The world according to PCT

W. T. Powers[1]

> *Beware of the man who works hard to learn something, learns it, and finds himself no wiser than before... He is full of murderous resentment of people who are ignorant without having come by their ignorance the hard way.* — Kurt Vonnegut, "Cat's Cradle"

In the 1930s electrical engineers discovered how to build systems of a kind that could do the following things:

1. Act on the world outside them to cause changes of temperature, speed, position, force, rate of rotation, altitude, and acidity, to mention at random only a few possible variables.

2. Bring any of those aspects of the world to a desired goal-state, starting from any previous state, by acting directly on it.

3. Keep any of those aspects in the goal-state (or it bring it back to the goal-state) even if unwanted variations occur so that different amounts and directions of action are required each time to achieve the same result, and do this even if the variations are unpredictable, and even if the causes of the variations are invisible and unknown.

4. When the specification for the goal-state of a variable is changed, alter the action on the physical world however necessary in the current environment to bring the sensed magnitude to the new value and keep it there.

Devices of this kind have been called, ever since the 1930s, negative feedback control systems. They are goal-seeking, intentional, purposive organizations of matter and energy.

However, by the time the nature of such devices was discovered, scientists who were concerned with a strictly scientific study of the behavior and functioning of living systems had reached a consensus that for any physical system, the behavior described above is impossible. And although their descendants in cognitive psychology and kindred fields, unlike their predecessors, may sometimes use words like *purpose* and *goal*, they are still uncomfortable with them because there seems to be no principled way to account for such phenomena. Crucial aspects of their theories and methods remain rooted in the old ‘strictly scientific’ view.

How was this counter-factual conclusion reached? Observably, cause precedes effect in physics, so it was assumed that, since living things must obey the laws of physics, behavioral actions are likewise preceded by causes (‘stimuli’) and followed by effects (‘responses’) in strict temporal order. It was further reasoned that since the causes of behavioral actions had to precede them, and since the actions in turn cause the observed results, there is no way for a goal — a specific result of the action — to reach back through time as a cause of the action from which it results. There was no way for a goal-state to ‘stimulate’ the behaving system into acting in just the way needed to produce that very goal-state as a ‘response’. It was concluded that the appearance of goal-directed, intentional, or purposive behavior had to be an illusion, and it was the duty of science to put that illusion aside along with all the other myths and illusions of an earlier time, and to search soberly for the correct explanation for these misleading appearances of behavior.

---

[1] Thanks to Bruce Nevin for comments and editorial assistance.

Many behavioral scientists nevertheless could not reconcile this conclusion with their experience. They rejected 'reflexology' and 'stimulus-response psychology' and studied behavior in a naturalistic way, accepting even unexplainable appearances at face value. They tried to adhere to the basic tenets of scientific research as much as they could, using statistics and mathematical modeling with some competence, but there remained an apparently unbridgeable gap between the empirical, functional understanding of the 'soft sciences' (as they came to be called) and the 'hard sciences' based on fundamental principles of biophysics and biochemistry. This discontinuity between the levels of analysis of the two camps persists today.

The soft sciences couldn't point to physical stimuli causing specific behaviors, but in the effort to be more scientific they could at least claim to study cause and effect by looking at patterns in the relations between environmental influences and general kinds of behavior. They studied phenomena that the stricter behavioral scientists could not, but their conception of linear causality from environment to behavior was the same, and their methodology was highly similar. They could study the relationship of childhood poverty to a subsequent life of crime and show that a link existed (for a significant portion of the population studied), even if the resultant 'general linear model' could give no insight into the form of the mechanisms involved. In this way, the same basic model of linear causation that was adopted by behaviorists has persisted in fields like cognitive psychology, sociology, and personality theory. Only prior causes can generate future effects, and goals or purposes represent future effects, therefore they cannot affect behavior. Either goals are illusory, or they exist in a subjective 'mental' world with little or no explainable connection to the physical world.

In the course of searching for the correct explanation to replace the illusion of purpose, behavioral scientists have proposed many theories, but first they had to change the way behavior is described. Since the appearance of goal-seeking is an illusion, a person who wants to be considered a scientist can't speak of what an animal or a person is 'intending', what it 'wants' or is 'trying to accomplish'. We observe that the environment stimulates the organism's sensors, generating neural signals that are routed by various known paths to the muscles and glands where they produce movements and secretions like those to which we give the names of behaviors. When the muscles contract, we see not simply *actions*, but *conditioned responses to stimuli* that produce effects radiating out into the environment. This, it was agreed by large numbers of scientists, is the only scientific way to describe behavior. Observable effects always come after observable causes. Goals are easily understood: they are simply outcomes of behavior determined by prior stimuli.

If a person doesn't happen to know of the existence of negative feedback control systems, this line of reasoning and its subsequent influence on interpretations may well seem inescapable. If one wishes to investigate living systems, and also to do this scientifically, environmental determinism seems inevitable. There is apparently no other path to take.
.
But in ordinary experience we do seem to have goals, intentions, wishes, desires, objectives, preferences, and targets for a great many things that we haven't achieved yet. Such goals and so on appear to exist; it is a very convincing illusion. When we read a description of what a negative feedback control system can do, it not only seems quite reasonable and familiar, but it sounds just like the things we can do ourselves. We can easily imagine needing milk, driving a car to the grocery store, buying milk, bringing it back home, and stowing it in the refrigerator. That looks very much like carrying out a series of intentions, like having a goal or goals and then achieving it or them. It's hard to imagine that this is an illusion. It seems completely real. [2]

---

[2] Imagining also appears really to occur,

It's even harder for a scientist to portray what happens without making it sound goal-directed and intentional. Yet that is exactly what large numbers of scientists have tried to do. Somehow, they had to describe what was observed so that the responses were always being caused by stimuli from outside, and show that while the outcome of an action might be some result beneficial to the organism, the action was produced without reference to the benefit, certainly not *in order to get* a benefit that did not yet exist. The very term ‘response’ carries the presupposition that it results from some prior stimulation. By careful use of language, scientists could bypass the illusion of purpose. This was not easy to do but they did it — or thought they had done it.

And all in vain. It was a futile exercise, a waste of effort, simply because negative feedback control systems are real physical systems subject to all the laws of physics and chemistry, yet they can do all those impossible things mentioned above, just as all organisms do them. All the arguments saying that such things can't be done are swept aside when we see them being done — more easily, of course, when we understand how they are done. We can stop trying to imagine how to describe behavior in any way but the right way, which — cross your fingers and knock on wood — we now know with as much certainty as today’s science permits.

In tracking down the place where the sciences of life went astray in the opening decades of the 20th Century, we are led to discoveries about the nervous system made early in the 1800s. It was discovered that there were two kinds of nerves, one carrying signals from sensory organs inward into the brain and the other carrying signals from the brain outward, from higher centers to lower centers and finally to the muscles (and glands). As the 20th Century started, the great neurologist Sherrington saw the outgoing signals as general commands formulated in the cerebral cortex and then elaborated, level by level, into the detailed commands necessary to cause complex and coordinated muscle actions. Also level by level in reverse order, the incoming signals were seen as stimulating the brain to formulate those commands.

We can be almost certain now that Sherrington was mistaken about the functions of those signals, particularly about the idea that the signals going from higher centers toward the muscles are commands to produce actions. The signals are not commands. Instead, they are specifications not for actions but for sensory perceptions of the results of actions — ’reference signals’ in the language of control engineers. The message they carry is, "Make the perception you are sending to me look like *this.*"[3]

The lower systems then, level by level, immediately alter their actions as required, the lowest system changing the environment (which includes muscles) so that the perceptual signals rising from sensory neurons toward the spinal cord are adjusted to match the lowest level of reference signals. At each more general or abstract level, perceptual signals are made to match the reference signals descending from the next level above.[4] So it is not a future event that acts at any level as a goal. Real

---

[3] Sherrington, a hundred years ago, realized that "the brain thinks in terms of movements, not muscles." He found that brain stimulation in a particular location could cause movement of a limb to a specific place from any starting position. Different efforts and even different muscles were required and only the final position repeated. But he couldn't explain how that could happen. Control theory can.

[4] At the top level, the reference signals are from some genetic or other fixed source. The values of fixed reference signals can also result from learning, a topic which is beyond the scope of this introductory survey. In broadest terms, learning progresses from ability to perceive a variable, through ability to influence its value, to establishing a preferred state for it and controlling the perception.

neural signals that exist here and now define a hierarchy of outcomes to be achieved, just as a real blueprint, here and now, specifies via a general contractor what the subcontractors and their employees are to build. Each system continuously compares perceptual signals (derived from a combination of perceptual signals from lower systems) with the reference signal it is receiving from above. They continue to act until the difference between reference and perception is close to zero. Most electronics engineers would recognize that organization as what they call a 'cascade' of negative feedback control systems. We call it a hierarchy of control.

To see the nervous system as a collection of control systems is thus no more controversial than it is to see the heart as a pump. The heart is a pump because it pumps. The nervous system is a (multidimensional, multilevel) control system because it controls. Of course a pump is easier to understand for someone who knows how pumps work than for someone who doesn't. But control systems have also been easy to understand since engineers figured out how they work. Unfortunately, they didn't figure it out until well after life scientists had committed themselves to some serious theoretical mistakes. This book is about those mistakes and how we can fix them.

Before any 'fixing' can be done, however, it has to be permitted to happen. Here is the barrier that prevented the discoveries by those control-system engineers of the 1930s from immediately revolutionizing all the life sciences. Even those who quickly approved of what the engineers had done failed to understand the disruption this would cause to established careers and their paradigms of work. All of us who hopped onto the cybernetics bandwagon and later took off in still newer directions naively expected these new ideas to be appreciated and adopted by traditionally open-minded scientists who preferred knowing the truth to being right.

The massive opposition from some quarters and the passive resistance from others came as a disappointing surprise, but perhaps it shouldn't have. Science has a social as well as an intellectual aspect. Scientists are not stupid. They can look at an idea and quickly work out where it fits in with existing knowledge and where it doesn't. And scientists are all too human: when they see that the new idea means their life's work could end up mostly in the trash-can, their second reaction is simply to think "That idea is obviously wrong." That relieves the sinking feeling in the pit of the stomach that is the first reaction. Being wrong about something is unpleasant enough; being wrong about something one has worked hard to learn and has believed, taught, written about, and researched, is close to intolerable. All scientists of any talent have had that experience. The best of them have recognized that their own principles require them to put those personal reactions aside or at least save them for later. They know that any such upheaval is going to be important and that they will never regret getting over their humiliation. But those who recognize and embrace a revolution in science are the exception. Most scientists practice 'normal science' within a securely established -- and well-defended -- paradigm.

That is what we are up against here, and have been struggling with since before most of you readers were born. We have spent that time learning more about this new idea and getting better at explaining it, but no better at persuading others to change their minds in a serious way when their career commitments are threatened by it. What we had thought would be a nutritious and deliciously buttered carrot has proven to function like a stick.The clearer we have made the idea, the more defenses it has aroused.

We are now facing reality. This is going to be a revolution whether we like it or not. There are going to be arguments, screaming and yelling or cool and polite. It's time to sink or swim.

As the father, or by now the grandfather, of Perceptual Control Theory, I have invited [eight?] colleagues to join me in making the case for PCT in this book. Since there are many more than [eight] people who could equally well be offered that invitation, I have asked members of the Control Systems Group around the world to nominate and vote for co-authors who are well-acquainted with some conventional branch of the life sciences as well as with PCT — who understand why the older theories were persuasive and even useful, and also what has to be changed about them if the realities of negative feedback control are to be introduced into them.

To make the contributor list more inclusive, I'm asking each selected colleague to serve as the organizer of a cooperative authorship of each chapter. Many of us have found that group editing of a common document[5] is an eye-opening experience that changes competition into cooperation, and self-interest into a desire that the document itself be better than one contributor could make it. If this part of the introductory chapter is still here to be read, you will know that this experiment worked. There are many other books and articles explaining PCT. Here you will see it explained in specific contexts. I will add brief comments after each chapter; aside from that, this is good-bye for now.

---

5 With many thanks to the authors of Google Docs.