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

Carey, Goldstein, Marken, Nevin, Powers, Taylor:
20100920 plus content from WTP20100916 (BN20100914, WTP20100915, TC20100915)

The phenomenon of control is integral to psychology.. Even a cursory glance through databases and journal contents pages 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 implicit references to the phenomenon are included, 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 undisputed. Curiously, though, while the phenomenon of control is often front and center in research programs, the same cannot be said for the mechanisms of control. It appears that investigations into the phenomenon of control have proceeded in the absence of a clear understanding of how control works.

In many ways, this seems to put the cart before the horse. Understanding the mechanisms that allow the phenomenon of control to occur should perhaps underpin any research investigating control. An accurate model of how control occurs could clarify current misconceptions about the phenomenon as well as explaining its widespread applicability including its relevance to other important concepts such as cooperation and conflict. Building a model that is capable of producing simulations is perhaps the most rigorous standard of model development. Knowledge acquired through functional models can sometimes be surprising, even counterintuitive, but conclusions can be made with a degree of confidence that is not justifiable with conceptual or statistical models of the same phenomenon.

I have my finger on a button beside a door. I ask myself: “What am I doing?”, and the answer seems simple: “I’m ringing the doorbell”. But is that why I am there? 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 an onlooker, 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. You probably would not propose that my objective was simply “to press my finger on a button”.

This illustration may seem to belabor the obvious—that in everyday life, people have purposes, that the purpose is what the person is “doing”, and that their actions are the means for doing it. Equally obvious is that one usually cannot determine a person’s purpose by observing their actions. What is perhaps less obvious is that there are three aspects to this, a “what”, a “why”, and a “how”.

In the doorbell illustration, the first "what" that an observer might guess is "he is ringing the doorbell". Although its "why" is obscure to an observer (but not to the doorbell ringer), its "how" is clearly "by pressing the doorbell button". However, even this "how" has its own "What-why-how" pattern. "What" is "seeing and feeling my finger pushing the button", "why" is "to make the bell ring" and "how" is "by moving my hand and arm to the appropriate place".

If our observer looks at the possible "why" of pressing the doorbell button, another pattern of "what-why-how" emerges, for which making the doorbell ring is the "how". "What" might be "to get someone to open the door", or put better "to see someone open the door". And so it goes. Every "what someone is doing" is part of a "what-why-how" structure. In every case, "why" is because some state of the world is not as the person would like it to be, and "how" is a means of making the world a little or a lot closer to what the person would wish.

All of this sounds self-evident, albeit anecdotal and not very scientific. But it can be scientific.
The "what-why-how" complex describes "control". This does not mean forceful dominance of people or the environment, it is a technical term of art that means bringing some particular condition toward a desired state and maintaining it there. That is the engineering definition of control, and the thesis of this paper is that control is what living organisms do. Indeed, it is what you are doing, on many levels and in many ways concurrently, as you read this paper.

This informal account of who-what-why points to a basic model of behavior that has been under development since the early 1950s. It was first published in 1960, and was named Perceptual Control Theory (PCT) during the 1980s by members of the interdisciplinary, international group of researchers and practitioners that have engaged with it. This paper is a summary of the PCT paradigm as it is presently understood, including methodology, results, and applications.

Behavior as the means of control

The basic thesis of PCT is not difficult to describe. The behavior of organisms—their observed activity—is not the final product of prior causes. Rather, it is understood to be a variety of means to ends. The ends are manifested to observers of organisms in the way the behavior stabilizes aspects of the local environment against disturbances. From the organism's standpoint, the ends are certain experiences (or cessation thereof) that are intended or preferred. PCT is about purposive behavior.

What, again, is "control" as used in PCT? Consider another example. I hear music. In the language of PCT, I have a reference value for how loud I like this kind of music. At the same time, I perceive the current loudness of the music. I compare the loudness I perceive with its reference value, and if there is a difference, I do something that changes the physical environment to alter the loudness I perceive. Maybe I put in earplugs, maybe I move to another room closer to or further from the source, maybe I ask someone to turn the volume knob or turn it myself to make the music I hear louder or softer.

At the same time, other things might influence the loudness I perceive of the music. Maybe someone closes the door of the room where the music is playing, or turns the volume knob. I continuously perceive the loudness of the music, and at any time that it differs from my reference value, which may change from moment to moment, I may behave in such a way as to bring the loudness that I perceive nearer to the reference value that I currently have for it.

What we are describing here is a feedback loop, as illustrated by the diagram of the canonical PCT control loop in Figure 1.

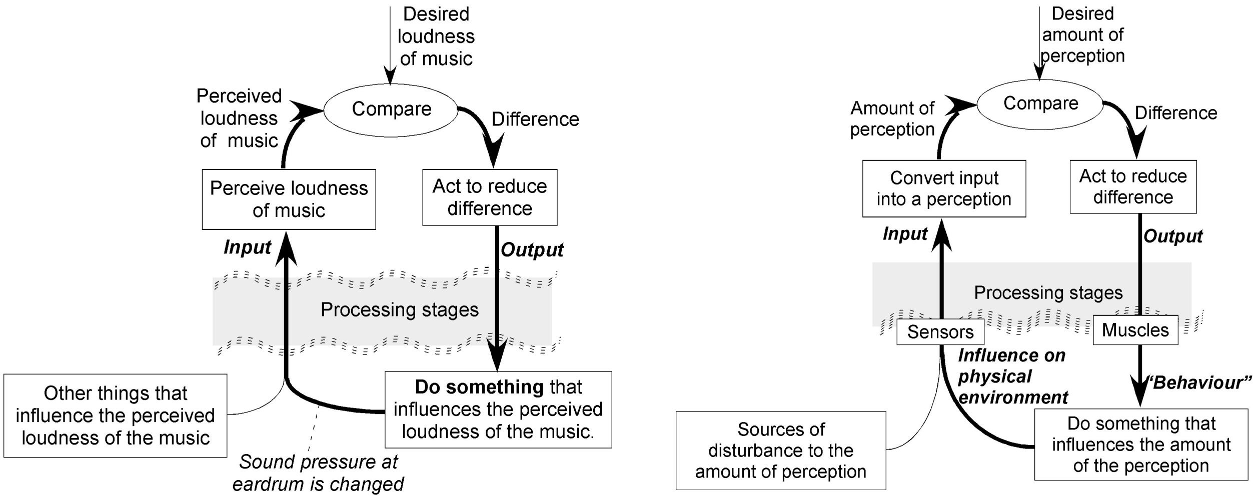

Figure 1. The canonical PCT control loop. (Left) controlling the loudness of music (Right) the generic loop. The key point is that what is controlled is the value of some perception, by means of the behaviour that influences the physical environment.

Not only is it a feedback loop, but the feedback is *negative*. In popular parlance "negative feedback" is equivalent to criticism, whereas "positive feedback" suggests encouragement.[1] In the original engineering meaning, however, positive feedback increases the difference between the reference (desired) amount of perception and the perceived amount, the opposite of what is needed for control, while negative feedback decreases the difference. A negative feedback control system can be designed so it can reduce the difference or error until it is at the limits of measurement (if even the smallest measurable differences are considered worth the energy needed to correct every little error). Only error correcting or negative feedback results in control.

This requires that the basic organization of a living system be of the kind that is capable of controlling.[2] A central concern of PCT from the outset has been to deduce the necessary properties of that internal organization by creating and testing generative working models of the actual behavior of individual organisms. Because behaviour resulting from changes is the means to create and stabilize specific conditions of the organism and its environment, causality in this kind of system is circular. What appear at first to be ordinary physical consequences of motor activities are recognized to be states of the world and of perceptions actively sought and defended against disturbances.

In the generic control loop of Fig. 1, we can discern the features of the two main concepts of behavior that preceded PCT. Following the path from "Desired amount of perception" through difference, output, and behavior, we have the same organization proposed by early neurologists and accepted by many modern neuroscientists and cognitive psychologists. A high-order plan or goal is converted,

---

[1] This usage stems from a misunderstanding when the term was taken up by the human potential movement, and has spread thence to fields such as counseling, education, and management.

[2] This is not the same as responding to inputs or generating patterned outputs, although both of these modes of action can be seen in the operation of different parts of the system.

step by step, into the simple or patterned behavior needed to achieve it. Following the path from sensors through intervening processes (omitting the comparison) to the output and the muscles, we have the organization known as a stimulus-response or cause-effect system, in which the environment causes behavior through causal paths connecting input to output. Both of these classical ideas omit the feedback path through the environment, although variations on the basic themes have been offered to take the feedback effects -- incorrectly -- into account. Both classical concepts solidified into schools of thought before engineers discovered the right way to analyze systems having this circular kind of causal organization.

PCT proposes a new answer to the question, what is it that distinguishes a natural arrangement of matter and energy that is alive from one that is not? The kind of control system described in PCT can have purposes of its own—that is, it can spontaneously select as goals future states of the world around it and alter its own behavior to achieve and maintain such goal-states. It can automatically, without external guidance or instruction, adjust its actions to oppose the effects of random and otherwise unpredictable disturbances (if they are not too powerful for it), quickly and accurately enough to prevent their having any important effects. It can control hierarchically; that is, it can adjust one set of goals as a means of achieving other, higher-order goals. It can control many different variables in parallel at the same level of the perceptual hierarchy, and by those means control multiple variables of a higher order at the same time. It can learn and adapt: it can alter aspects of its own organization in ways that matter to it less in order to control variables that matter to it more.

The biological, psychological, and social sciences have commonly studied organisms as simply one more possible arrangement of matter and energy, subject to the same laws of physics and chemistry as any other arrangement. PCT satisfies this requirement—control systems do not require any violation of the laws of physics and chemistry—but PCT recognizes additional laws that are emergent from the negative feedback control properties of suitably organized physical and chemical arrangements. This enables a systematic accounting for the behavior of organisms, individually and in groups, without which this recognition can only be treated statistically.

As has long been known, life is negentropic.[3] Organisms exploit the orderliness in the world around them as a means of increasing their own orderliness and stability. Control theory explains how an organism can impose order on its local world at the expense of order elsewhere. Control does not confer totally arbitrary intervention in the processes of the environment, but it often seems to do so, in that the organism and its world both behave quite differently when an organism is in control. A car left to steer itself would soon run off the road or collide with another car if only momentum, gravity, wind, and potholes affected it. But add a driver to the car, and it—along with a huge number of other cars with drivers—stays on the road, in its proper lane, for hundreds of miles, and travels to a destination with great reliability. This is a highly improbable outcome when a controlling agent is not present. With control added to the picture, the same outcome becomes highly probable.

Given the fundamental characteristics of negative feedback control, there follow significant differences from other life sciences in how PCT research is conducted; we will return to these later in this paper. At this point it is important to note a shift in perspective that comes with the recognition of the phenomenon of control. Without that recognition, behavior can be considered only from the point of view of an external observer, who as a scientist has little choice but to try to explain the activities of behavior as a mechanistic outcome of external forces acting on an otherwise inert “preparation.” Once

---

[3]Schrödinger, Erwin *What is Life - the Physical Aspect of the Living Cell*, Cambridge University Press, 1944; Brillouin, Leon: (1953) "Negentropy Principle of Information", *J. of Applied Physics*, v. **24(9)**, pp. 1152-1163

we recognize control as a phenomenon, and that it is the perceptual input that is controlled by means of behavioral activities, the relevant point of view becomes that of the organism, not that of an external observer. We cannot account for the how and what of the organism's activity until we have determined the why of it. How we do this in PCT will be explained in the section on methodology.

By this shift in perspective, PCT reconciles the objective approach of science and engineering with subjective experience. It provides a clearly mechanistic model of behavior that can be implemented and studied as a computer simulation, and which also explains how human beings can have goals, intentions, preferences, desires, and other experiences that have sometimes been thought to be figments of the imagination or simply errors of interpretation.

The question naturally arises: if PCT has been building into a coherent model for 60 years or so, with a vigorous research community gathered around it, why doesn't everyone know about it? Perhaps the most important reasons are found in an unfortunate development that occurred almost as soon as control engineers had elucidated the phenomenon of negative feedback control.

A discovery abandoned

Devices employing negative feedback control are documented as long ago as about 250 BC, but it was only in the 1930s that the principles were formalized by engineers. This was the basis of the wartime automation revolution of the 1940s. Recognizing the resemblance of electro-mechanical negative feedback control systems to living systems, Rosenbleuth, Weiner, Ashby, and others initiated the new field of cybernetics. A cybernetic revolution in the life sciences began to gather momentum in the late 1940s and early 1950s.

But the revolution came to a halt, essentially dead, in a decade. Negative feedback control was abandoned as a model of purposive living systems almost as soon as it was adopted by its main original proponent, the prominent cyberneticist W. Ross Ashby. In place of the negative feedback model, Ashby and others offered a different idea. Organisms, they proposed, analyze the environment, determine what actions would be needed to produce desired results, and then issue the commands necessary to make the muscles generate those actions. This represented a return to an idea of brain operation originally offered by Sherrington in 1906, in which the cerebral cortex formulated general commands that were then elaborated, level by level, into the detailed commands reaching the "final common pathway" to produce organized behavior.

From that time onward, negative feedback control has been regarded by many as old-fashioned. In 1960, Alfred Chapanis, then president of the Society for Engineering Psychologists, wrote "The servo-model, for example, about which so much was written only a decade or two ago, now appears to be headed toward its proper position as a greatly oversimplified inadequate description of certain restricted aspects of man's behavior." This was written in the same year that the first paper leading to PCT was published. Writing in *Purposive Systems*, the 1968 proceedings of the first annual symposium of the American Society for Cybernetics, Ralph Girard, a founder of the Society for Neuroscience and a contributor to the Macy Conferences on cybernetics, said "I have always regarded a drop of water sliding down a slightly inclined plain [sic] as showing all the manifestations of purposeful behavior."

Even in the book in which he first wrote about negative feedback control, Ashby had argued persuasively that this more complex design based on analyze-compute-execute processes would operate faster than the negative feedback control system, eliminating delays, and that it would be more accurate since it did not need to allow any errors to occur. It could even, he proposed, anticipate disturbances and generate actions to oppose them at the same instant they occurred. Since evolution would naturally have shaped organisms to operate in the best possible way, it was assumed that this model should also be used to explain the behavior of organisms, although this explanation would of necessity be limited to organisms that are sufficiently complex to carry out the required analyses and matrix inversions. (The cognitive capacities of primitive organisms have consequently been a perennial source of surprise, but puzzlement about this has not led to reconsideration.)

Unfortunately for this view, real organisms seldom behave as optimal control systems. It is, in fact, easy to design artificial control systems that control much better than people do, but that amounts to making a model of the behavior of a perfect robot, not of a human. To make a model that behaves as much as possible the way a real person does—in, for example, a tracking task—it is necessary to resurrect the negative feedback control model. Ashby had the right idea when he explained the importance of negative feedback control in the first part of his first book on the subject, *Design for a brain*. PCT would have been accepted long ago, at least in cybernetics, if he had not written the rest of the book.

What kind of theory is PCT?

The simplest form of theory in psychology, and the most prevalent, is a set of statements of what will be observed under certain experimental conditions, such as “Mothers hold their babies on the left side.” The only test of such a theory that is possible is to observe whether this is how mothers really behave. It has no generality, no necessary connection to any other observation. Either most mothers do this, or they don't. There are many other theories of this type. Another example is Piaget’s stage theory of the cognitive development of children, that is, a theory of changes in their cognitive abilities and processes. Piaget saw actions as the basis of early development, and mental operations coming in later. Out of this grew his notion of schemas, categories of knowledge supporting our interpretation and understanding of experience. A schema, according to Piaget, includes both a category of knowledge and the process by which it is established, and new experiences are fitted into existing schemas, adjusting or adding to them as needed. Thus, a child in a family with a pet chihuahua conceives of the category “dog” as short-haired, four-legged, and small, but after encounters with other dogs in the park this schema is modified to accommodate more varieties of dogs.

This theory is supported if in fact it is observed that cognitive development is first based on actions and later on mental operations. It must be tested, but the only test possible is to find that this pattern is repeated, or fails to repeat, when observed again. There is no hypothesis to test concerning how the child must be internally organized in order to show these patterns under the conditions described. The theory is purely descriptive.There isn’t even any statement about how categories exist as elements of behavior or experience—they just exist, along with the knowledge itself, as though given in the environment or innately in the child. So it seems that this sort of theory is about the information content of the brain, not the brain’s structure or organization.

PCT advances hypotheses, many of them testable even now, about the structure and organization of the nervous system, including the brain, at the level of functions if not anatomy. If Piaget’s “theory” (proposed observation) is correct, we would expect eventually to use PCT to try to explain how the child manages to do these things. From the PCT point of view, Piaget has offered some data that, if valid, require theoretical explanation. His theory is not a theory of the same sort as PCT. It describes a phenomenon, but does not explain it or provide any principled basis for it, and offers no justification beyond generalizations of descriptive observations. Generalizations can only be tested by further observation of instances. PCT neither supports nor denies theories of the type offered by Piaget: It accepts them to the extent that they are valid, and then, if all goes well, offers an explanation for why or how the phenomenon in question is generated -- and suggests other similar phenomena that might prove to be observable.

[I eliminated a paragraph bragging about how wonderful PCT is. WTP]

When we look in the behavioral sciences for a theory of the same type as PCT, probably the oldest, dating back to Descartes, is the proposition that stimuli acting on sensory nerves are the cause of motor behavior. Like PCT, this theory is an attempt to explain all behavior, not any particular one. The other main theory of this type, deriving from ideas such as Ashby's mistaken analyze-compute-execute hypothesis and Sherrington's map of the brain, is that such stimuli cause centers in the cerebral cortex to generate plans of action, which are then executed by lower systems to produce desired ends. It is with theories at this level of coverage and intent that PCT must be compared, theories of the kind that aim to explain how all behavior is produced, rather than attempting to describe or predict what specific behaviors will be observed under specific circumstances.

To recapitulate the basic principles of PCT: Behavior is not a linear result of prior causes, it is the variable means of achieving goals that the behaving organism specifies within itself. The activities of behavior are only one among many of the causes that affect some aspect of the environment about which the organism has some preferences; or more exactly, affect the organism's perception of that aspect. The difference between that perception (a neural firing rate) and the preference, an internally specified reference value for that same perception (another neural firing rate) is the cause of the behavioral activities (by way of the propagation of the resulting neural signal downward through the control hierarchy). By this control loop of circular causation, the organism does whatever works to maintain that difference at or near zero; behavioral actions vary precisely as needed to achieve consistent aims.

This answers philosophical objections that fortunately are not much heard since the fall of logical positivism into disfavor. Criticisms were made that any notion of goals or purposes must be disallowed in a science of behavior, because it would require the future to affect the present, or effects to become their own causes, or infallible predictions to be made. Those objections may be outdated, but they were responsible for a general rejection of the idea of purposive behavior at the time when important psychological theories were just starting to form. The consequences are still with us. The efficacy of PCT models demonstrates, however, that all that is needed to account for purposive behavior is continuous perception, comparison, and action, all of which go on simultaneously rather than in sequence, and each of which causes and is caused by the others.

Quantitative and qualitative theories: variables and categories

Following Ashby, the conventional ideas of control most widely accepted today propose that an organism achieves goals in steps, by first analyzing the environment, then calculating the actions and trajectories of action needed to bring the goal-state about, and finally by executing the actions. We have already noted two of the reasons this analyze-calculate-execute hypothesis cannot support an adequate model of behavior. Yet the evidence for this model certainly seems clear: the actions required to achieve a goal-state are indeed produced with the normal result of successful goal-attainment.

The evidence, however, is far richer and more informative when we measure the variability of behavior rather than counting instances of “behaviors”. Closer inspection shows that the actions are not as regular and repeatable as they seem at first, and that in fact repeated goal-seeking actions have regular effects precisely because they are not repeated exactly. The reason is that those regular effects are influenced by more than just the organism's actions; there are also independently varying influences in the environment, including past and present states of the organism itself. Results can be repeated only by varying the actions so that they precisely counter those unpredictable disturbances and changes in environmental conditions which simultaneously are also influencing the result. It is not just that many different actions *can* produce the same result, a qualitative observation that Skinner proposed in his definition of the “operant;” different actions *must* be employed, and just the right different actions each time. As we will see in demonstrations later in the paper, actions must vary quantitatively in exactly the right way if the same result is to recur.

The only reason that behavior (the observed activity) seems to repeat is that human observers tend to think qualitatively rather than quantitatively. Qualitative thinking is categorical, but behavioral activity does not leap discontinuously from one category to another, it is continuous. A driver making a left turn seems to be generating a stereotyped behavioral pattern that is qualitatively the same each time it is executed, as if it were a simple repetition of what has been done before. This has been taken to imply that repeating the result of the action means that the nervous system must be issuing the same commands to the muscles each time. But that implication is dissipated as soon as an engineer's or a physicist's eye is brought to the scene. The car never approaches the intersection of roads along exactly the same line or at the same speed as the last time; the tires distort, bounce, and slide by different amounts each time they encounter smooth or rough spots on a road that may be dry or slippery; crosswinds require more or less effort to be applied to the steering wheel to achieve the same turning path; the speed of the car influences the turning radius, as does the number of passengers in the car. Yet somehow, every time there is a left turn the steering wheel turns in just the manner required during that particular turn for the car to move in very nearly the same stereotyped fashion from the lane it is in to its proper place in the crossing lane. It is the result that is stereotyped, not the action that produces it or the neural commands that operate the muscles. Conversely, repeating precisely the same neural output signals or actions each time would not produce the same consistent result. That fact will prove below to have enormous consequences for the theory of reinforcement.

After sufficient quantitative observation of behavior, it becomes clear that it is not an organism's neural outputs or motor actions that repeat, but the consequences of those outputs and actions. The outputs and actions themselves vary exactly as required to keep the consequences the same. The small disturbances revealed by close inspection—as well as some large ones—have multiple independent causes that arise from different environmental sources on different occasions, at unpredictable times, in unpredictable directions, and to unpredictable degrees. Yet what we observe is exactly the kind of variation in behavior that is needed, given all the other influences acting at the same time, to make the critical consequences repeat.

By conventional ways of thinking this is impossible. But control systems do not operate in a simple input-output way. They can control consequences because they continually monitor the state of the consequences, and when that state differs, moment by moment, from what is expected or intended, the difference is used as the basis for altering the action in exactly the way that will keep the difference as small as possible. That is how the needed variations in behavior are produced, and why they do not need to be calculated in advance.

Symbolic vs. analogue computation in the nervous system

The idea of "computation" of outputs suggests that variables are converted into symbolic representations which are then manipulated according to the rules of mathematics, as in a digital computer, to generate a derived symbolic specification, which is then converted back to terms of action. But PCT models analog computation in the nervous system with continuous rather than discrete variables, and the mathematics involved in simulations is not intended to represent the physical processes taking place, but only to describe how variables change or to approximate their effects in the language of mathematics. The biochemical processes being modeled are direct physical interactions, not abstract symbolic computations.

An example is the construction of certain perceptual signals as weighted sums of raw sensory signals. In the symbolic approach, each sensory signal would be modeled as a discrete variable with a particular value; the weighted sum would be created by multiplying each signal by a weight and then adding together all the products to create the sum. The sum would then be converted into a magnitude of a neural signal.

The analog-computing version of this process has no need of the symbolic phase. Two or more signals reach synapses on a target neuron; each signal releases neurotransmitters which result in positive or negative changes in post-synaptic potentials; these changes contribute to the net setting of the firing threshold, which determines how fast the cell will send impulses into the cell's axon to provide input for the next cell in line. The relationship between incoming and outgoing impulse rates is a continuous function; output signals change as the input signals change. There is no pause for a computation phase; if we graph them, the output change curves are nearly simultaneous with the input change curves, and overlap in time.

Parallel computation

The simultaneity of all processes linking input and output emphasizes another fact about analog computation in the nervous system: all phases of the computation are occurring at the same time rather than one after another as in analytical mathematics. The cells in a nervous system function entirely in parallel, each converting its inputs into outputs at the same time that the others are doing the same thing. A control system made of neurons and muscles functions as a whole, not one part at a time in sequence. If there are time delays, the delays do not imply sequentiality of action; they mean only that the current inputs to some cells are the outputs from other cells as they were some milliseconds in the past. Continuous variations, even if delayed, are still continuous, and delays are subsumed in the rate of change as noted above.

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," 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, 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 anticipated cybernetics in many ways, and in the 1950s came to the same conclusion that was being developed in the ancestral theory that became PCT: behavior has to be multiordinal -- organized hierarchically, in layers. A simple observation 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 higher centers 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 way for 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 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. 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 casts new light on Sherrington's concept of a "final common pathway," which he took to consist of signals carrying commands telling the muscles how much to contract. In a control-system model of the reflexes, the muscles are operated not by reference signals or command signals, but by error signals. The signals from spinal motor cells carried by alpha efferent axons to muscles result from two inputs to the motor neuron: an excitatory input descending from higher centers, and an inhibitory input coming from sensors in the tendons measuring the force applied by the muscle. The net signal leaving the motor neuron represents the excess of excitation over inhibition, and the feedback loop at this level simply makes the sensed tension in the tendon (due to the force exerted by the muscle) match the constant or changing reference signal received from above. Thus the brain (or a system higher in the spinal cord) sends the motor neuron a signal saying, in effect, "Make the tension signal match *this* signal." The feedback loop alters the output to the muscle, in just a few milliseconds, until the match is achieved. The reference signal is not a command to produce a certain amount of action; it is a request for a certain amount of perceived force or tension.

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, not what action they should perform. What 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.

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, which normally calls for a loud voice, but they do not want to wake the baby. Control of the perception of warning the person wants to use a loud voice; control of the perception of the baby sleeping wants quiet. Two control loops are controlling the same environmental variable, the loudness of sounds in the room, trying to produce very different values of that variable. The person may resolve the internal conflict in this case by gesturing to get the person's attention and by whispering. The two controllers may be in different people. You approach an open doorway at the same time as someone else coming the opposite way. One may stand aside and wait, or they may each turn sideways to slip past each other.

Most conflicts are routinely resolved. When a conflict cannot be resolved, 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), which will be described in a later section.

When conflicts between control systems inside one organism or in different organisms are not resolved, the result can be a serious loss of function. Each system tries to make the same physical variable match a different reference condition. If the difference is moderately large, both systems will experience large control errors, and as a consequence at least one of the control systems will produce as much output action as possible, limited only by strength and endurance. A between-organisms example is the conflict between a rat in an operant conditioning experiment and the experimenter. The rat normally behaves in such a way as to maintain its body weight at the "free-feeding" level, but in this case the experimenter adjusts the available food (between experimental sessions) so as to keep the weight at 80% of the free-feeding weight. The experimenter, having complete control over the rat and a decisive strength advantage, does keep the weight at that low level, while the rat ends up simply pressing the lever at the fastest rate it can sustain in a vain attempt to increase the food intake. The rat has lost control of its own body weight. [Footnote: Abbott (unpublished) has shown that apparent changes in pressing rate when a fixed-ratio schedule is varied are due entirely to the time taken to consume the food after each reinforcement. The actual continguous pressing rate is simply the fastest rate the animal can maintain, and does not change as the schedule changes.]

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. In the case of conflict, the control actions of each are a disturbance to control by the other. In the case of cooperation, these conflicts are resolved as they arise. The mutual adaptations of this resolution are more complex and variable to model. Conflict is easy, all it takes is for control systems to control their inputs regardless of each other. For control systems to regard each other and accommodate their control of environmental variables to a common goal is more complex. Instrumental in this are learned reference values for perceptions that we think of as social expectations, mores, customs, rules, and laws. These have the effect of stabilizing an environment that includes other people so that fewer resources are tied up resolving external conflicts, at the cost of limiting the freedom of control and sometimes creating conflicts within the hierarchy. This explains why cooperation , even when highly valued, is difficult to put into practice.

The resolution of conflicts requires changes in some part of control systems that create behavior -- the perceptions or the actions must become different in the same environmental situation. In PCT such changes are described as a general process of reorganization.

Changes of organization

The final facet of PCT is concerned with ontogeny, how a mature control hierarchy grows out of the primitive organization of a new organism. In accord with the general principles of PCT, this process of changing control systems is seen itself as a control process, in which variables of basic importance, referred to by Ashby as "critical variables," are maintained near reference states by altering the organization of the organism.

The main alternative to the reorganization concept is the idea of reinforcement. When a behavior occurs that has a reinforcing effect, it is said that the probability of that behavior's occurring in the presence of the same discriminative stimuli is increased. This was Thorndyke's "law of effect," picked up and elaborated by B. F. Skinner and many others early in the 20th Century. Skinner summed up reinforcement by saying that behavior is controlled by its consequences.

There is one major flaw in this concept which has already been noted here. Repetition of a behavioral action is not likely to result in repetition of its consequences. Viewed qualitatively as countable categories, “behaviors” like pecking a key or pressing a lever do seem to repeat, and the repetition is what appears to result in more reinforcement of the same behavior[4]. But many experimentalists including Skinner noticed quickly that animals in conditioning cages do not actually repeat the same motor behavior again and again. They do succeed in making contacts beneath a key or lever close to deliver reinforcers, but the actual motor behavior used to do this can vary enormously. A rat can operate the lever by leaning on it, chewing it, sitting on it, or standing on it with a front or rear paw. The approach to the lever depends on immediately prior activities and many other factors. The categorization of diverse actions as “lever presses” by the observer conceals these differences, so that in fact the manner of recording data can attribute a specific rate of lever-pressing to the animal, the total number of presses divided by the total time of counting, even though the animal spent part of that time having a nap.

Taking this variability into account, we observe that a free-feeding animal at normal weight does whatever it takes to receive enough food in the artificial environment of the laboratory. It varies its motor behavior, without any particular repetition, in exactly the way required to make the same consequence occur under changed conditions. This can be done by a negative feedback control system without any reorganization being required once the system attains a final organization.

E. coli reorganization

[4]Reinforcement is an example of positive feedback. More reinforcement means more behavior of a certain kind; more of that behavior means more reinforcement. This is an unstable, error-increasing kind of organization which can only produce the maximum possible behavior or none at all.

Skinner explained the acquisition of the first successful behavior in conditioning experiments by saying that organisms spontaneously produce random variations of behavior. PCT adopts that idea but in a different form: the basic theory of reorganization is that behavior varies randomly when there is "intrinsic error." Starvation is an instance of such a challenge to the state of the organism. Deprivation is not just an "establishing condition." It causes control errors that bring reorganizing processes into action.

Intrinsic error means a difference between the state of some critical variable, such as blood sugar, and a genetically-determined reference condition. This difference results in random changes of organization. The kind of learning involved is fundamental, the kind that occurs when there is no systematic method available for higher levels of control to pursue, and when there is no prior experience to guide changes. Because the changes are unstructured, they are not constrained by anything but the existing organization, so the possibility of finding solutions to new control problems is maximized.

Clearly, if the random changes of organization produce behavior that eliminates 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 new patterns, just as if something had told the organism that the new pattern was a good one. But doubts about this idea are well justified; it doesn’t seem very likely to work.

This concept has been part of PCT since the first published paper in 1960, but it seemed too inefficient. 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 increased 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 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 fact that E. coli senses the rate of change of concentration of chemical substances that is 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 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 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 than down the gradient. For repellents, of course, the relationships are reversed. According to Koshland, E. coli can behave in this way in relationship to more than 20 different substances simultaneously.

To translate this principle into terms of reorganization, the spatial dimensions in which E. coli moves become parameters of control systems. Swimming in a straight line becomes 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.

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. This is because it makes use of information about the changing size of control errors. A 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. This 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. It is possible that this principle will provide the final rebuttal to arguments that natural selection with random variability of individuals in a population is unlikely to account for the facts of evolution. If evolution is actually carried out at the level of the genome by an organism-generated process of E. coli-type reorganization, it may easily prove to be as efficient as necessary. (The idea that organisms generate their own evolution is not entirely new.)

In the following demonstration of a negative feedback control model (Powers 2008), a person uses a mouse to make a cursor track a moving target for one minute. Data are sampled 60 times per second. The data for a single experimental run are shown in the upper plot of Fig. 1. The red trace shows the target movements; the green trace shows the mouse and cursor movements. The black trace shows the difference between target and cursor—the tracking error.

There is a consistent small time delay, hard to see in Fig. 1, between target movements and cursor movements (upper plot). The delay is not removed by anticipatory mouse movements as Ashby claimed would happen. In the upper part of Fig. 1 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.

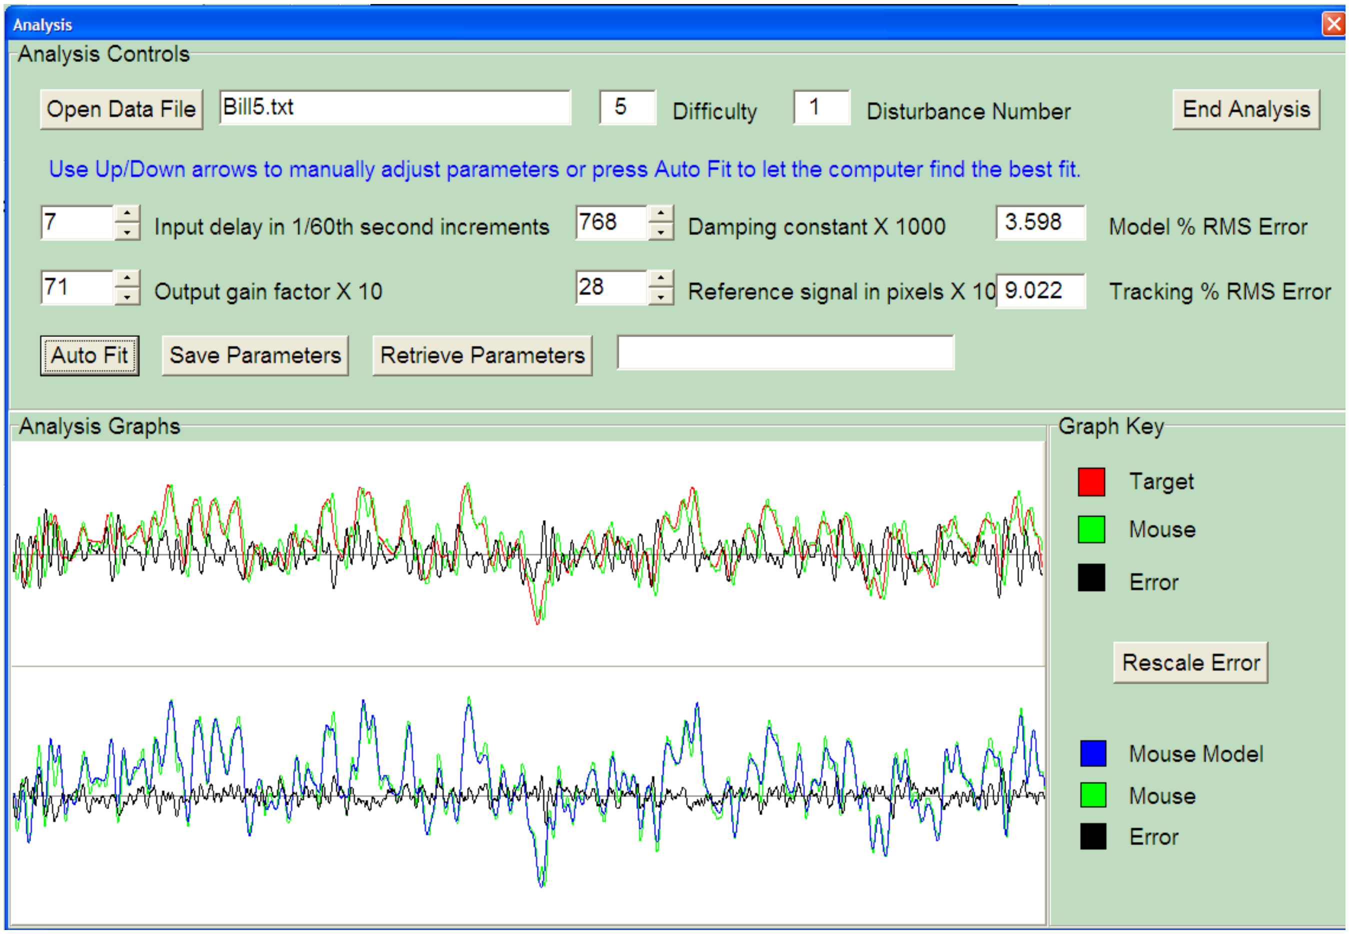

Fig. 1. Analysis of human tracking run and fit of negative feedback control model to the data.

The lower plot superposes the performance of a computer implementation of a negative feedback control system controlling 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 very nearly identical, with the same delay relative to the target movements. (The vertical scale is somewhat expanded, making this easier to see.) The mean 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; the model fits the real data well within the tracking error, showing that the model is making similar mistakes. This same model will work perfectly well with the delay set to zero. But it will work too well: with all the remaining parameters optimized, the mismatch with the real behavior rises from 3.6% to 6.0%. The delay is real.

Ashby's analyze-calculate-execute hypothesis is inadequate on at least two grounds; that it does not model actual behavior, and that it cannot model the behavior of organisms that lack the cognitive complexity required for such inverse kinematic/dynamic and planning computations. A homely example will suggest the difference in complexity. A certain apartment complex in Germany installed a building heating system which worked by measuring outside air temperature and calculating heat losses from the engineered insulative properties of the building envelope, then adjusting a furnace to counteract the losses. The same work is commonly done by negative feedback control loops, where the sensor and actuator in each space in the building is the simple analog device called a thermostat that doesn't even have a microcomputer in it..

In addition to Ashby's abandonment of classical control theory as a model of organisms, certain misconceptions about negative feedback control that have gained currency are an additional obstacle to its acceptance. Early in its history, various commentators noted that all real systems contain time delays. It was thought, apparently, that with any time delay at all, a negative feedback control system would have to become unstable. Error-correcting actions would start too late to prevent disturbances from having immediate effects, and would persist after disturbances disappeared, generating self-disturbance; and the time delay would convert negative error-opposing feedback into positive error-amplifying feedback, with the likely result that the whole system would oscillate violently.

While time delays can result in pathological behavior, all that is needed to correct it is to make the output driven by the error signal proportional to the time-integral of the error rather than to the error itself. This is equivalent to making the rate of change of output proportional to error. The constant of proportionality is adjusted so that during the time-delay that exists, the feedback effects from the output cannot change by more than the size of the perturbations caused by the disturbance. We will show this in greater depth in the next section.

This adjustment is sufficient to stabilize the system given any fixed or maximum time delay in its response. Even more important, as we have seen above and will further demonstrate presently, a working model of a control system incorporating this principle can reproduce experimental behavior of a human participant, including delays, with an accuracy of three to four percent of the range of variation of observed disturbances and responses, equivalent to a 25 to 30 sigma fit of model to data. There can be no practical possibility that this model fits the observations by chance, since $p < 1E-12$ or much less.

So while it is true that the success and stability of a control process depends on a number of static and dynamic aspects of the system and its environment, and that general treatments of the stability and accuracy of control systems can become very complex, nevertheless in applying control theory to organisms there is a shortcut to a solution: the living system's performance is observed to be stable and accurate, so a biological answer to the problem of stability, even if unknown, clearly exists. Given that observed performance is stable, and that we know of one way of stabilizing a model that accurately reproduces real behavior, we are assured that PCT gives a correct general picture of how control works without requiring that the exact method of achieving stability be known.

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.

Simulations and models

The very simple model that provided the illustration above exemplifies a method of analysis that originated in the "operations research" of World War II and the field of engineering psychology that grew up right after the war. Like PCT itself, it is basically a simple idea; but also like PCT, the power that it proves to have as an aid to understanding far exceeds what its simplicity seems capable of providing. It is important to understand both the simplicity and the power of a model constructed in this way, so we will take some time to study that here. Figure 2 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.

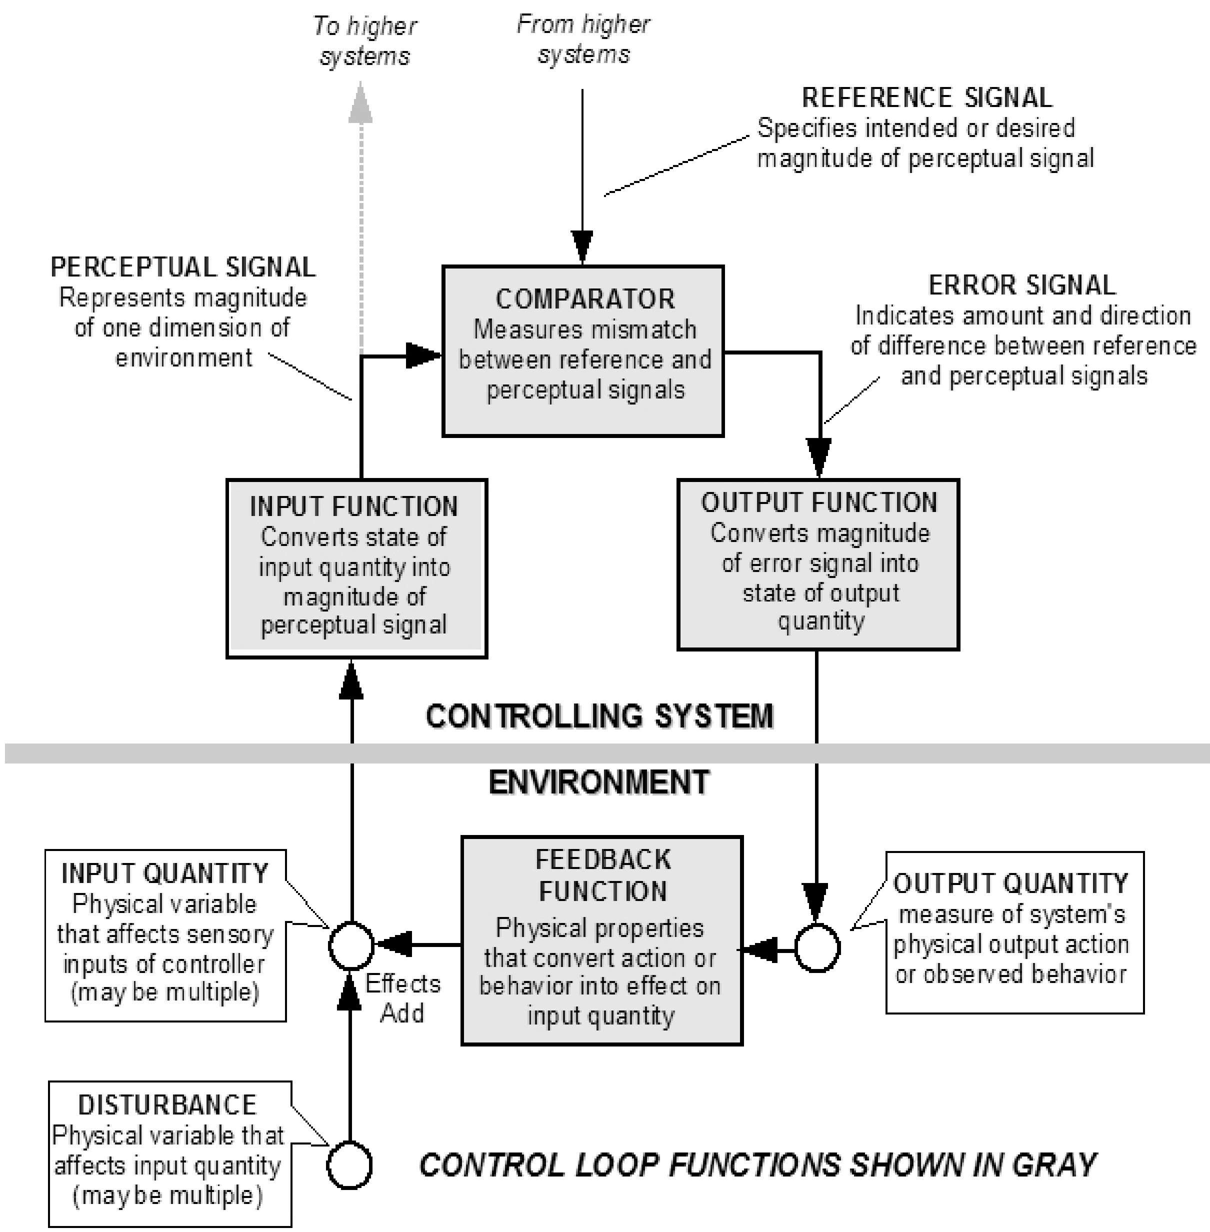

Fig. 2: The basic organization of a negative feedback control system.

The model assumes that inside the participant there is a perceptual signal, some kind of neural signal that literally and quantitatively represents (is an analog of) the input quantity. Applied to the tracking task, the input quantity is the vertical distance between the target position **T** and the cursor position **C**, and the random variation of the target position acts as a disturbance of that input quantity. This suggests that quantitatively the perceptual signal **p** is the cursor position **C** minus the target position **T**, as expressed in the equation $\mathbf{p=C-T}$.

Actually, of course, as we have noted, there is a delay involved in going from the perception of target and cursor to the signal representing the distance between them. (This delay is incurred at lower levels of the hierarchy that have been omitted from the present discussion for the sake of simplicity.) If the delay is τ seconds, the working perceptual signal at time **t** represents the target-to-cursor distance at a prior time, $\mathbf{t - \tau}$, so the correct equation as used in the model is

1. $\mathbf{p(t) = C(t-\tau) - T(t-\tau)}$

The basic negative feedback control system receives a reference signal (**r**) from elsewhere within the organism which specifies the currently intended or desired magnitude of the perceptual signal. The “comparator” emits an error signal **e** indicating the magnitude and sign of the difference between **r** and **p** (the time index is omitted but understood):

2. $\mathbf{e = r-p}$

Experiment has shown that in the best model for the output function the mouse velocity is proportional to the error signal. A positive error (perception less than reference) causes an upward velocity of the cursor that is proportional to the error by a gain factor **G** (that is, $\mathbf{V_{cursor} = G*e}$).

The next position of the cursor $\mathbf{C_{new}}$ is the current position $\mathbf{C_{old}}$ plus the velocity $\mathbf{V_{cursor}}$ times the duration **dt** of one iteration of the program. Making the substitution for $\mathbf{V_{cursor}}$ yields a third equation:

3. $\mathbf{C_{new} = C_{old} + G*e*dt}$

That is the totality of the simplest version of the model: a set of three simple equations or program steps which, evaluated over and over with the same pattern of target positions that the human participant experienced, duplicates the participant's actions in the tracking task above within 4.0% of their peak-to-peak range, in great detail. The model whose performance is illustrated in Figure 1 adds one more term to equation (3), a damping factor **d**, and that is what reduces the discrepancy between the model and the human participant to 3.6%, a small but consistent improvement. With this damping factor, the third equation (as it actually is implemented in the demonstration program) is

3′. $\mathbf{C_{new} = C_{old} + [(G*e)-(d*C_{old})]*dt}$

It is remarkable that these simple equations do so well in simulating real behavior, considering that we are ignoring possible nonlinearities such as the Weber-Fechner law, potential noise in the system, continuously varying angles at the joints, and many other possible causes of poor performance of a simple linear model. In this light, examine the lower plot of Fig. 1 again, showing the mouse/cursor positions of the real person and the model. The black trace representing the difference between model and person consists mainly of small high-frequency oscillations that are too fast for this system to suppress. Within the bandwidth of good control, the errors must be far smaller than the 3.6% to 4.0% of the range of target movement that is measured. There must be something fundamentally right about this hypothetical model.

A set of demonstration programs

Of course, 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.

One set of them can be downloaded from http://www.billpct.org/PCTDemo3.exe to run on a Windows computer. The reader is advised to do so now, because actually running the demonstration is probably the most effective way of understanding what PCT is about. Instructions for running the program are at that location.

“Responding” to an invisible stimulus

The first three demonstrations in the 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 3 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.

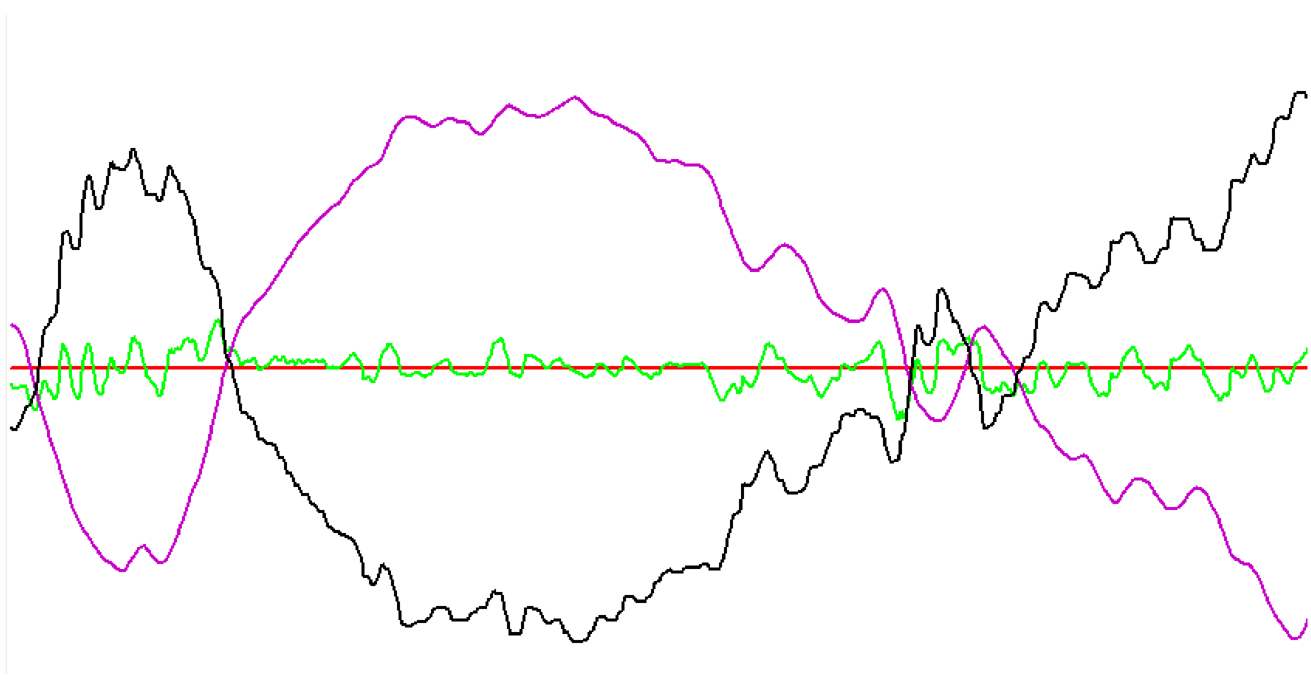

Fig. 3: Compensatory tracking. Black line shows mouse movements, green line shows cursor movements. Target position is horizontal red line. The purple trace shows an invisible disturbance that varied during the run: mouse position relative to the centerline is equal and opposite to the disturbance at 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. 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 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 in that the relationships between signals change (the functions change form), but there is no pattern of behavior to learn: what is learned is control.*

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 4 shows the appearance of the screen at about the 13-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.

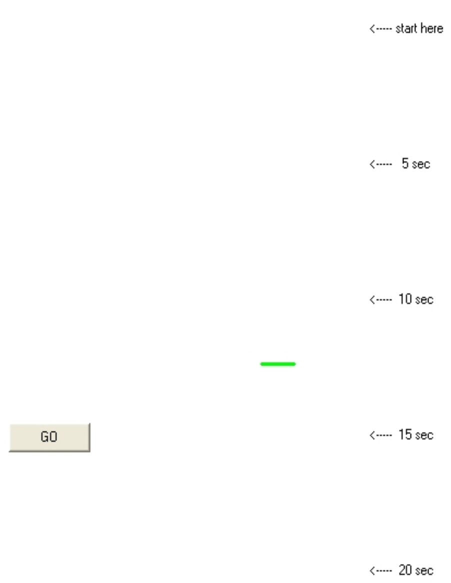

Fig. 4: Demonstration 5 at about the 12-second mark. The green cursor begins at the start line and is moved up and down by unseen disturbances as the subject, resisting these disturbances, attempts to move it smoothly downward so as to reach each successive mark at the indicated time.

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

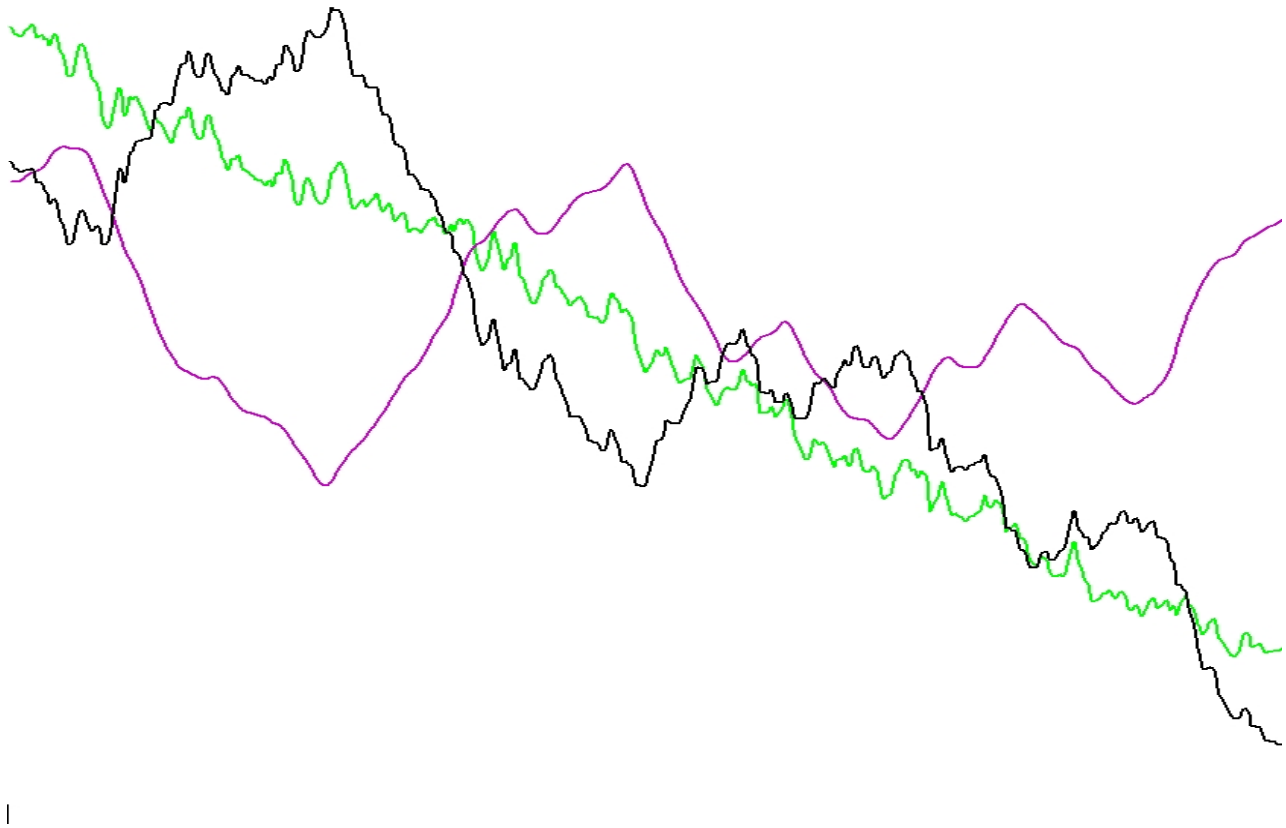

Fig. 4: One run of demonstration 5. The mouse cursor (green) approximates a steady progression from top to bottom (x axis) over the time elapsed (y axis). Mouse movements to accomplish this (black) are the sum of this descending movement plus the actions needed to counter the disturbance (purple). No stimulus corresponds to this activity, and no planning 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 actions 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 rates of change of position. This system generates a slowly-varying reference signal 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). 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.

These demonstrations actually are examples of the rigorous testing that PCT has undergone. As mentioned earlier in this paper, our attempts to prove that PCT is wrong are essential to doing science. Building and testing accurate models of individuals' behavior is at the heart of the theory and the experimental methodology of PCT. The remaining demonstrations illustrate some of the other principles that we have discovered and tested, and which, so far, have withstood all attempts to prove them wrong.

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. 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, 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 it could be used 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 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 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 a controlled perception, the true actions of the person often being markedly different.

In one later demonstration also available (see Resources at end), 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 traces out "hello" upside down and backward. The observed behavior of the participant is essentially unrelated to the control task being accomplished, even though overlaying the disturbance on the mouse movements would show that the mouse movements are very highly and negatively correlated (in the - 0.99s) with the disturbance.

The first demonstration of the LCS III set may be the most philosophically interesting. Here 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 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 merely side-effects..

"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 the north pole 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 they 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 correlate much better with the aspects that are *not* being controlled.

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, output amplification (gain) factor, 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 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. As observers of the 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 principle datum in PCT methodology is the controlled variable. All of the demonstrations that we have reviewed have clearly displayed three variables: the controlled variable (i.e. 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 what the mouse movements are. If the user, unbeknownst to us, is ignoring the moving 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). Since 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 even 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.

The fundamental step of PCT research, the Test for controlled variables, is the *gentle* application of control to a variable that the researcher surmises is already under control by the observed organism. If the organism resists the disturbance and restores that variable to the state that it desires, that is 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. And then it must be realized, in addition, that the perception of the environment by the observer is not the same as the perception of the "same" environment by the observed organism.

In order to build working generative models of behavior, like the simulations we have been exercising, there is one further requirement. We must be able to measure these influences affecting the state of the environmental variable that we have decided to test. 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, and others like it.

Our understanding of the inner workings of a hierarchical perceptual control system reorients our thinking in a number of fields.

PCT Applied to Psychotherapy

The field of psychological disorders and their treatment is an excellent example of the way in which PCT can provide a unifying framework to an otherwise fragmented area of research and practice. It is widely recognized, for example, that current classificatory systems of psychological disorders (such as the Diagnostic and Statistical Manual of Mental Disorders – Fourth Edition, Text Revision) do not easily map onto the lived experience of psychological distress. While concepts such as "comorbidity" have been invoked to explain the lack of correspondence between categories and experience, there is a growing awareness that this system of classification is unsatisfactory in important ways. Exploration of common or "transdiagnostic" processes has been a recent innovation that has attempted to explicate underlying pathways of symptom manifestation. PCT explains why this approach has merit and why categorizing symptoms is problematic.

We have already outlined in detail the fact that behavioral output varies in order to control perceptual input. There is a large amount of variation within current classificatory categories as well as a lack of clear differentiation between categories; internal problems do not give rise to recognizable, standardized, symptoms. Behavior must vary in the real world as a person repeats attempts to solve the same problem, according to PCT. In the same way that constellations in the night sky are arbitrary groupings of stars reflecting no underlying order or structure, categories of behavior -- control system outputs -- will not reveal the order or structure of internal malfunctions.

There is a developing acknowledgement that it is the distress associated with particular symptoms rather than the symptoms themselves that needs to be understood. Auditory hallucinations, for example, are experienced commonly in the general population. Who has never had cause to complain of not being able to stop a tune from replaying itself in imagination? Sometimes, auditory hallucinations are associated with debilitating distress and, at other times, they are a benign, perhaps even helpful, experience. A singer who can't mentally "hear" a note before singing it will not sing very well.

PCT contributes a useful perspective in understanding psychological disorders by first providing a model of satisfactory psychological functioning. Whereas current models of psychological dysfunction have been constructed by investigating one or more dysfunctional manifestations, PCT understands dysfunction by considering the way in which the process of control can be disrupted. As was previously discussed, conflict between control systems is problematic because it effectively removes the control abilities of both systems. While conflict of this nature is often transitory, on occasion it can become chronic. When this occurs the symptoms recognized as psychological disorder will become apparent and distress will result from an inability of an individual to control important experiences. Other problems can also arise (such as being overwhelmed by environmental forces) but a discussion of these other problems is beyond the scope of this paper.

Conflict, as it is conceptualized in PCT, occurs between two control systems at the same level. 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 biological markers 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 are, the more intense the conflict will be. Some kinds of mental illness, perhaps most, may be not a result 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 control systems are conflicted, it is the signals being sent from the next highest level that need to be altered.

The learning process of reorganization that was explained previously is, according to PCT, the change mechanism responsible for resolving conflict by modifying components of control systems that set reference signals for others. The therapeutic approach that is based on the principles of PCT is called the Method of Levels (MOL). Its remit is to help people redirect reorganization from the symptoms and the immediate efforts on both sides of the conflict to the control systems responsible for generating the conflict. It is hypothesized that reorganization and awareness are linked in such a way that it is the systems that are in awareness that will be the focus of reorganizing processes.

People tend to be most aware of painful or dramatic consequences of conflict: attention is drawn to the symptoms of loss of control 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. Consider wanting to stop smoking to avoid lung disease and at the same time wanting to continue it to relieve withdrawal symptoms caused by smoking; or wanting to leave a partner to avoid abuse and at the same time, 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, will not permanently alter the conflict. The conflict will be resolved 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. 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 an unusual blend of questioning about subjective experiences and selectively drawing a client's attention to seemingly tangential or peripheral comments the client might make -- comments that the therapist familiar with PCT recognizes as possibly indicating a higher-level system at work. In this way, the clients are showing 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 deficiency that causes 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. Details about MOL and its use in clinical practice are available elsewhere.

The idea of a common underlying process (conflict) as well as a common change mechanism (reorganization) might be particular significant for understanding the current situation in which psychotherapies based on quite different models of disorder can have similar effects. There has been an increasing call to move away from developing new techniques and strategies and instead to focus on underlying common principles and mechanisms. It is telling that despite the demonstrated effectiveness of psychotherapy there is still no generally accepted account of how these effects are achieved. The paradigm of perceptual control may provide the means to make sense of these otherwise puzzling results.

Recent research in neurobiology has indicated that psychotherapy can have effects in the brain that are similar to the effects that pharmacology achieves. Again, this result would come as no surprise from a PCT perspective. The hierarchy of PCT is a hypothesized neuronal architecture which is equally applicable to thoughts being explored or selective serotonin reuptake inhibitors being ingested.

It is certainly the case that, at this stage, PCT perhaps raises more questions for research than it answers in this field. Do conflicts at different levels of the hierarchy, for example, result in different types of distress? Does the rate of reorganization affect the experience of 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. 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.

In fact, full-time MOL practitioners, most of whom came from other schools of thought, agree that MOL is probably an explanation of why other therapies succeed when they are successful, and why they fail -- they fail to be consistent with MOL. Many therapists have independently developed methods that come close to MOL, simply by weeding out what doesn’t work. For some, such as Rogerians, a switch to pure MOL would involve only minor changes. For others, of course, such a switch would call for so many deviations from customary practice that it would be essentially impossible.

Exploring psychological disorders and their treatment from the perspective of perceptual control provides a new direction for psychotherapy researchers and practitioners. There is a growing possibility that it will enable a clearer understanding of the nature of psychological distress that is developed from a model of function rather than dysfunction. It may also promote the distillation of the important components of psychotherapy such that therapists can be clearer about their roles and treatments can become more efficient. Moreover it can, and already does, provide a guide regarding the purpose of psychotherapy. PCT, then, will have an impact on long standing debates such as the dodo bird hypothesis. 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 be easily 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.

Relation of PCT to conventional psychology

Mental Illness and Psychiatric Diagnostic Categories

Rewards and Punishments

Health and Disease

The future of PCT

PCT Research--how to get started