a steering wheel might produce exactly the wrong result if another car, a second later, changes direction by only a small amount. Planning all the turns of the steering wheel needed to drive from home to work couldn't conceivably get you to work the next day, no matter how precisely executed, even if exactly the same movements worked perfectly the day before. Think about other cars, traffic lights, pedestrians, weather, road repairs.

While planning clearly does take place, it can't operate by planning actions. We plan results, not actions, and that requires a new model of behavior. Even before cognitive science appeared, that new model was under construction.

4 Understanding Purpose

The new model was born in a parallel universe. Electronics engineers of the 1930s were using their new skills at designing electromechanical systems to automate tasks formerly done only by human beings. These tasks entailed a specification for some external condition to be brought about and maintained, even though it was impossible to predict or even detect all the events that might disturb that condition. The tasks included such things as aiming guns from the deck of a rolling ship; stabilizing the temperature of a room subject to opening and closing of doors and windows at unpredictable intervals on cool or cold days; adjusting the course of a torpedo to arrive at a moving target that made propeller-noises; keeping an airplane flying through rough air at constant altitude and speed, and on course.

To build such devices the engineers had to solve some basic problems. How could a (preferably) simple electromechanical device be given a specification for some effect that didn't yet exist, to be caused by a behavior that was not yet being carried out? How could this future state be made to cause an action in present time that would lead to that state? What if the effect of the action were disturbed *while* the device was producing the action? The engineers of the 1920s and 1930s, not knowing that the behavioral sciences had declared a device of this sort to be impossible (because future effects can't bring about their own causes), kept working at this problem until they solved it. The result was a new occupation called control system engineering, and (accidentally) a new theory of just about everything that lives.

These engineers had inadvertently discovered how purposive systems work. This discovery re-opened the door to the concept of intentional behavior directed by internal mental goals (which Watson, the founder of behaviorism, called a primitive superstition). The next logical step would have been to introduce this new understanding to the behavioral sciences. However, the sciences of life already had dozens of theories, all based on the idea that purpose is just causation misunderstood. They resisted mightily and that giant leap for mankind didn't happen.

5 Cybernetics *en Passant*

The Mexican physiologist Arturo Rosenblueth did notice the new ideas. He had been primed by studying under Walter B. Cannon, who worked to understand homeostasis, a process inside organisms that stabilizes critical variables such as nutritional state, body temperature, CO_2 level in the bloodstream, and other details of the life-support systems. Rosenblueth noticed that in the human body were many systems, behavioral systems, that appeared to work almost exactly in the way that the new artificial servomechanisms work. He communicated this discovery to Norbert Wiener, a mathematician at MIT where control engineering was rampant, and cybernetics was born.

Unfortunately, the main founders of cybernetics were not control-system engineers. They learned just enough about control systems to pattern cybernetic thinking around concepts like circular causation, but were more interested in subjects like communication, information theory, and (later) artificial intelligence and failed to carry the transformation to its ultimate conclusion.

That last step was not begun until the 1950s. That was when I learned of a recent school of thought called engineering psychology, and also started following the lead of W. Ross Ashby, a psychiatrist in the cybernetics movement who did have an understanding of control systems. With the help of R. K. Clark and R. L. MacFarland, I began to explore control systems with the idea of joining the cybernetics movement. After our first paper was published in 1960, we made overtures to psychology and cybernetics, but were put off by a general lack of interest. Clark and MacFarland went on to other things, and I kept working on PCT on my own. This led to my first book in 1973, then eventually to the formation of the interdisciplinary Control Systems Group in 1985, which in 1994 started a move toward becoming international by holding a meeting in Wales, and a few years later two meetings in Germany. The 22nd annual meeting of the CSG took place in 2006 at South China Normal University in Guangzhou, PRC, in collaboration with the Systems Society of China. PCT is part of the mainstream now. Almost.

6 A Scientific Revolution

The nature of a control system was almost understood by those who adopted behaviorism and cognitive science. There is something of each one in a control system.

The behaviorists realized, correctly, that behavior is based on perceptions that are caused by the physical events called stimuli. A driver can't keep a car on the road with both eyes closed. The kind of problem unsolved by behaviorism was how the stimuli could affect the driver's steering-responses in exactly the quantitative way needed to keep the car in its lane or steer it onto the correct exit ramp. This problem becomes worse when we realize that the driver also has to respond to *invisible* stimuli such as a crosswind. If the driver doesn't steer slightly into the wind by exactly the right amount, the car will drift into a ditch or into oncoming traffic. In general, stimuli as classes of happenings given names like "oncoming traffic" might lead to the right *consequences* of behavior ("avoiding collisions"), but are simply not the sort of thing that can produce the *quantitative amount and direction* of behavior needed.

Cognitive scientists realized, correctly, that behavior is the means an organism uses for achieving goals. An organism with a goal, they thought, must somehow figure out how to behave to achieve it. They noted, correctly, that the required behavior is not just a qualitative class of actions, but the quantitatively correct amount of action in exactly the right direction. The driver needs to perceive the environment to steer a car; the perceptions are supposedly the basis for the computations by which the organism calculates the actions that will achieve the goal. But it seems unbelievable that the driver could carry out all the repeated mental calculations required in the short time available, based on rather imprecise perceptions of what is going on out there.

In fact, neither behaviorism nor cognitive science hit on what now seems like the right explanation of behavior, though both hovered near it. The main mistake of both was to assume that the final product of brains was behavior, overcomplicated by the idea that behaviors must be exactly calculated.

7 The Solution: PCT

Here are the main questions unanswered by previous theories. How can stimuli produce not just responses, but *specifically appropriate* responses? What is a goal, that it can lead to just the behavior that will achieve it?

To answer these questions we have to look at things like perception and action a little differently. When someone steers a car, the perception that matters is the relationship of the car to the road as seen through the windshield. All the steering behavior has to be based on that perception—but not that perception alone.

It is also necessary for the driver to know, somehow, how that picture framed by the windshield *should* look if the car is to be properly located. This picture has to exist in the same place that the perception exists: in the brain. Without getting too neurological about this, we can say that whatever form the perception takes in the brain, the image of how the car and road *should* look must be in that same form, because the perception has to be compared with that image, the *reference image* ("goal:" goals are In Here, not Out There).

The difference between the imagined reference image and the real perception tells the driver how much steering error there is. "Error" just means the difference between reference and real. If the two coincide exactly, there is no error. If there is a mismatch in one direction, the driver should steer to the right. If in the other direction, to the left. That is basic control theory.

Now the cognitive scientist wakes up and says, "Yes, but exactly how much left or right? The brain has to calculate that, doesn't it?" The answer is yes, but. Yes, if there's a big error the brain should cause the steering wheel to turn a lot or if a small error, a little. But (and now we see the beauty of classical negative feedback control theory) the brain doesn't have to compute the exact amount because it can continuously adjust the action as the error changes, making smaller and smaller approximate adjustments as the error gets smaller until there is no error. Then no more changes in steering effort occur and the car is where it belongs in the lane. No complex computations. No planning. Just one swift simple process that converges smoothly to a final condition.

8 Behavior in the Real World

A driver traveling along a straight level road sees the picture in the windshield as exactly right; he steers neither to the right nor to the left. But is that true in the real world? Riding with a driver, we see endless little movements of the steering wheel, yet we don't feel or see the car moving left or right in its lane. The driver's steering efforts seem to be having no effect.

The reason is simple once you work it out. When the car starts drifting a little to either side for any reason, the driver immediately turns the wheel the other way as much as needed to keep the drift from getting larger, then a tweak more to eliminate it. If the driver can detect changes of the car's position as small as we can detect, or smaller, then we will never see or feel anything but tiny, barely-detectable, changes in position—if any at all. But the steering efforts can be quite large, in a gusty crosswind. It really looks as if the driver is responding directly to the crosswind, but of course in a closed comfortable car there is no way to detect the crosswind, except through effects on the car that the driver is mostly preventing. The result is that the deviations of the car are kept very small, especially in comparison to what would happen if the driver *didn't* make those steering movements. This is called negative feedback control—the same thing Descartes described.

So it seems that control means keeping disturbances from having much effect. But now, suddenly, the driver is turning the wheel so the car veers entirely out of its lane, a huge steering error. We immediately see why: it's an exit ramp. But why doesn't that steering control system act immediately to counteract the error? Because the reference image has been changed (one more time: reference image, reference perception, reference condition = GOAL). In fact, the driver's brain has smoothly changed the reference image from that of a car going straight in its lane to that of a car curving off to the right and up the ramp. The control system, still keeping the perception of the car's position matching the reference image, automatically alters the steering actions so as to keep the steering error close to zero. We see that simply by smoothly altering the goal of the behavior, the driver accomplishes the required change in behavior in an extraordinarily simple way, with no complex calculations.

9 Behavior: The Control of Perception

Behavior is the externally visible part of a process by which perceptions of various aspects of the experienced world are controlled. It is not the end-product of either the effects of stimuli or the goals sought by the organism. Behavior is simply the adjustable means by which an organism can keep its perceptions matching reference conditions. As disturbances come and go, behavior changes to have equal and opposite effects. As reference conditions vary, behavior changes to cause perceptions to vary in a matching way.

Behavior changes to cancel the effects of the disturbances on whatever the organism is controlling. The appearance is that the disturbances cause the actions, the observable behavior. But the real story is that the actions prevent the disturbances from significantly altering what the organism is concerned with: the perceptions it is controlling. This is how PCT explains the appearances that led to behaviorism.

When we make plans, the appearance is that we plan what behaviors will be needed to achieve what we want. But we can't predict what disturbances and changes in properties the environment is going to throw at us. What we can do is plan the perceived consequences we want to happen. We don't plan actions; planning successfully means planning perceptions. Higher levels in us tell lower control systems what perceptions to experience. The lower control systems adjust their actions to make their perceptions match the reference conditions they are given, and (without being told) enough more to cancel the effects of any disturbances that might be happening. This is how PCT explains the appearances that led to cognitive science. PCT does not require the brain to perform miracles of prediction and impossibly fast, complex, and accurate computations.

PCT thus encompasses the concepts of behaviorism and cognitive science, providing a single framework in which the observations of both can be understood. With one more added concept—levels of control—it expands to encompass all that human beings and perhaps all organisms experience.

10 Emotion

The control hierarchy can control perception at many levels by using actions from mild to strong, but there is something missing: feelings. This model doesn't suggest the physical feelings that accompany emotions, but one modification of the model can put feelings into relationship with the goals that go with them, to cover both the cognitive and feeling sides of emotion.

Disturbing higher control systems or changing goals causes errors that generate a cascade of changes in the reference signals passed down the hierarchy of control. We now divide this cascade into two branches. A behavioral branch goes to systems, mostly learned, that control using muscles. A somatic branch, primarily a product of evolution, goes into the amygdala, then the hypothalamus, and then the pituitary gland and autonomic nervous system which control the state of the body. This branch is where emotions supposedly originate, but in the PCT theory emotional feelings are effects, not causes.

Some control systems are inherited; most are learned. All act to adjust both the somatic systems and the action in the behavioral branch. The somatic branch adds sensations that generate the feeling component of the configurations we call emotions. Example: Either learned or innate systems can specify goals like escaping or attacking. If the perception differs from the reference, a "motivating" error signal is sent to multiple lower behavioral systems as reference signals. The effect of the error signal on the somatic branch provides the feeling part of the experience, the so-called fight-or-flight syndrome. The goals of attacking or fleeing distinguish fear from anger; the physiological states have been found to be identical in both emotions.

The feeling part of emotions often arises without any consciousness of the cause. This can happen if awareness is engaged at higher levels, and a disturbance occurs that affects lower-level control systems not currently in awareness. Those systems will react automatically by using the muscles and, according to this theory of emotion, will also adjust the physiological state of the body. The sensations arising from the physiological states will be processed level by level up the hierarchy, and when the perceptions reach a level accessible to awareness, will attract attention exactly as if they had occurred spontaneously, or had been caused from outside the body. An injection of adrenaline can be interpreted and experienced as fear *or* anger.

11 The Hierarchy of Control

The driver keeps the car in its lane, yes. But why? To stay alive, surely, but there are more immediate reasons. The driver has a destination in mind, and wants to get there. The reference perception: *I am at the entrance to the parking lot at the mall.* The actual perception: I am on 55th street a mile from the parking lot. So keep the car moving along in its lane. When the entrance appears, change the reference: *the car is following* ***this*** *path into the lot.*

The higher system is not telling the lower one what to do but showing it what to perceive. It does so by continuously varying the reference image, not by commanding steering wheel movements. The lower system automatically corrects the effects of disturbances and little steering errors on the car's path without having to be told to do it. The higher system needs only to alter the images that the lower system is to reproduce by turning the wheel. The lower system determines when, how much, and which way to turn the wheel.

The reason for going to the mall is to buy a dress shirt. The reason for buying the dress shirt is to look good at a wedding. The reason for looking good is to please the woman you're going to marry. The reason for pleasing her is that you want to show respect for her opinions. The reason you show respect for her opinions is that you want to make the marriage as ideal as you can, and see respect as an essential principle for making a good marriage.

Each level of control sets multiple goals for the next level down to perceive; that's how any higher system controls its own perceptions. The higher system's perception is built out of the perceptions that exist, some being controlled, at lower levels. There are many control systems at each level, and more than a few levels. The only systems that act on the environment directly are those at the first level. All the rest act by adjusting the perceptual goals for lower systems. All control their own perceptions, not their actions.

Now you know the essence of Perceptual Control Theory, which replaces the basic concepts of behavior in both behaviorism and cognitive science. A revolution, in progress.

Bill Powers,
Lafayette, Colorado, October 2009

This series continues with ***Reorganization and MOL***, an overview of how control systems may come into being, change, cause internal conflict, and ways to resolve internal conflict.