Perceptual Control Theory: A model for understanding the mechanism and phenomenon of control

Abbott, Carey, Goldstein, Kennaway, Mansell, Marken, Nevin, Powers, Robertson, Taylor, Yin

The phenomenon of control is integral to psychology. Even a cursory glance through databases and journals reveals a staggering number of references to control. Terms such as *perceived control*, *locus of control*, *cognitive control*, *subjective control*, and *vicarious control* speak directly to the phenomenon. If we include implicit references to the phenomenon, such as *self-determination*, *agency*, *learned helplessness*, and *emotional dysregulation*, the number of references grows exponentially. Recognition that control has an important place in the process of living, therefore, is ubiquitous.
But intuitive recognition of control has been divorced from any formal understanding of how control works. Without some kind of aid to understanding, the principles of control evade intuition.

The most common control mechanism in use is based on a principle with ancient roots. To maintain a constant water level in the supply reservoir of a water clock, Ktesibios of Alexandria devised a float-valve regulator—a negative feedback control system—in about 230 B.C.[1] For the next thousand years or more, that is how water clocks managed to keep reasonably accurate time. That principle of negative feedback control spawned many devices for the next twenty-odd centuries, from temperature regulators for furnaces to speed controls for steam engines to steering mechanisms for ocean liners to autopilots for aircraft to star trackers for spacecraft. However, it was only 70 or 80 years ago that engineers found a formal way to understand systems like this and named it *control theory*.

The problem with understanding control theory is that it violates intuitive ideas of cause and effect. In a float-valve regulator, any small increase in water level raises the float and closes the valve a little; a decrease opens the valve a bit. The inflow changes until it matches the outflow (adjusted as needed needed to operate the clock), so the water level remains constant . The difference between inflow and outflow determines how fast the water level rises or falls; simultaneously, the rise or fall of water level, by operating the float and valve, changes the inflow. Which is the chicken, which the egg? This chain of cause and effect runs in a circle, and therein lies the difficulty. The only simple and intuitive way to think of this chain is to break it up into a sequence of two cause-effect events. First A causes B, and when that process has produced a new value of B, B causes the next value of A. That is *too* simple; it is the wrong analysis of all but a very few kinds of very crude control systems. In most real control systems, A and B are affecting each other at the same time and in different, non-reciprocal, ways. Simple cause-effect logic can't handle that.

PCT is based on observations of human and animal behavior, but also on the mathematical analysis of artificial control systems worked out by engineers in the 1930s and during the war years that followed. However, the best introduction to PCT is to see some control systems in operation, with the

[1] Landels (1978).

mathematical analysis introduced only through simulations where it is easiest to grasp. In the following section we will assume that you have downloaded two suites of demonstrations into a PC (or into a Macintosh running Windows emulation software). Instructions and links are in the Resources section at the end of this paper. Downloading and running these programs will save you large amounts of time and difficulty in understanding PCT.

Simulations and models

The first demonstration is taken from the set called 'LCS3 Demos'. It is a simple pursuit tracking experiment like those studied in the early 1950s by engineering psychologists, but done using a modern desktop computer. In the menu of demonstrations it is called "Demo 4-1: TrackAnalyze." Instructions can be viewed by clicking a button with that name.

In this demonstration, a person uses a mouse to make a cursor track a moving target for one minute. After a few seconds for you to acquire control, data are sampled 60 times per second. The upper plot of figure 3 shows the data for a single one-minute experimental run. The red trace shows the target movements; the green trace shows the mouse/cursor movements. The black trace shows the difference between target and cursor—the tracking error. Note that the mouse movements are based on the difference between target and cursor position, while that difference depends, simultaneously, on the mouse movements. There are no stimulus events or response events.

There is a consistent small time delay, hard to see in Figure 3, between target movements and cursor movements (upper plot). The delay is not removed by anticipatory mouse movements. In the upper part of Figure 3 the results of fitting a negative feedback control model to the data are summarized; the best-fit delay in the model's response is 7/60ths of a second, which is 7 frames of the computer display running at 60 Hz, or 116 milliseconds. That is how far behind the target movements the participant is moving the cursor, on average. The delay does not get smaller with practice.

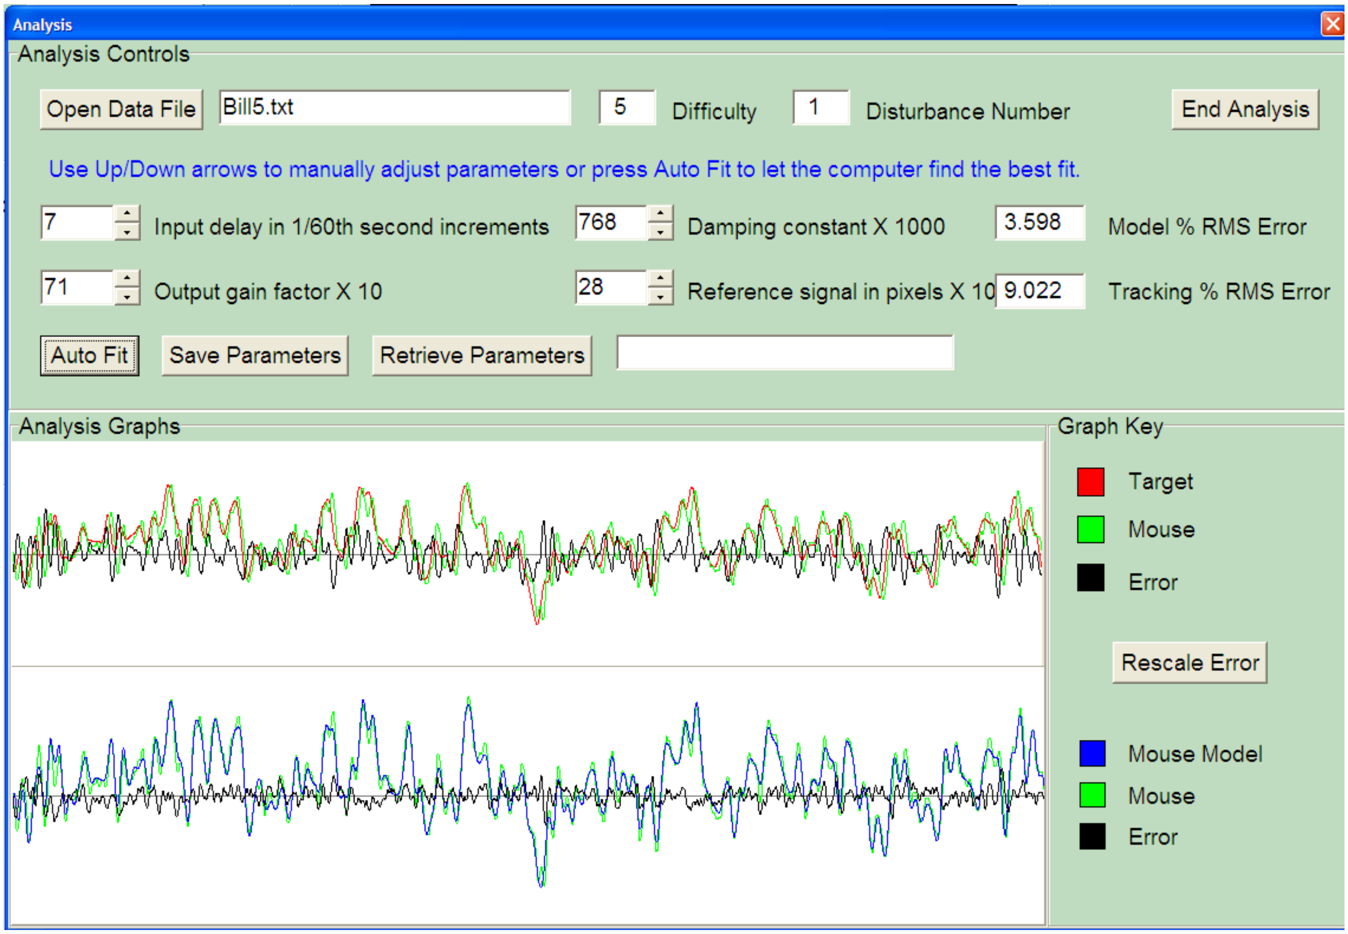

Figure 3. Analysis of human tracking run and fit of negative feedback control model to the data. Upper traces: experimental results; lower traces, match of model (blue) to the real mouse movements (green).

The lower plot compares real human performance with the performance of a computer implementation of a negative feedback system controlling the proximity of the cursor to the target, resisting unpredictable random disturbances to the target position just as the human subject did. The same target movements are used for model and human runs. The model's simulated mouse movements (blue) are compared with the person's real mouse movements (green). They are nearly identical, with the same delay relative to the target movements. The mean (RMS) difference between model and real behavior is 3.6% of the range of target movement. In this run, the target movements are rapid enough (maximum difficulty) that the tracking error is 9% of the target range, which to the subject seems like very poor tracking. The model fits the person's performance well within the tracking error, showing that the model is making similar mistakes. This same model will control even better with the delay set to zero, but it will work too well. With no delay and all the remaining parameters optimized, the mismatch with this sample of human behavior rises from 3.6% to 6.0%. The delay is real.

Soon after control theory started to appear in the psychological literature, various commentators noted that all real systems contain time delays. It was thought, apparently, that with any time delay at all, error-correcting actions would start too late to prevent disturbances from having immediate effects, and would persist after disturbances disappeared, generating self-disturbance, with the likely result that the whole system would oscillate violently. "Feedback is too slow" became a mantra for those resisting the new propositions. But this model contains a delay, and it is not only stable, but it matches real behavior closely.

Before we go on, some terminology has to be introduced, along with a diagram of the negative feedback control model. Figure 4 shows a generic model of a single control system, one system among many at one level among many: the building-block of the hierarchy of control systems that constitutes PCT.

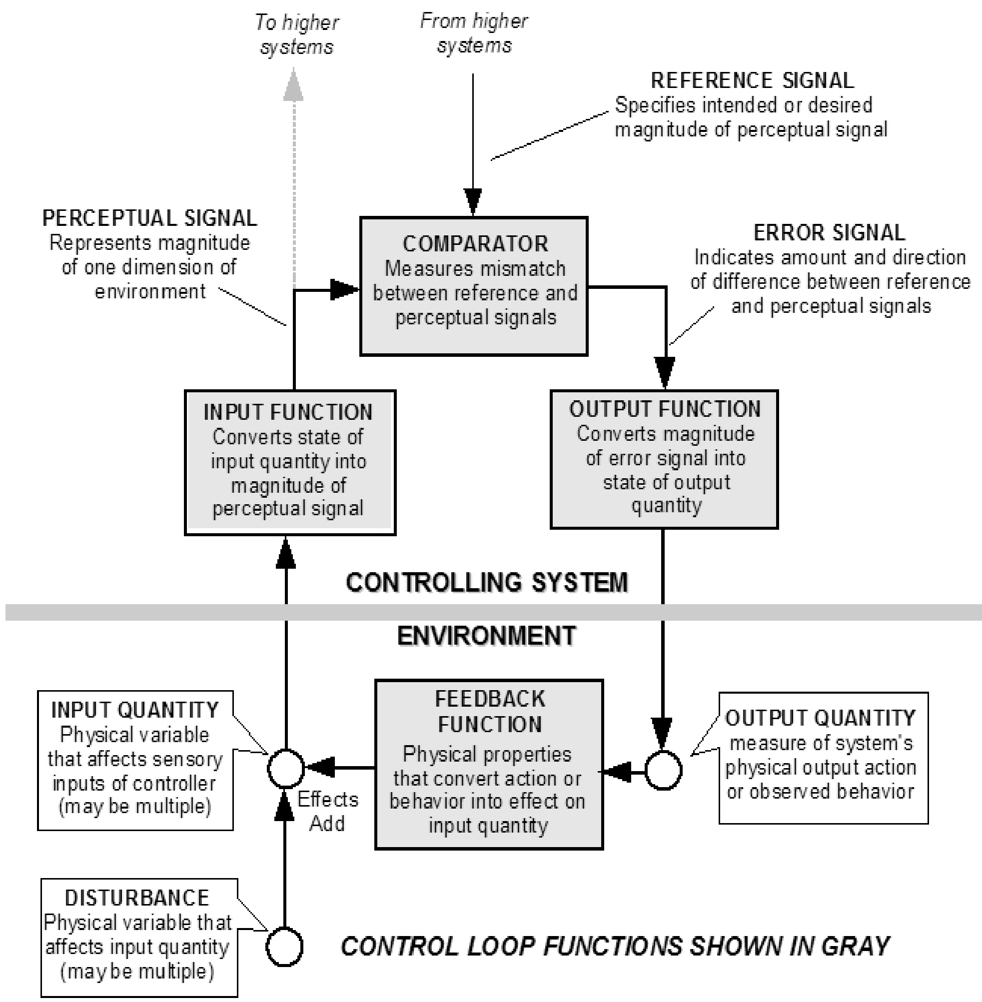

Figure 4:The basic organization of a negative feedback control system.

The what, why and how of control

Consider an ordinary, everyday situation. I have my finger on a button beside a door. If I pause to think about what I am doing, it seems simple enough: "I'm ringing the doorbell". But is that why I am here? Am I not trying to get someone to open the door? I am visiting Aunt Mary. That is why my finger is on the button. If you were a stranger passing by, you wouldn't know this. You might guess that I am trying to add to the expected vote total for my preferred candidate in an upcoming election, or that I am making some money by delivering pamphlets. Whatever motivation you imagined, you would probably not propose that my objective was simply “to press my finger on a button”.[2]

Multidimensional and multiordinal control

Any single control process can be modeled in isolation, as in the initial diagrams in this paper, but a model of the behavior of organisms must represent many control processes acting at once. In PCT, multidimensional control is modeled not as if complex signals or vectors were under control, but in the style called by Oliver Selfridge ‘pandemonium’,[3] in which many one-dimensional controllers are acting at the same time. Because each controller senses just one dimension of variation, complex control requires many one-dimensional controllers to be working in parallel. While this seems wasteful of neural resources, with considerable duplication of function (interactions between systems have to be handled by using the same input variables in many different control systems) , the resulting models are in fact computationally simple, and the bottom line is that they reproduce real behavior accurately, the *sine qua non* of model-based analysis.

The Russian physiologist Nicolas Bernstein[4] anticipated cybernetics in many ways, and in the 1950s came to the same conclusion that was being developed at the same time in the ancestral “feedback theory” that became PCT: behavior has to be multiordinal—organized hierarchically, in layers. The same simple problem led to this conclusion both in PCT and in Bernstein's work: if the spinal reflexes act to stabilize limbs against disturbances, they will prevent centers higher in the brain from using those limbs to carry out behavior. Any disturbance will cause a reflexive reaction against the disturbance. Since the brain obviously does use the spinal systems in producing behavior, there must be a principle that allows the higher systems to operate by incorporating the reflexes, not just by overcoming them or turning them off. This principle can be extended to higher feedback loops, each higher loop presenting the same problem to subsystems above it.

Bernstein never completely settled this problem. He was on the right track, but he lacked expert knowledge of the engineering principles of negative feedback control which inform PCT. The secret lies in the reference signal, the (variable) “Goldilocks standard” against which perceptual signals are judged as being too small, too large, or just right (though on a continuous scale). To use a reflex-type control system as means of control, all that the higher systems have to do is vary the reference signal.

This establishes a principle of hierarchical control that seems to apply equally well at many levels of organization. Higher systems act to control their own perceptual inputs by telling lower systems to produce a specific amount of the variable they are specialized to sense and control, not what action they should perform in order to produced the requested perception. How much to sense, not what to do. The lower systems autonomously act on their environments to make their own perceptual inputs match the specified reference condition of the moment.

---

[2]See Vallacher & Wegner (1985); Kozak, Marsh, & Wegner (2006); Marsh, Kozak, Wegner, et al. (2010). Links to these papers are at http://www.wjh.harvard.edu/~wegner/actid.htm.

[3]See http://en.wikipedia.org/wiki/Pandemonium_Architecture.

[4]Bernstein (1967).

Conflict and cooperation

The concurrent control of input variables by different controllers can result in conflict. An everyday example of conflict within the hierarchy occurs when a parent wants to warn someone of a hazard, but does not want to wake the baby. Warning the person calls for using a loud voice; letting the baby sleep calls for quiet. Two control loops are controlling the same perceived loudness of one's voice, trying to produce very different values of the same variable at the same time. The two controllers may be in different people: one person may focus on not waking the baby, the other on handling the hazard. While a conflict remains unresolved neither controller can achieve its goal; both are impaired, and one or both may effectively be removed from functioning. Psychological difficulties with this basis are addressed by the Method of Levels (MOL), a PCT-based method of therapy that will be described in a later section.

Both conflict and cooperation have the same formal description in the PCT model: two or more controllers are controlling their perceptions of one common variable in their environment. Cooperatioon requires resolution of conflicts. If two cooks interfere with each other, let one of them make the salad dressing while the other seasons the broth. When very skillful control of exactly the same variable is involved on both sides of a conflict, even small differences in the goals can cause large degrees of opposing efforts.This explains why cooperation , even when highly valued, can be difficult to put into practice.

The resolution of conflicts requires changes in the control systems that create behavior—in the given environmental situation, the perceptions or the actions (as physical forces or reference signals for lower systems) must become different.

Changes of organization

The final facet of PCT is concerned with adaptation and ontogeny. Both are accomplished by changing the structure of the control hierarchy, either to change the way it behaves (adaptation) or to develop capacities for control that it did not previously have at a more primitive stage (ontogeny). In accord with the general principles of PCT, this process of changing the structure of control systems is seen itself as a control process, in which variables of basic importance, referred to by Ashby as 'essential variables', are maintained near reference states by altering the organization of the organism. The resulting theory of *reorganization* incorporates one of Ashby's most important and still viable ideas, that of 'ultrastability' achieved not through systematic control or direct changes in behavior, but through random variations of the properties of a system.

The worst case for a learning theory is that in which the organism needs to learn some control process that has never been learned before, which cannot be extrapolated or generalized from past experience, and which cannot be worked out logically, either because it is unique or because the organism has not developed (or may never develop) the requisite capacities for using logic. This problem led Ashby to a control-system theory (adopted into PCT) of the most basic kind of learning, the kind that has to precede the learning of any systematic methods of learning. It is this sort of worst-case learning that is usually meant when the term *reorganization* arises in PCT discussions. The first version of this theory was not given much weight early in the development of PCT, but under the tutelage of a small bacterium it has developed into a testable model of basic learning.

E. coli reorganization

B. F.Skinner explained the acquisition of the first successful behavior in conditioning experiments by saying that organisms spontaneously “emit” random variations of behavior. PCT adopts that idea but in a different form: the basic theory of reorganization is that the organization of system parameters (and hence behavior) varies randomly at a rate that depends on the amount of *intrinsic error*. A condition for successful operant conditioning is that animals are deprived of whatever is to be used as a reinforcer. Conditioning with food reinforcers, for example, does not work on satiated animals. Deprivation is not just an ‘establishing condition’ as Skinner called it. It causes errors in control systems intrinsic to the organism that are organized to produce inputs of food, water, warmth, light, and so on and keep them at levels in the range necessary for comfort and even life. The theory of reorganization that is part of PCT proposes that when such basic control errors are large enough or protracted enough, they bring reorganizing processes into action.

Intrinsic error means a difference between the state of some essential variable, like one that reflects blood glucose concentration, and a genetically-determined reference condition. This difference is proposed to result in random changes of organization, trial-and-error learning that can work when there is no systematic method and no prior experience to guide changes.

Clearly, if the random changes of organization produce new behavior patterns that eliminate the deficit in blood sugar, the intrinsic error driving those changes will be eliminated and the changes will stop. That will leave the latest result of reorganization in effect, and behavior will show the new patterns from then on, just as if something—a ‘reinforcer’—had somehow told the organism that the new pattern was a good one for maintaining the right blood sugar levels. Reorganization, however, does not require that kind of knowledge or intelligence. It can produce organized behavior such as walking in a figure-eight pattern to correct a food deficit caused by a psychologist issuing rewards for walking that way -- a connection that is entirely unnatural.

This concept has been part of PCT since the first published paper in 1960, but it seemed at first too inefficient. Doubts about random reorganization are reasonable; at first blush it doesn’t seem very likely to work. Not until 1980 was it taken seriously. In that year, Daniel Koshland published a small book on bacterial chemotaxis which contained a principle that vastly increases the effectiveness of random reorganization.

The bacterium E. coli cannot steer, but it can make its way up and down chemical gradients very effectively. It does so by swimming in a (more or less) straight line and occasionally ‘tumbling’, changing direction in a way that Koshland reported was actually random. The explanation of the gradient-climbing is found in the observation (verified by perfusion experiments and biochemical analysis) that E. coli senses the time rate of change of concentration of chemical substances, change that is normally induced by its swimming in the gradient. If the rate of change of an attractant is positive, E. coli continues in a straight line. The attractant is diffusing radially from a source in the fluid medium, so the straight-line path of the bacterium may be visualized as a tangent line across concentric circles around a point, gradually reaching a closest approach to the source. As the path then starts to draw away from the source, the time rate of change of concentration goes negative, and E. coli immediately tumbles.

Since the tumbles change the direction of swimming at random, the result is just as likely to be worse as it is to be better. If the rate of change is still negative, however, another tumble ensues immediately, and tumbles keep repeating until the rate of change is once again positive. The bacterium does not swim far—a few body lengths—before tumbling again, so it does not travel much between successive tumbles. The result is that it travels much farther and faster up the gradient than down it. For repellents, meaning substances that E. coli avoids, the relationships are reversed. According to Koshland, E. coli can behave in this way in relationship to more than 20 different substances—simultaneously, apparently.

To translate this principle into terms of reorganizing a control system in a simulation, the spatial dimensions in which E. coli moves become multiple parameters of a control system. Swimming in a straight line becomes, in a simulation, adding small increments again and again to each parameter being varied, the direction of travel in parameter space being determined by the relative positive or negative amount of change per iteration in each dimension. A tumble corresponds to altering randomly the proportions in which different parameters are changing. To make sure the process does not overshoot its purpose, the amount added to each parameter is reduced as the control errors decrease.

In comparison simulations, the E. coli principle has proven to be over 50 times more efficient than a method based on random point-mutations of parameters, the proposal in the initial 1960 paper. This is because it makes use of information about the changing size of control errors. A slow progressive parameter change that continually reduces control errors simply continues as long as improvement continues. Only when the control error worsens does a ‘tumble’ take place, and then tumbles occur rapidly until the errors are declining again. The 50-fold gain in efficiency is seen when only two parameters are varying; the larger the number of parameters being reorganized, the greater is the gain in efficiency.

Reorganization theory replaces two other theories of change, Pavlovian or classical conditioning theory and Skinnerian operant conditioning theory. Reorganization—change in the internal organization of the behaving system—occurs only at special stages of these processes, and much of the behavior that looks like learning or adaptation can be produced by control systems without any change of internal organization.

Classical (Pavlovian) conditioning

Pavlovian conditioning begins, we propose, with an existing control process, either learned or inherited (a “reflex”). We will consider the example of thermoregulation. An animal subject to cold air blowing on its body will eventually start to shiver: seemingly, the stimulus of cold being sensed by skin receptors is causing the response of shivering. But bodily temperature control is most directly concerned with maintaining brain temperature close to a reference level. When loss of heat to a cold environment disturbs the core temperature, sensors in the brain report the reduced temperature as a reduced perceptual signal. When the perceptual signal is compared with a reference signal (or equivalent) that defines the current reference temperature, an error signal is generated, .which activates the shivering, which creates heat that opposes the effect of heat loss on the temperature, which reduces the magnitude of the error signal -- a simple negative feedback control loop. This would happen even if the skin temperature sensors were not involved.

The sensory effects occurring at the skin are not necessarily involved at first. However, there is a time lag between onset of cooling of the core and the buildup of shivering, and during that time the core temperature can drop significantly. This would be classed as a change in an essential variable, per Ashby, and that will activate a separate “reorganizing system”[5] that monitors essential variables, starting random changes in organization of existing control systems such as the one just defined. Those changes of organization will keep happening until they have the effect of reducing or eliminating the error in this basic inherited control system. Changes that do not accomplish that result simply allow reorganization to continue so further changes occur. What sort of change would correct the intrinsic temperature error detected by the reorganizing system and thereby stop the changes?

If shivering started earlier after the onset of heat-loss, the core temperature error would be counteracted sooner and perhaps might not occur at all. This is where the skin temperature, treated now as a stimulus that will become a CS, a conditioned stimulus, comes in. If the shivering could be started as soon as the skin temperature dropped significantly, it would start adding heat to the volumes where core temperature is sensed before the cooling can have much or any effect. This could happen if the input function of the core temperature control system were to be reorganized.

Reorganization changes the weights given to signals in other control systems at random, raising and lowering them at a rate that depends on the amount of “intrinsic error” detected in one of these reorganizing systems. If an input function is being reorganized, sooner many different input signals will become connected and then disconnected again as reorganization continues: the synaptic strengths will rise, then fall again; neural fibers might grow or atrophy. Let us suppose that one of the input signals that becomes connected to the functions that sense core temperature is the signal from the skin sensors that report a drop in skin temperature: “cold receptors”. The controlled variable will then become not just the core temperature, but the core temperature minus the “cold” signal[6]. If skin temperature now drops, the perceptual signal in the control system will decrease. As it falls below the reference temperature, shivering will commence and prevent the core temperature from dropping enough to make reorganization start. This new organization will therefor persist .

The same process can happen if there is any sensory signal that repeatedly precedes the core cooling and is capable of being connected with the correct sign to the input of this basic control system. The anticipatory signal could be a tone sounded before the experimenter causes cold air to blow on the organism, provided that the tone repeatedly stays in the correct temporal relationship to the core cooling.

Reorganization is not logical; no reasoning is involved. The tone ends up causing shivering because

[5]Ashby called his a “uniselector.” In a brain, this is probably a collection of localized reorganizing systems.

[6]The “cold” signal will have to make an inhibitory connection for reorganization to stop.

when that connection is made at some point during the random shuffling of connections, the shivering starts before the core temperature falls, and the reorganizing process stops.[7] If that did not happen, that connection would not have formed or persisted after forming in the first place. If it already existed, that connection would be reorganized away. That process is called extinction.

This model has not yet been simulated to see what other phenomena of classical conditioning it could reproduce, or to reveal any hidden flaws in the above proposal. That test will no doubt be done in the near future.

Operant (Skinnerian) conditioning theory

We can now see classical conditioning as a process of learning new perceptions to control. The same reorganizing process that creates the phenomena of classical conditioning can also explain operant conditioning. The main difference is that reorganization now works primarily on the output side of the control system rather than the input side.

In operant conditioning, the situation is more complicated because behavior affects one of the variables that is also considered a cause of the behavior. If that sounds like the kind of closed loop on which PCT is based, it should.

In all the basic forms of operant conditioning such as a fixed-ratio experiment, there is first a deprivation of something like food that is important to the organism. Again, we begin with an error signal in some basic and presumably inherited control system. Subjected to this ‘establishing condition’, the animal ‘emits’ whatever behaviors have already been learned that can counteract the deprivation. It is very likely that the animal already has learned or inherited some strategy for finding food, such as executing a search pattern that bring all parts of the local environment within range of food detectors employing vision, or more likely smell. During this search, the animal’s actions by chance actuate something in the experimental apparatus (depressing a lever, typically) to cause a bit of food to be dropped into a feeding cup, where the animal finds it and eats it. Typically, the strategy is organized to adjust so that the animal begins exploring in the area where the food was obtained. This increases the probability that the lever will be contacted again and more food will be delivered. So far, no learning, no change of organization has occurred. But now, gradually, the random-looking explorations begin to focus on the lever, and after some time the animal is pressing the lever in a way that looks organized and purposeful so as to provide for itself as much food as it needs (if the experimenter has been generous).

Two different processes appear to be working here. The first one is simply the search for food and the narrowing of the search for more to the area where food was found, an already-learned or perhaps inherited behavior. In the second process, the rat’s accidental and then purposeful use of the lever to obtain food, it is the progressive refinement of the behavior pattern that makes it instrumental—organized to produce a specific effect in the given environment. Only the second process requires any change of internal organization. Together, these two processes take place in what we may call the learning phase of a conditioning experiment. That phase is followed by a maintenance phase when the animal routinely uses the new technique to alleviate its hunger. The reorganizations in this kind

---

[7]If shivering starts too soon, core temperature will be increased above the reference level. If the increase is large enough, it could start reorganization and reduce the anticipatory shivering.Eventually, shivering will stop when just the right temperature is reached..

of conditioning are primarily on the output side, where errors give rise to changes in reference signals being sent to this or that lower-order control system that produces already-organized behaviors.

We have noted how in operant-conditioning experiments it is necessary to use 'establishing conditions'; if food is to be the reinforcer, the organism must first be deprived of food long enough to reduce its weight to 80 or 85 percent of its free-feeding weight. Of course, depriving the organism of water does not increase the 'reinforcing' properties of food—and anyone would say "of course"—but the theory of reinforcement can't explain why. Researchers in this area do—informally—understand that deprivation of something essential, like food, produces a 'drive' to obtain things that correct the deprivation. Removal of the deprivation reduces or removes the reinforcing value of the so-called reinforcer. When we realize that the food input is being controlled,[8] the reason for the deprivation effects becomes self-evident. When successful control is finally learned, that ends the deprivation and thus ends the random variations. The new organization of control systems can now maintain food intake close to whatever its reference level is. Thirst has its own intrinsic reference, and depriving the organism of water results eventually in its controlling the water intake, not the food intake, to bring it back to its reference level.

Reinforcement is said to increase the probability of the behavior that produced it. This has a descriptive basis in observations during the learning phase of an experiment. Observation of what happens in the operant cage shows, however, that it is the convergence of exploratory activities below, near, and above the lever that increases the probability of producing the reinforcement. The PCT alternative to reinforcement, up to this point, is simply to say that this is normal control behavior. Hunger could be a perceptual signal in some basic, perhaps inherited, control system. The reference signal would specify a level of zero for this perceptual signal. The error then causes already-learned organized exploratory behavior with the possible result of reducing the hunger toward its reference level. When the error is reduced, the tendency to go on exploring is decreased; when error is reduced enough, the exploring ceases.

This model leads us to expect, for different reasons, essentially the same series of events that the theory of reinforcement suggests, so for the initial learning, either theory accounts for the described facts. Simply having a plausible alternative to the theory of reinforcement, however, is useful in itself. It shows that reinforcement *is* a theory, not simply a description of a fact.

When we pass beyond the initial learning phase to later stages of performance, however, we encounter anomalies that the theory of reinforcement is unable to explain. We would expect the same relationship to continue to hold for further changes in reinforcement. If the rate of reinforcement increases, the behavior rate should also increase, and conversely if the rate of reinforcement decreases.[9] But that that is not what happens.

If the animal gets all its food from lever-pressing, it comes to press the lever often enough to generate and consume the food at the average rate which keeps the body weight constant for long periods of time. If food is consistently added, even by tube feeding, to what the behavior is producing, the behavior rate does not increase, but instead decreases. An animal offered high-calorie, tasty food in

[8]There is a *why* for this, of course: control of food input is the means of controlling a perception of hunger. Hunger is understood to be intrinsic, having no superordinate *why*.

[9]For simplicity, we will now assume a fixed-ratio schedule in our examples.

unlimited quantities will not, despite the implications of reinforcement theory, eat faster and faster until it explodes. Conversely, consistently removing food leads not to slower, but faster lever pressing, so that again the original rate of food intake is maintained. This was long ago confirmed in experiments with obesity. Rats that are obese because of certain hypothalamic lesions maintain a higher body weight, but they, too, defend it against disturbances in either direction, keeping their total food intake constant. Experiments with rats obtaining all their food by lever pressing (Collier, 19xx) showed that these animals maintain an approximately constant calorie intake when the apparatus varies the required ratio of presses to pellet deliveries over a range from 20:1 to1000:1, even if the reward size is varied, too. B. F. Skinner himself trained pigeons to peck thousands of time for each grain of food, not by increasing the reinforcement rate but by decreasing it (he concluded that intermittent reinforcement increases behavior, apparently not noticing that this entailed reducing the reinforcements) So the direct evidence shows that the effect of changes in reinforcement is actually the opposite of what reinforcement theory predicts.

In recent years, feedback phenomena have claimed more and more attention as researchers discover closed loops of causation as isolated phenomena at every level of organization in living systems. PCT shows how these observations fit into a systemic whole, but to grasp this we must begin with simple cases.

A set of demonstration programs

Tracking experiments involve only a very narrow range of behavior. They are a legacy from the engineering psychologists and physiologists of the 1950s, some of whom had worked in the war years on problems of aiming guns to track enemy targets. However, it must be acknowledged that we are in a position analogous to that of Galileo with his pendulums and inclined planes. As they are demonstrated and accepted, the principles of PCT can be applied to any behavior at all, but the most reliable experiments are still simple ones that can be implemented on a computer. Even so, many of the computer demonstrations of control processes that have been made publicly available by PCT researchers involve other kinds of behavior.

The most recent demonstrations are in a collection referred to as the ‘LCS3 set’, named for and explained in *Living Control Systems III* (Powers 2008). The program set, on a CD in the book, can also be downloaded from http://www.livingcontrolsystems.com/lcs3.html and runs on a Windows computer. The reader is advised to explore the set, because actually running the demonstrations is unquestionably the most effective way to learn what PCT is about (just as the best way to understand a float-valve regulator in a water clock is to see one in operation).

A second set of demos, also downloadable, are in a tutorial program called (for historical reasons) Demo3. The following refers to the demonstrations shown in various steps of this tutorial.

‘Responding’ to an invisible stimulus

The first three demonstrations in the “Demo 3” set explain how the mouse affects the cursor on the screen and the way numbers are used to determine positions. The first control task, step 4, is a tracking task: “compensatory” tracking in which the goal is to hold a cursor aligned with a stationary target and stabilize it against an invisible disturbance. After the 30-second experimental run is finished, a graph of the results appears. Figure 5 shows the result of one run. It differs from what you see when you exercise the program because the disturbance and the subject’s resistance to it both differ from one run to another.

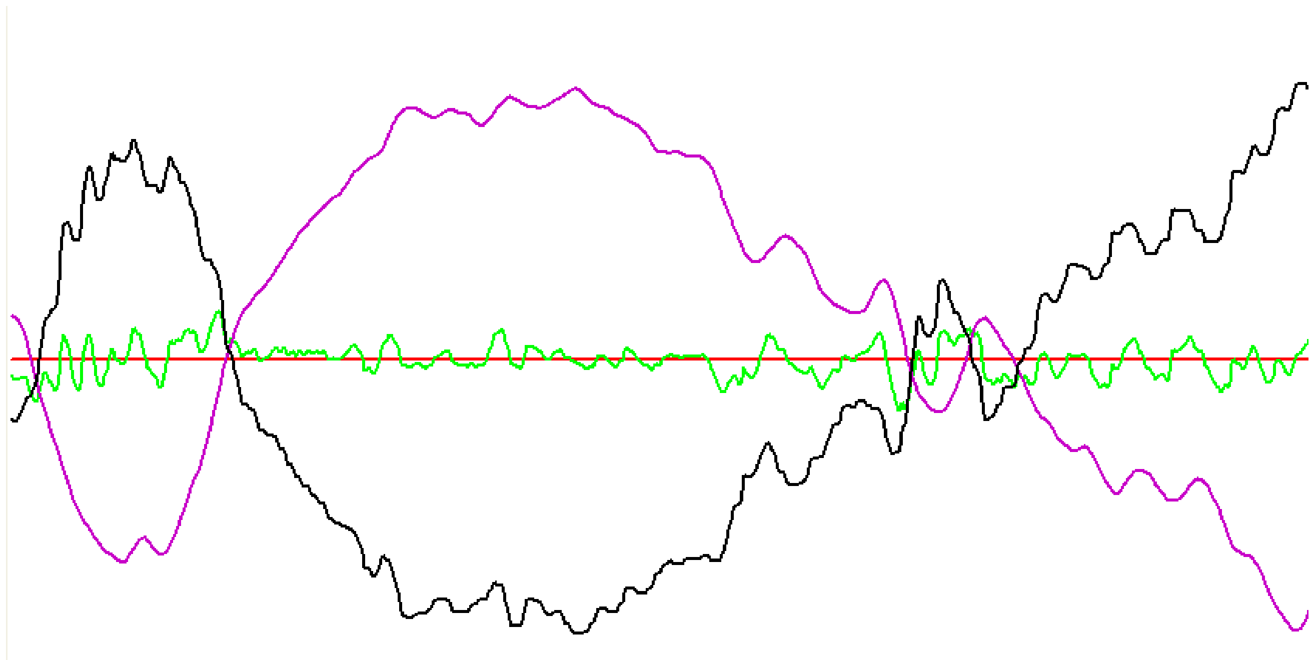

Figure 5: Compensatory tracking. Black line shows mouse movements, green line shows cursor movements. Target position is horizontal red line. The purple trace shows an invisible disturbance that varied during the run: mouse position relative to the centerline is equal and opposite to the disturbance at all times down to a moderate level of detail.

The main point of this demonstration is the way the participant moves the mouse so as to cancel the effects of an invisible disturbance (purple line) which, without these efforts, would move the cursor up and down so as to coincide with the disturbance trace. The green line shows the resultant cursor position during the 30 seconds of the experimental run (from left to right). There is no stimulus on the screen that corresponds to the purple disturbance plot, and clearly the green cursor line would be of no use in indicating the magnitude or direction of the disturbance. Thus there is no basis for claiming either that the mouse movements were a response to the cause of the perturbations of the cursor, or that the participant's brain was observing the disturbance and planning the actions needed to keep the cursor near the target. The information required to carry out either of those modes of action is simply not available in this demonstration. This is emphasized by the fact that one's performance improves over repeated exercise of these demonstrations, even though a new disturbance pattern is generated each time any step of the demonstrations is re-run. *Learning takes place, but there is no pattern of behavior* *to learn: what is learned is control.*[10]

Hierarchical control through reference signals

In demonstration 5, the participant is told to make the cursor descend from the top to the bottom of a range marked off in seconds, so that it passes each mark on schedule. An unseen disturbance is is still being applied to the cursor, so the participant must move the mouse so as to resist the effects of the disturbance and keep the cursor descending at a uniform rate. Figure 6 shows the appearance of the screen at about the 12-second mark. The participant is counting off the seconds, trying to make the green cursor move down so it passes each arrow at the time marked beside the arrow.

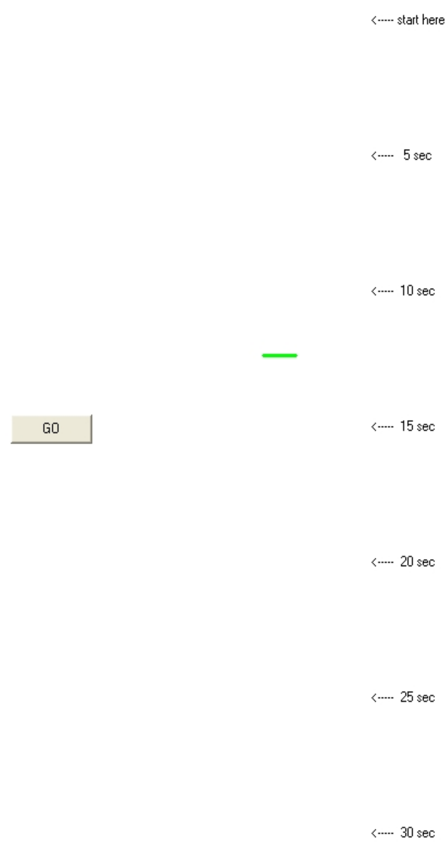

Figure 6: Demonstration 5 at about the 12-second mark. The green cursor begins at the start line and is moved up and down by unseen disturbances as the subject, resisting these disturbances, attempts to move it smoothly downward so as to reach each successive mark at the indicated time.

The graph for one run of demonstration 5 is shown in Figure 7.

---

[1]The control-system functions change form as the parameters are changed, and as a result the relationships between neural signals change. This is also how new control systems in the hierarchy develop out of collections of as-yet-unorganized neural nets, a process that largely involves ‘pruning’ of the surplus of connections with which the neonate begins.

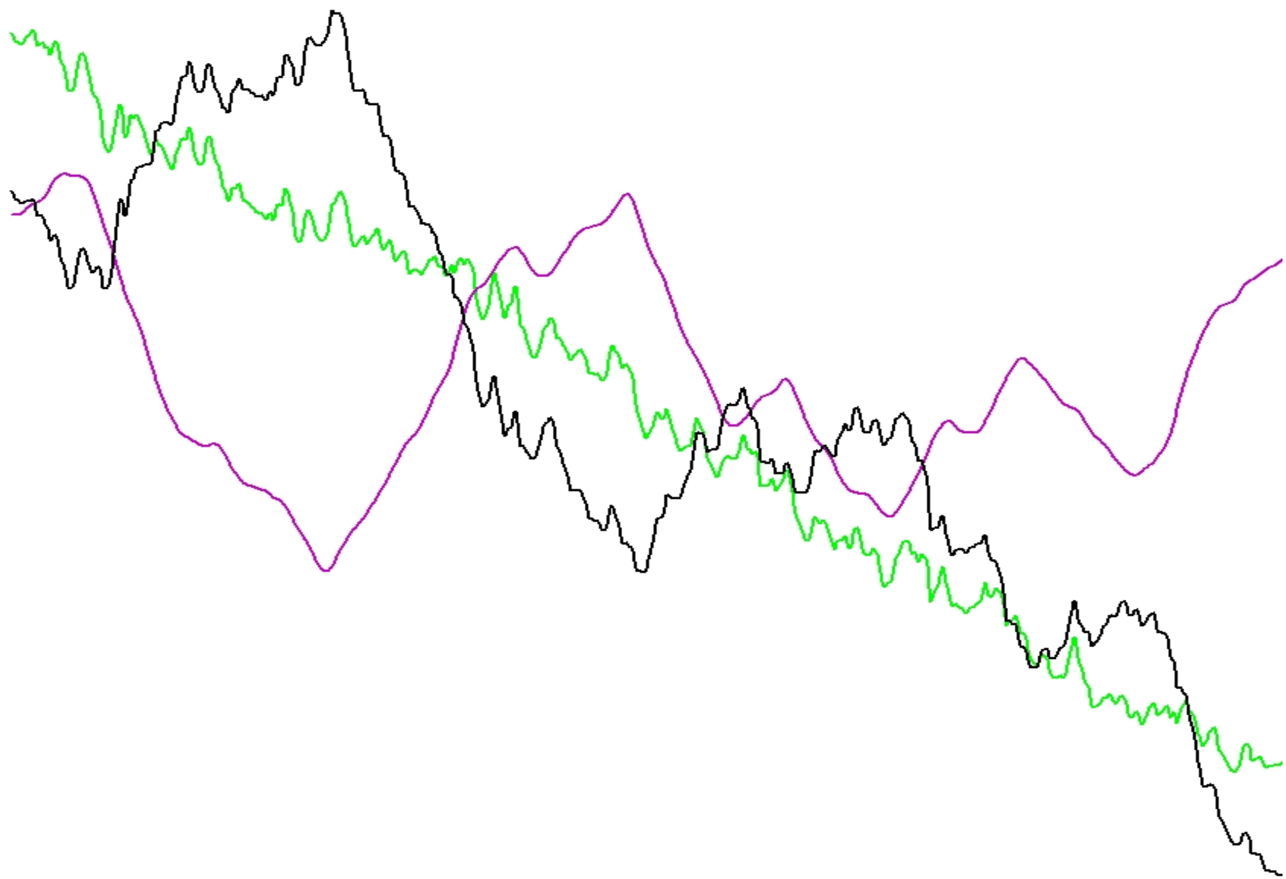

*Figure 7: One run of demonstration 5. The mouse cursor (green) approximates a steady progression from top to bottom (*y *axis) over the time elapsed (*x *axis). Effects of mouse movements (black) add to effects of the disturbance (purple) to accomplish this descending movement. No stimulus corresponds to the disturbance, and no planning of mouse movements is possible.*

The black trace showing the mouse movements executed by the subject's hand does not resemble either a mirror image of the disturbance pattern in purple or the pattern followed by the cursor in green. The result when the mouse movements are added to the disturbance is the requested slow movement of the cursor from high to low, as shown by the diagonal green line. But again, it is clear that this result cannot be accounted for in terms of responses to any visible stimulus, nor could the mouse movements have been planned in advance and then executed.

This demonstration shows what is meant by control through varying reference signals. The steady downward velocity of the green cursor bar is, according to PCT, the controlled variable for some system fairly high in the perceptual hierarchy, having to do with control of perceived rates of change of position. This system generates a slowly-varying reference signal[11] for a lower system concerned with maintaining the cursor in some particular position against disturbances. This lower system is just as in the previous demonstration, except that now the ‘particular position’ where the lower system is maintaining the cursor is being changed through time by the higher system. In both cases, the lower system acts to make the cursor position match the reference position at all times (as well as it can) by resisting disturbances. The difference is that in the previous demonstration the reference is stationary, but in this demonstration a higher-level system is changing the reference signal in the direction from positive toward negative, so that the lower system creates the requested perception of a slowly descending cursor—by, of course, using a still a lower level of organization to move the mouse up and down in whatever way works to make the cursor on the screen actually descend.

That the cursor is under positional control at all times is shown by the way it resists a disturbance that is trying to push the cursor up and down, away from its steady descent. The mouse position varies oppositely to the disturbance, not only canceling it as in the first demonstration, but also adding enough additional variation to maintain the steady downward velocity. This happens automatically at the level of position control. The higher control system concerned with downward velocity does not have to do much to resist residual effects of disturbances. Most of the resistance has been accomplished at the lower level.

In addition to illustrating principles of control, these demonstrations are one way of testing the theory through attempts to prove hypotheses wrong—attempts to falsify them. The remaining demonstrations illustrate and test additional principles of PCT.

Challenging PCT with experiments and simulations

The first book-length treatment of what is now known as PCT (Powers 1973, 2002) was finished before the advent of inexpensive desktop computers and the exponential growth of computing speed and memory storage. Some 12 years later, the first interactive computer demonstrations of the principles of PCT began to take shape, in time for the first meeting of the Control Systems Group in 1985.[12] At this meeting, a tracking experiment was shown in which a subject used a joystick to make a cursor on the computer screen track a moving target, the controlled variable being the separation of cursor from target and the reference condition (defined by instructions) being zero separation. Demonstration 4, the first one that we discussed above from the latest book in this series, recapitulates that demonstration.

This was also the first instance of a computer simulation of a PCT-type control system designed as a model of the person doing the tracking task. The parameters of the simulated control system were an integration sensitivity and a constant reference signal which were adjusted to make the performance of the model match the real person's performance with as little difference as possible. The RMS difference between modeled joystick movements and the real movements could be reduced to less than 10 per cent of the range of movement of the target. More recent versions have reduced the RMS error of fit to less than 4 per cent.

The most important aspect of this early simulation was that in a moving-target version it could be used

[1]The output function of this control system could be classified as a simple form of ‘central pattern generator’.

[1]This 1985 tracking experiment is included in DEMO 1, part of the series of DOS programs listed under Resources.

with either a single smoothed-random disturbance moving the target, or with a second uncorrelated disturbance added that made the cursor movements differ randomly and by large amounts from the joystick movements. With the second disturbance acting, the subject would move the joystick in a way that corresponded neither to the target movements nor to the second disturbance, but was exactly the movement needed to minimize the tracking error.

This demonstration illustrated the important point that the behavior observed in a control situation generates a regular result without itself being regular. This is the main feature of PCT that distinguishes it from the calculate-and-execute models of control behavior: it is not possible for the organism to calculate in advance the joystick movements that will be required, because the disturbances are being generated from random numbers during the experimental run, and are unknown in advance. A calculate-and-execute model necessarily fails in the presence of unpredictable disturbances. This is only noticeable if working models are made and tested.

Among the ‘LCS3’ demonstration programs introduced so far, of particular interest is the demonstration called “Square circle”. In this demonstration, a white dot is used by the participant moving a mouse to trace around the sides of a red square. At the end of one complete tracing, the path of the mouse is revealed: it is a circle. In a variant mode, the revealed path is a triangle—a bit more difficult to execute, but even more unexpected by the participant. The point is to show that what a person *experiences* as his or her own behavior is actually the perception that they are controlling, not their actions, which are often markedly different from it.

Other demonstrations are distributed with (Powers 2008) or can be downloaded from the Internet[13]. In one of these, the participant's task is to keep a small green circle aligned inside a slightly larger red circle in one corner of the screen. A white tracing shows the actual path of the mouse, which at the end of the run is seen to spell out in script the word “hello”. This is caused by a patterned disturbance of the green circle which (invisibly) traces out “hello” upside down and backward. Even though the observed behavior of the participant (the movement of the mouse) is essentially unrelated to the control task being accomplished—the behavior as the person experiences it—nonetheless, when we overlay the disturbance on the mouse movements they are very highly and negatively correlated (estimated to be in the -0.99s).[14]

The first demonstration of the LCS3 set may be the most philosophically interesting. A red ball is shown drifting left and right while it rolls vertically and changes shape from short and wide to tall and thin. Each aspect is affected by a smooth disturbance, the three disturbances being nearly uncorrelated. The mouse affects all three variable aspects of the ball—shape, orientation, and position—at the same time and by the same amount. The participant's task is to pick one of those aspects and keep it constant: shape as *round*, position as *centered*, or orientation as *level*. That this can be done at all is of considerable interest, but of equal interest is the fact that the computer can determine reliably which single aspect is being controlled and which two aspects are varying as side-effects. The computer deduces which effect of the action was intentional and which others were accidental. It is very seldom wrong even if control is poor.

---

[1]See Resources

[1]A similar result is seen with Richard Marken’s simulation of a bimanual control experiment described in Mechsner et al. (2001). The reader may exercise the demonstration at http://www.mindreadings.com/Coordination.html

Intention, in PCT, refers not to behavioral acts but to the consequences of those acts. The intended consequence of controlling the orientation of the red ball is to keep its axis pointing toward the viewer. Because each aspect of the ball is being influenced by a different pattern of disturbances, the same actions that stabilize orientation can't simultaneously stabilize position or shape; in fact they increase the variance of those two variables because the actions aren't systematically opposed to the relevant disturbances. The result is a rather puzzling combination of correlations: the actions that stabilize orientation correlate almost perfectly (-0.99) with the disturbance that tends to alter orientation, yet those actions and those disturbances show only a low correlation, close to zero, with the orientation that is being controlled. The mouse movements always correlate much better with the aspects that are *not* being controlled.[15]

To return to a subject at the beginning of this paper, a general-purpose demonstration called "LiveBlock" shows a basic control system as a 'live block diagram'. Here we have a control system with an adjustable transport lag, time constant, gain factor (output amplification), and environmental feedback factor, plus an adjustable reference signal and disturbance. The model runs continuously in the background so the effects of changing system parameters and independent variables can be seen as they occur. The method of stabilizing a system with time lags in it is illustrated, as are many other basic properties of a negative feedback control system. It is hoped that this demonstration can finally counteract many of the false ideas offered over the past 60 years about the limitations of negative feedback control as a model of behavior.

The discovery of the principles underlying a phenomenon for which there is a widespread naive (and incorrect) explanation is always of particular interest to many readers. Such is the discovery of what baseball players are really doing when they catch fly balls, as reported by McBeath, Shaffer, & Kaiser (1995), and the explication of their naive projection of perceptions as reported by Shaffer & McBeath (2005). The simulation at http://www.mindreadings.com/ControlDemo/CatchXY.html demonstrates this.

The methodology of PCT research

Near the beginning of this paper, we made note of an inherent difficulty of the experimental investigation of living things. An organism controls its own perception of some aspect of its environment, but that privileged point of view from inside the observed organism is unfortunately not available to scientific observers of other organisms. As observers of a different organism we do not have access to that perception, we only have our own perceptions from our own points of view, external to the organism. For that reason it has been crucially important to devise tests for determining which aspects of its perceived environment the organism is controlling.

The principal datum in PCT methodology is the controlled variable. All of the demonstrations that we have reviewed have clearly displayed three variables: the controlled variable (e.g. distance between the mouse cursor and the target), the disturbance (producing movements of the mouse cursor independent of the user's movements of the mouse), and the relevant behavioral actions (indicated by the changing mouse position). Obviously, the disturbance can't be identified until we know just what the controlled variable is and how mouse movements affect it. If the user, unbeknownst to us, is ignoring the moving

[1]Trying to explain this fact is a good test of one's understanding of PCT.

target and instead trying to draw a large circle with the mouse cursor, the measured and graphed results will not make sense to us. There will be no relationship between what is expected to be controlled (the position of the cursor relative to the target) and what the subject is actually controlling (following the outline of an imagined large circle). Even in this simplified, artificial, two-dimensional laboratory environment it is difficult to see what is actually under control; we would expect more naturalistic settings to present still more difficulties. Yet the technique for determining what perceptual variable is being controlled is essentially the same everywhere. The requirements are few. We must be able to make intelligent estimates of which aspects of the environment the organism can perceive and influence with its activities, and we must be able to also influence those aspects of the environment.[16]

The fundamental step of PCT research, the Test for controlled variables, is the *slow and gentle* application of disturbing influences to the state of a variable that the researcher surmises is already under control by the observed organism. If the organism changes its action and thereby prevents the disturbing influence from having the expected effect on that variable, that is strong evidence that the experimental action disturbed a controlled variable. It may take a number of variations of the disturbance to isolate just which aspect of the environmental situation is under control. This strategy gives a definition of the controlled variable in terms of the observer's way of perceiving the organism and its environment.

In order to build working generative models of behavior, like the simulations that we have been exercising, there is one further requirement. We must be able to measure the influences that affect the state of the environmental variable that we have decided to test. These measurements provide a numerical record of the disturbance and of the control actions that oppose it. Until a simulation produces very nearly the same numbers as were produced by measurement, it needs refinement; and when it does, we have a strong basis for the claim that the simulation models essential aspects of the unseen internal structure of the organism whose behavior we measured.

In sum, research methods in the PCT paradigm depend upon careful observation of individual specimens, followed by the construction of precise generative models under exacting standards of performance. The criteria for determining what is a good result are very different from what is usually accepted in the literature. The correlation between actual behavior and model behavior in the PCT work done with pursuit tracking is in the .90 and higher range. If a simulation does not perform at least this well, we do not have a correct functional model. This requires that research and refinement of the model continue until the correlation between theory and data moves into the .90 range and higher.

An important result is that this enables predictions about individuals, not just population averages.[17] Achieving this goal places us on much firmer ground when we turn to practical applications of psychology to personal and social ends. Clinical psychology is a good example to which we will next turn our attention.

PCT applied to psychotherapy

The field of diagnosing and treating psychological problems affords an excellent example of how PCT can provide a unifying framework to an otherwise fragmented area of research and practice. The alert

---

[1]Insofar as each level of perception is the environment for the next higher level in PCT, the test for the controlled variable may involve "environmental" variables that are available only to the person being studied (as in Robertson et al. 1999).
[1]Runkel (1990/2007), Gould (1996).

reader will anticipate from the foregoing presentation that there may be difficulties finding a good match to concepts that are organized around the commonplace assumption that perceptions (stimuli) cause behavior (responses). The juxtapositions in Table 1 must be viewed with this caveat in mind.

Table 1. Correspondences and differences of key terms in PCT, in relation to concepts in conventional psychological theories. The listings are not equivalences, but rather indications that the correspondent terms have different referents with different implications for theory and practice.

Conventional terminology	PCT terminology
Stimulus	1. Disturbance to a controlled variable (CV). 2. Effect of disturbance on a CV; when measured, the effect on the corresponding input quantity
Sensations, Perceptions, Concepts	Perceptual signals in at least three levels of the perceptual hierarchy (or more, depending on how the conventional terms are defined).
Motives	Reference signals.
Drives	Error signals
Responses; Behavior	Observable means of controlling perceptual input; when measured, output quantities; functionally, actions that are varied to control inputs.
Learning and Development	1. Control by existing systems: familiar means in unfamiliar circumstances. 2. Reorganization when error cannot otherwise be reduced: ranging from alteration of connections in existing control systems to formation of new control systems.[18]
Cognitive map	Awareness of many variables concurrently.

Some of the more radical differences in meaning are obvious, as for example *stimuli* and *responses*. Perceptual inputs (stimuli) are not independent of a person's actions on the environment (responses). Output quantities (responses) at one level become the reference signals at the next lower level. Some PCT terms, such as *error signals* and *reference signals*, refer to neural signals (firing rates in nerves) within the organism, and have no literal correspondence to the common parlance of psychology, though they have been corroborated in neurology at least at lower levels of the hierarchy. The observer may perceive and measure the reference *level* (the apparently desired value of a controlled variable) using objective behavioral experiments; the reference *signal* is an entity in a model that is meant to explain the observed reference *level*. As such, the reference level is an objective phenomenon with no

[1]Van de Rijt & Plooij (2008-2010), Powers (1979a).

clear correspondence to the subjective notion of *motives*.

Or again, because of physiological effects associated with preparation for action, error signals relate indirectly to terms used in conventional psychology such as *stress*, *distress*, and *emotions*.[19]

For the standard terminology of psychiatric diagnosis, however, we are unable to find PCT correlates. This is largely because diagnostic categories describe behaviors, not controlled variables. Our difficulty with this is corroborated by the growing recognition in the field that current classificatory systems of psychological disorders such as the DSM IV-TR (American Psychiatric Association, 2000) do not easily map onto the lived experience of psychological distress. Despite the invocation of concepts such as ‘comorbidity’,[20] there is a growing awareness that this system of classification is unsatisfactory in important ways. In recent years there has been great interest in processes that are said to be ‘transdiagnostic’ across DSM categories as an explication of underlying pathways by which diverse symptoms become manifested (Harvey, Watkins, Mansell, & Shafran, 2004). A PCT account of such processes explains why this approach has merit and why categorizing symptoms and behavioral actions is problematic.

We have already outlined and demonstrated in detail how behavioral output varies in order to control perceptual input. Consequently, internal problems of control do not give rise to recognizable, standardized patterns of symptoms. That is why there is so much variation within current classificatory categories of symptoms. This also accounts for the lack of clear differentiation between categories. In the real world, where unpredictable disturbances to control occur, behavior must vary as a person repeats attempts to solve a given problem. The appearance of symptom patterns and behavior patterns is analogous to the appearance of constellations in the night sky. They are arbitrary groupings in the eye of the beholder that reflect no underlying order or structure. Categories of behavior—that is, of variable control system outputs—cannot reveal the order or structure of goals that behaviors are intended to accomplish.

There is a developing acknowledgement that it is the causes of distress associated with particular symptoms rather than the distress or the symptoms themselves that need to be understood. Kazdin (1999) argues that functional impairment, rather than symptoms, is the main reason people seek psychotherapy. Large-scale population surveys, for example, have demonstrated that many people experience psychotic symptoms without requiring treatment (Bentall, 2009, p. 107). Auditory hallucinations are one type of symptom that can be experienced with or without the debilitating experience known as psychosis. Who has never had cause to complain of being unable to stop a tune from replaying itself in imagination?[21] Auditory hallucinations can even be helpful. A singer who can’t mentally ‘hear’ a note before singing it will not sing very well. However, when a person is distressed

---

[1]An account of emotions, an obviously important topic, is beyond the scope of this paper. Briefly, our present view is that an emotion perception is constructed from sensed physiological conditions—the sensations that we call feelings—in combination with the perceptions that are being controlled. An error signal in a control system has two coordinated effects. Firstly, it provides reference signals for the control actions which reduce the error signal. Secondly, it activates the endocrine and autonomic nervous systems to prepare the body for that activity. Thus, feelings are always potentially present when we are controlling, because there is always some body state present when a person controls. See Powers (2005), Chapter 17.

[2]A label meaning simply that one or more disorders or diseases are diagnosed in addition to a primary one.

[2]For a discussion of how imagination is modeled in PCT, see Powers (1989:277).

by auditory hallucinations, the hallucinations are problems. In PCT, the problems which are discussed in therapy are the ones which are distressful to the person and the focus of the discussion is the distress not the symptoms. There are situations where other people may be distressed by the person's symptoms, but this is considered to be more of a social problem than a psychological problem, to be settled by social processes like law, mediation, conflict resolution, or negotiation, not by therapy. PCT provides no way to decide which party in a social conflict needs therapy.

Each person's problems are understood to be unique to that person. There is no justification in PCT for applying what worked for someone else to this person, just because that other person had similar symptoms or DSM IV-TR diagnoses. The same symptomatic behavior can result from an entirely different set of internal conflicts; similar conflicts can lead to entirely different sets of symptoms.

Psychotherapy has focused, understandably, on pathology. PCT contributes a useful perspective in understanding psychological disorders by first providing a model of satisfactory psychological functioning. PCT portrays dysfunction in terms of disruption of successful control (Carey, 2006; Mansell, 2005). Conflict between control systems, as we noted earlier, is a problem because it effectively removes the control abilities of both systems. Conflict is usually transitory. It is when conflict is unresolved and becomes chronic that the symptoms recognized as psychological disorder become apparent. Distress is the experience that results from the person's inability to control important experiences.[22] It clearly can't be "treated" as if being distressed is itself the problem. Restore the ability to control and the distress will disappear.

Conflict, as it is conceptualized in PCT, occurs between two control systems at the same level (Carey, 2008). These control systems, however, are located within a hierarchical network of control systems so their conflicted arrangement will influence and be influenced by lower and higher level systems. This account of psychological distress may explain why no reliable biochemical correlates of mental illness have ever been discovered. From a PCT perspective, control systems that are in conflict are not dysfunctional or broken. In fact, it is quite the reverse. The better the control systems would control in isolation, the more intense the conflict will be when they oppose each other (Powers, 2005). Some kinds of mental illness, perhaps most, may be a result not of broken brains but of well functioning control systems locked in chronic conflict.

It is the hierarchy that provides a clue as to where treatments should focus to help conflicts resolve. Systems at one level receive their references from the next higher level. When two control systems are conflicted, it is the signals being sent to each from the next highest level that need to be altered. The senders, not the receivers, need reorganization. Otherwise competent control systems are being misused by higher systems.

Most therapies assume that attention to problems facilitates change, and people in general tend to be most aware of painful or dramatic consequences of conflict. This is seldom helpful in itself. Attention is drawn to the symptoms rather than the causes of loss of control, symptoms such as apathy, confusion, fear, or despair. Often a person will try to strengthen the 'good' side of a conflict, which usually just makes conflict more extreme because the other side resists the effort to change and starts to look good for other reasons.

---

[2]Other problems that can also arise, such as feeling overwhelmed by environmental forces, require more extended discussion than is possible within the scope of this paper.

PCT suggests that resolving conflicts requires the learning process of reorganization that was explained previously. Initially, reorganization can do this by modifying components of higher-level control systems that send conflicting reference signals to others at a lower level. That change may well result in other reorganizations being needed at higher and lower levels. The therapeutic approach that is based on the principles of PCT is called the Method of Levels (MOL; e.g., Carey, 2006, 2008). Its principle is to redirect attention to the higher level control systems responsible for generating the conflict, away from a preoccupation with the symptoms and the immediate efforts on both sides of the conflict. It is proposed that reorganization and awareness are linked in such a way that the systems in awareness become the particular focus of the reorganizing processes. Reorganization is an automatic response to intrinsic error; it can't be controlled voluntarily. But awareness can be redirected with the consequence being that the focus of the reorganizing process changes. .

Consider wanting to stop smoking to avoid lung disease and at the same time wanting to continue it to relieve withdrawal symptoms caused by stopping; or wanting to leave a partner to avoid abuse and, at the same time, wanting to stay with the partner for the sake of love. Ultimately, attempts to modify the actions of conflicted systems, or to give preference to one goal by will power, cannot permanently alter the conflict. The conflict is resolved only when awareness is shifted to the level above the conflicted systems so that reorganization can be directed to the systems creating the conflict—the systems that are establishing these conflicting goals (Carey, 2006, 2008). Hence the name, the Method of Levels.

For the person in therapy, MOL is an experience of describing in detail a current area of distress to a therapist who understands PCT. The therapist's approach is a blend of questioning about subjective experiences and selectively drawing a client's attention to seemingly tangential or peripheral subjects, usually on the basis of comments the client makes—comments that the therapist familiar with PCT recognizes as possibly reflecting involvement of a higher-level system. In this way, clients show the therapist what path to follow, and when the therapist helps them focus in the right place, their own reorganizing capabilities generate new perceptions and goals that may resolve the conflict, or uncover the deficiencies that cause trouble.

MOL has been used over a number of years by different clinicians in a variety of clinical settings. Evaluations have been conducted of the way in which MOL is experienced by routine clients in routine clinical contexts (e.g., Carey, 2005; Carey & Mullan, 2008; Carey, Carey, Mullan, Spratt, & Spratt, 2009) .[23]

It is telling that despite the demonstrated effectiveness of various psychotherapies there is still no generally accepted account of how these effects are achieved, and in fact, it has been shown that psychotherapies based on quite different models of disorder can have similar effects (e.g., Wampold, 2001). There has been an increasing call to move away from developing new techniques and strategies based on diagnosis and instead to focus on underlying common principles and mechanisms (e.g., Rosen & Davison, 2003). PCT provides a common underlying process (conflict) and a common change mechanism (reorganization) that might be particularly significant for understanding this peculiar situation. The paradigm of perceptual control may provide the means to make sense of these otherwise puzzling results.

[2]For details about MOL and its use in clinical practice, see Carey (2006).

Recent research in neurobiology has indicated that psychotherapy can have effects in the brain that are similar to the effects that pharmacology achieves.[24] Again, this result would come as no surprise from a PCT perspective. The hierarchy of PCT is a hypothesized neuronal architecture for the full range of control. At every level, brain chemistry and neural signal-handling functions are aspects of the same processes.

It is certainly the case that, at this stage, PCT raises more questions for research in this field than it answers. Do conflicts at different levels of the hierarchy, for example, result in different types of pathology? Does the rate of reorganization affect the experience of internal conflict? What influences the mobility of awareness such that some conflicts are resolved satisfactorily while others become chronic? The possibilities for new research, as usual with new ideas, proliferate.

While some of the propositions about the application of PCT principles to psychotherapy remain speculative, there is also indirect but strong evidence for this approach. Problems of control (such as behavioral control, impulse control, emotional control, and thought control) are widely recognized as important in psychological functioning. Many approaches to psychotherapy use conflict formulations to explain psychological distress (see Carey, 2008; in press for more details). Many approaches also discuss the importance of awareness in resolving problems as well as recognition of the need to consider problems from higher levels of thinking (such as important life values or belief systems). Finally, there is a growing body of literature that recognizes that the change involved in the resolution of psychological distress is not a linear or predictable process (e.g., Hayes, 2007).

Exploring psychological disorders and their treatment from the perspective of perceptual control provides a new direction for psychotherapy researchers and practitioners. An understanding of the nature of psychological distress that is developed from a model of function rather than dysfunction will help to clarify the function and purpose of treatments. The distillation of the important components of psychotherapy will allow therapists to be clearer about their roles and treatments to become more efficient. Moreover it can provide a guide regarding the purpose of psychotherapy. PCT, then, will have an impact on long standing debates such as the dodo bird hypothesis.[25] A unifying focus such as the one provided by PCT will allow a more consistent and coherent approach to emerge that will go a long way towards preventing the debilitating impact of psychological distress that is currently on the increase in many countries.

There may also be other implications of this approach that cannot easily be predicted at this stage. Perhaps the stigmatizing nature of mental illness will change with a more accurate explanation of these problems that is inherently psychological (yet firmly grounded in neurobiology) and intuitively optimistic and hopeful. The nature of the delivery of psychological treatments might also change as researchers and clinicians become more familiar with the reorganizing capabilities of individual systems. Perhaps we will learn to use both psychotherapy and pharmacotherapy more judiciously. While the outcomes may not be entirely obvious there seems to be sufficient justification at this stage to step into the paradigm of control and to build our knowledge of the mechanisms of psychotherapy from the foundations of these functional and rigorously tested models.

---

[2]For example, Schwartz, (1996) on OCT symptoms.

[2]This is the proposal that all therapies are equally effective, depending upon the practitioner, and “all must win prizes” (alluding to a pronouncement by the Dodo Bird in *Alice in Wonderland*). See http://en.wikipedia.org/wiki/Dodo_bird_verdict.

Directions in PCT research

Despite being more than 50 years old, and having the principles of negative feedback control and many known facts of neurology and physiology behind it, much of PCT is still provisional and hypothetical. The main lines of research that PCT theoreticians are concerned with have to do not with applications but with testing the core concepts of the theory. The demonstrations cited earlier were constructed as deliberate challenges to the basic concept of control, in which the assumptions of the theory were made explicit as properties of a simulated system. In some cases they were used to make quantitative predictions of behavior for experimental test. The fact that it is possible to give confident descriptions of what will happen when any randomly-selected person participates in the interactive demonstrations shows the progress that has been made at this basic level.

It is now quite clear that simple kinds of behavior are well described in great detail and under many variations in experimental conditions by a negative feedback control model. But the theory contains hypotheses about much more than a single control system operating in a simple experiment. Eleven levels of control have been proposed[26]. Do they really exist? If they exist, have they been correctly defined? Are there too many or not enough levels? It is proposed that reference signals are derived from remembered perceptions, a possibility raised by the way people can remember goals and act to match current perceptions to perceptions that have been experienced some time in the past. How long do remembered reference signals remain accurate? Are there differences at different levels of perception? Does imagining a perception interfere with present-time control? Does degree of awareness or attention affect the quality of control, and if so, in what way? Does reorganization really occur—can we detect it by measuring control parameters at frequent intervals? Will the changes prove to be random? Will they become smaller as an asymptote is approached? Does reorganization really follow awareness? Is it really driven by intrinsic errors, such as error in normal biochemical and homeostatic control systems?

The 'pandemonium' form of the model, in which control systems control only scalar variables in one dimension, is probably wrong in several regards, the main one being a lack of interaction between different perceptions. But there is a way of transforming this model into one with more realistic functions that correspond to sensory and motor nuclei or larger areas in the brain, and studies of interactions will give clues about what the groupings really should be.

Another large area for research concerns the way higher-level systems use lower-level systems. In PCT, the only link from higher to lower systems is the reference signal, but there are reasons to believe that higher systems can vary some parameters of lower-level control systems, for example the sensitivity of an output function to error signals (reducing the sensitivity to zero is equivalent to turning the control system off—but there are other ways to do that as well). It is also possible that higher systems can act to alter lower-level perceptual functions, changing the very nature of a controlled perception so the same perceptual signal represents a different aspect of the external world.

The inability of neural signals to change sign (a negative rate of firing is impossible) requires that

---

[2]Briefly: intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts.

bidirectional control be accomplished by pairs of control systems, each acting in only one direction. This implies that paired systems employing agonist-antagonist muscles to accomplish bidirectional motor control might have measurably different properties in the two directions. And more than two directions of control in one dimension are possible; muscles set at varying angles can contribute to X and Y control when they are shared by several control systems.[27]

Adaptation is handled in PCT by the basically random reorganizing process. However, control systems exist that provide more systematic ways of gaining control when it is difficult (that is, solving problems), and these have been explored to some extent. In theory, one result of random reorganization is to create control systems that are so much more effective than the random process that errors in a given area of experience no longer get large enough to activate reorganization. True or false?

If older theories had been subject to the kind of testing that PCT has undergone even at this early stage, most of them would have been discarded long ago. The flaws are simple and obvious (as they were for PCT in the beginning). When the premises of a theory are taken as unquestionable facts without even being tested (as in calling every action a “response”), flaws are perpetuated and go undetected, and the quality of science suffers. Much as we have tried to avoid that error in PCT, our success is no doubt partial at best. We have tried to be skeptical about what seem our best ideas, to demand demonstration of every claim. Whether that has worked remains to be judged.

The approach that has led through many years to this paper has entailed going through all the assumptions on which PCT is based and, one by one, testing to see if they can be supported or ruled out by experimental data. That task is far from finished. The current form of the model is the result of half a century of challenging the theory, with many failures and subsequent improvements. At first, just one person was doing this work, and in terms of actual computer modeling, never more, until recently, than half a dozen. One person or a small handful cannot develop or explore a new idea in all the necessary ways. Real progress and the building of a real science require the vast resources of a whole discipline, with thousands upon thousands of independent minds, each reorganizing in unpredictable ways, looking for difficulties with and improvements of the theory. PCT as it stands today is no more than a pilot study, a definition of a problem and a possible kind of solution. It is difficult to imagine what will happen when the full power of a scientific community is turned to developing it in all the directions that are possible. Will it ignominiously disappear? Or will it turn into the direction for all the life sciences before the 21st Century is out?

Resources

Computer simulations

The most recent set of simulation and demonstration programs designed for and included with (Powers 2008) can be downloaded at http://www.livingcontrolsystems.com/lcs3.html.
Earlier DOS and Windows programs by Powers can be downloaded at http://www.livingcontrolsystems.com/demos/tutor_pct.html

[2]Powers (1979b, Part 3); p. 10 of the reprint as byte_aug_1979.pdf.

Programs by R. Marken at http://www.mindreadings.com/demos.htm can be run in a web browser.

Reference websites

Introductions and discussions of Perceptual Control Theory can be found at several web sites. Some of the major reference sites are: http://www.livingcontrolsystems.com, http://www.pctweb.org/, http://www.mindreadings.com/, http://www.perceptualcontroltheory.org/

References

American Psychiatric Association. 2000. *Diagnostic and statistical manual of mental disorders DSM-IV-TR*, Fourth Edition (Text Revision). Arlington, VA: APA.

Ashby, W. Ross. 1952. *Design for a brain*. London: Chapman & Hall. [Repr. 1960, New York: Wiley. Online: http://www.archive.org/details/designforbrainor00ashb.]

Ashby, W. Ross. 1956. *An introduction to cybernetics*. London: Chapman & Hall. [Online: http://pcp.lanl.gov/ASHBBOOK.html & http://pespmc1.vub.ac.be/books/introcyb.pdf.]

Bennet, S. 1979. *A history of control engineering 1800-1930*. New York: Peter Peregrinus Ltd.

Bentall, R. P. 2009. *Doctoring the mind: Is our current treatment of mental illness really any good?* New York: New York University Press.

Bernstein, Nicolas. 1967. *Coordination and regulation of movements*. New York: Pergamon Press.

Black, Harold Stephen. 1934. "Stabilized feed-back amplifiers", *Electrical Engineering* 53:114-120.

Harold S. Black. 1977. "Inventing the negative feedback amplifier", *IEEE Spectrum* 14:54-60.

Brillouin, Leon: 1953. "Negentropy principle of information", *J. of Applied Physics*. 24.9:1152-1163.

Burke, Robert E. 2007. "Sir Charles Sherrington's *The integrative action of the nervous system*: a centenary appreciation". *Brain* 130.4:887-894. doi: 10.1093/brain/awm022

Carey, T. A. (2005). Can patients specify treatment parameters? *Clinical Psychology and Psychotherapy: An International Journal of Theory and Practice, 12*, 326-335.

Carey, T. A. (2008). Conflict, as the Achilles heel of perceptual control, offers a unifying approach to the formulation of psychological problems. *Counselling Psychology Review, 23*(4), 5-16.

Carey, T. A. (in press). Exposure and reorganization: The what and how of effective psychotherapy. *Clinical Psychology Review* DOI 10.1016/j.cpr.2010.04.004

Carey, T. A. (2008). *Hold that thought! Two steps to effective counseling and psychotherapy with the Method of Levels.* Chapel Hill, NC: newview Publications.

Carey, T. A. 2006. *The Method Of Levels: How To Do Psychotherapy Without Getting In The Way*. Hayward, CA: Living Control Systems Publishing.

Carey, T. A., Carey, M., Mullan, R. J., Spratt, C. G., & Spratt, M. B. (2009). Assessing the statistical and personal significance of the Method of Levels. *Behavioural and Cognitive Psychotherapy, 37,* 311-324.

Carey, T. A., & Mullan, R. J. (2008). Evaluating the Method of Levels. *Counselling Psychology Quarterly, 21*, 1-10.

Chapanis, Alphonse. 1961. "Men machines, and models". *American Psychologist* 16:113-131.

de Rosnay, Joël. 2000. History of cybernetics and system science. Principia Cybernetica Web: http://pespmc1.vub.ac.be/cybshist.html.

Gould, Stephen Jay. 1996. *The Mismeasure of Man*. New York: W. W. Norton & Company.

Harvey, A. G., E. Watkins, W. Mansell, & R. Shafran. 2004. *Cognitive behavioural processes across psychological disorders: A transdiagnostic approach to research and treatment.* Oxford: Oxford University Press.

Hayes, A. M., Laurenceau, J. -P., Feldman, G., Strauss, J. L., & Cardaciotto, L. A. (2007). Change is not always linear: The study of nonlinear and discontinuous patterns of change in psychotherapy. *Clinical Psychology Review, 27,* 715−723.

James, William. 1890/1950. *Principles of Psychology.* New York: Henry Holt/Dover Publications.

Kazdin, A. E. 1999. The meanings and measurement of clinical significance. *Journal of Consulting and Clinical Psychology, 67(3),* 332-339.

Ronald Kline. 1993. "Harold Black and the Negative-Feedback Amplifier", *IEEE Control Systems* 13.4:82-85.

Kozak, M., Marsh, A. A., & D. M. Wegner 2006. “What do I think you're doing? Action identification and mind attribution”. *Journal of Personality and Social Psychology*, 90.543-555.

Landels, J.G. 1978. *Engineering in the Ancient World.* Berkeley: University of California Press; London: Chatto & Windus.

Mansell, W. (2005). Control theory and psychopathology: An integrative approach. P*sychology and Psychotherapy: Theory, Research and Practice, 78*, 141-178.

Marsh, A. A., M. N. Kozak, D. M.Wegner, M. E. Reid, H. H. Yu, & R. J. R. Blair. 2010. “The neural substrates of action identification”. *Social Cognitive and Affective Neuroscience 2010*; doi: 10.1093/scan/nsq004.

Mayr, Otto. 1970. *The origins of feedback control*. Cambridge, MA: MIT Press.

Mechsner, F., K. Kerzel, G. Knoblich, & W. and Prinz. 2001. “Perceptual basis of bimanual coordination”. *Nature* 414:69-73.

McBeath, M. K., Shaffer, D. M., & Kaiser, M. K. 1995. “How baseball outfielders determine where to run to catch fly balls”. *Science*, 268:69–573.

Miller, GA, E. Galanter, & K.A. Pribram. 1960. *Plans and the structure of behavior*. New York: Holt, Rhinehart, & Winston.

Ozer, Mark N. (ed.) 1979. A cybernetic approach to the assessment of children: Toward a more humane use of human beings (A Westview special study). Boulder, CO:Westview Press.

Piaget, J. 1937/1954. *La construction du réel chez l'enfant* / *The construction of reality in the child*. New York: Basic Books.

Powers, William T. 1973. *Behavior: the control of perception*. 1st ed. Hawthorne, NY: Aldine de Gruyter. [See 2nd Exp. ed. 2005.]

Powers, William T. 1979a. "A cybernetic model for research in human development". Ozer (1979:11-66). [Repr. Powers (1989:167-220).]

Powers, William T. 1979b. *The nature of robots*. Part 1, “Defining behavior”. *BYTE*, 4.6:132-144 ; Part 2, “Simulated control systems”, *BYTE* 4.7:134-152 ; Part 3, “A closer look at human behavior”, *BYTE*. 4.8:94-116 ; Part 4, “Looking for controlled variables”, *BYTE*. 4.9:96-112. [Online at http://www.livingcontrolsystems.com/intro_papers/ as four PDF files.]

Powers, William T. 1989. *Living control systems*. New Caanan, CT: Benchmark Publications. [Selected papers 1960-1988.]

Powers, William T. 1998. *Making sense of behavior*. New Caanan, CT: Benchmark Publications.

Powers, William T. 2005. *Behavior: the control of perception*. 2nd expanded ed. New Canaan: Benchmark.

Powers, William T. 2008. *Living control systems III: The fact of control.* New Canaan, CT: Benchmark Publications. [Includes a CD with a set of demonstration programs.]

Robert A. Rescorla 1988. “Pavlovian conditioning; it's not what you think it is”. *American Psychologist* 43.3:151-160.

Robertson, R. J., D.M. Goldstein, M. Marmel, & M. Musgrave. 1999. “Testing the self as a control system:
Theoretical and methodological issues.” *International Journal of Human-Computer Studies* 50:571-580.

Rosen, G. M., & Davison, G. C. (2003). Psychology should list empirically supported principles of change (ESPs) and not credential trademarked therapies or other treatment packages. *Behavior Modification, 27,* 300−312.

Runkel, Philip J. 1990 / 2007. *Casting Nets and Testing Specimens: Two Grand Methods of Psychology*. New York: Praeger; Hayward, CA: Living Control Systems Publishing.

Schrödinger, Erwin. 1944. *What is life: The physical aspect of the living cell*. Cambridge: Cambridge University Press.

Shaffer, Dennis M., & Michael K. McBeath. 2005. "Naive beliefs in baseball: Systematic distortion in perceived time of Apex for Fly Balls". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31.6:1492–1501.

Sherrington, Charles S. 1906. *The integrative action of the nervous system*. New York: Scribners. [Repr. 1947 New Haven, CT: Yale University Press.]

Schwartz, J.M. (1996). *Brain lock*. New York, NY: HarperCollins Publishers.

Vallacher, R. R. & D. M. Wegner. 1985. *A theory of action identification*. Hillsdale, NJ: Lawrence Erlbaum Associates.

van de Rijt, Hetty, & Frans Plooij. 2010. *The Wonder Weeks. How to stimulate your baby's mental development and help him turn his 10 predictable, great, fussy phases into magical leaps forward*. Arnhem, The Netherlands: Kiddy World Promotions B.V. [First published in 1992 as Oei, ik groei! by Zomer & Keuning Boeken B.V., Ede and Antwerp. Updated with interactive participation of researcher-parents.]

Wampold, B. E. (2001). *The great psychotherapy debate: Models, methods, and findings.* Mahwah, NJ: Lawrence Erlbaum Associates.

Wiener, Norbert: 1948 [1961]. *Cybernetics or control and communication in the animal and the machine*. Cambridge, MA: MIT Press. [Repr. with new Introduction 1961.]