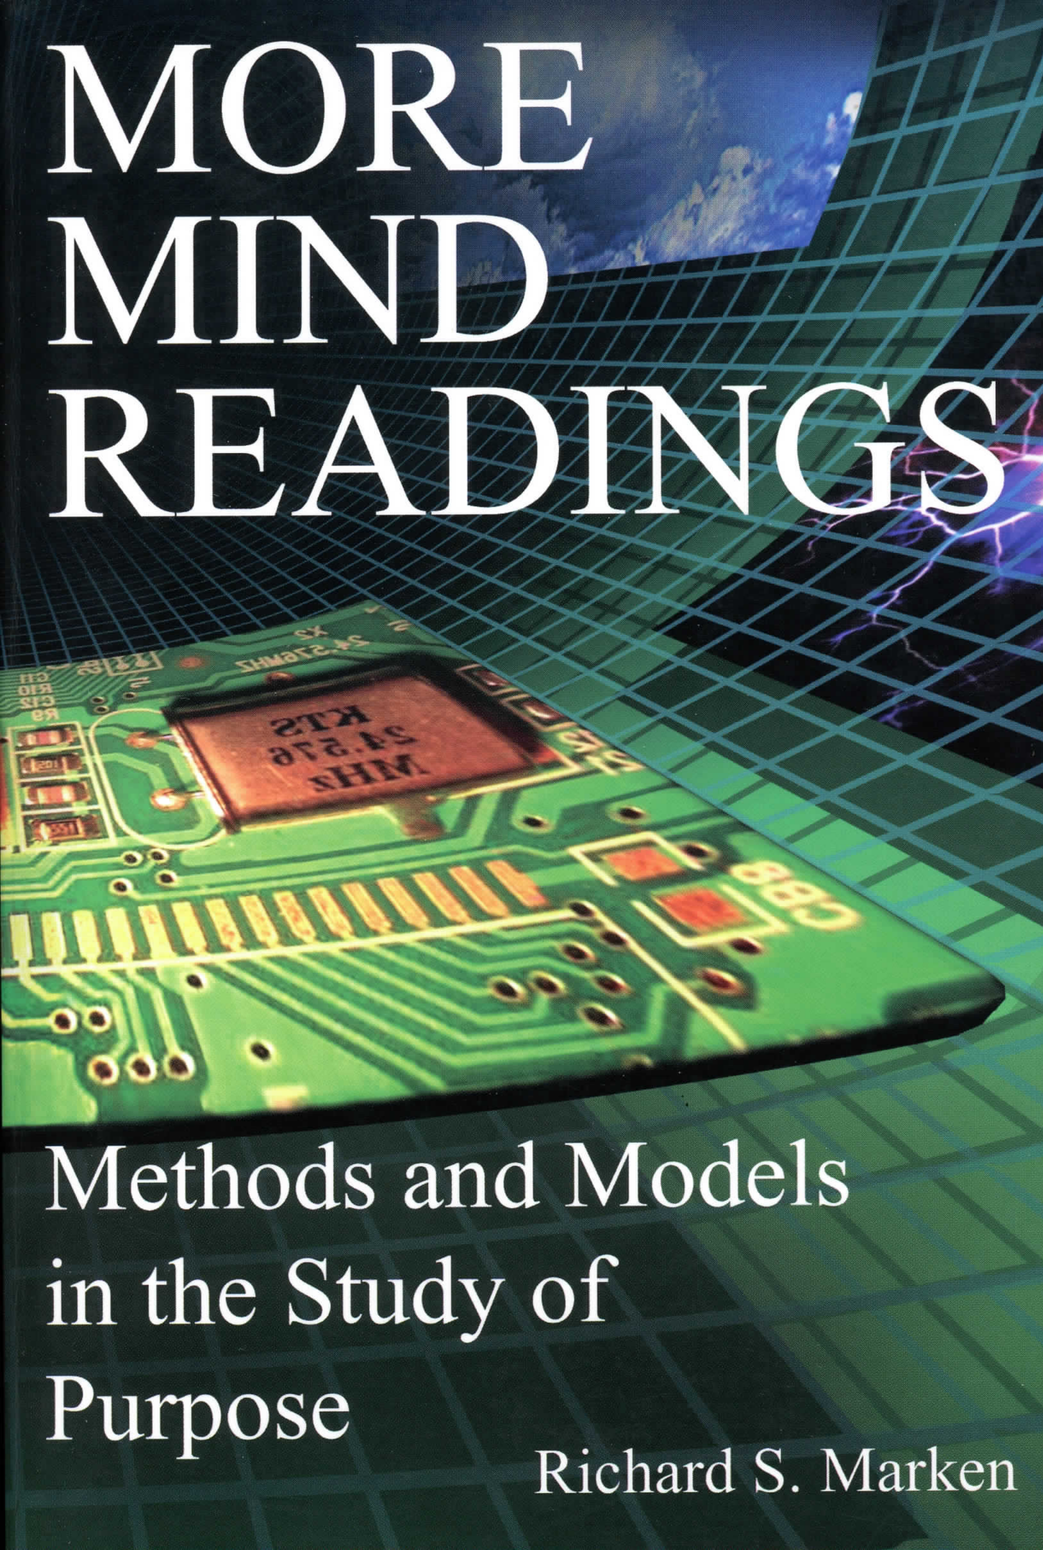

MORE MIND READINGS

Methods and Models in the Study of Purpose

Richard S. Marken

Dr. Marken is the author of two books, *Methods in Experimental Psychology* (1981) and *Mind Readings : Experimental Studies of Purpose* (1992), and over 40 papers on control system theory and psychology. He is currently a Senior Behavioral Scientist at the RAND Corporation where he does research on health care and national policy issues.

"You will find here the shape of a new psychology powerful in method and breathtaking in scope."

Philip J. Runkel
Casting Nets and Testing Specimens
Praeger Publishers, 1990

"... a delight to read. Marken writes lucidly and compellingly. He convincingly demonstrates the unprecedented power of negative feedback perceptual control models to describe purposeful behavior. Highly recomended."

Bruce Gregory
Harvard-Smithsonian
Center for Astrophysics

$18.00

ISBN 0-944337-43-0

9 780944 337431 51800

MORE MIND READINGS

MORE
MIND READINGS

Methods and Models in the Study of Purpose

Richard S. Marken

A MindReadings.com Book

A MindReadings.com Book

Library of Congress Catalog Number: NN-NNNNN

ISBN 0-9704701-7-7

Printed in the United States of America

Contents

1. Purpose in Perspective

2. Purpose in Research Methodology

3. Purpose in Psychology

4. Purpose in Biology and Economics

5. Purpose in Systems Engineering

Foreword

You have in your hands one of the foundation documents in the construction of a new scientific psychology. It is a firm addition to the steadily growing experimentation concerning Perceptual Control Theory. This book, like its predecessor, *Mind Readings: Experimental Studies of Purpose* (1992), contains both reports of experimentation and some theoretical comments.

Most books reporting psychological research tell you the conditions or circumstances under which you are likely to find a greater-than-likely frequency of this or that sort of behavior. The reports typically take this form: Twenty-three percent of the people of *this* sort (or in this condition) did what we predicted, whereas only eighteen percent of the people of *that* sort (or in that condition) did so.

Marken's experimentation departs radically from that sort of investigation. Marken does not predict *acts*. He predicts only that the person will do anything necessary to maintain a perception the person wants to maintain, *varying* action as necessary to maintain the perception. He predicts, too, that the person will do that continuously.

When I say that Marken predicts “only” the continuous control of perceptual variables, I do not mean that such a prediction is trivial or easy to investigate. To carry out this kind of investigation, Marken observes each subject continuously over a long enough period so that he can obtain *thousands* of data-points. In this feature alone, Marken's experimentation stands in startling contrast to typical psychological experiments – in which it is common, even usual, for each subject to yield only one or two data-points.

Marken predicts that the person will maintain a sensed variable (a condition or quantity) despite influences from the environment that would alter the variable were the person not there holding the variable steady. (And the person does so, in every case.) That is, the experiment does not “hold constant” all “unwanted” environmental variables. Marken does not even try to guess at what those unwanted variables might be like. He allows the environment to affect the variable the person wants to hold steady so that the behavior of the person in holding the variable steady can be observed.

The investigator's questions are “How can the person do this” and “How can we ascertain whether the person is indeed doing this?” The first question has been investigated in some piecemeal, fragmented ways by neurologists, but no one had investigated it in the behavior of the whole organism until the advent of Perceptual Control Theory. The second question has been rejected by many psychologists as impossible to test. Marken, using the radical assumptions of Perceptual Control Theory, shows both how the intentions of the person (of one person at a time) can be ascertained and how the person can be capable of maintaining the intended perception.

You will find here the shape of a new psychology powerful in method and breathtaking in scope.

Philip J. Runkel
Eugene, Oregon
January, 2001

Introduction

More Mind Readings is my second collection of papers describing research based on the control theory model of purposeful behavior developed by William T. Powers (1973). Since the publication of my first collection of papers (*Mind Readings*, 1992) Powers' theory has come to be known as perceptual control theory (PCT) to distinguish it from other applications of control theory in psychology. The name "perceptual control theory" describes what distinguishes Powers' use of control theory from all others: the idea that purposeful behavior is the *control of perception*.

The papers in the present collection describe what control of perception means, how it works and what can be learned from a theory that views organisms as controllers of their own perceptual experience. But it is more than a change in the name of Powers' theory that motivates this new collection of papers. In the introduction to *Mind Readings* I noted that the publication of that collection would mark "...the end of an era in which my research focused largely on what is wrong with current theories of behavior and the beginning of an era in which my research will focus almost exclusively on what is right with control theory." Nearly ten years after making that resolution I am finally able to put together a new collection of papers that describe research which shows what is right with the theory that is now dubbed PCT.

The papers collected in this book can be profitably read as a self-contained group; it is not necessary to have read *Mind Readings* before jumping into the present collection (though I believe that the papers in *Mind Readings* make a useful supplement to the papers in the present collection, and vice versa). The intended audience for both books is behavioral science researchers as inquisitive laymen who want to see what can be accomplished by looking at behavior as the control of perception. The present collection might actually be the best place to start since many of the papers in this volume show how PCT can be used to answer some familiar research questions in experimental psychology. For example, the paper entitled *Controlled Variables: Psychology as the Center Fielder Views It* shows how PCT can be used to answer the question “how do baseball outfielders catch fly balls?”

Purpose in Perspective

More Mind Readings begins with two papers that discuss the role of PCT in experimental psychology in particular and the behavioral sciences in general. The papers show that PCT is not just a theory of behavior; it is also a way to *look at* behavior. Experimental psychologists typically look at behavior as a cause-effect phenomenon. Behavior is seen as the end result of a causal process that begins in the environment or the brain of the organism. Psychologists have viewed behavior as a cause-effect process because it *looks like* a cause-effect process.

Placing food in a dog’s mouth seems to cause salivation; giving food to a rat after it presses a bar seems to cause the rat to press the bar again. But PCT shows that the cause-effect appearance of behavior is an illusion that results when the observer fails to notice the purpose of behavior.

In PCT, behavior is viewed as a purposeful phenomenon; the purpose of behavior is to control. An organism controls by bringing perceptual variables to reference states and maintaining them in those states, protected from disturbance. The perceptions that an organism controls are not always obvious to an outside observer. Indeed, an observer who is unfamiliar with PCT is likely to see behavior as caused when it is actually aimed at achieving purposes (controlling perceptions). For example, the salivation that appears to be caused by food is actually part of the process of controlling a perception of the “swallowability” of the food; the food that seems to cause the next bar press is actually part of the process of controlling a perception of the rate of food intake.

PCT occupies an unusual position relative to other theories of behavior. Its value as a theory of behavior can be seen only after one has learned to view behavior from the PCT perspective. The first two papers in *More Mind Readings* therefore aim at presenting the PCT view of purposeful behavior as the control of perception.

Purpose in Research Methodology

The methods used to study behavior from a PCT perspective differ from the familiar methods of experimental psychology. The methods of experimental psychology are designed to determine the causes of behavior; the methods of PCT are designed to determine the purpose of behavior. The basic method of PCT is the test for the controlled variable (or, simply “the test”). This method has been called “the nearest approach I know of to mind reading” (Powers, 1979). Indeed, it was this description of the test that provided the title for the first paper in this section as well as for this book and its predecessor. The “Mind Reading” paper shows how the test can be used to monitor a person’s intentions (or purposes) regarding the location of a line on a computer screen. Intentions regarding line position may not be as interesting as intentions regarding love or stock positions but the “Mind Reading” paper illustrates how the test can be used *in principle* to determine any intention (or purpose).

The second paper in this section describes how the PCT researcher approaches behavior from the point of view of the behaving system itself. The goal of PCT research is to understand what the “dance” of behavior looks like from the point of view of the behaving system itself; the “dancer”. The “dancer’s” perspective on behavior is described in the context of various examples of real, purposeful behavior. The “Dancer” paper shows how the test can be used to determine which of the dancer’s perceptions, when they are controlled, give rise to the dance of behavior that is perceived by an observer.

Purpose in Psychology

The two papers in this section describe psychological research studies based on the PCT model of behavior. The first paper describes studies of the human ability to produce complex behavioral sequences. The studies suggest that limitations on the speed with which people can produce various sequences of actions, such as sequences of trilled notes, spoken phonemes or typed letters, are the result of limitations on the ability to perceive such sequences, not on the ability to produce them. The relative time limits on the human ability to produce different types of sequences correspond to the relative position of these sequences in a perceptual hierarchy that is a part of the PCT model of behavior.

The second paper describes a detailed PCT model of a fielder running to catch fly balls. The model assumes that fielders control some visual variable and that the fielder runs with the purpose of keeping this perception in some desired state. The paper describes three possible perceptions that a fielder might be controlling; optical velocity, optical acceleration and linear optical trajectory. The model fits real fielder behavior best when it is assumed that the fielder is controlling optical velocity. This paper shows how PCT can explain familiar examples of purposeful behavior as the control of particular perceptions, such as the perception of optical velocity.

Purpose in Ethology and Economics

I have included the two papers in this section to give a glimpse of how PCT might be applied in fields outside the usual confines of experimental psychology. The first paper is a commentary on a theoretical paper that deals with species specific behavior, the concern of the field of ethology. The commentary points to the importance of taking the purpose of behavior into account when developing models of species specific behavior. The second paper shows how PCT can be used to explain economic behavior at the population level: macroeconomics. The PCT model described in this paper models the economy as a collection of input controllers, organized to produce goods and services for their own consumption. The model seems to hold promise for explaining some surprising macroeconomic phenomena, such as the effect of income distribution on economic growth rate and the effect of Federal Reserve discount rates on inflation.

Purpose in Systems Engineering

The final paper in this collection shows how PCT can be applied in the field of engineering. The paper describes the use of PCT as the basis for performing a task analysis, which is a basic component of the systems engineering process. The paper describes a PCT-based task analysis process called PERCOLATe (perceptual control analysis of tasks) which describes tasks in terms of controlled perceptions rather than responses to input. This approach to task analysis happens to be well suited to the design of the human-computer interface (HCI) to complex systems. The results of a PERCOLATe analysis tell the HCI designer what data should be displayed to the operator and what the computer must allow the operator to do to get the displayed data into the desired state.

I hope that the papers in the present collection give the reader an idea of the power and scope of the revolutionary new theory of behavior known as perceptual control theory (PCT). It takes some time to get used to looking at behavior through PCT glasses; from the unfamiliar perspective of the behaving system itself rather than from the more familiar perspective of the observer of that system. But once the PCT perspective is learned, it provides a basis for a surprising and satisfying new way of looking at the everyday mystery that is the behavior of living organisms.

1. Purpose in Perspective

A Science of Purpose

William James got scientific psychology off to an excellent start by describing, as well as anyone before or since, the nature of purposive behavior:

> Romeo wants Juliet as filings want a magnet; and if no obstacles intervene he moves toward her by as straight a line as they. But Romeo and Juliet, if a wall be built between them, do not remain idiotically pressing their faces against its opposite sides like the magnet and the filings with the [obstructing] card. Romeo soon finds a circuitous way, by scaling the wall or otherwise, of touching Juliet's lips directly. With the filings the path is fixed; whether it reaches the end depends on accidents. With the lover it is the end which is fixed, the path may be modified indefinitely. (James, 1890, p. 7)

A paragraph like this brings a tear to a control theorist's eye, not just for the beauty of the prose, but for the insight into the nature of behavior. James understood the difference between purposive and non-purposive behavior and tried, unsuccessfully, to launch psychology as the science of purposive behavior. He failed because, like other visionaries, his ideas were slightly ahead of their time.

James was trying to start a science of purpose at a time when "purpose" was a scientific profanity (it still is, to some extent, but we live in more permissive times). The tools that could make purpose scientifically respectable were not to appear for another 40 years. In the meantime,

---

From *American Behavioral Scientist*, 1990, 34 (1), 1-15.

psychologists abandoned any serious efforts to understand the nature of purpose and turned, instead, to the development of theories of behavior based on a model borrowed from the physical sciences, a model designed to explain the behavior of non-purposive systems. The model assumes that natural phenomena result from the operation of cause-effect laws. Events such as behaviors result from causes, not purposes.

The cause-effect model seems appropriate as long as behavior is seen in a way that ignores its purpose. This is done by a process that can be called the "objectification" of behavior. Objective behavior is "output" that is emitted by organisms (like heat from a stove or light from a filament). From the objective point of view, behavior is something organisms do, little more than "a show put on for the benefit of an observer" (Powers, 1978). If the observer happens to be a behavioral scientist, he or she can determine the cause of behavior in order to "predict and control" it, just as scientists in other fields predict and control other natural phenomena. But it is clear that objective behavior serves the purposes of the observer, not those of the behaving organism.

The objectification of behavior has not made purpose go away. It is difficult to avoid noticing, for example, that the results of behavior are often "good" for the organism. They seem to serve some purpose, such as providing nutrition or avoiding destruction. Thus a rat pressing a bar is producing an output (the bar press) but it is also feeding itself. The "objective" solution has been to admit that purpose exists, but to place it in the environment rather than in the organism (Skinner, 1981). Thus an apparently purposive behavior like feeding occurs because the environment selects the appropriate outputs (bar presses), not because the organism intends to eat. Purpose has come to be seen as a passive result of environmental selection rather than an active attempt to produce desired results.

James was not describing purpose as a passive process. Romeo is not pulled to Juliet like filings to a magnet. Rather, he intends to place his lips on hers and he actively adjusts his actions in order to produce this end. What makes active adjustment necessary is the unpredictable nature of the environment. James saw that the environment is as likely to be a hindrance as a help in achieving desired results. Thus Romeo must adjust to the environment (a wall, a balcony, or a Capulet) in order to produce the desired end. In purposive behavior, as James pointed out, it is the end that is fixed while the path is modified indefinitely. These modifications are made necessary by unpredictable changes in the environment. Romeo carries out his purpose by constantly adjusting his path to achieve his end. Unfortunately, James could not explain how Romeo could possibly behave in this way, but now we can. Control theory is a model of systems that act just like Romeo, that is, systems that produce fixed results in an inconsistent and often unhelpful environment.

Control theory, developed in the 1930s (Black, 1934), should have made it possible for psychology to reclaim James's title as the science of purpose. Indeed, there have been a number of attempts to apply control theory to the behavior of organisms. These efforts began during World War II with the pioneering work of Craik (1943) and continue to this day, notably in the study of “manual tracking” (see the review by Wickins, 1987). But there is still no psychology of purpose. I believe the difficulty can be traced to the objectification of behavior. Control theory has been applied to the wrong phenomenon (Marken, 1988).

Because psychologists "know" that behavior is output, control theory has been used to explain the causes of behavioral output. But control theory is designed to explain "control," a phenomenon more like the purposive behavior described so eloquently by James than the objectified output dear to contemporary behavioral science. Control is the process of producing consistent results in an inconsistent environment.

Figure 1

We can be a bit more precise (and less poetic) than James by describing control with a diagram of the variables involved in behavior (see Figure 1). The symbols R, B, and E represent variable aspects of the world. R is a response variable, something done by an actor (like tensing a muscle) to produce the results that we call behavior, B. "Response" and "behavior" are relative terms. A "response" at one level of observation can be a "behavior" at another. For example, muscle tensions are responses that cause a behavior called "leg movements"; similarly, leg movements are responses that cause a behavior called "walking." E is an environmental variable that also affects behavior. In the case of walking, E might be the changing slope of the terrain. A behavior such as walking is a variable consequence of the effects of responses, such as leg movements, and environmental variations, such as the changing slope.

Ordinarily, as R and E vary over time, so does B. This would happen if, for example, B were the distance between two pebbles, R were the wind, and E were an earthquake (of variable magnitude). The distance between the pebbles (B) will vary as the velocity of the wind (R) and magnitude of the earthquake (E) vary. If, however, B stays nearly the same while R and E vary, there is control. This is what happens with Romeo and Juliet. The distance between the two is a behavior (B) that stays about the same (zero distance) despite variations in the environment (E, consisting of walls and Capulets) and responses (R, such as Romeo jumping, running, and dodging).

Control occurs when variations in E and R are precisely equal and opposite over time. The result of this opposition is a stable behavior, B. Control theory explains how to design an agent that will produce variations in R that stabilize B. Such an agent is called a control system. The basic components of a simple control system are shown in Figure 2.

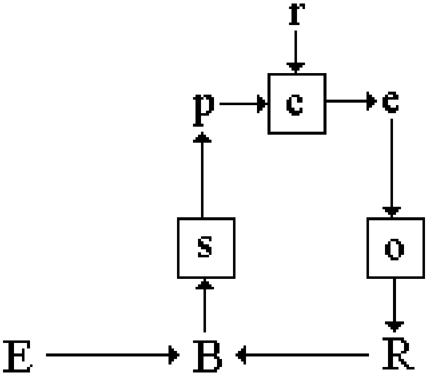

Figure 2

The control system consists of three components (s, c, and o) and three signals (p, r, and e). Signals and components are all in lower case to highlight the fact that they are part of the system. The three variables in upper case are outside the system in the environment, which, in a living organism, includes everything outside the nervous system.

The three components of the control system are the sensor (s), comparator (c), and output amplifier (o). The sensor converts the external behavioral variable into a perceptual signal. The comparator turns the difference between the reference (r) and perceptual signal (p) into an error signal (e). The output amplifier turns the error signal into a response variable, R, a process that could involve a huge amount of amplification, as when a weak neural signal is turned into a strong muscle tension. We have already seen that the response has an effect on the behavioral variable, B. There is also an effect of the behavioral variable on the response via the control system. This two-way connection between response and behavior is called a feedback loop. If the effect of the response on behavior is such that it tends to reduce the error signal, the feedback is negative.

Note that the control system senses only variable B. It does not sense E or R, the two environmental effects on behavior. It does not need to. If the system is set up for negative feedback, it will stabilize B by continuously varying R to counteract any effect of E on B. B will be stabilized at a value that corresponds to the value of the reference signal, r. If the system were to change the value of r, B would change as well. Thus, if B is the distance between

Romeo and Juliet, r (the reference signal in Romeo) determines what this distance will be. In the play, it seems to be set at a value that produces a B pretty close to zero. If the pair had survived and Romeo proved as fickle as he seemed before meeting Juliet, the r in Romeo might eventually have changed to make B considerably greater than zero (and stabilized at that new value against disturbances, such as the entreaties of well-meaning friends).

The control system model produces purposive behavior by controlling perception, not responses (Powers, 1973). A negative feedback control system keeps its perceptual signal, p, matching its reference signal, r. The responses of the system depend almost completely on environmental effects on behavior. If we ignore the behavioral variable (the purpose for producing responses), then it appears that variations in responses are caused by variations in environmental events. Indeed, much of modern psychology (and behavioral science in general) can be seen as an attempt to discover the laws relating environmental events, E (the independent variable), to response variables, R (the dependent variable). The assumption is that these laws (also called intervening variables, cognitive processes, or hypothetical constructs) will reveal something about the nature of the organism which is presumed to mediate the relationship between E and R. In fact, if the organism is a control system, the law that relates E and R depends almost completely on the nature of the environmental links between E, R, and B and tells little about the behaving organism (Powers, 1978).

An observer who is watching the behavior of a living control system can see only the variables outside of the system. The observer cannot see what the system perceives.

This suggests that it should often be difficult to see the purpose of behavior and, indeed, we sometimes find ourselves watching behavior and wondering what is happening. When we ask, “what is that person doing?” we mean, “what is that person's purpose?” We have been discussing behavior as though we always knew the actor's purpose. But we often do not, or we get it wrong. Is Romeo's purpose really to keep the distance between himself and Juliet at zero? If so, he is doing a rather poor job. A more accurate description of Romeo's purpose might be “perceiving myself close to Juliet 90% of the time.” This purpose is more complex and it is defined over time. Romeo could be consistently achieving this purpose (keeping the perceptual signal matching the reference signal) even while he is off achieving other purposes (like avenging his friend's honor and whatnot).

Control theory suggests a rigorous method for testing hypotheses about the nature of the perceptions being controlled by other control systems (that is, determining their purposes). The method begins with a clear description of a hypothesized controlled perception in terms of variables in the system's environment. Thus a hypothesized controlled perception, p, is described as a function of observable behavioral variables, B. The next step is to determine what the behavior of the hypothesized variable would be in response to various environmental factors (like E and R) if it were not controlled. The next step is to produce variations in these environmental factors and see if the hypothesized variable behaves as expected. If it does not (if, for example, it varies far less than expected), then the variable is under control.

Simple control systems have simple purposes. The basic control system model shown in Figure 2 could not produce the purposive behavior of a sea slug, let alone a human being. A single control system cannot be an adequate model of most living control systems, which achieve many purposes simultaneously. Some purposes are achieved as part of the process of achieving other purposes. A driver has the purpose of keeping the car in its lane. This purpose (along with others) is carried out to achieve the purpose of getting to work which, in turn, is carried out to achieve the purpose of making a living. Powers (1973) described a hierarchical arrangement of basic control systems that can achieve a hierarchy of purposes. In the hierarchy, higher level systems achieve their purposes by adjusting the reference signals (and hence the purposes) of lower-level systems; the error signals of the higher level systems become the reference signals of the lower-level systems. A single lower-level system can be part of the means used to achieve the purposes of several higher-level systems. Also, several lower-level systems can be used as the means for achieving the purpose of a single higher-level system.

The behavior of a hierarchy of control systems is rather amazing. When set up properly, all systems in the hierarchy are able to achieve their purposes, virtually simultaneously. There is no necessary conflict between systems, although it is always a dangerous possibility. Conflict occurs when one control system can achieve its purpose (that is, get its perception to match its reference signal) only by acting in a way that causes the perception of another system to move away from its reference. Conflict is the worst thing that can happen to a hierarchy of control systems because it prevents the systems from doing what they were designed to do: control. In humans, internal conflict (resulting from having purposes that work against each other) is recognized as a major cause of dysfunctional behavior. Control theory can help us understand the nature of conflict and possibly tell us how to avoid it or, at least, make it less frequent.

In a control hierarchy, higher-level systems control perceptions that are constructed from the perceptions of lower-order systems. The perception of a symphony is constructed from perceptions of musical phrases that are themselves constructed from perceptions of individual notes. Powers theorized that each level of perception represents a different class of perceptual variable. The perceptions in a class can be quite abstract, such as “relationships,” “principles,” and “system concepts.” Powers suggested that there may be just a small number of different classes of perceptual variable (ten at last count; Powers, 1979e) but this is enough to explain why we see people carrying out very complex purposes, like “being honest” or “being scientific.” These purposes are caused by reference signals that specify a particular value of a perceived principle (honesty) or system concept (science).

When it is said that something is done “on purpose,” it is implied that it is done consciously. However, control theory shows us that purpose and consciousness are two distinct phenomena. The hierarchical control model carries out rather complex purposes with no consciousness whatsoever. In fact, most of our daily purposes, simple and complex, are carried out quite unconsciously (when was the last time you were conscious of maintaining your balance or your sense of self?). Consciousness is part of the control model because it is a part of people but it is separate from the control hierarchy described so far. In the model, consciousness monitors the status of existing control systems. and, if necessary, changes them (a process called reorganization). Although it is a significant part of the model (and ourselves), consciousness is not likely to yield to coherent investigation until the nature of the purposive behavior generated by the control hierarchy is understood in some detail.

The Blind Men and the Elephant: Three Perspectives on the Phenomenon of Control

Behavior has been described as a response to stimulation, an output controlled by reinforcement contingencies and an observable result of cognitive processes. It seems like these are descriptions of three different phenomena but they are actually descriptions of three different aspects of the same phenomenon – control. Control is like the proverbial elephant studied by the three blind men; what one concludes about it, and how one tries to explain it, depends on where one stands. It is suggested that the best place to stand is where one has a view of the whole phenomenon, be it elephant or control.

The behavior of living organisms (and some artifacts) is characterized by the production of consistent results in an unpredictably changing environment, a phenomenon known as control (Marken, 1988). Control can be as simple as maintaining one's balance on uneven terrain or as complex as maintaining one's self-esteem in a dysfunctional family. Control is a pervasive aspect of all behavior yet it has gone virtually unnoticed in psychology. What has been noticed is that behavior appears to be a response to stimulation, an output controlled by reinforcement contingencies or an

---

From *Closed Loop*, 1993, 3, 37-46.

observable result of cognitive processes. Each of these appearances is what would be expected if people were looking at control from different perspectives. The situation is similar to that of the three blind men who were asked to describe an elephant; the one near the tail described it as a rope, the one near the leg described it as a tree trunk and the one near the trunk described it as a snake. Each description gives an accurate picture of some aspects of the elephant, but a false picture of the elephant as a whole. If behavior involves control then psychology, too, has given an accurate picture of some aspects of behavior but a false picture of behavior as a whole.

Closed-Loop Control

The basic requirement for control is that an organism exist in a negative feedback situation with respect to its environment. A negative feedback situation exists when an organism's response to sensory input reduces the tendency of that input to elicit further responding. Negative feedback implies a closed-loop relationship between organism and environment; sensory input, s, causes responding, r, that influences the sensory cause of that responding, s, as shown in Figure 1.

Figure 1 shows the closed-loop feedback relationship that exists between an organism, represented by the rectangle, and its environment, represented by the arrows outside of the rectangle. A sensory variable, s, influences responding, r, via the organism function, k.o. Responding influences the sensory variable via the feedback function, k.f. The sensory variable is also influenced by an environmental variable, d, via the environmental function, k.e. There is feedback in this closed-loop relationship because the effect of responding is “fed back”, via the environmental feedback function, k.f, as an influence on the sensory cause of that responding.

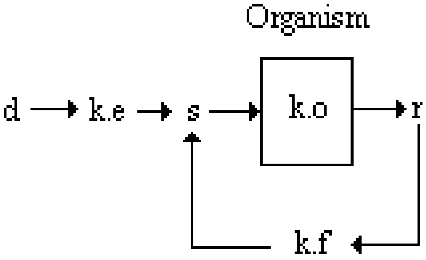

Figure 1

It is hard to imagine an organism that does not exist in such a closed-loop situation because all organisms are built in such a way that what they do affects what they sense. Eyes, for example, are located on a head that moves so that what the eyes see depends on what the head does. To the extent that what the head does depends on what the eyes see (such as when the head turns in response to an attractive passer-by) there is a closed loop; sensory input causes responding (head movement) which affects the cause of responding (sensory input). The feedback in this loop must be negative because behavior is typically stable (organisms, for example, do not normally exhibit the “run away” behavior that characterizes positive feedback loops, such as the “feedback” from a microphone that amplifies its own output).

The fact that organisms exist in a closed negative feedback loop means that two simultaneous equations are needed to describe their relationship to the environment. These are given as equations (1) and (2), below. For simplicity we will assume that all functions are linear and that all variables are measured in the same units. Equation (1) describes the effect of sensory input on responding so that:

$$(1) \qquad r = k.o\,(s^*-s)$$

This equation says that responding, r, is a linear function of sensory input, s. The sensory input is expressed as a deviation from the value of input, s^*, that produces no responding; s^* defines the zero point of the sensory input. Equation (2) describes the effect of responding on sensory input. For simplicity it is assumed that responding, r, adds to the effect of the environment, d, so that:

$$(2) \qquad s = k.f\,(r) + k.e\,(d)$$

The variables r and d have independent (additive) effects on the sensory input, s. The nature of the environmental effect on sensory input is determined by the environmental function, $k.e$. The feedback effect of responding on the sensory cause of that responding is determined by the feedback function, $k.f$.

Equations (1) and (2) must be solved as a simultaneous pair in order to determine the relationship between stimulus and response variables in the closed loop (the derivation is shown in the Appendix). The result is:

$$(3) \qquad r = l/((l/k.o) + k.f)\, s^* - k.e/((l/k.o) + k.f)\, d$$

Equation (3) can be simplified by noting that the organism function, k.o, transforms a small amount of sensory energy into a huge amount of response energy (such as when a pattern of light on the retina is transformed into the forces that move the head). In control engineering, k.o is called the "system amplification factor" or "gain" and it can be quite a large number. With sufficient amplification (such that k.o approaches infinity) the (l/k.o) terms in equation (3) approach zero, so equation (3) reduces to:

(4) $r = s^*/k.f - (k.e/k.f)\ d$

Equation (4) is an input-output equation that describes the relationship between environmental (stimulus) and response variables when an organism is in a closed-loop, negative feedback situation with respect to its environment. The result of being in such a situation is that the organism acts to keep its sensory input equal to s*, which is called the reference value of the input. Equation (4) shows that the organism does this by varying responses, r, to compensate for variations in the environment, d, which would tend to move sensory input away from the reference value; this process is called control.

Three Views of Control

All variables in equation (4), with the possible exception of s*, are readily observable when an organism is engaged in the process of control. The environmental variable, d, is seen as a stimulus, such as a light or sound. The response variable, r, is any measurable result of an organism's actions, such as bar pressing or speaking. The reference value for sensory input, s*, is difficult to detect because an observer cannot see what an organism is sensing. But s* is the central feature of control since everything an organism does is aimed at keeping its sensory inputs at reference values. Because these reference values are difficult to detect it will not be obvious to an observer that an organism is engaged in the process of control. What will be obvious is that certain variables, particularly the environmental and response variables and the relationship between them, will behave as described by equation (4). Thus, equation (4) can be used to show what control might look like if one did not know that it was occurring. It turns out that there are three clearly different ways of looking at control depending on which aspect of the behavior described by equation (4) one attends to.

1. *The stimulus - response view.* This view of control sees behavior as a direct or indirect result of input stimulation. An example of stimulus-response behavior is the so-called “pupillary reflex” where changes in a stimulus variable (illumination level) lead to changes in a response variable (pupil size). The stimulus-response view is the basis of several current approaches to understanding behavior, such as the “synergistic” or “coordinative structure” theory of motor coordination. Warren, Young and Lee (1986), for example, describe a synergistic model of running in which “vertical impulse is directly modulated by the optical variable Δt.” (p.264). The behavior of running is seen in stimulus-response terms; a stimulus variable, Δt, determines (“modulates”) the value of a response variable, vertical impulse. The stimulus-response view is also the basis of a re-

cent theory of attention (Cohen, Dunbar and McClelland, 1991) in which connections between printed word stimuli and verbal responses in the Stroop effect are modulated by connections in a neural network.

Equation (4) shows that behavior will look like a stimulus-response process when the reference value for sensory input, s*, is a constant; for simplicity assume that it is zero. Then responding is related to environmental stimuli as follows:

(5) $r = - (k.e/k.f)\ d$

Equation (5) shows that, when there is a fixed reference level for sensory input, it will look to an observer of behavior as though variations in an environmental stimulus, d, cause variations in a response, r. This is what we see in the pupillary reflex where pupil size, r, is proportional to illumination level, d. Of course, this relationship between pupil size and illumination level is precisely what is required to keep a sensory variable (sensed illumination) at a fixed reference value (s* = constant).

One's inclination when looking at an apparent relationship between stimulus and response is to assume that the nature of that relationship depends on characteristics of the organism. Equation (5) shows, however, that when an organism is engaged in control, this relationship depends only on characteristics of the environment (the functions k.e and k.f); the organism function, k.o, that relates sensory input to response output, is rendered completely invisible by the negative feedback loop. This characteristic of the process of control has been called the “behavioral illusion” (Powers, 1978).

2. *The reinforcement view*. This view of control sees behavior as an output that is shaped by contingencies of reinforcement. A reinforcement contingency is a rule that relates outputs (like bar presses) to inputs (reinforcements); in equation (4) this contingency is represented by the feedback function, k.f, that relates responses to sensory inputs. The reinforcement view is the basis of at least one influential theory of generalization and discrimination (Shepard, 1987). In a connectionist implementation of the theory, a reinforcement contingency is used to shape the formation of generalization gradients (Shepard, 1990). The reinforcement view is also the basis of modern theories of operant behavior. According to Domjan (1987) the contemporary perspective on operant behavior focuses on how contingencies “restrict freedom of action and ... create redistributions of various types of activities” (p. 562). In other words, contingencies shape (redistribute) responses (activities).

Equation (4) shows that it will look like contingencies (the feedback function) control responses when s*, d and k.e are constants, as they are in the typical operant conditioning experiment. In these experiments, s* is the organism's reference value for the sensory effects of the reinforcement. The environmental variable, d, is the reinforcement, which, if it is food, is typically a constant size and weight. The sensory effect of a reinforcement can be assumed to be directly proportional to its size and weight, making k.e = 1. So, for the operant conditioning experiment, equation (4) can be re-written as

(6) $r = S^*/k.f - D/k.f$

where S* is the constant reference value for sensed reinforcement and D is the constant value of the reinforcement itself.

The only variable in equation (6) is the feedback function, k.f, which defines the contingencies of reinforcement. One simple contingency is called the “ratio schedule” in which the organism receives reinforcement only after a certain number of responses. The ratio corresponds to the function k.f in equation (6). When the ratio is not too demanding it is found that increases in the ratio lead to increased responding. More demanding ratios produce the opposite result; increases in the ratio lead to decreased responding (Staddon, 1979). Either of these results can be produced by manipulating the relative values of S* and D in equation (6). The important point, however, is that the apparent dependence of responding on the feedback function, k.f, is predicted by equation (6). To an observer, it will look like behavior (responding) is controlled by contingencies of reinforcement. In fact, the relationship between behavior and reinforcement contingencies exists because the organism is controlling sensed reinforcement; responding varies appropriately to compensate for changes in the reinforcement contingency so that sensed reinforcement is kept at a constant reference value, S*.

3. *The cognitive view*. This view of control sees behavior as a reflection or result of mental plans or programs. This kind of behavior is seen when people produce complex responses (such as spoken sentences, clever chess moves or canny investment decisions) apparently spontaneously; there is

often no visible stimulus or reinforcement contingency that can be seen as the cause of this behavior. The cognitive view is the basis of numerous psychological theories that propose mental algorithms to explain the appearance of cognitive behavior. Examples of such theories include the ACT (Anderson, 1983) and SOAR (Newell, 1990) models of cognition, hierarchical models of the generation of movement sequences (Rosenbaum, Kerry and Derr, 1983), connectionist models of speech production (Jordan, 1989) and schema models of expertise in problem solving (Lesgold, A., Robinson, H., Feltovitch, P., Glaser, R., Klopfer, D. and Wang, Y., 1988).

Cognitive behavior is most obvious when environmental factors (such as stimulus variables and environmental and feedback functions) are held constant. When this is the case, equation (4) becomes

$$r = s^*/F + K \tag{7}$$

where F is the constant feedback function and K = (k.e/k.f) d, a constant. Since s* is typically invisible, equation (7) shows that there will appear to be no obvious environmental correlate of cognitive behavior. An observer is likely to conclude that variations in r are the result of mental processes – and, indeed, they are. But it is actually variations in s*, not r, that are caused by these processes; variations in r being the means used to get sensory inputs equal to s*. Thus, chess moves are made to keep some sensed aspect of the game at its reference value. When the environment is constant, r (the moves) may be a fair reflection of changes in the reference value for sensory input. However, under normal circumstances r is only indirectly related to s*, variations in r being mainly used to compensate for variations in the environment that would tend to move sensory input from the reference value, s*.

Looking at the Whole Elephant

The blind men never got a chance to see the whole elephant but if they had they would have instantly understood why it seemed like a rope to one, a tree trunk to another and a snake to the third. Psychologists, however, can take a look at control and see why the appearance of behavior differs depending on one's perspective. What is common to the three views of behavior discussed in this paper is the reference for the value of sensory input, s*. Organisms behave in order to keep sensory inputs at these reference values (Powers, 1989). They respond to stimulation in order to keep the sensory consequences of this stimulation from moving away from the reference value; so it appears that stimuli cause responses. They adjust to changes in reinforcement contingencies by responding as needed in order to keep the sensory consequences of reinforcement at the reference value; so it appears that contingencies control responding. And they change their responding in order to make sensory input track a changing reference value for that input; so it appears that responding is spontaneous.

What appear to be three very different ways of describing behavior can now be seen as legitimate ways of describing different aspects of one phenomenon – control. Each is just a different way of describing what an organism must do to keep its sensory inputs at their reference values. Indeed, once you know that the appearances called “behavior” are merely the visible consequences of an organism's efforts to control its sensory inputs, the problem of explaining behavior changes completely, from an attempt to build models that simulate the appearance of behavior (S-R, reinforcement or cognitive) to an attempt to build models that control the same sensory inputs as those controlled by real organisms. In order to build the latter type of model it is necessary to learn what sensory variables are actually being controlled by organisms. This type of investigation cannot be done by simply looking at the appearance of behavior. Methods based on control theory can be used to test which sensory variables an organism might be controlling at any time (Marken, 1992). These methods make it possible to take off the blindfolds and see the whole elephant – the phenomenon of control.

Appendix

Given the two system equations:

(1) $r = k.o (s*-s)$ and

(2) $s = k.f (r)+ k.e (d)$

we want to solve for r as a function of s. First, substitute equation (2) for s in equation (1) to get:

(A.l) $r = k.o (s*-(k.f (r)+ k.e (d)))$

Multiply through by k.o to get:

(A.2) $r = k.o (s*) - k.o\ k.f (r) - k.o\ k.e (d)$

Move all terms with r to the left side of the equation to get:

(A.3) $r + k.o\ k.f (r) = k.o (s*) - k.o\ k.e (d)$

Factor r out of the left side of the equation to get:

(A.4) $r (1 + k.o\ k.f) = k.o (s*) - k.o\ k.e (d)$

Divide both sides of the equation by (1 + k.o k.f) to get:

(A.5) $r = k.o/(l + k.o\ k.f)\ s* -k.o\ k.e/(l +k.o\ k.f)d$

Finally, divide k.o out of the numerators on the right side of (A.5) to get equation (3):

$$r = 1/((l/k.o)+ k.f)\ s^* - k.e/((l/k.o)+ k.f)\ d \tag{3}$$

2. Purpose in Research Methodology

“Mind Reading”: A Look at Changing Intentions

Methods adapted from control engineering can be used to discriminate intended from unintended consequences of an organism’s actions. By continuously monitoring a quantity called the stability factor it is possible to observe changes in intentions, which are not visible in overt behavior.

We often find ourselves wondering what people are doing although their behavior is completely visible. This is a problem for those who consider behavior an objective phenomenon. What is not always obvious about behavior is its purpose. To understand the purpose of behavior one must determine the intentions of the actor. This involves 'mind reading,' something current methods in experimental psychology are not equipped to handle. The solution has been to ignore intentions; behavior (what a person is doing) is defined as whatever the experimenter says it is.

Control theory provides an alternative approach. Intentions are viewed as internal models of the desired or reference states of perceptual inputs (Powers, 1973). The purpose of action is to keep perceptions matching reference states, a process called control. Organisms that control perceptions are control systems. To understand the behavior of a control system one must determine what percept-

---

From *Psychological Reports*, 1983, 53, 267-270. Reprinted with permission of Psychological Reports.

tions are being controlled. Although we cannot see what a system perceives we can measure physical variables in the system's environment (such as the room temperature near a thermostat) which may correspond to controlled perceptions. Methods adapted from control engineering can be used to determine which variables are controlled. These methods have been dubbed the test for controlled variables or simply *the test* (Powers, 1973; Marken, 1982).

The test is based on the fact that control systems act to resist disturbances to a controlled variable. For example, a thermostat acts to resist the disturbance to room temperature produced by a suddenly opened window. If we suspect that some variable is under control, we can test this suspicion by applying disturbances (called driving functions in control engineering) to the variable. If the disturbances produce less than the expected effect (expected on physical grounds) and this lack of effect can be traced to the actions of the organism, then the variable (actually, the perception of the variable) is under control.

One version of the test involves computation of a quantity called the stability factor (Powers, 1978), which is the ratio of expected to observed variance of a suspected controlled variable. For a controlled variable which is being disturbed, expected variance will be considerably greater than observed variance; the organism's actions tend to reduce variance created by the disturbance and the stability factor will be greater than 1.0. For an uncontrolled variable, expected and observed variance will be equal; the stability factor will be close to 1.0.

The use of the stability factor is illustrated in a simple experiment. The subject is seated in front of a video monitor that shows a small square above a small diamond. The subject is asked to move either symbol back and forth in an arbitrary manner using the handle of a joystick. Moving the joystick to the right moves *both* symbols to the right; moving it to the left moves both symbols to the left. The movement of the symbols is also influenced by slowly varying random disturbances, a different disturbance for each symbol. At random times during a 10-minute experimental run, a tone is played, signaling the subject to “change his mind” and move the symbol that is not currently being moved.

At any point in the experiment it is impossible for an observer to tell whether the subject is moving the square or the diamond. The subject is actually moving both symbols all the time; movements of both the square and the diamond are behaviors of the subject. To determine which symbol is being moved intentionally a stability factor is computed under the hypothesis that the position of the square is the controlled variable. The position of the square is determined at each instant by the sum of the values of two variables, one corresponding to the position of the joystick and the other to the value of the disturbance. Treating these as independent random variables, the expected variance of square position is the sum of the variances of joystick and disturbance values. The observed variance of square position is the sum of these variables. The variance measures are based on 240 equally spaced samples of joystick and disturbance, which are updated during each screen refresh of the video animation. The stability factor is computed continuously throughout an experimental run. The value of the stability factor should be significantly greater that 1.0[2] when the subject is intentionally moving the square (in practice it was always greater than 1.2) and close to (or less than) 1.0 when the subject is intentionally moving the diamond.

The results of one experimental run are shown in Figure 1. The lower two traces show the movements of the square and diamond symbols over time. Nothing about the movements of these symbols suggests which is being moved intentionally at a particular time. The vertical lines show the occurrence of the tones signaling the subject to change from one symbol to the other. Again, there is nothing about the movements of the symbols after each tone, which suggests that a change has occurred.

The upper trace in the figure shows the value of the stability factor over time. The stability factor changes markedly after each tone. The value of the stability factor suggests that this subject started by moving the square; the stability factor is significantly greater than 1.0 just before the first tone. After the tone the stability factor gradually approaches 1.0, suggesting that the subject has switched to the diamond. (Actually, this only means that the subject is not moving the square but it is assumed that the subject followed instructions and was moving the diamond when not moving the square). The stability factor often goes

---

[2] The stability factor is a random variable distributed as F. In this experiment the probability that the stability factor will be greater than 1.2 when the subject is not intentionally moving the square is 0.05.

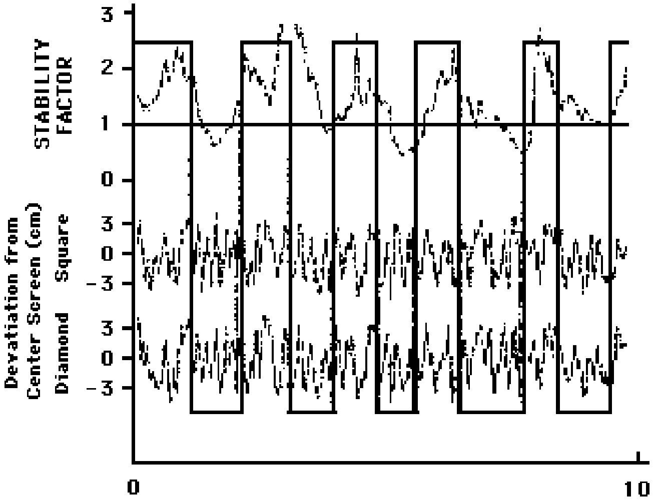

Figure 1. Movements of the square and diamond symbols during one 10-minute experimental run. The top trace shows changes in the stability factor during the run. Vertical lines indicate the occurrence of the tone signaling the subject to change symbols.

below 1.0 when the subject is not moving the square and indicates that actions destabilize the uncontrolled variable. The stability factor changes regularly after each tone (vertical line) indicating that the subject is intentionally moving a different symbol each time.

Table 1 shows the value of the stability factor before the onset of each tone signaling the subject to change intentions. For the two subjects tested the stability factor alternates regularly from high (greater than 1.2) to low (less than 1.2) values at the end of successive time periods. This is expected if the subjects (as instructed) are moving alternate symbols after each tone. The results in the rightmost column (labeled control) show the value of the stability factor before the onset of each tone when a subject is moving the joystick but not looking at the video screen. This is a control condition to show that the stability factor does not alternate regularly if the subject is not controlling either symbol.

TABLE 1

Value of the Stability Factor at the End of Seven Successive Time Periods Terminated by Occurrence of a Tone Signaling the Subject to Move a Different Symbol.

Time	Subject		
Period	IW	RM	Control
1	1.43	1.86	0.94
2	0.96	1.14	1.10
3	1.20	2.01	0.95
4	0.86	0.96	0.92
5	1.33	1.96	0.93
6	1.13	0.73	0.92
7	1.23	1.62	1.26

Note. The control condition shows the values of the stability factor at the end of these time periods when a subject is moving the joystick but not looking at the video screen.

The stability factor was used to monitor continuously the purpose of a person's actions (joystick movements). The purpose of action is never self-evident. Purpose corresponds to intended perceptions, not objective results of action. To know what a person is doing an observer must discover what a person is trying to perceive. The test for controlled variables is the basis for a new methodology in psychology, one aimed at the experimental detection of purpose.

The Dancer and the Dance: Methods in the Study of Living Control Systems

This paper describes methods for studying behavior when organisms are viewed as living control systems. These methods are aimed at determining the variables that organisms control when they are engaged in various observable behaviors. Controlled variables are perceptual representations of the environment that are protected from the effects of disturbance by the responses of the organism. It is possible to detect controlled variables by applying disturbances to aspects of the environment that might be under control and looking for *lack* of an effect of these disturbances. This *Test for Controlled Variables* (TCV) makes it possible to see the "dance" of behavior from the perspective of the behaving organism (the "dancer") rather than from that of the observer, who sees only the dance.

Powers (1989) has suggested that what we call "behavior" is actually a visible side-effect of a process called "control". Control is what organisms do to keep their own perceptual experiences in preferred states. This approach to behavior is based on control theory which views behavior as the control of perception (Powers, 1973). From the control theory perspective, behavior is a

---

From *Psychological Methods,* 1997, 2, 436-446. Reprinted with permission of the American Psychological Association.

visible side-effect of an organism's efforts to control its own perceptions. For example, a dance is the visible side-effect of the dancer's efforts to control perceptions of muscle tension, limb position and body movement (among many others). The dance that is seen (and controlled) by the dancer is quite different than the one seen by an observer. The skilled dancer knows how to control her own perceptions in order to produce the dance that the observer sees.

Conventional psychological research has been aimed at the discovery of variables that control the “dance” of behavior as seen from the perspective of observers (who, in this case, consist of behavioral researchers) while ignoring the perceptual variables that are being controlled by the organism’s behavior (Marken, 1993). For example, operant researchers (e.g. Staddon, 1979, Timberlake, 1984) study how variables such as food delivery schedules, constraints and set points control behaviors, such as response rates, while ignoring variables, such as the perception of stomach fullness, that might be controlled by these behaviors.

The control theorist, on the other hand, sees the dance of behavior as a side issue; an interesting side issue but beside the point from the perspective of the behaving organism. A control theorist studying operant behavior wants to identify the perceptual variables that are controlled by behavior rather than the behavioral variables that are controlled by food delivery schedules, constraints, set points and the like.

Proximal and Distal Cause

The conventional approach to the study of behavior is based on a *causal model* of the relationship between an organism and its environment. According to this model, variations in environmental (independent) variables (IVs) cause variations in response (dependent) variables (DVs). The goal of research is to discover the environmental variables that actually do have an effect on response variables. This is done by manipulating an IV and looking for concomitant changes in a DV. If changes in the IV are associated with changes in the DV when all other variables are held constant we can conclude that variations in the IV are the cause of variations in the DV (e.g. Bordens and Abbott, 1991; Marken, 1981). The causal model assumes that these IV-DV relationships reveal something about the nature of the organism under study. It is, therefore, the basis of even the most advanced, "state of the art" approaches to the study of behavior (e.g. Shadish, 1996).

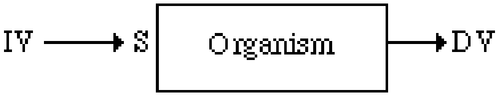

Figure 1. The causal model of the relationship between an organism and its environment (IV).

In most behavioral experiments it is clear that the IV cannot be the immediate cause of behavior since it is out in the environment, some distance from the organism. Thus, an implicit assumption of the causal model of behavior is that IVs have sensory effects (S) that are the immediate cause of the organism's behavior (Figure 1). The IV is a "distal" cause while the sensory effects are the immediate or “proximal” cause of behavior. For example, in an experiment where the IV is the schedule of reinforcement, the different schedules are the distal cause while the sensed effect of each schedule (such as the sensed rate of food delivery) is the immediate or proximal cause of behavior (rate of pecking). Research based on the causal model of behavior is based on the assumption that variations in the distal cause produce concomitant variations in the proximal cause of behavior. This can be called *the assumption of concomitant variation*. This assumption justifies the use of the IV as a legitimate “surrogate” for the immediate or proximal cause of behavior (S). Only if the assumption of concomitant variation holds can we conclude that any observed relationship between an IV and a DV reflects the actual causal relationship between proximal cause and behavioral effect.

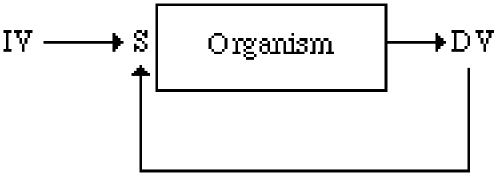

Figure 2. Closed-loop relationship between organism and the proximal cause (S) of its behavior (DV).

Powers (1978, 1979) has shown that the assumption of concomitant variation breaks down when organisms exist in a closed-loop relationship with respect to the proximal causes of their behavior (Figure 2). In a closed-loop, the proximal (sensory) cause of behavior (S) is influenced *simultaneously* by both distal variables (IVs) and the organism's own responses (DV). So an IV is not the only influence on the proximal cause (S) of behavior; the other influence on S is behavior itself. In the schedule of reinforcement experiment, for example, the sensed rate of food delivery (S) is influenced by both the schedule of reinforcement (IV) and pecking rate (behavior). The effect of the organism's own behavior on the proximal causes of that behavior reduces or eliminates any concomitant variation between the distal and proximal causes of behavior.

Closed-Loop Control

If there is a closed-loop relationship between organisms and the proximal causes of their behavior then a fundamental assumption of the causal model of behavioral research – the assumption of concomitant variation – is violated. The fact that organisms do exist in such a closed-loop relationship is evident from inspection; the proximal causes of an organism’s behavior occur at sense organs (eyes, ears, nose, skin etc.) that are located on the surface of the organism itself. Any behavior caused by stimulation of a sense organ changes the orientation of the organism with respect to the distal cause of that stimulation. This necessarily affects the way the stimulation is represented at the sense organs. The effect of behavior (DV) on the proximal cause of that behavior (S) is immediate and strong; the DV affects S *while* S affects the DV. For example, when the image of an attractive passer-by (S) moves across the eye it causes head movements (DV) that affect the movement of the image *while* the image is causing the head movements.

One might imagine that failure of the assumption of concomitant variation would have become evident soon after the first behavioral experiments were completed; this failure would have shown up as an inability to find consistent IV-DV relationships in behavioral experiments. But behavioral researchers do find fairly clear relationships between distal causes (IVs) and behavioral effects (DVs). While most of these relationships are statistical in nature (which is what one would expect if there were no concomitant variation of the proximal and distal causes of behavior) reliable relationships between IV and DV are also found. For example, highly reliable relationships between reinforcement schedule and response rate are found in the study of operant behavior (Skinner, 1971). Such reliability does not necessarily mean, however, that the assumption of concomitant variation holds in these cases. Reliable relationships between IV and DV are expected when organisms are in a closed-loop *negative feedback* relationship with respect to the proximal cause of their behavior.

There is negative feedback in a closed-loop when the sensory input to the loop causes responses that reduce the tendency for that input to cause further responses; responses suppress (have a negative effect on) the cause of responses. This negative feedback process stabilizes sensory input (S), keeping it at a *reference value* (a value that causes no further responses), protected from the effects of *disturbances*. Disturbances are environmental variables (IVs) that influence sensory input variables. The relationship between variables in a negative feedback loop is shown in Figure 3. The position of the passer-by is the IV or disturbance variable; the angle of the head is the DV or response variable; the image of the passer-by on the retina is the sensory input variable (S). Changes in the IV (position of the passer-by) lead to changes in the DV (head angle); the result is little or no change in S (the position of the passer-by on the retina).

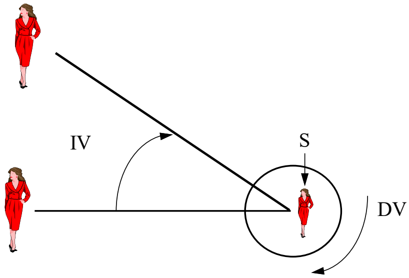

Figure 3. Relationship between IV, DV and S in a negative feedback loop. As the IV (the position of the passer-by) changes, the DV (head angle) changes so that S (the position of the passerby on the eye) remains approximately constant.

Organisms in a closed negative feedback loop act to bring their sensory inputs to reference values and protect them from disturbances; this process is called *control*. Organisms that act to control their sensory inputs are *living control systems*. A living control system controls its sensory inputs by varying its responses so as to oppose disturbances. This means that variations in the control system's responses (DV) will tend to be strongly related to variations in environmental disturbances (IV) to the controlled sensory input. This relationship between IV and DV exists even though the proximal cause of responses (the controlled sensory input, S) has little or no relationship to variations in the IV or DV (Powers, 1979); there is, thus, no causal path from IV to DV via the organism. When an organism is in a closed-loop, negative feedback relationship with respect to the proximal causes of its behavior, the causal path from IV to DV runs through the environment and is completely external to the organism. However, an external observer, seeing the relationship between IV and DV, is likely to see the variations in IV as acting on the organism to cause the variations in the DV. This is called the *behavioral illusion* (Powers, 1978).

Behavioral Illusion

The behavioral illusion is seen when an environmental variable appears to be acting on an organism to cause its responses. For example, we see this illusion when a person's head turns to see an attractive passer-by (Figure 3). It looks like the passer-by causes the person's head to turn but the person is actually controlling the image of the passer-by, keeping it centered on the fovea. The movement of the passer-by is a disturbance to the position of the image on the eye and head-turning is the response that counters this disturbance. It looks like movements of the passer-by (IV) act on the person to cause proportional movements of the head (DV). In fact, the causal path from IV to DV runs through the environment, not the person; the proportional changes in the DV in response to changes in the IV are a result of the laws of optics (that relate the passer-by to the image on the eye) not the laws of behavior (that relate the image on the eye to head turning).

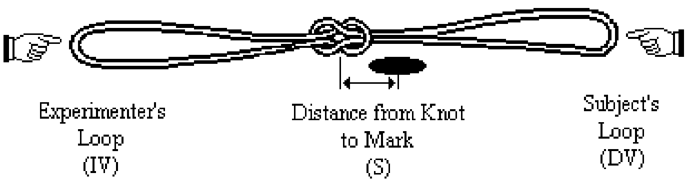

Figure 4. Variables involved in the rubber-band demonstration of the behavioral illusion.

The behavioral illusion has important implications for the study of living control systems. It shows that the IV-DV relationships seen in conventional behavioral research can be quite misleading when the object of study is a living control system. The problem can be illustrated with a simple demonstration involving an experimenter, a subject and a pair of rubber-bands (Powers, 1973, p 241). The rubber-bands are knotted together as shown in Figure 4 and placed on a tabletop. The experimenter puts a finger in one rubber-band loop; the subject puts a finger in the other. The subject is asked to keep the knot next to some target mark on the table; so the subject is asked to control the distance from knot to target mark. This distance, which is the controlled sensory input, S, is to be kept at zero. The experimenter applies disturbances to S by pulling on his loop of the rubber-bands; the subject responds by pulling on her loop.

The experimenter's pull on the knot is the IV in the experiment; the subject's pull is the DV. The IV and DV can be measured in terms of the distance from each finger to the target mark; the greater this distance, the greater the pull on S. The results of this experiment, for three different values of the IV, are shown in the left panel of Figure 5. As expected, the subject's pull is proportional to the experimenter's pull, as it must be if the distance from knot to target mark is to be kept at the reference value (zero distance). The slope of the relationship presumably reflects the subject's responsiveness to changes in the proximal cause of behavior (S). This is the conventional interpretation of the IV-DV relationship in Figure 5. It looks like there is a “behavioral law” of the form: one unit of change in the IV (and, by assumption, a proportional change in the proximal cause of behavior) results in one unit of change in the DV. This behavioral law presumably shows the subject's responsiveness to the stimulation caused by the experimenter’s pulls (IV).

Because the subject's behavior occurs in a negative feedback loop, the IV-DV relationship seen on the left of Figure 4 does not reflect the subject’s responsiveness to stimulation. Rather, it reflects characteristics of the environment that link IV and DV to the controlled sensory input, S. That is, the IV-DV relationship shown on the left of Figure 5 reflects the relative elasticity of the two rubber-bands. The results of this experiment tell us about the subject's environment, not the subject. This can be shown by replacing one of the rubber-bands with another that is far more elastic. If, for example, the subject's rubber-band is replaced with a large, thick one we get the IV-DV relationship shown on the right in Figure 5. It looks like the subject has become far less responsive to changes in the IV. But there has been no change in the subject; only a change in the environment. The appearance that the IV-DV relationship observed in this experiment reveals something about the subject's tendency to respond to pulls on the rubber-band – or to the proximal effect of those pulls – is the behavioral illusion.

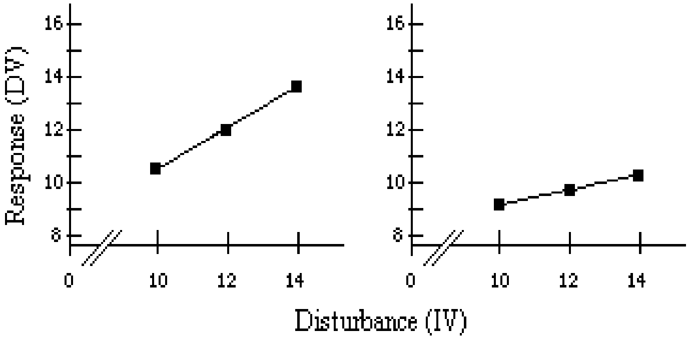

Figure 5. The behavioral illusion. The left panel shows the relationship between IV and DV when the elasticity of the experimenter and subject's rubber-bands are nearly equal. The right panel shows the relationship between IV and DV when the elasticity of the experimenter's rubber-band is less than that of the subject. It looks like the subject becomes less responsive to the IV. In fact, the results reflect an increase in the responsiveness of the environment, not a decrease in the responsiveness of the subject.

The rubber-bands in this demonstration make visible what is typically invisible in studies of behavior: the environmental connections between IV, DV and S. These connections (like the laws of optics in the example of watching an attractive passer-by) are what determine the form of the relationship between IV and DV when studying living control systems. Surprisingly, when dealing with living control systems, conventional IV-DV relationships tell us almost nothing about the system's behavior but a great deal about the nature of the environment in which the system behaves.

Controlled Variables

The rubber-band demonstration illustrates the problems of the conventional IV-DV approach to the study of living control systems but it also points to a useful new approach to the study of these systems. The behavior in the rubber-band demonstration (the DV in Figure 5) becomes completely transparent once we know that the subject is trying to keep the knot on the target mark; the visible “dance” of behavior makes sense as soon as we know what sensory input (S) the subject is controlling. Controlled sensory inputs are called *controlled variables*. So we can understand the behavior of living control systems if we can identify the variables they are controlling: controlled variables. Controlled variables are aspects of the system's own sensory experience that it is keeping under control. Once we know what sensory inputs the system is controlling, we can deduce IV-DV relationships (if we care to do so) by determining the nature of the environmental connection between the IV, the DV and the controlled variable (S).

In the rubber-band demonstration we knew that the subject was controlling the distance from knot to target mark because we asked the subject do this; assuming that the subject complied, we knew that the controlled variable was the distance from knot to target mark. But under ordinary circumstances the variables that organisms control are by no means obvious. Controlled variables are variable aspects of the organism's own experience; they are what the organism perceives, like the location of the image of the passer-by on the eye or the sweetness of coffee in the mouth. We see the dance of behavior – the movement of the head or the sugar cubes placed in the cup – but not the behavior of the perceptions controlled by the dancer – the centered image of the passer-by or the sweetness of the coffee.

The fact that organisms are controlling their own perceptions explains why it is often difficult to tell what an organism is doing, even though we can see its every action. We see the dance of behavior (the actions that keep perceptions under control) but not the reason for the dance (the perceptions that the organism controls). For example, we see a person bolting across the street but we don't see why (to catch the bus that is pulling away from the stop). It is difficult to tell what an organism is doing because it is difficult to see what an organism is controlling. Controlled variables are perceptual variables and we can't get inside an organism to see what it perceives. But we can determine what aspects of our own perceptions of the environment correspond to the perceptions the organism is controlling. In order to be able to tell what an organism is doing (controlling) we have to be able to perceive the world as the organism perceives it.

Controlled variables can be viewed as perceptual “maps” of the environment in which behavior takes place. The perception of the distance between knot and target mark, for example, is one mapping of the environment in which the rubber-band demonstration takes place. But there are other ways to map the same external environment into perceptions. For example, it is possible to see the knot and target mark as part of a pattern, such as a triangle. This pattern perception can be controlled in the same way as the distance perception is controlled – by pulling on the rubber-band. So it would be difficult for an observer to tell, by looking at the subject's responses (pulls on the rubber-band), whether the subject was controlling the distance (knot to target mark) or the pattern (triangle with the knot and target mark as components).

We can tell what variables an organism is controlling if we can determine what aspects of the environment the organism maps into controlled perceptions. Of course, what we see as the organism's environment is our own perceptual map of the “real” environment beyond our senses; we, like the organisms we study, have no privileged access to what's really out there. So the problem of determining an organism's map of the environment is really a matter of determining which one of our own perceptions corresponds to the perception the organism is controlling. We can determine this mapping through the use of a formal procedure called *The Test for Controlled Variables* or simply the *TCV*.

The TCV

The TCV is based on the fact that living control systems resist disturbances to the variables they control. This disturbance-resisting characteristic of living control systems was seen in the rubber-band demonstration where the subject resisted experimenter produced disturbances (pulls on the experimenter's loop of the rubber-bands) to the distance between knot and target mark. The subject resisted these disturbances by pulling in the opposite direction on her loop of the rubber-bands. The result was that disturbances had far less of an effect on the controlled variable (distance between knot and target mark) than expected.

The TCV starts with a guess or *hypothesis* about a perceptual variable that an organism might be controlling. This is really a hypothesis about the aspects of the tester's own perceptions that might correspond to the perception that the organism is controlling. The hypothesis about the controlled variable is aimed at explaining some observed behavior, such as catching a baseball. The tester tries to think of a perceptual variable, which, if it were being controlled, might explain the observed behavior. For example, it has been suggested that outfielders catch fly balls by controlling the optical acceleration of the ball on the eye (Babler and Dannemiller, 1993). This hypothesis is based on the assumption that the observed movements of the outfielder are aimed at keeping optical acceleration under control.

Hypotheses about controlled variables are tested by applying disturbances to these variables and looking for *lack* of an effect of these disturbances. A disturbance is any change in the environment or the organism's relationship to the environment that would change the value of the hypothesized controlled variable if that variable were *not* under control. For example, suppose that it has been hypothesized that outfielders do control optical acceleration. This can be tested by hitting fly balls in different trajectories relative to the outfielder; the different trajectories are disturbances to the optical acceleration. If optical acceleration is under control, the outfielder will *resist* these disturbances (by running relative to the ball at the appropriate rate) so that optical acceleration remains constant; there would be little or no effect of the disturbances (ball trajectories) on the controlled variable (optical acceleration). If optical acceleration is not under control, however, the different trajectories (disturbances) will meet with no resistance; optical acceleration will differ for each trajectory.

The TCV can be viewed as a version of conventional IV-DV research where the disturbances are the IV and the hypothetical controlled variable is the DV. In conventional IV-DV research, a successful experiment is one where we find a large effect of the IV on the DV. In IV-DV research using the TCV, a successful experiment is one where we find little or no effect of the IV on the DV. When there is little or no effect of the IV on the DV and we can trace this lack of effect to the subject's actions (such as the outfielder's movements when running to catch a fly ball), then we can be reasonably sure that the hypothesized controlled variable (DV) is the actual controlled variable (S) (see Marken, 1982, pp. 53-56).

Doing the TCV

In order to determine whether or not disturbances have an effect on the hypothesized controlled variable, it is necessary to monitor the state of both disturbances and the hypothesized controlled variable itself. It is often easier to monitor disturbances, which are typically environmental events created by the observer, than possible controlled variables. In the case of catching fly balls, for example, it is clearly easier to monitor the trajectory of the fly ball (disturbance) than the optical acceleration of the ball on the retina (hypothesized controlled variable). But it is possible to develop clever measurement systems, such as the shoulder mounted video system used by McBeath, Shaffer and Kaiser (1995), to monitor the state of hypothesized controlled variables. By plotting video measures of the optical position of fly balls as seen from the outfielder's shoulders, McBeath et. al. were able to reject the hypothesis that optical acceleration is the variable controlled when catching fly balls; optical acceleration was found to be strongly affected by the different trajectories.

If disturbances have a pronounced effect on the hypothesized controlled variable, as they did in the McBeath et. al. study, we proceed to the next step in the TCV, which is to revise the hypothesis about the controlled variable and test again (apply disturbances to the newly hypothesized controlled variable and look for lack of effect). McBeath et. al. were actually ready with an alternative hypothesis for the controlled variable at the start of their study of catching fly balls. They had guessed that outfielders might be controlling the linear optical trajectory (LOT) rather than the optical acceleration of the fly ball; that is, they hypothesized that the controlled variable is the projection of the ball onto an imaginary straight line on the eye. McBeath et. al. found that LOT is, indeed, a variable controlled by the outfielders; regardless of the trajectory of the fly ball (disturbances) the optical projection of the ball followed a LOT (the controlled variable), as shown in Figure 6.

The McBeath et. al. approach to the study of the variables controlled when outfielders catch fly balls was a particularly good one because the hypothetical controlled variable (LOT) was monitored directly during disturbance (movement of the ball) using a video camera as a stand-in for the retinal surface. Other efforts to determine this variable have been based on inferences from observation of the relationship between responses, such as the player's movements on the field, and disturbances, such as the trajectory of the fly ball (Babler and Dannemiller, 1993; McLeod and Dienes, 1996). The accuracy of such inferences depends on the correctness of the assumptions

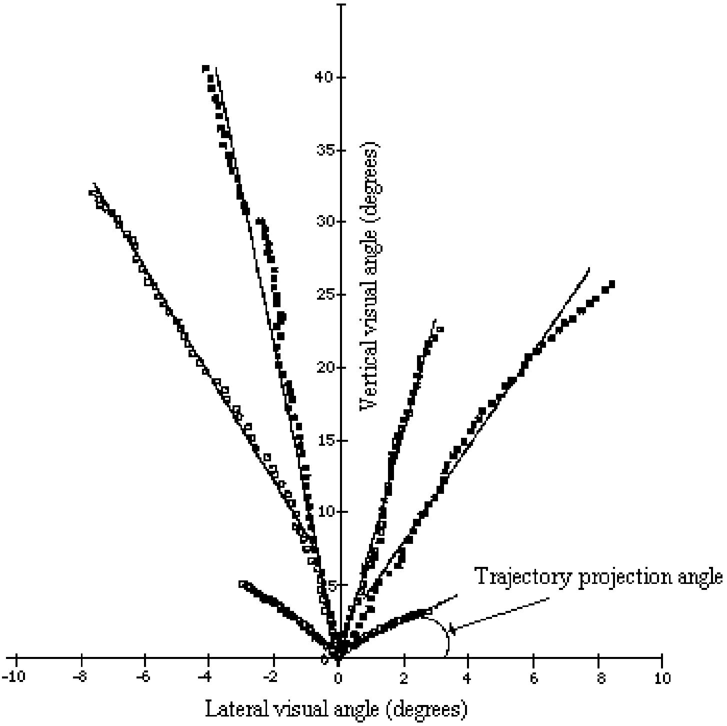

Figure 6. Optical projection of fly ball trajectories recorded by video camera mounted on the outfielder's shoulder; the origin of the plot corresponds to the fielder's view of home plate. The

fly balls are maintained in linear optical trajectories (LOTs) which are straight-line patterns on the optical surface. (Adapted from Figure 4 in McBeath, Shaffer and Kaiser, 1995; reprinted with permission)

that are made about the functions relating disturbance and response to the hypothetical controlled variable. In many cases these functions are assumed to be linear when, in fact, they are almost certain to contain complex non-linearities. The McBeath et. al. results show that the relationship between response (movement on the field) and disturbance (trajectory) can appear to be consistent with the optical acceleration hypothesis even when there is no evidence that optical acceleration is being stabilized on the retina. When done properly, the TCV lets us see past the dance of behavior to the perceptual variables controlled by the dancer.

When to Use the TCV

The TCV is used to study the behavior of control systems because it is the only way to identify, with any confidence, the perceptions that these systems are controlling. But it is not necessary to know, *a priori*, that the system under study is a control system before using the TCV. Indeed, it is recommended that researchers use the TCV whenever there is even the slightest possibility that the behavior under study is that of a control system. This is because there is no penalty for using the TCV to study the behavior of a non-control system but there is a large penalty for using conventional IV-DV methodology to study the behavior of a control system.

If the TCV is used to study the behavior of a non-control system, the results of such a study will be precisely equivalent to those that would have been obtained using conventional IV-DV methodology. For example, suppose that the TCV were used to study the behavior of a metal ball rolling down an inclined plane. It might be hypothesized that the ball is controlling its perception of acceleration down the plane. Disturbances (IV) to this hypothetical controlled variable are applied in the form of variations in the inclination of the plane; measures of the hypothetical controlled variable (DV) are obtained for each inclination. If perceived acceleration is under control there should be little or no effect of the IV on the DV. What will be found, of course, is a strong effect of the IV on the DV; the effect of inclination (IV) on acceleration (DV) will be precisely what is expected on the basis of a causal model of the system (Newton's laws). The TCV reveals that the metal ball is not a control system. At least, it is not a control system with respect to the hypothetical controlled variable: acceleration.

If, on the other hand, conventional IV-DV methodology is used to study the behavior of what is actually a control system, the results of such a study will *not* be equivalent to what would have been obtained using the TCV. What would be missing is evidence regarding the variables the system is controlling. If the results of the conventional IV-DV study show a lack of effect of the IV on the DV we cannot, then, conclude that the DV is a controlled variable because the system's actions may not be responsible for this lack of effect. Similarly, if the results of the conventional IV-DV study show a strong effect of the IV on the DV we cannot conclude that this effect is the system’s response to a disturbance to a controlled variable (as it was in the rubber-band demonstration) because we did not determine whether the hypothetical controlled variable actually remained stable while it was being disturbed.

The results of conventional IV-DV research can provide strong *hints* about the perceptual variables that are being controlled by a living control system. But the only way to determine whether or not these variables are actually under control is by doing the TCV for the Controlled Variable.

Modeling Control

Once we know what an organism is controlling, it is possible to develop models that reproduce the organism's behavior and predict how the organism will behave in new situations (Marken, 1992). One of the main difficulties involved in modeling the behavior of organisms is determining the variable or variables that the organism (and, hence, the model) controls. Once we know what variable is being controlled in a particular situation, it is possible to model the behavior in this and future situations with great accuracy (Bourbon, 1996).

Modeling can actually be used as part of the process of determining the variables that are being controlled by a living control system. Powers (1971), for example, used modeling to determine one of the variables controlled by rats in a shock avoidance study. The model was set up to control either the average rate of shock occurrence or the probability of a shock in a fixed time interval. Both versions of the model fit the data extremely well (model behavior deviated from actual behavior by less than one press per minute) but the model that controlled shock probability gave a significantly better fit to the data than the one that controlled shock rate. So it was possible to conclude that the rats in this experiment were controlling a perceptual variable that was more like shock probability than shock rate.

The modeling approach to determining controlled variables requires data that is collected in situations where disturbance and response measures have a known relationship to possible controlled variables. Modeling could be applied to the shock avoidance experiment, for example, because the data record included measures of shock delivery schedule (IV) and response rate (DV) that could be related to the actual rate and probability of shock experienced by each rat. That is, measures of the IV and DV could be related to the actual state of the possible controlled variables. Unfortunately, the relationship of IV and DV to possible controlled variables is not always this clear in behavioral research.

Data are usually expressed as averages over subjects, which eliminates the ability to compute the variables that might have been controlled by an individual subject. Moreover, measures of the IV and DV are often arbitrary (in terms of their mutual effect on possible controlled variables) so they may have no obvious link to a common variable. In this case, it is necessary to test for controlled variables before modeling can be usefully applied to the behavior observed in the experiment.

Reference Values

Controlled variables can be maintained at any one of many different reference values. For example, the distance between knot and target mark in the rubber-band demonstration could have been maintained at many values besides zero: 1 cm, 2 cm etc. Indeed, the reference value at which a controlled variable is maintained can be changed continuously over time by the organism itself. This is what happens when we move a limb, for example. The reference values of the variables controlled when a limb is moved from one position to another are gradually changed along a continuum.

Changing reference values present a problem when doing the TCV because they can be confused with the effects of disturbances to the controlled variable. If, for example, the subject in the rubber-band demonstration decides to continuously change the reference distance between knot and target mark while the experimenter is applying continuously changing disturbance (pulls on the loop of the rubber-band) the observed changes in the value of the controlled variable could be seen as evidence that the disturbances are having an effect; the experimenter could mistakenly conclude that the distance between knot and target mark is not the controlled variable.

Changing reference values are taken into account by looking for lack of any *systematic* effect of disturbances on a possible controlled variable. While it is true that any particular change in the controlled variable may appear to be a result of the disturbance when it is actually a result of a change in reference value, it is highly unlikely that there will be a systematic relationship between disturbances and changes in reference values in the long run.

The use of the TCV when reference values are varying was demonstrated in an experiment where subjects were asked to control a particular variable (the position of one of five objects on a computer screen) but to vary the reference value of this variable by moving the object in a temporal pattern on the screen (Marken, 1992, p. 41-46). The subject's responses influenced all five objects simultaneously so it was impossible to tell which of the five objects was under control by simply looking at object movement. Therefore, the TCV was used to determine which object was actually under control. The disturbances were filtered random waveforms that were added to the subject's effects on each object. The controlled object was reliably detected as the one for which the correlation between variations in disturbance and object position was closest to zero.

The detection of the object controlled in the "five objects" study was done by computation of a running correlation between variations in the disturbance to each object and the hypothesized controlled variable – the position of each object. A high correlation between the disturbance and the hypothetical controlled variable occurs if there is an effect of the disturbance that is greater than the effect of the changing reference for the controlled variable. So a high correlation between disturbance and hypothetical controlled variable means that an object's position is *not* under control. A low correlation between the disturbance and the hypothetical controlled variable occurs if there is no systematic effect of the disturbance that is greater than the effect of the changing reference for the controlled variable. So a low correlation between disturbance and hypothetical controlled variable means that an object’s position *is* very likely to be under control; the disturbance has no systematic effect on the hypothesized controlled variable (the object’s position) because it is being resisted by the subject.

This statistical approach to the TCV is useful whenever it is likely that a large portion of the variation in the value of a possible controlled variable is a result of the organism itself changing the reference value for the variable. If a possible controlled variable varies when all disturbances seem to be constant, it is likely that the variation in the value of this variable is the result of changing reference values.

The “five objects” study shows that it is not necessary, in principle, to know the reference level of a controlled variable in order to successfully perform the TCV. It was possible to determine which of the five objects was under control without guessing the momentary reference value for the position of each object at every instant during the TCV.

The Perceptions People Control

Controlled variables can be simple, like the position of an object on the computer screen, or complex, like a person's position on a political issue. But simple or complex, controlled variables are always perceptions; living control systems control perceptual representations of various aspects of the external environment. The position of an object on the computer screen is a perceptual representation of just one of many possible aspects of the object, such as its size, color, shape etc. Similarly, a person's position on a political issue is just one of many possible aspects of the words on the sign that the person is waving in front of you. Living control systems control many of these perceptions, simple and complex, simultaneously. We test for controlled variables one at a time; but that doesn't mean that living control systems are like thermostats, controlling only one variable at a time.

One way to get a sense of the variety of perceptions people control is by playing the “coin game” described by Powers (1973). The game is played with two people, an experimenter (yourself) and a subject. The subject is asked to place four coins on a table so that “...they satisfy some specific condition or exemplify some specific pattern. The experimenter is to discover what the subject has in mind [what the subject wants to perceive] without any verbal communication” (Powers, 1973, p. 235). The experimenter uses the TCV to do this, applying disturbances by changing the coins in some way. If a disturbance changes the coins so that they are no longer in the desired (reference) condition or pattern, the subject corrects the error. If the disturbance leaves the coins in the reference condition, the subject says “no error”. The experimenter has determined what perception of the coins the subject is controlling when he can make three changes that will be corrected and three that will not be corrected by the subject.

If you play this game with several different subjects you will see that people will often control perceptual aspects of the coins that you may not have anticipated. Some subjects may try to control the spatial pattern of the coins. In this case a change in the position of any coin could be an error producing disturbance; the only change that is not a disturbance is turning the coins over (from heads to tails, say). But some subjects may try to control more complex aspects of the coins. For example, a subject might try to control a relationship between the coins showing heads and those showing tails; the reference perception could be something like "two coins showing heads; two showing tails". In this case, a change in the position of any coin would not produce an error; the disturbance to the pattern of coins would be fully effective.

The coin game shows that preconceptions about what a person is "doing" (controlling) can interfere with the ability to determine the variables that are actually under control. If the experimenter can only imagine that the subject is controlling something about the pattern of the coins, then the absence of a corrective response to pattern disturbances will prove quite puzzling and frustrating. This puzzlement and frustration will eventually disappear if the experimenter perseveres and tries testing for different types of controlled variables. When the experimenter eventually guesses that the subject is controlling "two coins showing heads, two showing tails" he will be relieved to find that every disturbance to this condition of the coins results in corrective action.

It is possible to conceive of many different perceptual aspects of the coins that could be controlled. For example, the subject could control for two coins being on one side of an imaginary line and two being on the other. Or the subject could control for having any *one* of the four coins on the table. Or the subject could control for a contingency between the coins such as "if the nickel is heads then any two other coins must show tails otherwise all the other coins must show heads". All of these perceptions are relatively complex and it can be very difficult to tell when a subject is controlling perceptions like these. But people control perceptions at least this complex every day.

The coin game shows how difficult it can be to determine what perceptions another organism is controlling. The TCV requires creativity (in generating hypotheses about possible controlled variables) skill (to know which disturbances to apply to test these hypotheses) and persistence (to keep testing as hypothesis after hypothesis is rejected). The TCV is not easy; but it's the only way to learn what is most important about the behavior of a living control system: what perceptions it is controlling.

Testing In Everyday Life

The TCV can be carried out using formal or informal procedures. The McBeath et. al. study of the variables controlled when catching a fly ball is an example of a formal application of the TCV; there is precise measurement of disturbances (the trajectory of the fly ball) and the state of the hypothetical controlled variable (the optical position of the ball). The experimental test to determine which one of five objects was being moved around the computer screen was also an example of a formal application of the TCV: the computer continuously calculated the correlation between program generated disturbances and measures of the positions of the objects on the screen.

Formal applications of the TCV are used when it is possible to obtain quantitative measures of the variables involved in control, in particular, the disturbance variable and the hypothetical controlled variable. But it is still possible to test for controlled variables in situations where it is difficult or impossible to obtain such measures (Runkel, 1990, p. 151). The essence of the TCV is the development of hypotheses about controlled variables. Once these hypotheses are developed it is usually possible to think of various ways to disturb possible controlled variables. These disturbances do not need to be quantitatively precise; if a variable is under control the organism will clearly act to restore the situation after the disturbance. If the variable is not under control, the organism will do nothing about the disturbance at all.

A familiar example of an everyday use of the TCV (familiar to moviegoers) is determining whether you are being followed. The hypothesis is that the driver in the other car is controlling a perceived relationship: “behind your car”. You test this hypothesis by applying disturbances; randomly turning up one street and down another. If the car remains behind yours after a number of these disturbances have been applied, it's a good bet that you are being followed. If one disturbance is effective, however (that is, if you turn and the car that was following is no longer behind you) then the apparent “following” was a coincidence; the perception of being behind you was not under control.

More common (and useful) examples of informal uses of the TCV occur in counseling and therapy situations. The goal here is to discover, without explicitly asking, what the client or patient wants. Again the process starts with a hypothesis about what perception is under control. Disturbances can be applied with words. For example, you might guess that the patient is controlling for being seen as a "good son". This can be tested by suggesting (verbally) that the client has not been nice to his mother. If the client is controlling for the hypothesized perception, the verbal disturbance will be resisted; the client becomes "defensive". If the client is not controlling for the hypothesized variable the disturbance will be ignored or, at least, not resisted.

The TCV can also be used in therapeutic situations to reveal the existence of internal conflicts. Conflict exists when two control systems try to keep the same perception at two different reference values; for example, a conflict exists when you want to stay at a party but you also want to leave; you are trying to keep the same variable (perception of where you are) at two different reference values ("at the party" and "not at the party"). Conflicts like this make life difficult and stressful. But it is relatively easy to resolve such conflicts (Powers, 1992, p. 41-54) once their existence has been revealed by the TCV.

Conclusion

The results of conventional IV-DV research are misleading when the objects of study are living control systems. It seems like the relationship between IV and DV reveals something about the nature of the organism; in fact, the relationship between IV and DV reveals something about the nature of the connection between the organism and the aspects of the environment that it controls. Living control systems are organized around the control of their own perceptual experiences of the environment. Thus, one of the main goals of the study of living control systems is to determine the perceptual variables that these systems control. The TCV is the basic methodology in the study of living control systems. The TCV provides a systematic means for determining the perceptions a system is controlling, even when the reference states of these perceptions are changing. The TCV makes it possible to see the dance of behavior from the point of view of the dancer rather than from that of an observer of the dance.

Note

Demonstrations of some of the research discussed in the article are available on the World Wide Web at http://home.earthlink.net/~rmarken/demos.html.

3. Purpose in Psychology

The Hierarchical Behavior of Perception

This paper argues that the coincidental development of hierarchical models of perception and behavior is not a coincidence. Perception and behavior are two sides of the same phenomenon – control. A hierarchical control system model shows that evidence of hierarchical organization in behavior is also evidence of hierarchical organization in perception. Studies of the temporal limitations of behavior, for example, are shown to be consistent with studies of temporal limitations of perception. A surprising implication of the control model is that the perceptual limits are the basis of the behavioral limits. Action systems cannot produce controlled behavioral results faster than the rate at which these results can be perceived. Behavioral skill turns on the ability to control a hierarchy of perceptions, not actions.

Psychologists have developed hierarchical models of both perception (e.g. Bryan and Harter, 1899; Palmer, 1977; Povel, 1981) and behavior (e.g. Albus, 1981; Arbib, 1972; Greeno and Simon, 1974; Lashley, 1951; Martin, 1972; Keele, Cohen and Ivry, 1990; Rosenbaum, 1987). This could be a coincidence, a case of similar models being applied to two very different kinds of phenomena. On the other hand, it could reflect the existence of a common basis for both perception and behavior. This paper argues for the latter possibility, suggesting that perception and behavior are two sides of the same phenome-

---

From *Closed Loop*, 1993, 3, 33-54.

non: control (Marken, 1988). Control is the means by which agents keep perceived aspects of their external environment in goal states (Powers, 1973). It is argued that the existence of hierarchical models of both perception and behavior is a result of looking at control from two different perspectives; that of the agent doing the controlling (the actor) and that of the agent watching control (the observer). Depending on the perspective, control can be seen as a perceptual or a behavioral phenomenon.

From the actor's perspective, control is a perceptual phenomenon. The actor is controlling his or her own perceptual experience, making it behave as desired. However, from the observer's perspective, control is a behavioral phenomenon. The actor appears to be controlling variable aspects of his or her behavior in relation to the environment. For example, from the perspective of a typist (the actor), typing involves the control of a dynamically changing set of kinesthetic, auditory and, perhaps, visual perceptions. If there were no perceptions there would be no typing. However, from the perspective of someone watching the typist (the observer), perception is irrelevant; the typist appears to be controlling the movements of his or her fingers in relation to the keys on a keyboard.

These two views of control have one thing in common; in both cases, control is seen in the behavior of perception. For the actor, control is seen in the behavior of his or her own perceptions. For the observer, control is seen in the behavior of his or her own perceptions of the actor's actions. (The observer can see the means of control but can only infer their perceptual consequences as experienced by the actor). If control is hierarchical then it can be described as the behavior of a hierarchy of perceptions. Hierarchical models of perception and behavior can then be seen as attempts to describe control from two different perspectives, those of the actor and observer, respectively. This paper presents evidence that hierarchical models of perception and behavior reflect the hierarchical structure of control.

A Perceptual Control Hierarchy

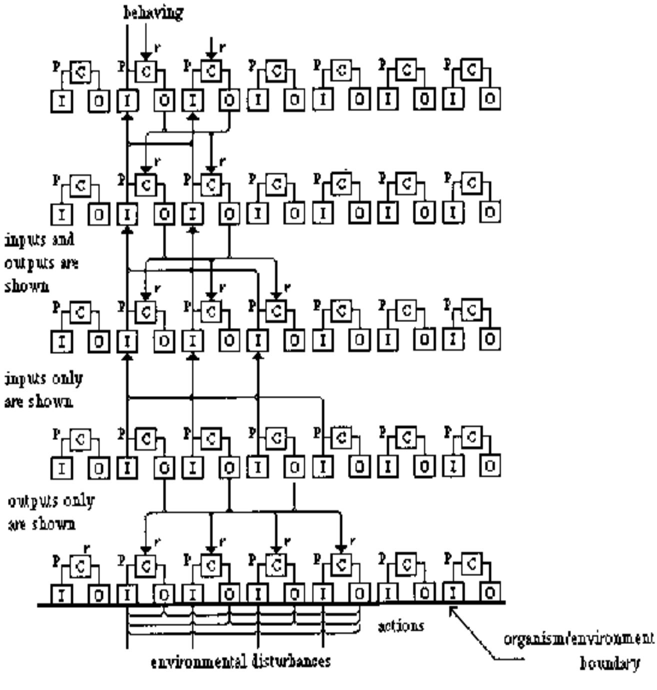

Figure 1. Hierarchy of perceptual control systems.

The concept of control as the behavior of perception can be understood in the context of a hierarchical control system model of behavioral organization (Powers, 1973;

1989). The model is shown in Figure 1. It consists of several levels of control systems (the figure shows four levels) with many control systems at each level (the figure shows seven). Each control system consists of an input transducer (I), comparator (C) and output transducer (O). The input transducer converts inputs from the environment or from systems lower in the hierarchy into a perceptual signal, p. The comparator computes the difference between the perceptual signal and a reference signal, r. The output transducer amplifies and converts this difference into actions which affect the environment or become reference signals for lower level systems.

The control systems at each level of the hierarchy control perceptions of different aspects of the external environment. However, all systems control perceptions in the same way; by producing actions that reduce the discrepancy between actual and intended perceptions. Intended perceptions are specified by the reference signals to the control systems. The actions of the control systems coax perceptual signals into a match with reference signals via direct or indirect effects on the external environment. The actions of the lowest level control systems affect perceptions directly through the environment. The actions of higher-level control systems affect perceptions indirectly by adjusting the reference inputs to lower level systems.

The hierarchy of control systems is a working model of purposeful behavior (Marken, 1986; 1990). The behavior of the hierarchy is purposeful inasmuch as each control system in the hierarchy works against any opposing forces in order to produce intended results. Opposing forces come from disturbances created by the environment as well as interfering effects caused by the actions of other control systems. The existence of disturbances means that a control system cannot reliably produce an intended result by selecting a particular action. Actions must vary to compensate for varying disturbances. Control systems solve this problem by specifying what results are to be perceived not how these results are to be achieved. Control systems control perceptions, not actions. When set up correctly the control systems in the hierarchy vary their actions as necessary, compensating for unpredictable (and, often, undetectable) disturbances, in order to produce intended perceptions. Indeed, the term "control" refers to this process of producing intended perceptions in a disturbance prone environment.

Levels of Perception

Powers (1990) has proposed that each level of the hierarchy of control systems controls a different class of perception. These classes represent progressively more abstract aspects of the external environment. The lowest level systems control perceptions that represent the intensity of environmental input. The next level controls sensations (such as colors), which are functions of several different intensities. Going up from sensations there is control of configurations (combinations of sensations), transitions (temporal changes in configurations), events (sequences of changing configurations), relationships (logical, statistical, or causal covariation between independent events), categories (class membership), sequences (unique orderings of lower order perceptions), programs (if-then contingencies between lower level perceptions), principles (a general rule that exists in the behavior of lower level perceptions) and system concepts (a particular set of principles exemplified by the states of many lower level perceptions; see Powers, 1989, pp. 190-208). These eleven classes of perception correspond to eleven levels of control systems in the hierarchical control model. All control systems at a particular level of the hierarchy control the same class of perception, though each system controls a slightly different exemplar of the class. Thus, all systems at a particular level may control configuration perceptions but each system controls a different configuration.

The rationale for hierarchical classes of perceptual control is based on the observation that certain types of perception depend on the existence of others. Higher level perceptions depend on (and, thus, are a function of) lower level perceptions. For example, the perception of a configuration, such as a face, depends on the existence of sensation (color) or intensity (black/white) perceptions. The face is a function of these sensations and intensities. The lower level perceptions are the independent variables in the function that computes the higher level perception. Their status as independent variables is confirmed by the fact that lower level perceptions can exist in the absence of the higher level perceptions, but not vice versa. Color and intensity perceptions can exist without the perception of a face (or any other configuration, for that matter) but there is no face without perceptions of intensity and/or color.

The Behavior of Perceptions. From the point of view of the hierarchical control model, “behaving” is a process of controlling perceptual experience. Any reasonably complex behavior involves the control of several levels of perception simultaneously. For example, when typing the word “hello”, one controlled perception is the sequence of letters “h”, “e”, “l”,”l” and “o”. The perception of this sequence is controlled by producing a sequence of key press event perceptions. Each key press event is controlled by producing a particular set of transitions between finger configuration perceptions. Each finger configuration is controlled by a different set of force sensations that are themselves controlled by producing different combinations of intensities of tensions in a set of muscles.

The perceptions involved in typing “hello” are all being controlled simultaneously. Transitions between finger configurations are being controlled while the force sensations that produce the configuration perceptions are being controlled. However, the typist is usually not aware of the behavior of all these levels of perception. People ordinarily attend to the behavior of their perceptions at a high level of abstraction, ignoring the details. We attend to the fact that we are driving down the road and ignore the changing muscle tensions, arm configurations and steering wheel movements that produce this result. Paying attention to the details leads to a deterioration of performance; it is the opposite of “Zen” behavior, where you just attend to the (perceptual) results that you intend to produce and let the required lower level perceptions take care of themselves (Herrigal, 1971). However, while it violates the principles of Zen, attention to the detailed perceptions involved in the production of behavioral results can provide interesting hints about the nature of the perceptual control hierarchy.

The Perception of Behavior. The behavior of an actor who is organized like the hierarchical control model consists of changes in the values of variables in the actor's environment. An observer cannot see what is going on inside the actor; he or she can only see the actor's actions and the effect of these actions on the external environment. The effect of these actions is to cause purposeful behavior of certain variables in the environment; the variables that correspond to perceptions that the actor is actually controlling. The purposefulness of the behavior of these variables is evidenced by the fact that consistent behaviors are produced in the context of randomly changing environmental disturbances. Thus, a typist can consistently type the word “hello” despite changes in the position of the fingers relative to the keyboard, variations in the push-back force of the keys or even a shift from one keyboard arrangement to another (from QWERTY to Dvorak, for example).

Since the actor controls his or her own perceptions, the observer cannot actually see what the actor is “doing”; the actor's “doings” consist of changing the intended states of his or her own perceptions. All the observer sees is variable results of the actor's actions; results that may or may not be under control. For example, the observer might notice that a click occurs each time the typist presses a key. The click is a result produced by the typist and the observer is likely to conclude that the typist is controlling the occurrence of the click. In fact, the click may be nothing more than a side effect of the typist's efforts to make the key feel like it has hit bottom. There are methods that make it possible for the observer to tell whether or not his or her perceptions of the actor's behavior correspond to the perceptions that are being controlled by the actor (Marken, 1989). These methods make it possible for the observer to determine what the actor is actually doing (i.e. controlling).

Hierarchical Control

The hierarchical nature of the processes that generate behavior would not be obvious to the observer of a hierarchical control system. The observer could tell that the system is controlling many variables simultaneously but he or she would find it difficult to demonstrate that some of these variables are being controlled in order to control others. For example, the observer could tell that a typist is controlling letter sequences, key press events, finger movements and finger configurations. But the observer would have a hard time showing that these variables are hierarchically related.

The observer could make up a plausible hierarchical description of these behaviors; for example, finger positions seem to be used to produce finger movements that are used to produce key presses that are used to produce letter sequences. But finding a hierarchical description of behavior does not prove that the behavior is actually produced by a hierarchical process (Davis, 1976; Kline, 1983).

Hierarchical Invariance

Hierarchical production of behavior implies that the commands required to produce a lower level behavior are nested within the commands required to produce a higher level behavior. For example, the commands that produce a particular finger configuration would be nested within the commands that produce a movement from one configuration to another. Sternberg, Knoll and Turlock (1990) refer to this nesting as an invariance property of hierarchical control. Lower level commands are like a subprogram that is invoked by a program of higher level commands. The invariance of hierarchical control refers to the assumption that the course of such a subprogram does not depend on how it was invoked from the program (low level invariance); similarly, the course of the program does not depend on the nature of the commands carried out by the subprograms (high level invariance).

Convergent and Divergent Control. The hierarchical control model satisfies both the low and high-level invariance properties of hierarchical control. The commands issued by higher level systems have no effect on the commands issued by lower level systems and vice versa. It is important to remember, however, that the commands in the control hierarchy are requests for input, not output. Higher level systems tell lower level systems what to perceive not what to do. This aspect of control system operation solves a problem that is either ignored or glossed over in most hierarchical models of behavior: How does a high level command get turned into the lower level commands that produce results that satisfy the high level command? If commands specify outputs then the result of the same command is always different due to varying environmental disturbances. The high level command to press a key, for example, cannot know which lower level outputs will produce this result on different occasions. This problem is solved by the hierarchical control model because intended results are represented as a convergent function rather than a divergent network.

Most hierarchical models of behavior require that a high level command be decomposed into the many lower level commands that produce the intended result. In the hierarchical control model, both the high level command and the intended result of the command are represented by a single, unidimensional signal. The signal that represents the intended result is a function of results produced by many lower level commands. But the high level command does not need to be decomposed into all the appropriate lower level commands (Powers, 1979). The difference between the high level command and the perceptual result of that command is sufficient to produce the lower level commands that keep the perceptual result at the commanded value (Marken, 1990).

Levels of Behavior

The hierarchical invariance properties of the control hierarchy provide a basis for determining whether its behavior is actually generated by hierarchical processes. Hierarchical control can be seen in the relative timing of control actions. In a control hierarchy, lower level systems must operate faster than higher level systems. Higher level systems cannot produce a complex perceptual result before the lower level systems have produced the component perceptions on which it depends. This nesting of control actions can be seen in the differential speed of operation of control systems at different levels of the control hierarchy. Lower level systems not only correct for disturbances faster than higher level ones; they carry out this correction process during the higher-level correction process. The lower level control process is temporally nested within the higher-level control process.

Arm Movement. Powers, Clark and McFarland (1960) describe a simple demonstration of nested control based on relative timing of control system operation. A subject holds one hand extended straight ahead while the experimenter maintains a light downward pressure on it. The subject is to move his or her arm downward as quickly as possible when the experimenter signals with a brief, downward push on the subject's extended hand. The result of this simple experiment is always the same: the subject responds to the downward signal push with a brief upward push followed by downward movement of the arm. An electromyograph shows that the initial upward push is an active response and not the result of muscle elasticity.

The arm movement demonstration reveals one level of control nested within another. The subject's initial upward push (which cannot be suppressed) is the fast response of a lower level control system that is maintaining the perception of arm position in a particular reference state (extended forward). The behavior of this system is nested within the response time of a higher level system that moves the arm downward. The higher level system operates by changing the reference for the arm position control system. The downward signal push causes the brief upward reaction because the signal is treated as a disturbance to arm position. This is particularly interesting because the signal is pushing the arm in the direction it should move; the lower level reaction is “counter productive” with respect to the goal of the higher level system (which wants to perceive the arm down at the side). The reaction occurs because the lower level system starts pushing against the disturbance to arm position before the higher level system can start changing the reference for this position.

Polarity Reversal. More precise tests of nested control were carried out in a series of experiments by Marken and Powers (1989). In one of these experiments, subjects performed a standard pursuit-tracking task, using a mouse controller to keep a cursor aligned with a moving target. At intervals during the experiment the polarity of the connection between mouse and cursor movement was reversed in a way that did not disturb the cursor position. Mouse movements that had moved the cursor to the right now moved it to the left; mouse movements that had moved the cursor to the left now moved it to the right.

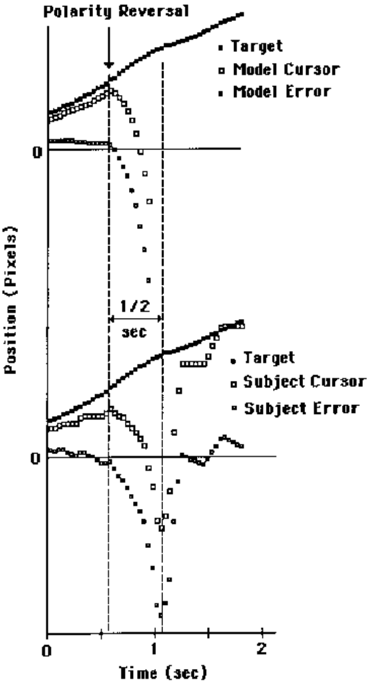

Figure 2. Hierarchical response to polarity change.

A sample of the behavior that occurs in the vicinity of a polarity reversal is shown in Figure 2. The upper traces show the behavior of a control system model and the lower traces show the behavior of a human subject. When the reversal occurs, both the model and the subject respond to error (the deviation of the cursor from the target) in the wrong direction, making it larger instead of smaller (any deviation of the error trace from the zero line represents an increase in error).

The larger error leads to faster mouse movement that causes the error to increase still more rapidly. A runaway condition ensues with error increasing exponentially.

About 1/2 second after the polarity reversal the subject's behavior departs abruptly from that of the model. The subject adjusts to the polarity reversal and the error returns to a small value. The model cannot alter its characteristics and the error trace quickly goes off the graph. These results provide evidence of two nested levels of control operating at different speeds. The faster, lower level system controls the distance between cursor and target. This system continues to operate as usual even when, due to the polarity reversal, this causes an increase in perceptual error. Normal operation is restored only after a slower, higher level system has time to control the relationship between mouse and cursor movement.

Levels of Perception

The arm movement and polarity shift experiments reveal the hierarchical organization of control from the point of view of the observer. The hierarchical control model suggests that it should also be possible to view hierarchical organization from the point of view of the actor. From the actor's point of view, hierarchical control would be seen as a hierarchy of changing perceptions. One way to get a look at this hierarchy is again in terms of relative timing; in this case, however, in terms of the relative timing of the perceptual results of control actions rather of the actions themselves.

Computation Time Window. The hierarchical control model represents the results of control actions as unidimensional perceptual signals. A configuration, such as the letter “h”, is a possible result of control actions, as is a sequence of letters, such as the word “hello”. The model represents these results as perceptual input signals, the intensity of a signal being proportional to the degree to which a particular result is produced. This concept is consistent with the physiological work of Hubel and Wiesel (1979) who found that the firing rate of an afferent neuron is proportional to the degree to which a particular environmental event occurs in the “receptive field” of the neuron.

Many of the higher level classes of perception in the control hierarchy depend on environmental events that vary over time. Examples are transitions, events, and sequences. The neural signals that represent these variables must integrate several lower level perceptual signals that occur at different times. Hubel and Wiesel found evidence of a computation time window for integrating perceptual signals. Certain cells respond maximally to configurations (such as “lines”) that move across a particular area of the retina at a particular rate. These are “motion detector” neurons. The neuron responds maximally to movement of a configuration that occurs within a particular time window. Movement that occurs outside of this time window is not included in the computation of the perceptual signal that represents motion.

Levels by Time. The hierarchical control model implies that the duration of the computation time window increases as you go up the hierarchy. The minimum computation time window for the perception of configurations should be shorter than the minimum computation time window for the perception of transitions, which should be shorter than the minimum computation time window for the perception of sequences. I have developed a version of the psychophysical method of adjustment, which makes it possible to see at least four distinct levels of perception by varying the rate at which items occur on a computer display. A computer program presents a sequence of numbers at two different positions on the display. The presentation positions are vertically adjacent and horizontally separated by 2 cm. The numbers are presented alternately to the two positions. The subject can adjust the rate at which the numbers occur in each position by varying the position of a mouse controller 1.

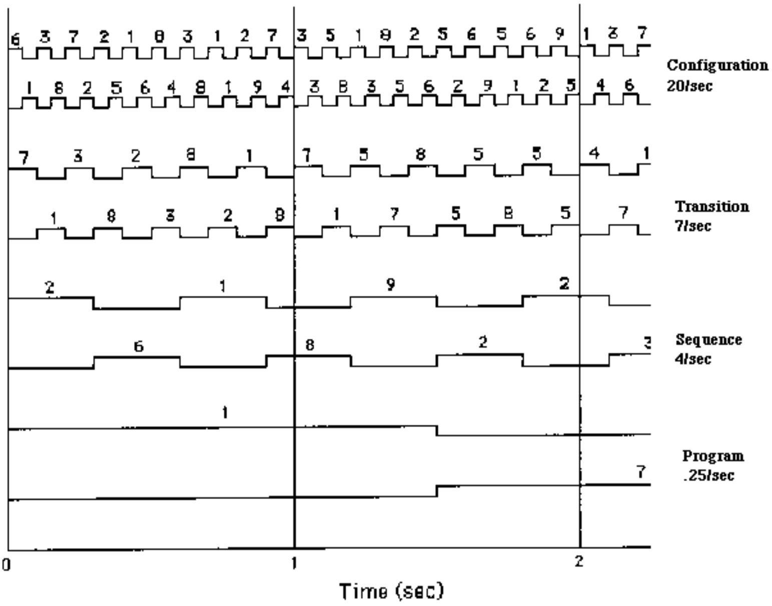

Figure 3. Temporal constraints on pattern perception.

The results of this study are shown schematically in Figure 3. At the fastest rate of number presentation subjects report that the numbers appear to occur in two simultaneous streams. The fact that the numbers are presented to the two positions alternately is completely undetectable. However, even at the fastest rate of number presentation subjects can make out the individual numbers in each stream. At the fastest rate, there are approximately 20 numbers per second in each stream. This means that there is a 50 msec period available for detecting each number. This duration is apparently sufficient for number recognition suggesting that the computation time window for perception of configuration is less than 50 msec. Studies of the “span of apprehension” for sets of letters suggest that the duration of the computation time window for perception of visual configuration may be even less than 50 msec, possibly as short as 15 msec (Sperling, 1960).

As the rate of number presentation slows, the alternation between numbers in the two positions becomes apparent. Subjects report perception of alternation or movement between numbers in the two positions when the numbers in each stream are presented at the rate of about 7 per second. At this rate, an alternation from a number in one stream to a number in another occurs in 160 msec. This duration is sufficient for perception of the alternation as a transition or movement from one position to the other suggesting that the computation time window for transition perception is on the order of 160 msec. This duration is compatible with estimates of the time to experience optimal apparent motion when configurations are alternately presented in two different positions (Kolers, 1972).

The numbers presented in each stream are always changing. However, subjects find it impossible to perceive the order of the numbers as they alternate from one position to another even though it is possible to clearly perceive the individual numbers and the fact that they are alternating and changing across positions. The rate of number presentation must be slowed considerably, so that each stream of numbers is presented at the rate of about two per second, before it is possible to perceive the order in which the numbers occur. At this rate numbers in the sequence occur at the rate of four per second. These results suggest that the duration of computation time window for the perception of sequence is about 0.5 seconds. This is the time it takes for two elements of the sequence to occur: the minimum number that can constitute a sequence.

The numbers in the rate adjustment study did not occur in a fixed, repeating sequence. Rather, they were generated by a set of rules: a program. The sequence of numbers was unpredictable unless the subject could perceive the rule underlying the sequence. The rule was as follows: if the number on the right was even then the number on the left was greater than 5, otherwise the number on the left was less than 5. (Numbers in the sequence were also constrained to be between 0 and 9). Subjects could not perceive the program underlying the sequence of numbers until the speed of the two streams of numbers was about .25 numbers per second so that the numbers in the program occurred once every two seconds. The perception of a program in a sequence of numbers requires considerably more time then it takes to perceive the order of numbers in the same sequence.

The perception of a sequence or a program seems to involve more mental effort than the perception of a configuration or a transition. Higher level perceptions, like programs, seem to represent subjective rather than objective aspects of external reality; they seem more like interpretations than representations. These higher level perceptions are typically called “cognitions”. Of course, all perceptions represent subjective aspects of whatever is “out there”; from the point of view of the hierarchical control model, the location of the line separating perceptual from cognitive representations of reality is rather arbitrary. Behavior is the control of perceptions which range from the simple (intensities) to the complex (programs).

Perceptual Speed Limits. The hierarchical control model says that all perceptions of a particular type are controlled by systems at the same level in the hierarchy. This implies that the speed limit for a particular type of perception should be about the same for all perceptions of that type. The 160 msec computation time window for perception of transition, for example, should apply to both visual and auditory transition. There is evidence that supports this proposition. Miller & Heise (1950) studied the ability to perceive an auditory transition called a “trill”. A trill is the perception of a temporal alternation from one sound sensation or configuration to another. The speed limit for trill perception is nearly the same as the speed limit for visual transition perception found in the number rate adjustment study – about 15 per second. As in the visual case, when the rate of alternation of the elements of the auditory trill exceeds the computation time window the elements “break” into two simultaneous streams of sound; the perception of transition (trill) disappears even though the sounds continue to alternate.

There is also evidence that the four per second speed limit for sequence perception found in the number rate adjustment study applies across sensory modalities. Warren, Obusek, Farmer, & Warren (1969) studied subjects' ability to determine the order of the component sounds in a sound sequence. They found that subjects could not perceive the order of the components until the rate of presentation of the sequence was less than or equal to four per second. This was a surprising result because it is well known that people can discriminate sequences of sounds that occur at rates much faster than four per second. In words, for example, the duration of the typical phoneme is 80 msec so people can discriminate sequences of phoneme sounds that occur at the rate of about 10 phonemes per second. But there is reason to believe that the phonemes in a word are not heard as a sequence; that is, the order of the phonemes cannot be perceived. Warren (1974) showed that subjects can learn to tell the difference between sequences of unrelated sounds that occur at rates of 10 per second. However, the subjects could not report the order of the sounds in each sequence; only that one sound event differed from another. A word seems to be a lower order perception – an event perception – that is recognized on the basis of its overall sound pattern. There is no need to perceive the order in which the phonemes occur; just that the temporal pattern of phonemes (sound configurations) for one word differs from that for other words.

The Relationship between Behavior and Perception

Configurations, transitions, events, sequences and programs are potentially controllable perceptions. An actor can produce a desired sequence of sounds, for example, by speaking sound events (phonemes) in some order. An observer will see the production of this sequence as a behavior of the actor. The hierarchical control model suggests that the actor's ability to produce this behavior turns on his or her ability to perceive the intended result. Since perception depends on speed, it should be impossible for the actor to produce an intended result faster than the result can be perceived. The observer will see this speed limit as a behavioral limit. An example of this can be seen in the arm movement experiment described above. In that experiment it appears that the time to respond to the signal push is a result of a behavioral speed limit; the inability to generate an output faster than a certain rate. But a closer look indicates that the neuromuscular “output” system is perfectly capable of responding to a signal push almost immediately, as evidenced by the immediate upward response to the downward signal push. The same muscles that produce this immediate reaction must wait to produce the perception of the arm moving downward. The speed limit is not in the muscles. It is in the results that the muscles are asked to produce; a static position of the arm (a configuration perception) or a movement of the arm in response to the signal push (a relationship perception).

Sequence Production and Perception. Some of the most interesting things people do involve the production of a sequence of behaviors. Some recent studies of temporal aspects of sequence production are directly relevant to the hierarchical control model. In one study, Rosenbaum (1989) asked subjects to speak the first letters of the alphabet as quickly as possible. When speed of letter production exceeded four per second the number of errors (producing letters out of sequence) increased dramatically, indicating a loss of control of the sequence. The speed limit for sequence production corresponds to the speed limit for sequence perception – four per second.

The letter sequence study does not prove that the speed limit for letter sequence production is caused by the speed limit for letter sequence perception. It may be that the speed limit is imposed by characteristics of the vocal apparatus. However, in another study Rosenbaum (1987) found the same four per second speed limit for production of errorless finger tap sequences. The speed limit for finger tap sequence production is likely to be a perceptual rather than a motor limit because we know that people can produce finger taps at rates much higher than four per second. Pianists, for example, can do trills (alternating finger taps) at rates that are far faster than four per second. Further evidence of the perceptual basis of the finger tap sequence speed limit would be provided by studies of finger tap sequence perception. When a subject produces a sequence of finger taps he or she is producing a sequence of perceptions of pressure at the fingertips. A perceptual experiment where a pressure is applied to the tip of different fingers in sequence should show the four per second speed limit. Subjects should have difficulty identifying the order of fingertip pressures when the sequence occurs at a rate faster than four per second.

Confounding Levels. It is not always easy to find clear-cut cases of behavioral speed limits that correspond to equivalent perceptual speed limits. Most behavior involves the control of many levels of perception simultaneously. People control higher level perceptions (like sequences) while they are controlling lower level perceptions (like transitions). This can lead to problems when interpreting behavioral speed limits. For example, Rosenbaum (1983) presents some finger tapping results that seem to violate the four per second speed limit for sequence perception. When subjects tap with two hands they can produce a sequence of at least 8 finger taps per second. But each tap is not necessarily a separate event in a sequence. Some pairs of taps seem to occur at the rate at which sequences are experienced as events. A sequence of finger taps is an event in the same sense that the sequence of muscle tensions that produce a finger tap is an event; the order of the components of the sequence cannot be perceived. These finger tap events are then unitary components of the sequence of finger tap perceptions.

The fact that certain pairs of finger taps are produced as events rather than ordered sequences is suggested by the errors made at each point in the finger tap sequence. Errors occur most frequently at the point in the sequence at which a fast pair is being initiated. Errors rarely occur for the second element of a fast pair. This suggests that the errors occur at the sequence level rather than the event level. The subject's attempts to produce a key press sequence too rapidly apparently interferes with sequence rather than event production. Events are already produced at a fast enough rate and an increase in the speed of sequence production has little effect on the ability to control the component events.

Changing Perception Can Change Behavior: "Going Up A Level". The relationship between perception and behavior can be seen when a person learns to perform a task by controlling a new perceptual variable. An example of this can be seen in simple pursuit tracking tasks. In the typical tracking task the target moves randomly. When, however, a segment of target movement is repeated regularly the subject's tracking performance improves markedly with respect to that segment (Pew, 1966). According to the hierarchical control model, control is improved because the repeated segment of target movement can be perceived as a predictable event. With the random target the subject must wait to determine target position at each instant in order to keep the cursor on target. With the repeated target, the subject controls at a higher level, keeping a cursor movement event matching a target movement event. The fact that the subject is now controlling a higher level perception (an event rather than a configuration) is evidenced by the longer reaction time when responding to a change in target movement. When controlling the target-cursor configuration the subject responds almost immediately to changes in target position. When controlling target-cursor movement events it takes nearly 1/2 second to respond to a change to the same change in target movement pattern.

An experiment by Robertson and Glines (1985) also shows improved performance resulting from changed perception. Subjects in the Robertson and Glines study performed a learning task where the solution to a computerized game could be perceived at several different levels. Subjects who were able to solve the game showed three distinct plateaus in their performance. The level of performance, as indicated by reaction time measurements, improved at each succeeding plateau. Because the same outputs (key presses) were produced at each level of performance, each performance plateau was taken as evidence that the subject was controlling a different perceptual variable.

Behavior/Perception Correlations. Few psychologists would be surprised by the main contention of this paper: that there is an intimate relationship between perception and behavior. However, most models of behavior assume that the nature of this relationship is causal: behavior is guided by perception. This causal model provides no reason to expect a relationship between the structure of perception and behavior: no more than there is to expect a relationship between the structure of computer input and output. This does not mean that there might not be such a relationship; it is just not demanded by the causal model.

The control model integrates perception and behavior with a vengeance. Behavior is no longer an output but, rather, a perceptual input created by the combined effects of the actor and the environment. Behavior is perception in action. From this point of view, behavioral skills are perceptual skills. Thus, it is not surprising to find some indication of a correlation between behavioral and perceptual ability. For example, Keele and his colleagues (Keele, Pokorny, Corcos and Ivry, 1985) have found that the ability to produce regular time intervals between actions is correlated with ability to perceive these intervals. These correlations were fairly low by control theory standards but they are expected if the production of regular time intervals involves control of the perception of these intervals.

Conclusion

This report has presented evidence that human behavior involves control of a hierarchy of perceptual variables. There is evidence that the behavior of non-human agents, such as chimpanzees, also involves the control of a similar hierarchy of perceptions (Plooij and van de Rijt-Plooij, 1990). A model of hierarchical control shows how studies of perception and behavior provide evidence about the nature of control from two different perspectives. Perceptual studies provide information about the ability to perceive potentially controllable consequences of actions. Behavioral studies provide information about the ability to produce desired consequences. The factors that influence the ability to perceive the consequences of action should also influence the ability to produce them. In both cases we learn something about how agents control their own perceptions.

The hierarchical control model shows that limitations on the ability to produce behavior may reflect limitations on the ability to perceive intended results. The speed at which a person can produce an errorless sequence of events, for example, is limited by the speed at which the order of these events can be perceived. But not all skill limitations are perceptual limitations. Controlled (perceived) results are produced, in part, by the outputs of the behaving agent. The ability to produce certain outputs can limit the ability to control certain perceptions. For example, it is impossible to perceive oneself lifting a 300 pound barbell until the muscles have been developed to the point that they are able to generate the output forces necessary to control this perception.

Perception and behavior are typically treated as two completely different types of phenomena. Perception is a sensory phenomenon: behavior is a physical phenomenon. But the concept of control as the behavior of perception suggests that this separation is artificial. Perception and behavior are the same phenomenon seen from two different perspectives. In order to understand how this phenomenon works, it will be necessary to understand how agents perceive (perception) and how they act to affect their perceptions (behavior). Studies of perception and behavior should become an integral part of the study of a single phenomenon: control.

Controlled Variables: Psychology as the Center Fielder Views It

Perceptual Control Theory (PCT) views behavior as the control of perception. The central explanatory concept in PCT is the *controlled variable*, which is a perceived aspect of the environment that is brought to and maintained in states specified by the organism itself. According to PCT, understanding behavior is a matter of discovering the variables that organisms control. But the possible existence of controlled variables has been largely ignored in the behavioral sciences. One notable exception occurs in the study of how baseball outfielders catch fly balls. In these studies it is taken for granted that the fielder gets to the ball by controlling some visual aspect of the ball's movement. This paper describes the concept of a controlled variable in the context of research on fly ball catching behavior and shows how this concept can contribute to our understanding of behavior in general.

The publication of John B. Watson's (1913) *Psychology as the Behaviorist Views It* signaled the beginning of an era of psychological research dominated by the search for *controlling variables*; the variables that control behavior. Behavioral psychologists started looking for these variables in the organism's environment. Cognitive psychologists are now looking for these variables in the organism's mind (or brain). But in both cases the search is for controlling vari-

From *American Journal of Psychology*, 2001, 114(2), 100-110.
Reprinted with permission of University of Illinois Press

ables; the variables that cause organisms to behave as they do. This preoccupation with controlling variables may be one reason psychologists have paid so little attention to a theory of behavior that focuses on *controlled variables*; the variables that are controlled *by* behavior. The theory is called Perceptual Control Theory (PCT) and its basic assumption is that behavior is organized around the control of perceptual variables (Powers, 1973a).

Purpose and Control

PCT was developed to explain the purposeful behavior of organisms (Marken, 1990). Purposeful behavior involves the production of consistent results in a world where unpredictable disturbances make such consistency highly unlikely (Old Faithful notwithstanding). For example, a person sipping tea is producing a consistent result – the sips – despite unpredictable disturbances, such as head movements, that change the relative location of cup and lips. Sipping tea is a purposeful behavior.

PCT is based on the realization that purposeful behavior, like that of the tea drinker, is equivalent to the *controlling* done by artificial *control systems*, such as a thermostat. In both cases, a consistent result is produced despite unpredictable disturbances that should produce inconsistency. The thermostat produces a consistent room temperature despite unpredictable changes in outdoor air temperature; the tea drinker produces consistent sips despite unpredictable changes in the relative location of cup and lips.

Control systems act to bring variable aspects of the environment to pre-selected states while protecting these variables from the effects of disturbance. This process is called *control*. The variable aspects of the environment that a control system controls are called *controlled variables*; room temperature and distance from cup to lips are controlled variables. The purposeful behavior of a control system, whether it is living (like the tea drinker) or artificial (like the thermostat), is organized around controlled variables. Control theory explains how a control system acts to keep these variables under control. PCT is the application of control theory to understanding the purposeful behavior of control systems in general and living control systems in particular.

Control Theory

Control theory describes the organization of systems that can control variables like room temperature or the distance from cup to lips. A basic control system is shown in Figure 1. The upper part of the figure represents the control system itself; the lower part of the figure represents the system's environment. The most important aspect of the environment is the variable controlled by the control system: the controlled variable (q_i) The thermostat's controlled variable is room temperature; the tea drinker's controlled variable is distance from cup to lips.

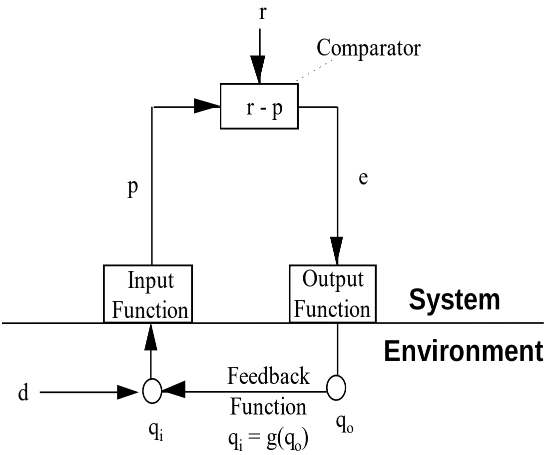

Figure 1. A basic control system

The state of the controlled variable is influenced by other variables in the environment, which include independent influences on q_i , called disturbances (d), and the actions of the control system itself, called outputs (q_o). Room temperature is influenced by disturbances, such as variations in outdoor air temperature and by the outputs of the thermostat itself – the variations in the amount of heat generated by the furnace. The distance from cup to lips is influenced by disturbances, such as the variations in the location of the mouth, and by the outputs of the tea drinker herself – the forces exerted on the cup as it is lifted to the lips.

The controlled variable is represented inside the control system as a perceptual signal (p). This perceptual signal is the output of a transducer, called the *input function,* which converts variable aspects of the control system's environment into signals inside the control system itself. The perceptual signal is continuously compared to a reference signal (r) that specifies the target or intended value of the perceptual signal. The *comparator* computes the difference between the reference and perceptual signals (r-p); this difference is the error signal (e). The error signal causes, via the *output function*, the outputs that have effects on the controlled variable.

The variables in a control system trace out a *closed loop* of cause and effect. Variables in any part of this loop have effects that *feedback* on themselves. Moreover, the feedback effects of variables in this loop tend to cancel themselves out; a process called *negative feedback*. The result of this negative feedback process is that the perceptual signal is bought to and maintained at a fixed or variable *reference state* by the outputs of the system, protected from the effects of disturbance. So the behavior of a negative feedback control system is properly described as *control of perception* (Powers, 1973a). When a perception is controlled, the environmental correlate of that perception – the controlled variable – is also controlled.

Controlled Variables and Behavior

A control system can produce some rather complex-looking behavior. For example, the thermostat turns the furnace on and off for varying amounts of time, producing a complex pattern of "furnace actuating" behavior. One approach to understanding this behavior is to try to discover its *causes*. But control system behavior cannot be understood in cause-effect terms because cause-effect relationships ignore the possible existence of variables that the system might be controlling (Marken, 1993; Powers, 1973b). For example, the apparent cause-effect relationship between window opening and furnace actuating does not reveal the fact that the thermostat is controlling room temperature. The same cause-effect relationship would be seen if the thermostat were controlling some other variable, such as relative humidity. Indeed, this cause-effect relationship would be seen even if the thermostat were controlling nothing at all; furnace actuating behavior could be caused by a switch that is turned "on" when the window is opened and "off" when it is closed.

Control system behavior can be properly understood only in terms of the variables the system controls: controlled variables. Once you know that the system is controlling a particular variable you can predict its behavior with great accuracy. Much of the apparent complexity of control system behavior results from the fact that the system's outputs *mirror* the effects of disturbances to the controlled variable (Powers, 1978). Complex behavior will be seen in environments where disturbances produce complex effects on controlled variables. For example, a complex pattern of "furnace actuating" behavior is seen in an environment where disturbances (such as windows opening and closing, people entering and leaving the room, etc.) produce a complex pattern of effects on the variable the thermostat is controlling: room temperature. The discovery of controlled variables (such as room temperature) can, therefore, provide a simple and elegant explanation of what may appear to be very complex behavior.

Noticing Controlled Variables

Controlled variables are the central feature of purposeful behavior but they have gone largely unnoticed in the behavioral sciences. This may be because controlled variables are difficult to notice under ordinary circumstances. For example, if you did not already know that a thermostat controls room temperature it would be difficult to notice that room temperature is being controlled by the thermostat's behavior. What is noticed is the thermostat's reaction to stimuli like the cold blast of air from the opened window. What is not noticed is the controlled variable itself, the room temperature, which is relatively unaffected by the cold air.

Controlled variables are hard to notice under ordinary circumstances precisely because these variables *do not* react to stimuli (disturbances). It is harder to notice something that does not happen (such as the almost non-existent response of the controlled variable to disturbances) than something that does (such as the control system's marked response to any disturbance to the controlled variable, a response that prevents the disturbance from having much effect on the controlled variable).

Although it is hard to notice controlled variables, it is not impossible. A good approach to noticing controlled variables is to try to look at behavior from the point of view of the control system itself. For example, if you did not know that the thermostat was controlling room temperature you might be able to figure it out by asking yourself what the thermostat might be trying to perceive by turning the furnace on shortly after the window is opened and off shortly after it is closed. Putting yourself in the thermostat's shoes (or housing) might help you realize that the thermostat is trying to perceive a constant room temperature; room temperature is a controlled variable.

Similarly, you could put yourself into the paws of Pavlov's dog and ask what you might be trying to perceive by salivating when dry food is placed in your mouth. Perhaps you are trying to feel food that is wet, smooth and easy to swallow rather than food that is dry, sticky and impossible to swallow; the texture of the food might be a controlled variable. Or you could put yourself behind the nose of Skinner's rat and ask what you might be trying to perceive by quickly repeating the bar press that just netted you a food pellet. Perhaps you are trying to perceive food pellets arriving as quickly as possible; the rate of food pellet delivery might be a controlled variable.

The Test for the Controlled Variable (TCV)

When behavior is viewed in terms of controlled (rather than controlling) variables an important question immediately presents itself: How do you determine whether or not a variable that *seems* to be under control actually is under control? How do you know, for example, whether or not the thermostat actually is controlling room temperature? How do you know whether or not Pavlov's dog actually is controlling the texture of the food placed in its mouth? How do you know whether or not Skinner's rats are controlling the rate of food pellet delivery? The problem, of course, is that the variables that are being controlled are *perceptual* variables. In order to know what a control system is controlling it seems that you would have to be able to perceive what the control system is perceiving, which is obviously impossible.

Fortunately, there is a very simple procedure, based on PCT, that can be used to determine whether or not a variable that can be perceived by an observer corresponds to a variable that is being controlled by a control system (Marken, 1997). The procedure, called the Test for the Controlled Variable (TCV), involves applying a disturbance to a possible controlled variable and looking for *lack of effect* of the disturbance. For example, we can test to determine whether or not a thermostat is controlling room temperature by applying a disturbance, such as a blast of cold air, and looking to see if it has the expected effect – a lowering of the room temperature. If room temperature, which is perceived by the observer as the reading of a thermometer, is under control, the disturbance will have little or no effect – the room temperature reading stays about the same.

The same type of test can be used to determine whether or not a tea drinker is controlling the distance from cup to lips. The test is done by applying a disturbance, such as a gentle push on the cup, and looking to see if it has the expected effect – increasing the distance between cup and lips. If distance between cup and lips is under control, the disturbance will have little or no effect – the distance between cup and lips stays about the same.

A properly conducted TCV involves the application of *many* different disturbances, all of which *would* have an effect on the hypothetical controlled variable if that variable were *not* under control. While it is impossible to prove that a variable is unquestionably under control, it is possible to conduct tests until one becomes very confident that the variable is under control. If *every* disturbance that should have an effect on the variable doesn't, then one can be almost certain that the variable is under control.

The View from Center Field

The PCT approach to understanding behavior, which is based on the TCV, is rarely seen in psychological research. A notable exception occurs in research aimed at determining how baseball outfielders catch fly balls. The behavior under study is quite familiar; when the ball is hit in the air the outfielder runs to the spot where the ball will land and (usually) catches it. The conventional approach to understanding this behavior would be aimed at finding its causes; it would try to answer the question "what variables guide the fielder to the spot where the ball lands?" The PCT approach to understanding fly ball catching behavior is aimed at finding controlled variables; it tries to answer the question "what variables, if controlled, would result in our seeing the fielder move to the spot where the ball lands?"

Chapman (1968) proposed that an outfielder can get to the spot where a fly ball lands by running "...so as to maintain a constant speed of increase of tan α [the tangent of the optical angle of the ball relative to home plate as seen by the fielder]" (p. 870). Though he did not describe it this way, Chapman was actually proposing a hypothesis about a perceptual variable that the fielder might be controlling in order to get to the spot where the ball can be caught.

Chapman's hypothesis was that the rate of change in tan α (a variable called "optical velocity") is a controlled variable. This was an ingenious proposal, based on the observation that optical velocity is constant (and positive) when a fly ball is hit directly to the fielder. The fielder's behavior (running towards or away from the ball as it flies through the air) is, according to Chapman, a side effect of the control of perception; the perception being controlled is the rate of change of the projection of the image of the ball on the eye.

Chapman was making a guess about what fly ball catching behavior might look like from the fielder's, as opposed to the observer's, perspective. From the observer's perspective, fly ball catching looks like a pattern of running movements that bring the fielder to the spot where the ball lands. From the fielder's perspective, according to Chapman, fly ball catching looks like a ball that is rising at a constant rate (constant positive optical velocity).

A proper test of Chapman's hypothesis requires use of the TCV, which involves applying disturbances to the hypothetical controlled variable and looking to see whether or not these disturbances have an effect. If optical velocity is, indeed, under control then disturbances will be seen to have little effect; optical velocity will remain nearly constant despite disturbances that should cause it to change. In order to perform such a test it is necessary to apply disturbances to optical velocity *while* monitoring this variable to see whether or not these disturbances have an effect.

It is possible to disturb optical velocity by hitting fly balls in different trajectories relative to the outfielder. Each trajectory will result in a constantly changing optical velocity unless the fielder does something (runs in the appropriate direction) to keep optical velocity constant. Most tests of Chapman's hypothesis do use different fly ball trajectories as disturbances but nearly all of these tests have failed to apply these disturbances *while* monitoring the hypothesized controlled variable itself – optical velocity. The first test of Chapman's hypothesis where the hypothetical controlled variable was directly observed while disturbances were applied was done by McBeath, Shaffer and Kaiser (1995).

Optical Velocity, Optical Acceleration and LOT

McBeath et al. hit fly balls to a fielder who caught them while carrying a shoulder mounted video camera. The videotaped view of the fly balls provided a record of what the fielder saw while running to catch the ball. The record is a plot of the optical position of the ball at 1/30-second intervals during each catch. An analysis of this record shows that optical velocity (size of the change in the optical elevation of the ball relative to home plate during each interval) remains relatively constant regardless of the trajectory of the fly ball. So optical velocity passes the TCV; this hypothetical controlled variable remains nearly constant despite disturbances (the different fly ball trajectories) that should cause it to change.

Although optical velocity passes the TCV, it is not necessarily the variable fielders control when catching fly balls. There are alternatives that are also consistent with the data. McBeath et al. discovered one of these alternatives when they noticed another constancy in the videotape record: the linearity of the optical pattern traced out by the ball during each catch (see McBeath et al., 1995; Figure 4, p. 572). The plot of the optical elevation of the ball for each fly ball trajectory was always close to being a straight line, which McBeath et al. called a linear optical trajectory (LOT). LOT also passes the TCV since this variable remains nearly constant (a straight line) despite disturbances that should cause it to change (curve away from a straight line).

There is at least one other possibility regarding the variable fielders control when catching fly balls: optical acceleration. Many studies of fly ball catching behavior are based on the hypothesis that fielders control the rate of change in optical velocity (e.g. Babler and Dannemiller, 1993; Peper, Mastre, and Bakker, 1994; Dienes and McLoed, 1993). These studies assume that fielders act to cancel (zero out) variations in the optical acceleration of the fly ball seen by the fielder. Thus, the hypothesis that fielder's control optical acceleration has been called the optical acceleration cancellation (OAC) hypothesis.

Optical acceleration is the mathematical derivative of optical velocity so one might imagine that controlling optical acceleration is equivalent to controlling optical velocity. Indeed, keeping optical acceleration at zero is equivalent to keeping optical velocity constant. But an infinite number of different constant velocities (corresponding to the different possible values of the constant of integration that is part of the integral that transforms acceleration into velocity) are consistent with an acceleration of zero. So control of optical acceleration (keeping acceleration at zero for *all* trajectories) is not equivalent to control of optical velocity (keeping velocity at the same constant value for all trajectories); optical acceleration and optical velocity represent two different hypotheses about the variable fielders control when catching fly balls. The McBeath et al. results are consistent with both of these hypothesis; the hypothesis that fielders maintain optical velocity at some constant, non-zero value and the hypothesis that they maintain optical acceleration at zero.

Thus, there are three plausible hypotheses about the variable fielders might be controlling when they catch fly balls: optical velocity, optical acceleration and linear optical trajectory (LOT). The videotape records from the McBeath et al. study do not rule out any of the hypotheses about the controlled variable in fly ball catching behavior. There is evidence that optical velocity and acceleration remain constant and that LOT remains straight despite disturbances (different fly ball trajectories) that should have an effect on these variables. These hypotheses represent three different views of the *same* behavior (fly ball catching) from the point of view of the fielder. When we watch a fielder catch a fly ball we are either watching a side-effect of the fielder's efforts to keep optical velocity at some reference speed, to keep optical acceleration at zero or to keep the optical projection of the ball falling on a straight line (LOT). We could even be watching the fielder control some combination of these variables (McBeath, Shaffer and Kaiser, 1996).

Choosing Between Different Hypotheses about Controlled Variables

What is needed is a way to choose which one of several different hypotheses about the controlled variable is the best representation of the variable that is actually under control. The most straightforward approach is to systematically produce disturbances that should affect one variable at a time and watch to see if the disturbances have an effect on each of the variables. For example, it may be possible to disturb optical velocity without disturbing optical acceleration or LOT. If the disturbance has an effect on optical velocity then that variable can be eliminated as a hypothetical controlled variable; if not, then optical velocity remains in the pool of possible controlled variables.

The next step is to disturb optical acceleration without disturbing optical velocity or LOT. If the disturbance has an effect on optical acceleration then that variable can be eliminated as a hypothetical controlled variable; if not, then optical acceleration remains in the pool of possible controlled variables. This process continues until all but one of the hypotheses about the possible controlled variable is eliminated.

In practice, it is not always possible to disturb one hypothetical controlled variable without disturbing others. If a disturbance has no apparent effect on a hypothetical controlled variable it could be because this variable is, indeed, under control; but it could also be because this variable is related to the actual controlled variable. Variations in the hypothetical controlled variable could be *confounded* with variations in the actual controlled variable. For example, a disturbance to optical acceleration is also likely to be a disturbance to optical velocity. If the disturbance has no effect on optical acceleration it may be because optical acceleration is under control. But it also may be because optical velocity is under control and the actions that protect optical velocity from the effects of the disturbance happen to protect optical acceleration from these effects as well. The variations in optical acceleration are confounded with variations in optical velocity.

Modeling Control

One way to deal with the problem of confounding variables is to compare the behavior of the control system under study to that of a model of the system. One such model was developed by Dannemiller, Babler and Babler (1995) in order to demonstrate problems with the LOT hypothesis. Dannemiller et al. built a model of fly ball catching to show that controlling LOT does not guarantee that the fielder will get to the ball in situations where fielders actually do get to the ball. The model shows that two different straight line running paths, one of which does not get the fielder to the ball (an erroneous path), will produce a LOT. This result suggest that variations in LOT are confounded with variations in other variables which, when controlled, allow the fielder to get to the ball consistently.

The Dannemiller et al. analysis of fly ball catching is an example of the use of a *descriptive model* of control system behavior. The model is a mathematical *description* of the feedback function (the function g (), relating q_o to q_i in Figure 1) that relates fielder behavior (running path and speed) to a possible controlled variable (LOT). The mathematical equations that make up the model show how a variable (LOT) *would be* affected by other variables (fly ball trajectory) *if* the control system (fielder) behaved in a particular way (ran in a straight-line path at a particular angle and speed relative to the ball). The problem is that the real control system (fielder) may not actually behave in a way that is consistent with the assumptions of the model. And there is evidence that fielders do not run in the straight line paths assumed by the Dannemiller et al. model (McBeath et al., 1995, Jacobs, Lawrence, Hong, Giordano, and Giordano, 1996). This suggests that the Dannemiller et al. model may not give a complete description of the confounding variables that exist in the study of control of LOTs.

The Dannemiller et al. model describes how a system *might* control a hypothetical controlled variable (like LOT) but it does not actually control that variable. Another approach is to build a model that actually controls a hypothetical controlled variable. Such a model is called a *generative model* because it generates the behavior that keeps the hypothetical controlled variable under control. A generative model of a fielder controlling LOT, for example, would generate the behavior (such as running at a particular angle and speed relative to the ball) that keeps LOT under control. Like the descriptive model, the generative model includes a mathematical representation of the environment in which model behavior occurs and the feedback function that relates model behavior to the aspect of that environment that is being controlled (LOT). But the generative model also includes a model of the behaving system itself; a model that actually controls the proposed controlled variable.

Generative models handle confounding variables by including them in the description of the model's environment. If the model is an accurate representation of the real system, then confounding variables will have the same effect on the behavior of the model as they have on the behavior of the real system. If the behavior of the generative model matches that of the real system then the model is an accurate representation of the real controller that automatically takes the effect of confounding variables into account.

A Model Center Fielder

A generative model of fly ball catching behavior is shown in Figure 2[1]. The model represents a fielder playing straight away center who catches fly balls by controlling optical velocity. This optical velocity control model was developed only to illustrate the *concept* of a controlled variable in the context of a generative model of behavior. Nevertheless, the results of this modeling effort suggest that more complex hypotheses about the controlled variable – optical acceleration and LOT – may not be necessary.

In the optical velocity control model, the fielder is modeled as two control systems; one system controls *vertical* optical velocity, q_V, which is the time change in the vertical elevation of the ball relative to home plate. The other system controls *lateral* optical velocity, q_l, which is the time change in the lateral position of the ball relative to the fielder's direction of gaze (the model fielder always looks straight ahead). In both cases, the projection of the ball is onto a retinal plane that is normal to the fielder's direction of gaze.

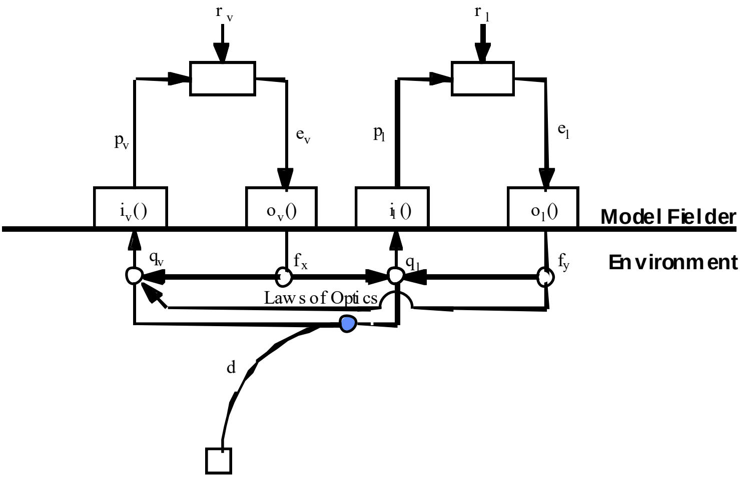

Figure 2. A generative model of fly ball catching behavior. The model is composed of two separate control systems; the system on the left controls vertical optical velocity; the system on the right controls lateral optical velocity.

The optical velocities, q_V and q_l, are environmental variables that are converted into one dimensional perceptual signals, p_V and p_l, by the input functions, $i_V()$ and i_l (), respectively. These functions are linear, based on the assumption that the psychophysical function relating actual to perceptual velocity is relatively linear, at least in the range of velocities experienced by a fielder. So p_V is a time vary-

---

[1] A Java simulation of the optical velocity control model is available on the World Wide Web at http:// home.earthlink.net/~rmarken/ControlDemo/CatchXY.html.

ing signal that is proportional to q_V and p_l is a time varying signal that is proportional to q_l. The input functions also add noise to each perceptual signal to simulate the fielder's less than perfect sensitivity to differences in optical velocity. Low pass filtered random noise was added to each perceptual signal; the average amplitude of the noise was approximately 2% of the average amplitude of each perceptual signal.

The reference signals in the model, r_V and r_l, represent the target values of p_V and p_l , respectively. Model behavior gave the best qualitative match to available fly ball catching data (ground running patterns and optical projections) when r_V was set to .4, indicating an intention to see the vertical projection of the ball constantly rising at the rate of .4 units/sec and r_l was set at 0.0, indicating an intention to the see the lateral projection of the ball remain stationary. The model included a 100 msec transport lag to represent the time it takes for a perceptual signal originating at the retina to get to the point in the central nervous system where it is compared to the reference signal. So time-lagged versions of the perceptual signals, p_V and p_l, were continuously compared to the reference signals, r_V and r_l.

The difference between the time-lagged perceptual signals and the reference signals in each system are the error signals, e_V and e_l; e_V is transformed, via the output function, $o_V()$, into the forces that move the fielder forward or backward (depending on the sign of e_V); e_l is transformed, via the output function $o_l()$, into the forces that move the fielder leftward or rightward (depending on the sign of e_l).

Output forces were limited to those that produce running speeds that are no greater than those that can be produced by a real fielder (a maximum of 6 meters/sec). The control system outputs move the fielder to different positions on the field (in a coordinate system where the x-axis is a line from home plate to straight away center and the y-axis is perpendicular to the x-axis). The fielder's x position, f_X, is determined by the forward or backward motions caused by the outputs of the system controlling vertical optical velocity, q_V. The fielder's y position, f_Y, is determined by the leftward or rightward motions caused by the outputs of the system controlling lateral optical velocity, q_l.

The fielder's position at any instant affects both controlled optical variables, q_V and q_l, via the *laws of optics*. These laws, which determine how the fielder's position relates to the optical projection of the ball on the fielder's eye, are represented by the lines connecting system outputs to system inputs in Figure 2. The lines show that the fielder's x position, f_X, affects both optical velocities, q_V and q_l, as does the fielder's y position, f_Y. This means that movements in the x dimension that are aimed at control of q_V are also a disturbance to q_l; similarly, movements in the y dimension that are aimed at control of q_l are also a disturbance to q_V. This relationship between the fielder's inputs and outputs is the feedback function that connects the outputs of both control systems to the variables controlled by these systems.

Model Behavior

The model's ability to catch fly balls was tested by simulating fly balls hit in different trajectories relative to the model. The model started each catch from the same field position (defined by the initial values of f_x and f_y) which was located in straight away center field, about 50 meters from home plate. The balls were launched at a variety of realistic angles (vertical and lateral) and initial velocities (Adair, 1994). Vertical angles ranged from 42 to 48 degrees with respect to horizontal. Lateral angles ranged from +20 to -20 degrees with respect to the line connecting the fielder to home plate. Initial velocities off the bat ranged from 21-25 meters/sec. The trajectories were limited so that most balls could be caught, given the limits on the model fielder's output (running rate) capabilities.

The behavior of the model as it catches fly balls hit in four different trajectories is shown in Figure 3. The left panel shows a top view of the fielder's running path to the ball (changes in f_x and f_y over time). The right panel shows the "model fielder's view" of the ball during each fly ball trajectory.

The running paths on the left in Figure 3 are qualitatively similar to those observed for real fielders (see McBeath et al., 1995, Fig 3, p. 571; Jacobs et al., 1996, Fig 2A, p. 258). Like the real fielder, the model fielder moves laterally, in a curved path, before moving forward to intercept balls hit to the left or right of the fielder's starting position. The model fielder also shows an initial reverse movement for some trajectories that is also seen in some of the running patterns of the real fielder observed by McBeath et al. The exact shape of the model fielder's path to the ball depends on control system parameters, particularly the relative gains of the systems controlling q_V and q_l, and the reference (r_V) for vertical velocity (the larger the value of r_V the less the model fielder backs up prior to running to the intercept point).

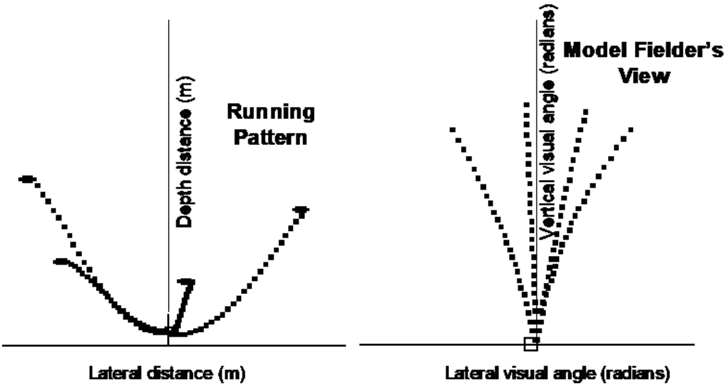

Figure 3. Behavior of the control model of fly ball catching. Data on the left are running paths to four different fly balls; data on the right are the optical projection of these four different fly balls as seen by the model.

The "model fielder's view" of the fly ball trajectories on the right in Figure 3 are qualitatively similar to the "fielder's view" of the fly ball trajectories observed by McBeath et al. (1995, Fig 4, p. 572). The model fielder's views of the trajectories are basically LOTs as are the fielder's views of such trajectories in the McBeath et al. data. These LOTs are a surprising side effect of controlling for optical velocity. The model fielder's views of the trajectories curve slightly away from a straight line but careful inspection of the McBeath et al. data shows that the fielder's views of the trajectories curve in the same way. Again, the degree of curvature seen in the model's views of the trajectories depends on the model parameters, principally the relative gains of the systems controlling q_V and q_l. When these gains are high the model fielder's views of the trajectories have very little curvature; they are basically straight lines.

Comparing Models

The optical velocity control model (Figure 2) is presented as an example of a model that generates behavior (running paths, model fielder's views of trajectories) that is qualitatively consistent with the available data on catching fly balls. The model is very similar to a model of fly ball catching developed by Tresilian (1995). The main difference between the models is the controlled variable, which in the present model is optical velocity and in Tresilian's model is optical *acceleration*. A comparison of equivalent versions of the models (equivalent in terms of running speed, gain, transport lag and other relevant parameters) showed that the behavior of a model controlling optical velocity is not the same as the behavior of a model controlling optical acceleration. The comparison showed that the acceleration control model (which was controlling for a vertical optical acceleration of zero) often over- or under- ran fly balls that were caught by the equivalent optical velocity control model (which was controlling for a vertical optical velocity of .4 units/sec).

The problem with the acceleration control model is that its behavior depends strongly on the ball's initial optical velocity (the fielder's view of the ball's velocity as it leaves the bat). Since the model is trying to keep optical acceleration at zero, the ball's optical velocity remains near its initial value. For example, if optical velocity during the first sampling interval (when the ball leaves the bat) is .5 units/sec then the model will act to keep the velocity in the next sampling intervals near .5 units/sec (constant velocity equals zero acceleration). So the optical velocity that results from controlling optical acceleration at zero can be too fast (as it would be if it were kept at .5 units/sec) or too slow relative to the optimal optical velocity, which corresponds to the reference value (.4 units/sec in this case) for optical velocity. The result is that the ball sometimes falls in front of or behind the model fielder controlling acceleration even when the model fielder is successfully keeping optical acceleration at zero. These results are strong evidence that optical velocity, not acceleration, is the variable controlled by fielders catching fly balls.

The optical velocity control model also generated straight line plots (LOTs) of the fielder's view of the fly ball trajectories like those predicted by the LOT control model. These LOTs were generated for *every* fly ball trajectory as a side effect of controlling optical velocity. This result suggests that the LOTs observed in the McBeath et al. study may not have been a controlled variable. The observed LOTs may have been a side effect of the fielder's control of optical velocity. But the fact that LOTs are generated by the optical velocity control model is not proof that the observed LOTs are a side effect. In fact, the observed LOTs may occur because LOT is, indeed, a controlled variable; fielders may try to keep the optical projection of the ball moving in a straight line. What is needed is a way to test whether the observed LOTs are a controlled variable or a side effect of control of optical velocity.

The Answer is Blowing in a Wind

It is possible to distinguish LOT as a controlled variable from LOT as a side-effect of optical velocity control by introducing disturbances to the optical projection of the fly balls that will be resisted only if LOT is a controlled variable. One such disturbance is a slowly varying lateral wind that causes the ball to move away from its otherwise straight course. Such disturbances were applied to fly balls that were hit to the optical velocity control model. The results are shown in Figure 4.

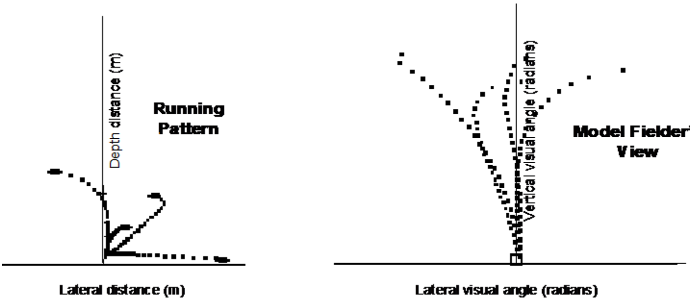

Figure 4. Same as Figure 3 except that a time varying lateral wind was added to each of the four fly ball trajectories.

The left panel of Figure 4 shows a top view of the model fielder's paths to four different fly balls; the right panel shows the "model's view" of the optical trajectory of each fly ball. The model was able to catch all fly balls but, as can be seen from the model fielder's view data, the path of the model fielder's view of the ball is no longer a straight line (LOT). In several cases, the path of the model fielder's view of the ball is quite curved. The model results show that a lateral wind will be a fully effective disturbance to LOT (it makes the LOT curve) if the fielder is controlling optical velocity.

Figure 4 constitutes a prediction of what will be seen as the "fielder's view" in a videotape record of fly balls caught in a lateral wind if the fielder is controlling optical velocity. If the fielder is actually controlling LOT then the "fielder's view" of fly balls caught in a lateral wind will still be a straight line; the fielder will move so as so keep the wind-disturbed projection of the ball falling on a straight line. If the fielder is controlling optical velocity then the fielder's view of the ball will look like the results on the right in Figure 4.

Some Lessons From Center Field

The relatively simple, PCT-based optical velocity control model handles, at least qualitatively, the existing data on fly ball-catching behavior. So it seems appropriate to see what lessons about behavior can be learned from the model. The first lesson comes from watching the behavior of the model in real time. An animation of model behavior shows running patterns that appear quite realistic. The model fielder seems to be *anticipating* where the ball will land, *planning* the best path to the ball and *calculating* the forces necessary to get there. It looks like the model is doing complex anticipatory computations of the movements it should make to get to the ball. In fact, the model is just continuously calculating the difference between perceived optical velocity and the reference for this velocity. The behavior of the optical velocity model suggests that the appearance of anticipation in purposive activities (e.g. Wing and Lederman, 1998) may be an illusion. The model shows the importance of identifying controlled variables and developing models that generate behavior by controlling these variables before concluding that any behavior is the result of anticipatory plans and calculations.

A second lesson from the model concerns reference states for controlled variables. An often unnoticed assumption that has been made since the earliest applications of control theory to human manual tracking behavior (e.g. McRuer and Krendel, 1959) is that the only possible reference state for controlled variables is zero. In studies of tracking this shows up as the assumption that any non-zero distance between the cursor and the target is an “error”. Of course, this is only true if the controller intends to keep the controlled variable – the distance between the cursor and the target – at zero. If the controller intends to keep the distance between cursor and target at some value other than zero (a non-zero reference state) then a non-zero distance between the cursor and the target will not be an “error”. In fact, this non-zero distance will be the distance that results in zero error.

The option of selecting a non-zero reference for the controlled variable is made explicit in PCT because the reference signal is shown in the diagram of the control system model of performance (see Figures 1 and 2). Control system diagrams that are not based on PCT typically leave out the reference signal (see, for example, Tresilian, 1995, Figure 1, p. 692) which seems to encourage the assumption that the control system is an input-output transfer function with a nominal offset (reference signal value) of zero. This is, indeed, the assumption made by the OAC model of fly ball catching (Babler and Dannemiller, 1993; Tresilian, 1995). The OAC model keeps optical velocity non-zero when its reference for optical acceleration to set to zero. But the non-zero optical velocity produced by the OAC model is not always the one that gets the fielder to the ball. The PCT-based optical velocity control model is able to catch balls that are missed by the OAC model because it explicitly sets a specific non-zero reference for vertical optical velocity.

PCT assumes that reference signals are set by higher level control systems in the organism itself; the non-zero reference for vertical optical velocity is presumably set by a higher level system that is controlling some other perception, like the perception of the ball being caught. This is *hierarchical control* (Powers, 1979). The higher level systems in the hierarchy control perceptions by setting an appropriate (in this case, non-zero) reference for the perceptions controlled by the lower level systems (like those controlling optical velocity). The fact that the optical velocity control model can only catch the ball when its reference for vertical optical velocity is set close to .4 units/sec suggests that real fielders have *learned* to set the reference for the lower level controlled variable (optical velocity) at .4 units/sec in order to control the higher level perception: catching the ball.

Finally, the PCT-based optical velocity control model shows how behavior can be modeled using multiple independent control systems, each controlling a different variable. Most applications of control theory to behavior have used a single control system to model the behaving system. For example, Tresilian's (1995) model of fly ball catching behavior uses a single control system to model the fielder. The model controls only one input variable even though the fielder must move in two dimensions to control this variable. The second dimension of fielder movement in Tresilian's model is derived by open loop computations based on information about the direction of the ball's movement relative to the fielder.

PCT immediately suggests a simple solution to the problem of simulating behavior that involves controlling in two dimensions; simply have separate control systems control different dimensions of *perception*. It might seem that these systems could work at cross-purposes since the outputs of each system disturb the perceptual variable controlled by the other. But separate control systems can control their inputs successfully as long as the inputs controlled by each system represent relatively independent perceptual degrees of freedom (Marken, 1992, pp. 185-206). The control systems that make up the optical velocity control model successfully control two independent perceptual variables (vertical and lateral optical velocity) while automatically compensating for the effects their outputs have on the inputs to the other system. Output coordination is an automatic consequence of the control of several independent perceptual inputs.

Beyond Baseball: Principles of a PCT-Based Psychology

The fly ball catching model shows how one familiar behavior – catching fly balls – can be seen as an observable side effect of the process of keeping a set of perceptual variables (vertical and lateral optical velocity) under control. According to PCT, *all* behavior is aimed at the control of perceptual variables. So understanding the behavior of living organisms means knowing what variables the organism is controlling. The PCT approach to understanding fly ball catching behavior gives an idea of what a PCT-based psychology – one based on the principles of PCT – would look like.

One principle of a PCT- based psychology is that behavior must be viewed in terms of the perceptual variables that an organism *might* be controlling. The behavior we see is presumed to be an observable side effect of the organism's efforts to keep controlled variables under control. This was the case in the fly ball catching studies, where it was assumed that the observed running patterns are a side-effect of the fielder's efforts to control an optical perception. An analogous situation exists for other observable behavior. For example, a PCT-based approach to understanding the observed egg rolling behavior of the greylag goose (Lorenz, 1981) would be based on the assumption that the observed behavior is a side effect of the goose's efforts to control some perceptual variable, such as sensed pressure on the inside of the bill. Complex movements of the bill, such as those that occur when the egg is surreptitiously removed, could then be understood as outputs aimed at protecting the pressure perception (the controlled variable) from disturbance.

A second principle of a PCT-based psychology is that evidence of controlled variables must come from studies where individuals are tested *one at a time*, as in the fly ball catching research. Runkel (1990) calls this approach to research the *method of specimens*. It should be possible to find evidence of many types of controlled variables in existing psychological studies of behavior – even if these studies were not intentionally designed to expose the existence of controlled variables. But this is true only if these studies used the method of specimens. Studies of the average behavior of groups of individuals cannot reveal what variables each individual is controlling, unless all individuals happen to be controlling exactly the same variables. Indeed, the behavioral laws revealed by group averages can be precisely the opposite of the laws that actually characterize the behavior of the individuals in the group (Powers, 1990).

One field of psychological research where the use of the method of specimens is common is the study of operant behavior. The goal of the study of operant behavior is ostensibly to discover the variables that control behavior. But it is possible to look at operant behavior from the perspective of PCT and see evidence of variables that are controlled by behavior: controlled variables. For example, studies of the effects of reinforcement schedules on behavior suggest that organisms will vary their responding appropriately in order to keep rate of reinforcement relatively constant (Staddon, 1979). The rate at which reinforcements are received from the apparatus may be a variable that organisms control.

Another field of research where the use of the method of specimens is common is the study of cognitive processes. Atwood and Polson (1976), for example, studied problem solving by looking at the behavior of individual subjects solving water jar problems. Their research suggests that the relative amount of water in the jars is the controlled variable. The reference state for this variable seems to be the relative amounts of water in the jars in the solution state of the problem.

Evidence that people try to make the current problem state look as much as possible like the solution state comes from the fact that, given a choice between a legal move that makes the relative amounts of water look more like the solution state and one that makes the relative amounts look *less* like the solution state, a person nearly always chooses the former rather than the latter. Trying to achieve this reference state for the controlled variable makes problem solution difficult because, in order to solve a water jar problem, the person must select moves that make the relative amounts of water in the jars look *less*, rather than more, like the solution state.

A third principle of a PCT-based psychology is that evidence of controlled variables must come from studies where individuals are actually able to control the possible controlled variable. In the fly ball catching studies, fielders were able to control all possible controlled variables because impossible trajectories (leading to fly balls that could not be caught) were not used. But possible controlled variables are often found to be uncontrollable in otherwise appropriate studies of purposeful behavior. For example, in many operant scheduling experiments, the organism has very little influence over possible controlled variables, such as the amount of reinforcement received. The typical operant schedule is arranged so that the organism can get only a fraction of the reinforcement that it wants even when it is responding at its maximum rate. This means that it is impossible to tell whether any failure to keep a variable (like reinforcement rate) at a particular level (reference state) occurs because the organism is not controlling the variable or because it *cannot* control the variable.

A fourth principle of a PCT-based psychology is that evidence of controlled variables must come from studies that use some version of the TCV to determine whether or not a particular variable is actually under control. The fly ball catching studies did use a version of the TCV where possible controlled variables were monitored while disturbances (varying trajectories) were applied. Unfortunately, the steps involved in doing the TCV are rarely carried out in most psychological research studies. For example, the TCV is rarely applied in operant research. In particular, there is no monitoring of possible controlled variables while disturbances are applied.

It would be easy to replicate many operant experiments using the TCV. All that is needed is a way to produce disturbances and monitor the state of possible controlled variables. A possible controlled variable like reinforcement rate, for example, could be disturbed by delivering extra reinforcers after certain responses. Something like this was done by Teitelbaum (1966) who added extra food pellets randomly after certain lever presses in an operant situation. Although Teitelbaum did not monitor the state of a hypothetical controlled variable while these disturbances occurred, the results suggest that organisms do adjust their response rate in a way that would keep some variable related to reinforcement rate under control.

Finally, a fifth principle of a PCT-based psychology is that the study of purposeful behavior must include the development of generative models that produce behavior by controlling the variables that are presumably being controlled by the real system. Generative models, like Tresilian's (1995) model of fly ball catching behavior, make it possible to determine whether observed behaviors, such as the running patterns seen when fielders catch fly balls, would actually be produced by a system controlling the presumed controlled variables. Moreover, generative models make it possible to see whether the observed behaviors are produced in an environment like that in which the real systems must behave.

Conclusion

This paper has described an approach to understanding the purposeful behavior of living systems in terms of controlled variables. The approach, based on PCT, was illustrated in the context of studies aimed at determining how outfielders catch fly balls. The central feature of this approach is research aimed at discovering the variables organisms control when performing various purposeful behaviors. This research involves testing for controlled variables using the TCV and developing generative models that produce the behavior under study by controlling these variables. In the case of research on fly ball catching behavior, the TCV suggested three variables that fielders might be controlling when they run to intercept the ball: optical velocity, optical acceleration and linear optical trajectory

(LOT). A generative model of fly ball catching suggests that the most likely controlled variable is optical velocity. The model also suggests a way to eliminate LOT as a possible controlled variable by looking for the effect of lateral wind disturbances on the optical path of the fly ball.

Once a controlled variable has been correctly identified it is possible to understand many aspects of the behavior under study. The optical velocity control model (Figure 2) shows how correct identification of a controlled variable (optical velocity) makes it possible to understand many aspects of the behavior observed when fielders catch fly balls. The model also makes clear (and, thus, clearly falsifiable) predictions of other behaviors that should be observed (Figure 4) if fielders actually do control optical velocity.

According to PCT, all *purposeful* behavior involves the control of some perceptual aspect of the environment. This means that it should be possible to explain all purposeful behavior in terms of controlled variables. In order to do this, it will be necessary to begin the search for controlled variables in earnest. The aim of this paper is to encourage a systematic search for the controlled variables that underlie various purposeful behaviors. This search can be centered on any purposeful behavior. All that is required is some evidence that these behaviors are purposeful; that there is an effort to produce a pre-selected state of some variable while protecting this variable from the effects of disturbance. Other aspects of these behaviors, such as learning and memory, can then be understood in terms of their relationship to the basic process of purposeful behavior: the control of perception (Powers, 1973a).

4. Purpose in Biology and Economics

The Ethology of Purpose

In their target *Behavioral and Brain Sciences* article, R. A. and R. T. Gardner (1988) propose a "unified feedforward model of the learning of adaptive and maladaptive behavior" that emphasizes the importance of ethology in operant learning. Unfortunately, the "model" turns out to be little more than a restatement of the observations on which it is based. However, the observations themselves are of interest. First, we are shown evidence that the responses typically used in operant conditioning experiments (key-pecking, bar-pressing) are actually variants of species-specific behaviors evoked and maintained by food and water without any contingency at all. Second, we find that increases in deprivation can lead to decreases in the effectiveness of contingency-based training – surprising if one imagines that deprivation increases the "positiveness" of the consequences that shape behavior. Taken together, these observations are inconsistent with the notion that "arbitrary" behaviors can be trained by making positive consequences contingent on production of the behavior (Skinner 1981).

Gardner and Gardner (1988) have taken aim at one of the most venerable principles of psychology and fired with both barrels. As a practicing non-fan of reinforcement theory (Marken, 1985) I applaud the goal and feel that some

---

From *Behavioral and Brain Sciences*, 1988, 11(3), 460-461. Reprinted with permission of Cambridge University Press

solid hits have been scored. The target article provides strong evidence that the role traditionally ascribed to response consequences is wrong. In operant conditioning experiments it appears that consequences select or strengthen responses. But, according to Gardner and Gardner, this only happens when "the to be-conditioned response is similar to or compatible with the obligatory responses evoked by S* [reinforcement]. Since bar-pressing is similar to the obligatory response evoked by food, it has appeared that contingent presentation of food strengthens this "arbitrary" behavior. The law of effect, then, is something of an illusion.

Gardner and Gardner have set out to do more than repeal the "law of effect". They also offer an alternative explanation of operant learning. I find their cure no better than the disease. Actually, I had some difficulty understanding the theory that is offered. As far as I can tell, Gardner and Gardner propose that learning results when the to-be-conditioned response becomes part of a recursive loop that is initiated by an event that evokes an obligatory species-specific response. If we ignore the problem of implementing a recursive behavioral loop, the animal presumably learns a behavior, such as pressing a bar, because "inputs indicating the presence of food" (I presume this is the sight of food) evoke prefeeding responses that get directed at the bar. The bar goes down to produce food, which (l) gets eaten and (2) provides more "inputs indicating the presence of food". A great deal is also made of the importance of ethology. Unfortunately, little is said about the mechanism by which ethology produces learning.

The problem with this "feedforward" model is the same as that with the reinforcement model it is meant to replace.

Both models assume an extremely intelligent and helpful environment. The “feedforward” model holds that behavior is “poked out” by prior events rather than “pulled out” by subsequent events. But how do these events, whether prior or subsequent, know what the animal wants to do? Of course, many would object to talk about “wants” especially in animals. But it is difficult to avoid noticing that animals do seem to achieve goals (Rescorla, 1987). In fact, one distinctive feature of operant behavior is that it is goal-oriented; animals consistently avoid shocks (Powers, 1971) and maintain food intake (Timberlake, 1984) under the most unpropitious natural (ethological?) and artificial conditions. How does the environment know how to guide the animal to these goals? How, for example, does the “food stimulus” know when to start and stop evoking responses so that food intake remains stable?

The fact that animals produce consistent results in the face of environmental variation means that these results are under control (Marken, 1988). The source of this control can be shown to be the animal itself – not the environment. The only known explanation of how control occurs is control theory. Gardner and Gardner are aware of the existence of control theory (they present an incorrect description of the operation of a control system in section 4.3) and seem to be contrasting their feedforward model with the feedback model of control theory (though it's anybody's guess what section 4.3 is really about). I get the uncomfortable feeling that Gardner and Gardner (like many other operant learning theorists) imagine that the concepts of feedback control and reinforcement are equivalent. They are not. Reinforcement theories hold that responses are *controlled* by consequences; control theory shows that responses *control* consequences. It may be years before learning theorists can understand that distinction, let alone accept it.

The ability to control certain results is not built into the organism; nature can't anticipate the details of the environment into which an individual animal is born. Thus, animals must learn to control. Gardner and Gardner have pointed out some interesting constraints on that learning process. “Obligatory” pre-feeding responses, for example, can interfere with acquisition of the means to control food intake. But even here, it is not the pre-feeding behavior that is “obligatory” but the result we see as prefeeding. The animal does not repeat exactly the same actions to do prefeeding; it can't and still produce a consistent result: prefeeding. Small changes in the environment and the orientation of the organism relative to that environment require different actions to produce the same result. What is apparently built-in is the “want” to produce the result seen as prefeeding.

The “feedforward” model described in Gardner and Gardner's target article aims to orient our attention to the neglected reality of species-specific behavior (ethology) and this is certainly worthwhile. Unfortunately, the “new” theory (it looks a lot like S-R theory to me) is based on the same fundamental (and mistaken) conviction as the “old” theories, namely, that goal-oriented, operant behavior can be explained by environmental guidance. The “feedforward” model is pretty much business as usual. Researchers will continue to search for *controlling variables* but now with the emphasis on variables that evoke species-specific behaviors. The model lets us continue to ignore the existence of *controlled variables* – environmental events controlled by the animal – thus missing the purpose of the animal's every move.

H. Economicus: A Perceptual Control System Model of the Economy

This paper describes a model of the economy that is based on Treval C. Powers' (1996) historical analysis of economic data found in the *Statistical Abstract of the United States*. T. C. Powers' analysis is surprisingly (and unintentionally) consistent with the perceptual control theory model of individual behavior developed by his son, William T. Powers (1973). Powers *père* views the economy as a circular flow of money between producers and consumers. The behavior of this circular flow can be explained in terms of a perceptual control system model such as that described by Powers *fils*. A perceptual control system controls a perceptual representation of some aspect of environment that is shaped by the outputs of the system itself. The perceptual control theory model of the economy that is described in this paper controls a monetary representation of aspects of the economy that are shaped by the outputs of the model itself.

Circular Flow Analysis

T. C. Powers' analysis of the economy assumes that

Paper presented at the 2000 Meeting of the Control Systems Group, Boston University, Boston. MA the basic economic process is a *circular flow* of money. Money flows between two composite entities: a *composite producer* and a *composite consumer*. The composite producer consists of all the people in the economy who are contributing to the production of goods and services. The composite producer is, thus, made up of everyone in the economy except those who are unable to contribute to the production of goods and services – the very young, the very old and the infirm. The composite consumer consists of all the people in the economy who are contributing to the consumption of the goods and services that are being produced. The composite consumer is, thus, made up of everyone in the economy since everyone must consume at least some portion of all goods and services (such as the food portion) in order to survive.

To a first approximation, the composite producer and composite consumer are the same group of people: the entire population of people that makes up an economic entity such as the United States. The terms *composite producer* and *composite consumer* simply refer to two different roles – production and consumption – that are being carried out by the population of people that makes up this economic entity. Moreover, at the composite level, production and consumption can be thought of as occurring *simultaneously*: the composite producer is making cars while, as the composite consumer, it is driving home; the composite producer is growing and packaging food while, as the composite consumer, it is eating; the composite producer is teaching computer science while, as the composite consumer, it is applying this knowledge to internet innovation.

The simultaneous interaction between composite producer and consumer is shown in Figure 1. The composite producer is producing goods and services, Q’, and paying *itself*, in the role of composite consumer, for the work that produces Q’. The amount of dollars the composite producer pays itself, as composite consumer, is the average price (P) of all goods and services times the amount of goods and services produced (Q’). This payment is seen leaving the composite producer as an *output of dollars*, PQ’. The composite consumer receives these dollars as income; wage income (W), which includes both wages and profits, and capital income (K), which includes interest income and rent. Consumer income (W+K) represents buying power, B. The composite consumer uses this income to buy the goods and services that it was paid for producing them. This payment is shown being returned to the composite producer as *spent* purchasing power.

Figure 1 shows a circular flow of money going from composite producer to composite consumer and back to composite producer. It is important to understand that money is flowing in all parts of this circle simultaneously; money is flowing from the composite producer to the composite consumer as income *while* it is flowing from the composite consumer to the composite producer as payment. This flow is in equilibrium when the money flowing into the composite producer, in the form of payments for consumption of Q’, is equal to the money flowing out of the composite producer in the form of wages and capital payments for production of that same Q’.

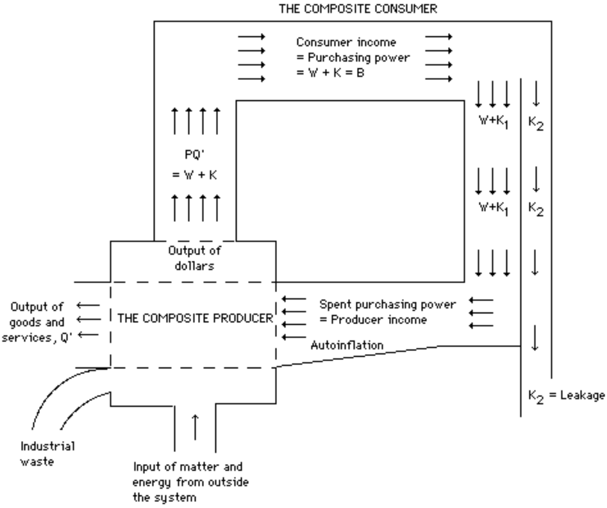

Figure 1. Circular flow analysis of the economy.

In circular flow analysis, there is a new law of supply and demand that operates at the composite level. PQ’, the dollars paid for production of Q’, is supply; W + K, the dollars available to purchase Q’, is demand. The circular flow keeps the economy going only when supply equals demand such that the composite producer is being repaid exactly what it spent (and handed to the composite consumer as income) for producing Q’. Supply will not equal demand if money is lost from the circular flow. Money will be lost from the circular flow if it is received as income by the composite consumer but not used to purchase goods and services (Q’). Money that is lost from the circular flow in this way is called *leakage*.

Figure 1 shows that some proportion (K_2) of the total income (W + K) received by the composite consumer is being leaked away through the path on the right side of the figure, leaving only a portion of total income ($W + K_1$) to be returned to the composite producer as payment for Q'. Leakage creates a mismatch between supply and demand. This mismatch will reduce the circular flow because the composite producer is being forced to produce less (and, thus, spend less) in order to make up for the reduced income it receives. So the reduced circular flow resulting from leakage will produce an economic slow down (a recession or depression) unless money is pumped into the economy. This pumping process is called *autoinflation* in Figure 1 because the composite consumer is now spending more of its income (more than $W + K_1$) for the same quantity of goods and services, Q'.

The amount of money that is being paid by the composite producer for production of Q' (PQ' or W + K) is equivalent to a very important measure of economic performance that is provided regularly in the *Statistical Abstract of the United States* (2000): *Gross National Product* or GNP. T. C. Powers' (1996) historical analysis of the economic data available in the *Statistical Abstract* shows that, over the last 100 years, the composite consumer has failed to return, on average, about 8% of its yearly income (GNP) to the composite producer. This means that the composite consumer spends 8% less on consumption than it receives as income. This unspent money is *not* savings. At the composite level, the amount of money being put into savings for future use is about equal to the amount being withdrawn from savings for current consumption. Powers (1966) presents evidence that this unspent income tends to be negatively related to economic growth and positively related to inflation, as predicted by circular flow analysis. In other words, unspent income acts like leakage.

T. C. Powers (1996) presents a considerable amount of evidence to support the circular flow analysis of the economy. What T. C. Powers does not do is propose a mechanism that keeps the circular flow flowing. For example, circular flow analysis describes no mechanism that keeps the outflow of dollars (PQ’) from the composite producer matching the inflow of dollars (autoinflated W + K_1) into the composite produced when there is leakage. Nor does the analysis describe a mechanism that can account for the effect of leakage on economic growth (the growth of Q’). What is needed is a mechanism that will produce the behavior predicted by circular flow analysis. The mechanism that will produce this behavior turns out to be the control system model of individual behavior developed by T. C. Powers’ son, W. T. Powers.

H. Economicus

A control model that produces most of the economic behavior predicted by the circular flow analysis is shown in Figure 2. The model, called *H. Economicus,* consists of two control systems: the *composite manager* and the *composite GNP controller*. Each system controls a variable in the economic *environment*. The composite manager system controls the difference between the amount of money paid out for production of Q’ (PQ' or GNP) and the amount returned as payment for Q’, which is called *Producer income* in Figure 1 and P’Q’ in Figure 2. P’Q’ is what it costs the composite consumer to buy Q’. So P’Q’ can be thought of as GNP seen from the point of view of the composite consumer while PQ’ can be thought of as GNP seen from the point of view of the composite producer. PQ’ is what it costs the composite producer to produce Q’ (which is the same as the GNP measured by government economists); P’Q’ is what it costs the composite consumer to purchase Q’. The composite GNP controller system controls just P'Q’, which is GNP from the composite consumer’s perspective.

P’Q’ is a new variable that is not found in the circular flow analysis. What is new about P’Q’ is P’, which is the average cost of Q’ to the composite consumer. P’ is to be distinguished from P, which is the average cost of Q’ to the composite producer. P and P’ are not always the same because the cost of Q’ to the consumer must sometimes be increased to make up for any loss of income to the producer due to leakage. So there are two controlled variables in the model of *H. Economicus*: PQ’-P’Q’, which is controlled by the composite manager and P'Q' which is controlled by the composite GNP controller.

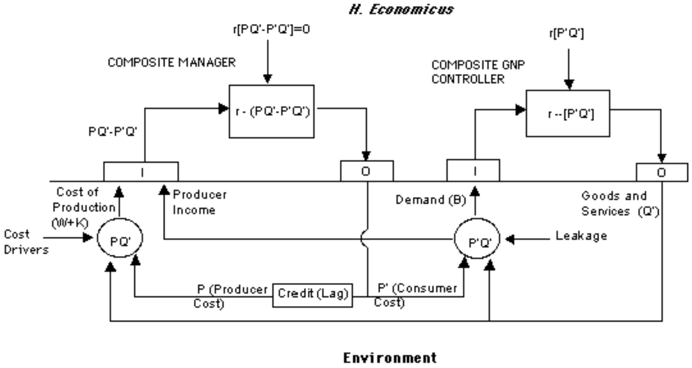

Figure 2. *H. Economicus,* a two control system model of the economy. One system controls PQ’-P’Q’ and the other controls just P’Q’. PQ’ is the cost to the composite producer of producing Q’; P’Q’ is the cost to the composite consumer for purchase of Q’.

The composite manager component of *H. Economicus* does part of the job of the composite producer in circular flow analysis; it produces income (PQ') by paying the cost of production. But the composite manager does something that is not done in the circular flow model; it balances the books. The composite manager acts to keep the difference between production costs (PQ’) and income from sales (P’Q’) equal to a reference value (r[PQ’-P’Q’]) that is set equal to zero. So the composite manager keeps output (PQ’) matching input (P’Q’). It does this by raising or lowering the cost of goods and services (P’) to compensate for disturbances to the controlled variable, PQ’-P’Q’. There are actually two disturbances to PQ’-P’Q’: cost drivers, such as unpredictable variations in the availability of natural resources, that influence the cost of production (PQ’) and leakage (variations in the amount of income that is not used for consumption) that influences the per item cost of goods and services.

The composite GNP controller component of *H. Economicus* fills the role of both the composite producer and composite consumer in circular flow analysis. It acts as a producer by working to make the goods and services (Q') it consumes; it acts as a consumer by consuming P'Q', the goods and services it produced. It works to produce Q’ in order to keep the amount of goods consumed (P’Q’), perceived in terms of their dollar value, equal to a reference for the amount of goods and services desired (r[P'Q']). The composite GNP controller can compensate for disturbances to P'Q' (caused by leakage) only by varying the amount of goods and services (Q') produced. In fact, most of the composite GNP controller's efforts (in terms of production of Q') are aimed at keeping perceived GNP, P'Q', equal to an ever increasing reference level.

The reference for P'Q' (r[P'Q']) is equivalent to the composite GNP controller's *demand* for GNP as perceived in terms of its cost (P'). In the *H. Economicus* model, increases in the reference for perceived GNP are the driving force behind economic growth. With P' relatively constant, the composite GNP controller must continuously increase Q' in order to keep P'Q' equal to a constantly increasing r[P'Q']). However, if there is leakage, increases in P' will offset increases in Q', leading to lower levels of production of goods and services; the composite GNP controller gets growth in P'Q' but this growth is a result of increases in P' (cost) as well as Q' (production of actual goods and services). The composite GNP controller's efforts to counter the effects of leakage on P'Q' by reducing output, is the mechanism that accounts for the effect of leakage on the productive capacity of *H. Economicus*.

The Behavior of *H. Economicus*: Leakage and Inflation

The *H. Economicus* model was implemented as a dynamic spreadsheet simulation. When the simulation was run, the reference for P'Q' (r[P'Q']) was automatically increased at a rate equivalent to 13% per year. The reference for PQ'-P'Q' remained equal to zero. At the beginning of each simulation run the user could enter a value for the rate of leakage. At the end of a simulation run (which lasted the equivalent of 20 years) the spreadsheet calculated three measures of economic performance: the relative output of goods and services produced by the economy (Q'/Q), the rate of economic growth (dQ'/dt) per year and an index of inflation. Relative output is the ratio of actual economic output (Q') to the economic output that *would have* been produced if there were no leakage. Rate of growth is the percentage change in Q' in one year. The index of inflation is the ratio of the actual average cost of consumer goods (P') to what the average cost of consumer goods *would have* been without leakage.

Table 1. Effect of leakage on Q'/Q, growth rate and inflation for circular flow analysis and *H. Economicus model.*

Leakage %	Relative Output Q'/Q %		Rate of Growth % per year		Index of Inflation	
	CF	*H. Econ.*	CF	*H. Econ.*	CF	*H. Econ.*
0	100	100	13	13	100.0	100.1
2	98	98	11	13	102.0	102.1
4	96	96	9	13	104.2	104.3
5	95	95	8	13	105.3	105.9
6	94	94	7	13	106.4	106.4
7	93	93	6	13	107.5	107.3.
8	92	92	5	13	108.7	108.4
9	91	91	4	13	109.9	109.5
10	90	90	3	13	111.1	111.2
11	89	89	2	13	112.3	112.8
12	88	88	1	13	113.6	113.6

13	87	87	0	13	114.9	115.1
14	86	86	-1	13	116.3	116.2
15	85	86	-2	13	117.6	117.4
16	84	84	-3	13	119.0	118.5

The control systems that make up *H. Economicus* were not designed to produce particular values of Q’/Q, growth rate or inflation. The values obtained are side effects of the operation of the control systems. They are the values of Q’/Q, growth rate and inflation that result when the control systems act to protect the variables they are controlling (PQ’-P’Q’ and P’Q’) from disturbances, in this case, from the disturbance caused by different levels of leakage.

The fact that the *H. Economicus* model produces values of Q’/Q and inflation that are very close to the predictions of circular flow analysis for all values of leakage is a reassuring indication that the mechanism of the *H. Economicus* model is able to capture the important aspects of the behavior of the circular flow analysis. The small differences between the predictions of the circular flow and *H. Economicus* model result from the fact that the P'Q' variable in the *H. Economicus* model includes a portion of Q' that represents unsold inventory.

The glaring difference between circular flow analysis and the *H. Economicus* model occurs in the results for rate of growth (columns 4 and 5 in Table 1). The circular flow analysis predicts a large effect of leakage on rate of growth but the rate of economic growth produced by *H. Economicus* is not affected by leakage at all. The reason for this discrepancy is clear when one takes a closer look at the circular flow analysis. In circular flow analysis the dependence of growth rate on leakage is simply assumed to exist; it is not derived from the interaction of variables in the circular flow, as was the case with Q/Q' and inflation. Rather, the effect of leakage on growth rate is taken as an axiom in circular flow analysis. That is, the effect is *assumed* (Powers, 1996, equation 2-29, p 101) rather than *predicted*.

Circular flow analysis does not suggest a mechanism to explain the dependence of growth rate on leakage and such dependence does not exist in *H. Economicus*. However, simulations using the *H. Economicus* model did reveal a surprising dependence of growth rate on *rate of change* in leakage (increasing rates of change in leakage lead to decreasing rates of change in economic growth). But it is difficult to tell whether this aspect of the behavior of *H. Economicus* is consistent with the economic data.

The effect of leakage on inflation that is found in both *H. Economicus* and the circular flow analysis in Table 1 (columns 6 and 7) was also found in the economic data. One of the main causes of leakage is the Federal Reserve's discount rate policies; the higher the discount rate the greater the leakage. Based on the results in Table 1, one would expect to find a positive relationship between discount rate (leakage) and inflation rate in the economic data. This is precisely the opposite of the effect expected by conventional economists. But Figure 3 shows that the data fit the prediction of circular flow analysis as implemented in the *H. Economicus* model rather than the conventional expectation. In fact, the results shown in Figure

3 are exactly what would be expected if the Federal Open Market Committee (which sets discount rates with the aim of keeping inflation in check) is in a positive feedback relationship with respect to the variable it is trying to control: inflation. The positive feedback comes from the fact that discount rates, which are being raised with the aim of decreasing inflation, are actually increasing it. Small fluctuations in inflation would be emphasized by this positive feedback process resulting in the large swings in inflation rate seen in Figure 1.

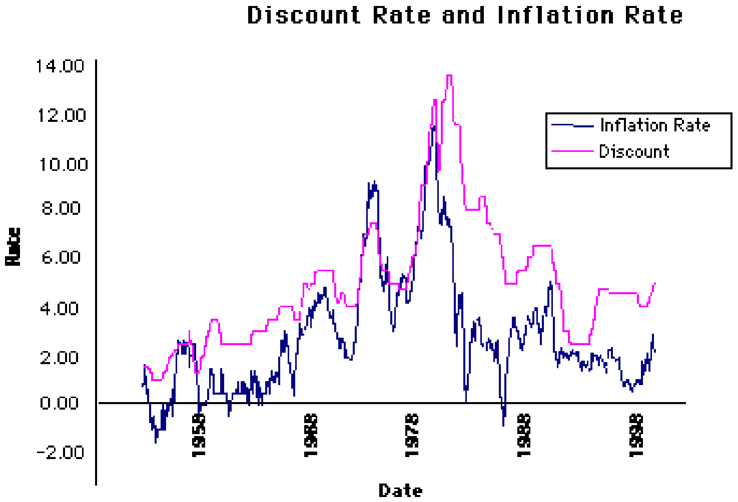

Figure 3. Relationship between Fed discount rate (leakage) and inflation over time.

Apparently, economists at the Federal Reserve who formulate the policy in which discount rate is used to control inflation are aware of the relationship shown in Figure 3 (Canterbery, 2000). Nevertheless, increases in discount rate are still thought to decrease inflation, but after a long delay. Unfortunately, the facts contradict even this hopeful interpretation of the data. The correlation between discount rate and inflation is still positive (and large) even when the discount rate from as much as a year earlier is correlated with current inflation rate. Apparently, the belief in the negative effect of discount rate on inflation persists because there is no way, using conventional economic models, to explain why increases in the discount rate (which decreases the amount of money in circulation) would lead to increases in inflation. Current economic models say that increases in the discount rate *should* decrease inflation. The *H. Economicus* model described in this paper does explain why increases in the discount rate would lead to increased inflation. Now that this phenomenon is explained, perhaps economists at the Federal Reserve will accept the data in Figure 3 as a representation of real phenomenon and act accordingly.

Conclusion

The *H. Economicus* model of the economy explains some of the most important observed economic phenomena in terms of collective control of monetary variables. The model suggests that some of our basic assumptions about what makes an economy function well may have to be revised. In the *H. Economicus* model, the economy functions best (low inflation, high productivity) when leakage (unspent consumer income) is low. Leakage is influenced by the monetary policies of the Federal Reserve Bank and by the distribution of income received by the composite consumer (when a small portion of the composite consumer receives a larger share of GNP than it can use to purchase Q' there is unspent income; leakage). If these influences on leakage can be controlled, a well functioning economy – one that works best for all its members – can be readily achieved and maintained.

5. Purpose in Systems Engineering

PERCOLATe: Perceptual Control Analysis of Tasks

This paper describes a new approach to task analysis based on perceptual control theory (Powers, 1973; 1990; Marken, 1992). Conventional task analysis (e.g. Kidd and Van Cott, 1972; Kirwin and Ainsworth, 1992) views the operator as an input-output device. The goal of conventional task analysis is to describe operator inputs, outputs and the rules that relate them. Perceptual control theory-based task analysis views the operator as a perceptual control system; the goal of the analysis is to determine the variables that the operator is to keep under control and the means the operator must have to effect this control.

The approach to task analysis described in this paper is called perceptual control analysis of tasks (PERCOLATe). PERCOLATe is an interview procedure that is designed to extract information from domain experts about how tasks are performed. PERCOLATe is based on the idea that all tasks involve control. A task involves control if the operator has some goal to achieve by carrying out the task. Thus, tasks like “searching” and “monitoring”, which are not typically seen as control tasks, fit into the PERCOLATe

---

From *The International Journal of Human-Computer Studies*, 1999, 50 (6), 481-487. Reprinted with permission of Academic Press perspective because they are done in order to achieve some goal; in search tasks the goal is to find a target; in monitoring tasks the goal is to respond to alarms.

PERCOLATe views the operator doing a task as a controller who is trying to keep perceptual representations of physical variables in preferred or reference states (see Figure 1). In the tasks to be discussed in this paper, the physical and informational variables to be controlled are typically represented on a computer display screen. In PERCOLATe these variables are called "displayed variables." The preferred or reference states of these variables exist in the operator's brain. In PERCOLATe these reference states are called "task objectives". Task objectives are the operator's representation of the intended or goal states of the displayed variables.

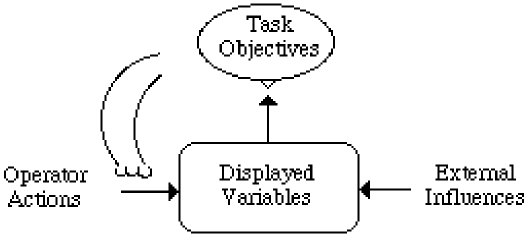

Figure 1. PERCOLATe model of a task.

The operator achieves task objectives by taking actions that bring the displayed variables to their goal states. These actions might include turning dials or throwing switches on a control panel or typing commands and clicking icons on a computer display. The effect of the operator's actions on displayed variables depends, to some extent, on disturbances to the variables that are represented as the displayed variables on the computer display. In PERCOLATe these disturbances are called "external influences." To a large extent, these external influences are the reason why an operator must be present. The operator must be available to take actions that may be necessary to prevent external influences from moving the displayed variables from their goal states (task objectives).

PERCOLATe recognizes that there is not one "right" set of actions that will achieve the task objectives. What the operator must do to achieve task objectives typically depends on the behavior of unpredictable external influences on displayed variables. This means that the same task objectives will be achieved by a different set of actions each time. This characteristic of task performance is not well captured in conventional task analysis, which represents tasks as though there were only one correct set of actions to be used to achieve task objectives. The result of a PERCOLATe task analysis is not a description of the particular set of actions that the operator must take to achieve task objectives. Rather, it is a description of the range of actions that the operator must have available in order to achieve a task objective in the context of all possible external influences to the displayed variables.

Satellite Control

PERCOLATe was developed as part of an effort to identify human-computer interface (HCI) requirements for satellite control. These requirements were to be based on an analysis of satellite control tasks as performed by satellite control operators in the Air Force Satellite Control Network. The PERCOLATe interview system was used to gain a generic, high-level picture of the tasks involved in satellite control from experts in several different areas of the satellite control process.

The PERCOLATe analysis was begun after satellite control experts had already decided on a decomposition of the satellite control process into a set of three satellite control task components. These task components were actually descriptions of high level task objectives. At the highest level, satellite control has three task objectives: 1) prepare for satellite support; 2) perform support; and 3) perform post-support analysis. These three task objectives were further broken down into the actual "tasks" to be analyzed using PERCOLATe. For example, preparation for a satellite support was broken down into three tasks 1) prepare a support plan; 2) configure ground resources; and 3) verify contact.

Task Objectives. The first step in the PERCOLATe process was to identify task objectives associated with each task. The task objectives were described as the results that had to be produced in order to complete the task. In most cases, the experts identified only two or three task objectives that had to be achieved in order to complete a task. For example, the experts identified three task objectives that had to be achieved in order to complete the "state of health (SOH) data collection/ verification" task: 1) request SOH data; 2) perform SOH data collection; and 3) process telemetry data.

Displayed Variables. The experts were then asked to think of all the variables that the operator controls in order to achieve a particular task objective; these are the variables that must be displayed to the operator. It was not always easy for the satellite control experts to think of the required displays as variables. For example, one task involved the creation of a plan for activities during a satellite pass – a pass plan. The satellite control experts found it difficult, at first, to think of the pass plan itself as a variable that had to be controlled. These experts had to learn that, in PERCOLATe, a variable is anything that can be in different states at different times; a pass plan is a variable that has to be displayed to the operator. A pass plan is a variable because it can be in several different states ranging from not completed to almost completed to completed. The experts eventually became comfortable with the idea that a controlled (or displayed) variable was anything (data, information, switch positions etc.) that the operator would have to be able to see and operate on in order to bring it to the state that corresponds to the task objective.

External Influences. Once a set of displayed variables have been identified, the experts were asked to identify external influences that might keep these variables from remaining in the state required to meet the task objective. The external influences on the variables displayed to satellite controllers include ground equipment failures, satellite anomalies, schedule conflicts and radio frequency interference. The identification of external influences is a unique and important aspect of the PERCOLATe analysis. It explains the problems the operator might encounter in the process of achieving task objectives and it provides a rationale for giving the operator the means of dealing with these problems. These means are the actions the operator can take to counter the effects of these external influences.

Operator Actions. Once the external influences on displayed variables were identified, the experts were able to describe the kinds of actions that the operator could take to compensate for the effects of these influences. The experts identified all the ways in which the operator must be able to affect the variables represented by the displayed variable in order to prevent external influences from interfering with achievement of the task objectives.

Results. The results of a PERCOLATe analysis of one satellite control task ("Verify Configuration") are shown in Table 1. The task involves verification of the equipment that has been reserved for a satellite support. The "task objective" column describes the intended or goal state of the displayed variables – the state that these variables should be in for successful accomplishment of the task. The "displayed variable" column is a list of all the variables that must be controlled – that is, the variables that must be brought to the state described in the task objectives. The "external influences" column is a description of all factors that might cause any of the displayed variables to deviate from the state described in the "task objectives" column. Finally, the "operator actions" column lists all of the ways in which the operator should be able to affect the state of the displayed variables.

Table 1. PERCOLATe Analysis of the "Verify Configuration" task.

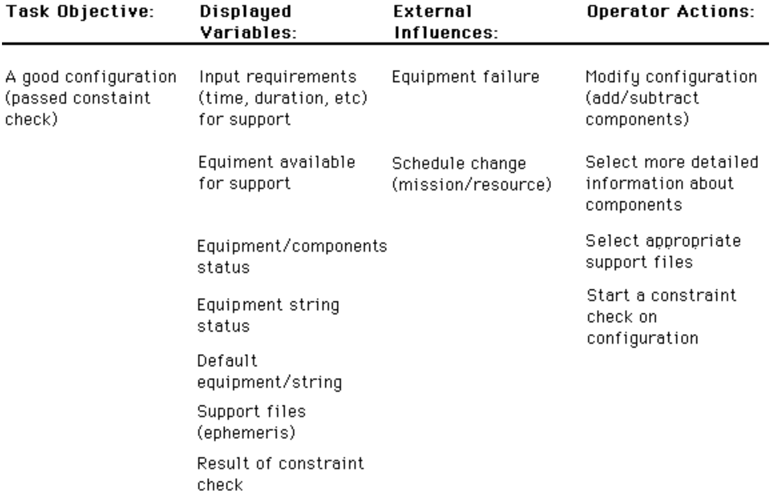

Task Objective:	Displayed Variables:	External Influences:	Operator Actions:
A good configuration (passed constaint check)	Input requirements (time, duration, etc) for support	Equipment failure	Modify configuration (add/subtract components)
	Equiment available for support	Schedule change (mission/resource)	Select more detailed information about components
	Equipment/components status		Select appropriate support files
	Equipment string status		Start a constraint check on configuration
	Default equipment/string		
	Support files (ephemeris)		
	Result of constraint check		

Level of Analysis

The PERCOLATe task description shown in Table 1 reflects a level of analysis that was appropriate for the purpose of developing HCI requirements for satellite control. If required, the analysis can be much more detailed. For example, in the analysis of the “verify configuration” task, the “task objectives” column could have been a list of all variables that make up a “configuration” and the exact states that these variables should be in to make the configuration to be controlled “good”. There could also have been a more detailed description of the “operator actions” used to control the variables listed in the “displayed variables” column.

The level of detail of a PERCOLATe analysis depends on how the analysis will be used. The PERCOLATe analysis of satellite control tasks was aimed at identifying general categories of HCI requirements unique to satellite control; a high level PERCOLATe analysis proved sufficient for this purpose. A far more detailed analysis would be needed if the goal of the analysis were to provide a basis for detailed HCI design.

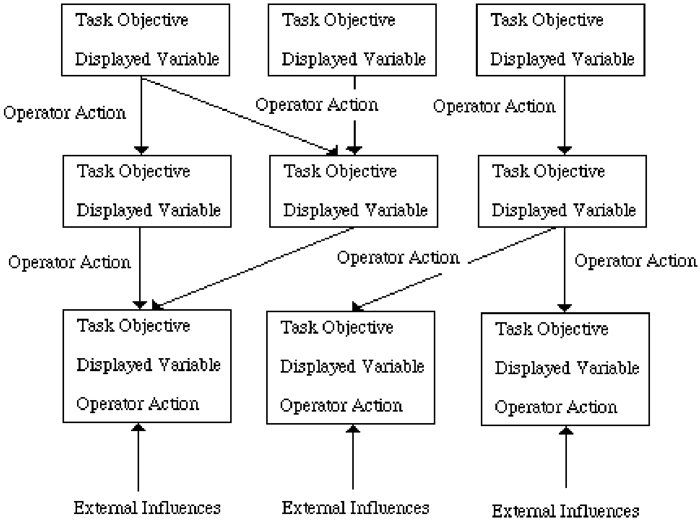

Figure 2. Hierarchical PERCOLATe model of a task.

Hierarchical Control. One useful way to increase the level of detail achieved in a PERCOLATe analysis is by determining "how" each task objective is accomplished; this gives a *hierarchical* picture of the control processes involved in performing a task. How a task objective is accomplished can be determined by a PERCOLATe analysis of the actions that the operator must take to achieve it. The actions that achieve this "higher level" task objective are themselves task objectives for a lower level task. There is a hierarchical relationship between task objectives because the same set of lower level task objectives may be used as the means (actions) to accomplish several different higher level tasks.

One possible hierarchical mapping between tasks is shown in Figure 2. Each node in this hierarchy can be thought of as a PERCOLATe task description like that in Table 1. Figure 2 shows that different lower order task objectives can be the “operator actions” that are used to achieve higher level task objectives. For example, one of the three highest level task objectives at the top of Figure 2 can be thought of as the objective of the “Verify Configuration” task: a good configuration. The two arrows coming out of this “task objective” box can be thought of as two of the “operator actions” that can be used to achieve the “good configuration” objective: 1) modify configuration and 2) select more detailed information. These two operator actions become “task objectives” at the next lower level in the hierarchy. The “operator actions” that are used to achieve these task objectives become task objectives at the next lower level of the hierarchy. The goal of the hierarchical PERCOLATe analysis is to “drill down” to determine the lower level “task objectives” that are the “operator actions” used to accomplish the higher level “task objectives.”

Each of the lower level PERCOLATe task descriptions will include the displayed variables, external influences and operator actions involved in achieving the lower level task objective. Only the lowest level task objectives are achieved by actually taking action on the environment (such as typing data with a keyboard or selecting icons with a mouse). Higher level task objectives (like verifying a configuration) are achieved by selecting various lower order task objectives (like producing a display of the configuration requirements).

HCI Design

The results of a PERCOLATe analysis are directly relevant to HCI design. In particular, the description of displayed variables and operator actions that influence those variables tell the HCI designer what variables must be represented on the computer display and what kinds of computer inputs (actions) must be available to influence the state of these variables. The HCI designer chooses how to represent the variables (graphics, text etc.) and the kinds of computer inputs the operator can use to influence these variables (mouse selection, text input etc.). The lower levels of the hierarchical PERCOLATe analysis specify the variables that are controlled via navigation techniques. These are the lower level displayed variables that must be controlled as the means of controlling higher order displayed variables.

The external influences column lets the HCI designer know why certain actions must be made available to the operator. In some cases the HCI designer might decide to inform the operator of the nature of the external influences that are thought to be affecting a displayed variable.

The PERCOLATe task objectives describe display states that the operator must adopt as goal states for the displayed variables. A description of these task objectives should be built into the HCI by the HCI designer as a reminder to the operator. But, ultimately, the operator must learn these task objectives; this is the job of training. The operator must learn what variables to control (the displayed variables), the goal states of these variables (task objectives) and the actions that can be taken to achieve these task objectives (operator actions).

A PERCOLATe task analysis provides a basis for the design of a training program since it specifies what the operator should know (task objectives), what the operator should monitor to determine whether these objectives are met (displayed variables) and what actions to take of if these objectives are not met (operator actions).

Summary

PERCOLATe is an approach to task analysis that is based on the notion that the operators performing a task are perceptual control systems. A perceptual control system acts to bring perceptual (displayed) variables to reference states (task objectives) while protecting these variables from the effects of disturbances (external influences). PERCOLATe can be performed at various levels of analysis to address a range of task analytic needs. A PERCOLATe analysis describes tasks in terms of the main components of this control process: task objectives, displayed variables, operator actions and external influences. A PERCOLATe analysis provides a very practical blueprint for HCI design. It tells the HCI designer what variables to display to the operator and what actions the operator should be able to take to influence this display. PERCOLATe analysis also provides a basis for the development of training.

References

Albus, J. *Brains, behavior and robotics.* Petersborough, NH: Byte Books

Adair, R. K. (1994) *The physics of baseball.* New York: HarperCollins

Albus, J. (1981) Brains, behavior and robotics. Petersborough, NH: Byte Books

Anderson, J. R. (1983) *The architecture of cognition*. Cambridge, MA: Harvard University

Arbib, M. (1972) *The metaphorical brain*. New York: Wiley

Atwood, R. and Polson, P. (1976) A process model of water jar problems, *Cognitive Psychology*, 8, 191-216

Babler, T. and Dannemiller, J. (1993) Role of acceleration in judging landing location of free-falling objects, Journal *Experimental Psychology: Human Perception and Performance*, 19, 15-31

Black, H. S. (1934, January). Stabilized feed-back amplifiers. *Electrical Engineering*, pp.114-120

Bordens, K. S. and Abbott, B. B. (1991) *Research design and methods: A process approach*, Mountain View: Mayfield

Bourbon, W. T. (1996) On the accuracy and reliability of predictions by perceptual control theory: five years later. *The Psychological Record*, 46, 39-47

Bryan, W. L. and Harter, N. (1899) Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6, 345-375

Canterbery, E. Ray (2000) *Wall Street Capitalism: The Theory of the Bond-Holding Class*, Singapore: World Scientific

Chapman, S. (1968) Catching a baseball, *American Journal of Physics*, 36, 868-870

Cohen, J. D., Dunbar, K. and McClelland, J. L. (1991) On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332-361

Craik, K. J W. (1943) *Psychological and physiological aspects of control mechanisms with special reference to tank gunnery. Part 1.* Medical Research Council (Great Britain), Military Personnel Research Committee, BPC 43/254.

Dannemiller, J. L., Babler, T. G. and Babler, B. L. (1996) Technical comment, *Science*, 273, 256-257

Davis, W. J. (1976) Organizational concepts in the central motor network of invertebrates. In R. M. Herman, S. Grillner, P.S.G. Stein, & D. Stuart (Eds.) *Neural control of locomotion*, New York: Plenum (p. 265)

Dienes, Z. and McLoed, P. (1993) How to catch a cricket ball, *Perception*, 22, 1427-1439

Domjan, M. (1987) Animal learning comes of age. *American Psychologist*, 42, 556-564

Greeno, J.G. and Simon, H.A. (1974) Processes for sequence production. *Psychological Review*, 81, 187-197

Herrigal, E. (1971) *Zen in the art of archery*. New York: Vintage

Hubel, D. H. and Wiesel, T. N. (1979) Brain mechanisms of vision. In J.M. Wolfe (Ed.*) The mind's eye: Readings from Scientific American*, New York: Freeman (pp. 40-52)

Jacobs, T. M., Lawrence, M. D., Hong, K., Giordano, N. and Giordano, N. (1996) Technical comment: On catching fly balls, *Science*, 273, 257-258

James, W. (1980) *The principles of psychology*. New York: Dover

Jordan, M. I. (1989) Serial order: A parallel, distributed processing approach. In J. L. Elman and D. E. Rumelhart (Eds.) *Advances in connectionist theory: Speech*. Hillsdale, NJ: Erlbaum

Keele, S. W., Cohen, A. & Ivry, R. I. (1990) Motor programs: Concepts and issues. In M. Jeannerod (Ed.) *Attention and performance XIII: Motor representation and control*. New Jersey: Erlbaum. (pp. 77-110)

Keele, S. W., Pokorny, R. A., Corcos, D. M. & Ivry, R. I. (1985) Do perception and production share common timing mechanisms: A correlational analysis. *Acta Psychologica*, 60, 173-191

Kline, R. (1983) Comment on Rosenbaum et al. Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 834-36

Kolers, P. (1972) The illusion of movement. In R. Held and W. Richards (Eds.) *Perception: Mechanisms and models*, San Francisco: Freeman

Kidd, R. and H. Van Cott (1972) Task Analysis. In Van Cott, H. P. and Kinkade, R. G. (Eds.) *Human Engineering guide to equipment design* (Wash. D C: American Institute for research)

Kirwin, B. and Ainsworth, L. K. (1992) *A guide to task analysis*. London: Taylor and Francis

Lashley, K. S. (1951) The problem of serial order in behavior. In L. A. Jeffress (Ed.) *Cerebral mechanisms in behavior*. New York: Wiley

Lesgold, A., Robinson, H., Feltovitch, P., Glaser, R., Klopfer, D. and Wang, Y. (1988) Expertise in a complex skill: Diagnosing X-ray pictures. In M.T.H. Chi, R. Glaser and M.J. Farr (Eds.) *The nature of expertise*. Hillsdale, NJ: Erlbaum

Lorenz, K. Z. (1981) *The Foundations of Ethology*. Springer-Verlag: New York.

Marken, R. S. (1981) *Methods in experimental psychology*. Monterey: Brooks/Cole

Marken, R. S. (1982) Intentional and accidental behavior: a control theory analysis. *Psychological Reports,* 50, 647-650.

Marken, R. S. (1985) Selection of consequences: Adaptive behavior from random reinforcement. *Psychological Reports,* 56, 379-383

Marken, R. S. (1986) Perceptual organization of behavior: A hierarchical control model of coordinated action. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 267-76

Marken, R. S. (1988) The nature of behavior: Control as fact and theory, *Behavioral Science,* 33, 196-20

Marken, R. S. (1989) Behavior in the first degree. In W. Hershberger (Ed.) *Volitional action: Conation and control*. Amsterdam: North Holland (pp. 299-314)

Marken, R. S. and Powers, W. T. (1989) Levels of intention in behavior. In W. Hershberger (Ed.) *Volitional Action:*

Conation and Control, Elsevier Science Publishers: North-Holland

Marken, R. S. (1990) Spreadsheet analysis of a hierarchical control system model of behavior, *Behavior Research Methods, Instruments, & Computers*, 22, 349 - 359

Marken, R. S. (1990) A science of purpose, *American Behavioral Scientist*, 34, 6-13

Marken, R. S. (1991) Degrees of freedom in behavior, *Psychological Science*, 2, 92 - 100

Marken, R. S. (1992) *Mind readings: Experimental studies of purpose.* Gravel Switch, KY: CSG Press

Marken, R. S. (1993) The blind men and the elephant: Three views of the phenomenon of control, *Closed Loop*, 3, 37-46

Marken, R. S. (1997) The dancer and the dance: Methods in the study of living control systems, *Psychological Methods*, 2,436-446

Martin, J. G. (1972) Rhythmic (hierarchical) versus serial structure in speech and other behavior. *Psychological Review*, 79, 487-509

McBeath, M. K., Shaffer, D. M, and Kaiser, M. K. (1995); How baseball outfielders determine where to run to catch fly balls, *Science*, 268, 569-573

McBeath, M. K., Shaffer, D. M, and Kaiser, M. K. (1996) Technical comment: response, *Science*, 273, 258-260

McLeod, P. and Dienes, Z. (1996) Do fielders know where to go to catch the ball or only how to get there? *Journal Experimental Psychology: Human Perception and Performance*, 22, 531-543

McRuer, D. T. and Krendel, E. S. (1959) The human operator as a servo system element. *Journal of the Franklin Institute*, 267, 381-403

Miller, G. A. and Heise, G. A. (1950) The trill threshold, Journal of the Acoustical Society of America, 22, 637-638

Newell, K. (1990) *Unified theories of cognition*. Cambridge, MA: Harvard University Press

Palmer, S. E. (1977) Hierarchical structure in perceptual representation. *Cognitive Psychology*, 9, 441-474

Peper, L. Bootsma, R. J., Mastre, D. R. and Bakker, F. C. (1994) Catching balls: How to get the hand to the right place at the right time. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 591-612

Pew, R. W. (1966) Acquisition of hierarchical control over the temporal organization of a skill. *Journal of Experimental Psychology*, 71, 764-771

Plooij, F.X. and van de Rijt-Plooij, H.H.C. (1990) Developmental transitions as successive reorganizations of a control hierarchy. *American Behavioral Scientist*, 34, 67-80

Povel, D-J (1981) Internal representation of simple temporal patterns. *Journal of Experimental Psychology: Human Perception and Performance,* 7, 3 - 18

Powers, T. C. (1996) Leakage*: The Bleeding of the American Economy*, New Canaan, Connecticut: Benchmark

Powers, W. T. , Clark, R. K. and McFarland, R. L. (1960) A general feedback theory of human behavior: Part II. *Perceptual and Motor Skills*, 11, 309- 323

Powers, W. T. (1971) A feedback model of behavior: Analysis of a rat experiment. *Behavioral Science*, 16, 558-563

Powers, W. T. (1973a) *Behavior: The control of perception*, New York: Aldine-DeGruyter

Powers, W. T. (1973b) Feedback: Beyond behaviorism, *Science*, 351-356

Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, 85, 417-435

Powers, W. T. (1979a). The nature of robots, Part 1: Defining behavior. *Byte*, 4(6), 132-144

Powers, W. T. (1979b). The nature of robots, Part 2: Simulated control system. *Byte*, 4(7), 134-152

Powers, W. T. (1979c). The nature of robots Part 3: A closer look a human behavior. *Byte*, 4(8), 94-116

Powers, W. T. (1979d). The nature of robots, Part 4: Looking for controlled variables. *Byte*, 4(9), 96-112

Powers, W. T. (1979e). A cybernetic model for research in human development. In M. Ozer (Ed*.). A cybernetic approach to the assessment of children: Toward a more humane use of human beings* (pp.11-66). Boulder, CO: Westview

Powers, W. T. (1989) Living *control systems*. Gravel Switch, KY: CSG Press

Powers, W. T. (1990) Control theory and statistical generalizations, *American Behavioral Scientist*, 34, 24-31

Powers, W. T. (1992) *Living control systems II: Selected papers*, Chapel Hill, NC: New View

Rescorla, R A. (1987) A Pavlovian analysis of goal directed behavior, *American Psychologist*, 42, 119-129

Robertson, R. J. and Powers W. T. (1990). *Introduction to modern psychology: The control-theory view*. Gravel Switch, KY: Control Systems Group, Inc.

Rosenbaum, D.A. (1987) Hierarchical organization of motor programs. In S.P.Wise (Ed) Higher brain functions: Recent explorations of the brain's emergent properties, New York:Wiley

Rosenbaum, D. A., Kerry, S. and Derr, M. A. (1983) Hierarchical control of rapid movement sequences. *Journal of*

Experimental Psychology: Human Perception and Performance, 9, 86 - 102

Runkel, P. J (1990) *Casting nets and testing specimens: Two grand methods of Psychology*, New York: Praeger

Shadish, W. R. (1996) Meta-analysis and the exploration of causal mediating processes: A primer of examples, methods and issues. *Psychological Methods*, 1, 47-65

Shepard, R. N. (1990) Neural nets for generalization and classification: Comment on Staddon and Reid (1990). *Psychological Review*, 97, 579 - 580

Shepard, R. N. (1987) Toward a universal law of generalization for psychological science. *Science*, 237, 1317 - 1324

Simon, H. A. (1992) What is an “explanation” of behavior? *Psychological Science*, 3, 150 - 161

Skinner, B. F. (1971) *Cumulative record*. (Revised edition) New York: Appleton

Skinner, B. F. (1981) Selection by consequences. *Science*, 213, 501-504

Staddon, J. E. R. (1979) Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General*, 108, 48-67

Teitelbaum, P. (1966) The use of operant methods in the assessment and control of motivational states. In W. K.

Honig (Ed.) *Operant behavior*, New York: Appleton-Century-Crofts

Timberlake, W. (1984) Behavior regulation and learned performance: Some misapprehensions and disagreements. *Journal of the Experimental Analysis of Behavior*, 41, 355-375

Tresilian, J. R. (1995) Study of a servo-control strategy for projectile interception, *The Quarterly Journal of Experimental Psychology*, 48A, 688-715

U. S. Census Bureau (2000), *Statistical Abstract of the United States*

Warren, W. H. Jr., Young, D. S. and Lee, D. N. (1986) Visual control of step length during running over irregular terrain. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 259 - 266

Watson, J. B. (1913) Psychology as the behaviorist views it. *Psychological Review*, 20, pp. 158-177.

Wickins, C. D. (1987) The effect of control dynamics performance. In K. R. Boff, L. Kaufman and J. P. Thomas (Eds.), *Handbook of perception and human performance. Vol. 2, Cognitive processes and performance* (chap.39). New York: Wiley.

Wing, A. M. and Lederman, S. J. (1998) Anticipating load torques produced by voluntary movements. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1571-1581