Without Miracles

8 Adapted Behavior as the Control of Perception

- The Insights of James, Dewey, and Tolman
- An Introduction to Control Systems
- Perceptual Control Theory

What we have is a circuit, not an arc or broken segment of a circle. This circuit is more truly termed organic than reflex, because the motor response determines the stimulus, just as truly as sensory stimulus determines movement. Indeed, the movement is only for the sake of determining the stimulus, of fixing what kind of a stimulus it is, of interpreting it.

--John Dewey[1]

In chapters 3 and 7 we examined several of attempts by philosophers, biologists, ethologists, and psychologists to provide accounts for the fit of organisms' behavior to their environment. We first examined instinctive behavior and showed how attempts to explain its fit progressed from the providential theories of Aquinas and Paley to the instructionist theories of Lamarck and Erasmus Darwin, and then to the selectionist theories of Charles Darwin, Wallace, and Lorenz. We considered learned behavior and saw how the Pavlovian conditioning theory of Pavlov and Watson and the operant conditioning theory of Thorndike and Skinner attempted to explain how behavior can adapt to environmental conditions over the lifetime of an organism through the formation of new stimulus-response connections.

But we also noted one serious omission in all of these theories--they provide no adequate explanation for the goal-oriented, *purposeful* or *intentional* nature of adapted behavior. In addition, Pavlovian conditioning emphasizes a one-way transmission of instructions from stimulus to response, and both Pavlovian and operant conditioning see the role of the environment as a sort of behavioral conductor that orchestrates all of the adaptive changes in behavior. In this chapter we will consider these issues in greater detail and introduce a radically different approach to understanding the adaptive nature of behavior.

The Insights of James, Dewey, and Tolman

It was over 100 years ago in 1890 that the influential American philosopher and psychologist William James (1842-1910) strikingly characterized the purposeful nature of the behavior of living things in contrast with the movements of inanimate objects:

> Romeo wants Juliet as the [iron] filings want a magnet; and if no obstacles intervene he moves toward her by as straight a line as they do. 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.[2]

To make it clear that such purposeful behavior was not uniquely human, James provided an example from the amphibious world:

> Suppose a living frog . . . at the bottom of a jar of water. The want of breath will soon make him also long to rejoin the mother-atmosphere, and he will take the shortest path to his end by swimming straight upwards. But if a jar full of water be inverted over him, he will not, like the bubbles, perpetually press his nose against its unyielding roof, but will restlessly explore the neighborhood until by re-descending again he has discovered a path round its brim to the goal of his desires. Again the fixed end, the varying means![3]

Even though James's *Principles of Psychology*, from which these passages are taken, was considered to be the most important psychological work of its day, the stimulus-response, conditioning theories of Pavlov, Watson, Thorndike, and Skinner all totally discounted purpose as having a role in a scientific account of behavior. In fact, Skinner repeatedly used the analogy of biological evolution to argue against purpose in behavior. For example:

> Evolutionary theory moved the purpose which seemed to be displayed by the human genetic endowment from antecedent design to subsequent selection by contingencies of survival. Operant theory moved the purpose which seemed to be displayed by human action from antecedent intention or plan to subsequent selection by contingencies of reinforcement. A person disposed to act because he has been reinforced for acting may feel the condition of his body at such time and call it "felt purpose," but what behaviorism rejects is the causal efficacy of that feeling.[4]

Darwin showed how natural selection leads to adapted complexity in the structure and instinctive behavior of organisms, without purpose either on the part of the organism or on the part of a supernatural provider. So Skinner is arguing here that the selection of behavior by the environment can similarly explain the adapted complexity of learned behaviors without recourse to purpose. This is accomplished by forming new stimulus-response connections strengthened by environmental reinforcement.

But such a stimulus-response view of learning is both seriously incomplete and misleading. In addition to the problems mentioned in the previous chapter, American philosopher and educator John Dewey (1859-1952) provided another. He observed that the stimulus-response interpretation of behavior was flawed since it recognizes that stimuli influence responses, but it neglects the equally important fact that *responses also influence stimuli*. Consequently, he criticized the concept of the stimulus-response reflex arc by noting that

> what we have is a circuit, not an arc or broken segment of a circle. This circuit is more truly termed organic then reflex, because the motor response determines the stimulus, just as truly as sensory stimulus determines movement. Indeed, the movement is only for the sake of determining the stimulus, of fixing what kind of a stimulus it is, of interpreting it.[5]

To understand how behavior can adapt to its environment we have to consider just what Dewey meant by his statement that behavior determines stimulus just as much as stimulus determines behavior, an insight that was almost totally ignored until the middle of this century. As we have seen, Pavlovian and operant conditioning theories view all adapted behavior as responses to external stimuli, including those caused by the behavior of other organisms. Acid is placed in a dog's mouth and it salivates. A hungry rat is placed in a familiar cage where in the past food was made available whenever it pressed the bar within two seconds after a bell sounded, and so it immediately proceeds to push the bar after hearing yet again the familiar peal of the dinner bell. To use the more technical terms of psychological jargon, the stimulus is considered to be the independent variable and the response the dependent variable; that is, the response *depends* on the stimulus, and the stimulus is *independent* of the response.

Another way to conceptualize this one-way view of the relationship is to see the stimulus as the sole determining *cause* of the response, and the stimulus to be isolated from any effects of the response. But Dewey took exception to this view, stating that the *response also causes the stimulus*. But how can this be? How can the dog's salivating influence the presence of acid in its mouth? How can the rat's pressing of the bar in any way cause the bell to be rung, the sounding of which is controlled by the experimenter?

To appreciate Dewey's important insight, we have to do something that is rarely done in experimental psychology. We have to abandon the point of view of an objective outside observer and instead attempt to imagine how things appear from the organism's point of view. This shift in perspective brings with it a realization that a stimulus can have an effect on an organism only insofar as it is experienced or *perceived* by the organism. A dog who cannot perceive that acid has been (or is about to be) placed in its mouth will not salivate when the acid is so placed (or about to be). Once we imagine how the world appears to the organism, Dewey's point about response influencing stimulus begins to make sense. Actually, it is not necessary to consider the world as it may appear to dogs and rats. A cursory look around our own world will quickly and clearly reveal that behavior causes changes in perception as much as perception causes changes in behavior.

As you move your eyes across this page you cannot fail to notice that as you do so your perception of the page changes. We might reasonably conclude that we move our eyes for the purpose of bringing into view that which we want to read next, and that this in fact is the very reason for the behavior. So response does influence stimulus and so behavior does influence perception. But this is not to deny that stimulus also exerts an influence on response. A poorly written sentence or complex sentence may well lead to your returning to it in a second attempt to decipher its meaning. And a loud sound coming from behind your head will likely have you quickly turning around to see what has happened. So proper understanding of behavior has to take into account the reciprocal give-and-take relationship of stimulus and response, something neither Pavlovian conditioning nor operant conditioning proposed to do.

The purposeful nature of animal behavior was clearly demonstrated by the research conducted by psychologist Edward C. Tolman (1886-1959) and his students at the University of California at Berkeley from the 1920s to the 1950s. Among the best known of these studies was one conducted by Tolman's student D. A. Macfarlane in which rats learned to swim through a maze to obtain a food reward.[6] After they had learned to do this well, a raised floor was installed in the maze so that the rats now had to wade through the maze to get to the goal box. It was hypothesized that if the rats' learning consisted of acquiring specific swimming behaviors (that is, specific responses to specific stimuli), they would have to relearn the maze in the wading condition, as the movements and stimuli involved in wading are very different from those involved in swimming. It was found instead that after a very brief period of adjustment to the new situation (just one "run" through the maze), the rats performed as well in the new wading condition as they had in the old swimming condition. This was a clear demonstration that what the rats had learned while swimming the maze could not be described as the acquisition of stimulus-response connections but rather as more general knowledge about the location of the goal box, since it made little difference to the rats whether they swam or waded to their destination. Similarly, once a person knows how to get to a specific location by driving a car, he can also get there by bicycle (if he knows how to ride one) or by walking (if not too far), regardless of the fact that the stimuli and responses differ greatly from one mode of transportation to another.

Regardless of these findings and many others like them, Tolman was never able to eliminate the concept of stimulus-response connections from the very core of his theory of purposeful behavior. Indeed, his attempt to explain how behavior can vary and yet reach a consistent goal involves imagining long, complicated, invisible chains of such connections existing within the organism in the form of *intervening variables*, and conceiving of responses not as specific muscular contractions but rather as a *performance*. With respect to the latter Tolman wrote:

> It is to be stressed . . . that for me the type of response I am interested in is always to be identified as a pattern of *organism-environment rearrangements* and not as a detailed set of muscular or glandular activities. These latter may vary from trial to trial and yet the total "performance" remains the same. Thus, for example, "going towards a light" is a *performance* in my sense of the term and is not properly a response (a set of muscular contractions).[7]

But substituting the word "performance" for "response" does nothing to explain how an organism is able to accomplish a repeatable "organism-environment rearrangement" by responding to stimuli; it simply states that it somehow happens. If "behavior may vary from trial to trial and yet the total `performance' remains the same," how is it that the organism is able to vary its behavior to arrive at a desired goal?

Nonetheless Tolman made an important initial step toward solving this problem in his realization that sensory *feedback* was important; that is, the rat's behavior changed the stimuli it perceived and this feedback was essential in guiding the organism toward the final goal.[8] But Tolman never provided an explicit model for just how such a system would work, and so he never broke out of the behaviorist tradition of considering stimuli as causes of behavior. The first successful attempt to develop a working model of purposeful behavior would have to await the development of negative-feedback control systems in engineering and their application to the life sciences by a few bold pioneers with interests and expertise in both engineering and psychology.

An Introduction to Control Systems

Macfarlane's swim-first-wade-later experiment provided clear evidence that the rat's behavior can be purposeful. Careful observation of the naturally occurring behavior of animals and people is all that is necessary to lead us to the same conclusion. Environments do not normally remain cooperatively still as behavior takes place, particularly not when they consist of other competing organisms. No fixed, predetermined pattern of muscular contractions and tongue movements will guarantee the frog's success in hunting flies. Nor will some unvarying pattern of wing movements ensure the fly's success in avoiding the frog's lunging tongue. To facilitate an animal's survival and reproduction, behavior must somehow take account of an ever-changing and unpredictable environment. A particularly striking example is the nest-provisioning behavior of the solitary wasp:

> Sometimes she drops the fly behind her, and then turning around, pulls it in [the nest] with her mandibles. In other cases, where a longer portion of the tunnel has been filled with earth, the fly is left lying on the ground while the wasp clears the way. The dirt that is kicked out sometimes covers it so that when the way is clear the careless proprietor must search it out and clean it off before she can store it away. In one instance, in which we had been opening a nest close by, the tunnel was entirely blocked by the loose earth which we had disturbed, and the wasp worked for ten minutes before she cleared a way to her nest.[9]

Countless examples of purposeful human behavior are easy to find. One with which many readers have daily personal experience is driving a car from home to work. Such behavior requires a remarkably complex, coordinated pattern of behavior of the fingers, hands, and arms operating on the steering wheel, shift lever, and turn indicator lever, while the legs and feet operate the accelerator, clutch, and brake pedals. Many patterns of behavior will get the driver and car safely to work (depending on the speed of the car and the particular route taken), but just slight changes in any one may have fatal consequences. And yet, no one set pattern of behavior will always be successful. Traffic may be heavy or light, fast moving or slow. The road surface may be dry and firm or wet and slippery. The engine may be responsive or balky. Road construction and traffic accidents may make deviations from the normally preferred route necessary. Because of these changeable conditions and other unpredictable disturbances, an exact replay of the driving behavior that was successful in getting you to work on Monday would certainly not get you to work on Tuesday. What is true for the person driving a car to work is also the case for the person walking from bedroom to bathroom, the bee finding and collecting nectar, the fox pursuing the hare and the hare avoiding the fox, the monarch butterfly migrating from Mexico to Canada, and the greylag gosling closely following the steps of its mother (or Konrad Lorenz). It is difficult to see how any complex behavior can remain adapted to some function if the organism does not continually make adjustments to it *while it is performing the behavior*.

But how is this possible? Can we imagine the functioning of an organism, made up of sense organs connected by a nervous system and brain to muscles, that is able to pursue a goal by continually modifying its behavior to adjust for environmental disturbances that would cause any fixed pattern of behavior to miss its mark? Yes, we can. The explanation comes to us not from the life sciences of biology, ethology, or psychology, but from electrical and mechanical engineers who in the 1930s began to make devices that could duplicate the purposeful behavior of humans.

These devices are known as *control systems*, and they work using the same type of stimulus-response circuit, or *loop*, that Dewey first mentioned in 1896. They also behave in the same variable-path but fixed-goal manner James described over a century ago as characteristic of animal and human behavior. Such devices are now commonly found in a multitude of electronic devices from the simple thermostat of a home heating system to the highly complex guidance system of antiaircraft missiles. To see how they work, we will consider one that is now commonly found on automobiles--the cruise control that automatically maintains a steady driving speed with no assistance from the driver.[10]

The cruise control system is engaged by first turning it on and then pushing the "set" button after the car has reached the desired speed. This speed then becomes the system's goal, or what control system engineers refer to as the reference level, and the system will then increase or decrease the amount of fuel it delivers to the motor as necessary to maintain this speed. If the car begins to climb a hill, the cruise control system will sense a reduction in speed (being equipped with a speedometer that measures the speed of rotation of the wheels) and will provide more fuel to the engine through a mechanical linkage to the throttle, causing it to increase its power output to maintain the speed despite the hill. As the car begins to descend the other side of the hill, the cruise control system will sense the increased speed, which will cause it to close the throttle, reducing the amount of fuel delivered to the engine so that again the desired speed is maintained. Because the system responds to too-high speeds by reducing the amount of fuel delivered to the motor and to too-low speeds by increasing the flow of fuel, it is referred to as a *negative feedback* system.

A clearer idea of the nature and functioning of a simple control system can be obtained by examining figure 8.1. The *sensor* converts some variable aspect of the environment (for example, light, sound, or speed) into a sensor signal (s), which varies from zero to some higher positive value. This sensor signal (s) is then compared with a *reference signal* (*r*) in the *comparator*, which subtracts *s* from *r* yielding an error signal *e*. This error signal is then amplified and converted by the *activator* into behavior (*o* for output). This behavior then acts on the *environment*, changing it in the intended direction, which again provides input (*i*) to the *sensor*, thereby closing the loop. However, it is not only the control system's output that influences the input to sensor, but also disturbances (*d*) emanating from the environment. So the feedback resulting from the control system's own behavior and the current disturbance from the environment combine to provide the input to the sensor.

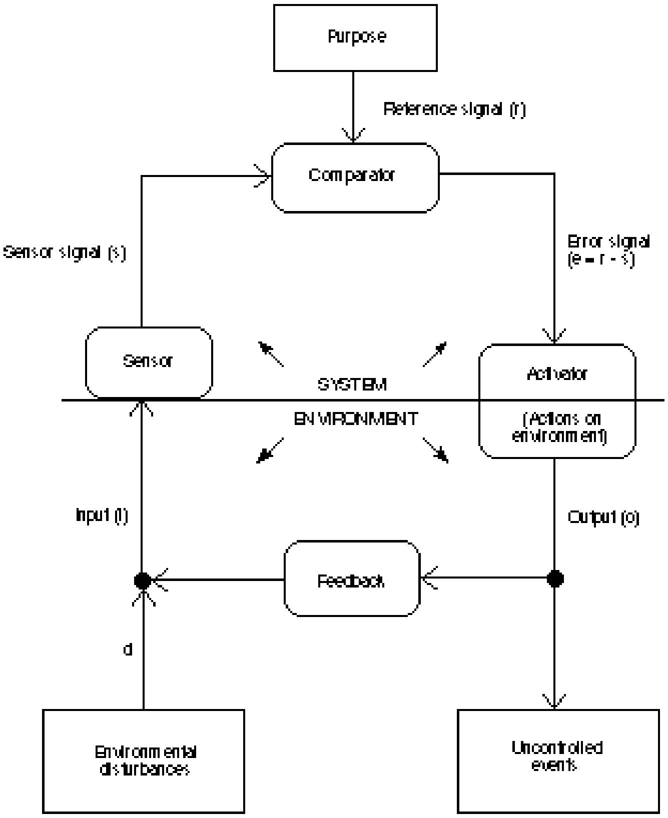

A cruise control system acts very much like a human driver, with the goal of maintaining a given speed. To do this, the driver must attentively monitor the speedometer. If the speed drops, the driver must press down on the accelerator pedal. If the speed rises, the driver must reduce pressure on the accelerator. While the car is moving at the desired speed, no action is called for. It should come as no surprise that the cruise control system mimics the functioning of a human driver so well, as this is exactly what it was designed to do.

The cruise control system has a number of intriguing aspects that are shared by all properly functioning control systems. First, it does not perceive the actual disturbances for which it must compensate. It has no way of determining whether the road is climbing or descending. It cannot tell if there is a stiff headwind or tailwind. It cannot know if a heavy trailer was attached to the car at the last stop, if a tire is losing air and offering steadily increasing rolling resistance, or if a spark plug has fouled causing a cylinder to fail and the engine to lose power. All it can sense, and therefore control, is the car's speed. Yet despite its complete ignorance of the multitude of interacting influences, it does a very good job of maintaining the desired speed.

Second, *a control system does not control what it does; it controls what it senses*. The word *control* is used here in its precise technical sense of *maintaining some variable at or near specified fixed or changing values* regardless of the disturbances that would otherwise influence it to vary. The cruise control system can only control what it senses to be the speed of the vehicle, and it does so by changing its output as required, that is, by delivering varying amounts of fuel to the engine. The only way it can maintain its sensing of the car's speed close to the reference level speed in the face of disturbances is to vary its output (change its behavior) as necessary. So we see that it controls its *input* (what it senses) and not its *output* (or behavior). Consequently, using a cruise control system to maintain a constant speed on a trip will allow you to predict accurately how long it will take to cover a certain distance. It not let you predict how much fuel will be used in getting there, because fuel consumption is not controlled, varying as it must to compensate for unpredictable disturbances. Since a control system controls what it senses, and since an organism's sensing of the environment is generally referred to as perception in the behavioral sciences, the application of control theory to the behavior of living organisms is known as *perceptual control theory* to distinguish it from the control theory applied by engineers and physicists to the inanimate world.

Finally, it is important to realize that whereas a control system's behavior is clearly influenced by the environment, it is not determined solely by the environment. Rather, its behavior is determined by what it senses (or perceives) *in comparison with its internal goal or reference level*. And it is here that we find a crucial difference between the nonliving control systems designed by engineers and the living ones fashioned by biological evolution. An engineered control system is usually designed so that its reference level can be manipulated by the operator, for example, by pushing the "accelerate" button of the cruise control system or by turning up the house thermostat. No such direct manipulation of the reference levels of living control systems is usually possible. We can certainly ask our taxi driver to drive more slowly or our child to be home before midnight, but there is no way to guarantee that the person will comply with our wishes.

Control systems were specifically designed to replace human operators in jobs calling for the control of important variables (for example, steam pressure in a boiler) and have been in wide use since the 1930s. They were almost completely ignored by the behavioral sciences, however, until the appearance of Norbert Wiener's groundbreaking book *Cybernetics* in 1948, which was quickly followed by important related works by W. Ross Ashby.[11] William T. Powers, an American control system engineer, was also struck by how the behavior of such systems resembled the purposeful behavior of living organisms, and it is primarily due to his work that a control system theory of human and animal behavior exists today.[12]

Perceptual Control Theory

Perceptual control theory does what no behaviorist stimulus-response theory has ever been able to do--it provides an explicit, working model that accounts for goal-oriented, purposeful behavior. Behavior must be purposeful if it is to enable an organism to survive and reproduce despite the unpredictable disturbances the organism continually encounters. Indeed, we could consider behavior to be adaptedly complex only insofar as it is able to achieve its purposes regardless of the continual challenges posed by environmental disturbances. The perspective provided by perceptual control theory, however, has not been received enthusiastically by behavioral scientists. The principal reason for its neglect appears to be that it turns the traditional analysis of behavior on its head. Instead of the still-dominant view of seeing stimuli (both past and present) controlling responses (or perceptions controlling behavior, or environment instructing organism), the theory offers the unorthodox view of behavior as controlling perception through the organism's *control of its environment*. Hence the title of Powers's seminal book, *Behavior: The Control of Perception*.[13]

Diagrams can be useful in providing some basic understanding of the functioning of control systems, but they have some serious limitations. For one thing, they may easily lead one to interpret the functioning of a control system as a series of sequential steps, with each step waiting for the completion of the previous one. Instead, in functioning control systems, both nonliving and living, all parts are active simultaneously so that both perception and behavior are happening *at the same time*.[14] Also, such diagrams may give the impression that control systems are limited to quite simple variables. In fact, complex control systems composed of many simple control systems can be designed to control quite complex variables, that is, variables that are computed composites of values of lower-order perceptual variables.

Powers proposes that a complex hierarchy of control systems underlies human and animal perception and behavior. Currently, 11 levels are envisaged. From lowest to highest they are intensity, sensation, configuration, transition, event, relationship, category, sequence, program, principle, and system concept. These levels will not be described in detail here,[15] but it should at least be noted that higher-level perceptions such as a configuration (for example, the visually perceived printed letter A) depend on a particular set of combinations of lower-order sensations, which in turn depend on a set of particular combinations of still lower-order intensities. But whereas higher-level perceptions depend on lower-order ones, the control of higher-level perceptions is achieved by manipulating lower-order reference signals. Thus as you write a note to a friend, you are controlling for the appearance of certain letters on the page. But to produce those letters, you must vary the reference levels for the positions and movements of your arm, hand, and fingers. In the same way, a higher-level reference level for reading results in varying the lower-level reference levels for the movements and positions of your eyes. That higher-level perceptions are made up of combinations of lower-order ones and that higher-order systems control their perceptions by varying the reference levels of subordinate systems is illustrated in figure 8.2. Notice, however, that this hierarchy is not a typical chain of *command*, since higher-order control systems do not tell lower-order ones what to *do*, but what to *perceive*. (We will return to the control system hierarchy at the end of chapter 12 in our discussion of education.)

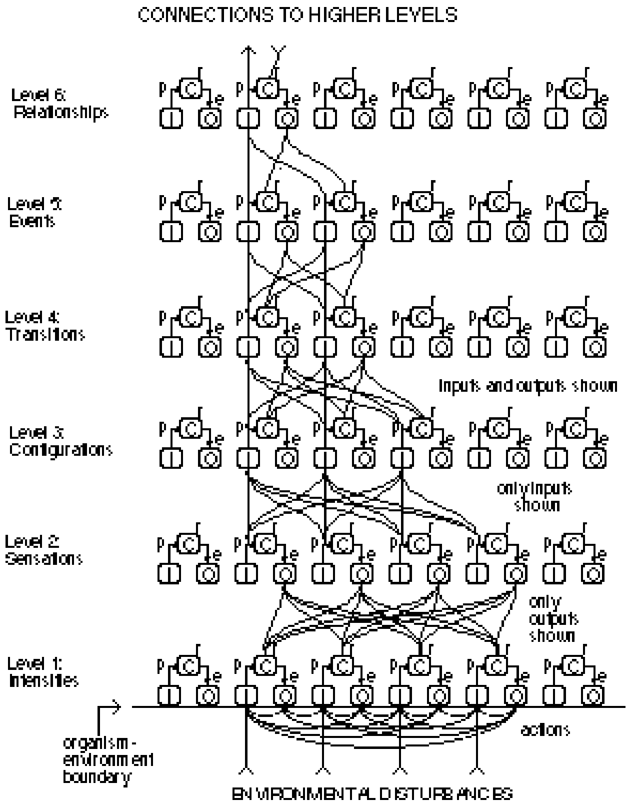

Figure 8.2 A hierarchy of control systems (after McClelland, 1991).

Let us now turn away from the complexities of perceiving and controlling higher-order perceptions and consider the relatively simple variable of distance to see how perceptual control theory can be used to understand the behavior of the greylag gosling. As described in chapter 3, Lorenz discovered that the greylag gosling will imprint on the first large object it sees after hatching, which in natural settings is its mother. Thereafter it will maintain close contact with this object throughout its goslinghood. From the perspective of perceptual control theory, we would say that the gosling develops a control system that permits it to maintain a relatively fixed distance between it and its mother despite the disturbances caused by the mother goose's own walking, wading, and swimming, and regardless of obstacles such as bushes, trees, rocks, and other geese that may come between them. Such a living control system is a great deal more complex than the one used in an automobile cruise control system. We can nevertheless see how all of its properties can apply to this situation, the two most important being the presence of a reference level specifying the distance to be maintained and the negative-feedback loop connecting perception to action and action back to perception.

Powers and sociologists Clark McPhail and Charles Tucker developed a computer program that uses control systems to simulate just this type of collective behavior.[16] Figure 8.3 shows the results of one such simulation in which four individuals (each indicated by the letter G) maintain close proximity to another individual (M) who is moving toward a goal location indicated by a large circle. Obstacles for all individuals to avoid are indicated by small circles. By examining the paths of the individuals (shown by the meandering lines), it is seen that the four Gs are successful in both maintaining close contact with M and avoiding the obstacles as M moves to its destination. Although these computer simulations were developed to model a particular type of collective human behavior, they also provide a striking simulation of the behavior of a group of goslings (Gs) maintaining close contact with their mother (M) as she herself moves to some destination, with all of them avoiding obstacles along the way.

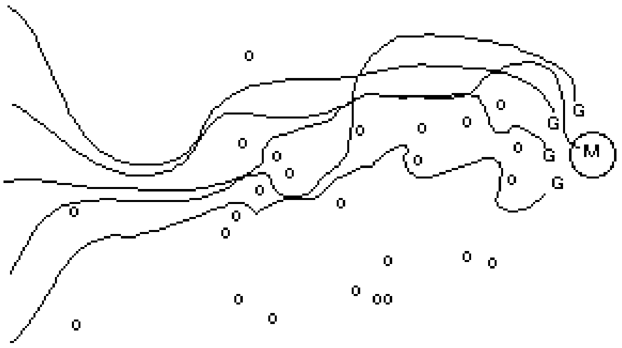

Figure 8.3 Simulation of four individuals (G) following another (M) (after McPhail, Powers, & Tucker, 1992).

A strong argument can therefore be made that it is not behavior in the form of fixed motor patterns as conceived by Lorenz that evolves and becomes instinctive. Rather, the basis of instinctive behavior is an interacting network of control systems that permits organisms to maintain certain perceptual goals (reference levels) despite the disturbances that they encounter. In much the same way that we instinctively know how to keep our body temperature at or near a constant 37deg.C through such automatic, unlearned behaviors as shivering and sweating, organisms also appear to inherit other control systems that underlie adaptive, species-specific behaviors. It is not that the spider is programmed with a fixed behavior pattern that will result in the construction of a web, but rather the spider is able to control its perception of its environment to match an inherited internal goal specification (reference level) by varying its behavior as necessary to construct its web. And in keeping with

Lorenz's original insights concerning the central role of natural selection in the evolution of instinctive behavior, it is only through selecting those organisms with fit control systems and eliminating those with less fit ones that the adapted complexity of instinctive behavior can be explained. As Powers remarked:

> . . . it's the capacity to perceive and control that evolves, not the specific acts by which control is effected. Behavioral acts achieve repeatable results only if they change appropriately with every disturbance, every change in initial conditions. There's no way to inherit behavioral outputs, because the outputs must remain adjustable to current circumstances, which never repeat exactly. All that can be inherited are control systems, and at the highest existing level perhaps some reference signals.[17]

But what about learned behavior? Can perceptual control theory account for adaptive changes in behavior resulting from experiences during the lifetime of an organism? To see how the theory views learning, we first must take a fresh look at both Pavlovian and operant conditioning. With respect to the former, we must reconceptualize the unconditional response to an unconditional stimulus not as the functioning of a stimulus-response connection but rather as that of a control system.[18] To take the example of Pavlov's dog into whose mouth acid (the unconditional stimulus) is introduced leading to salivation (the unconditional response), we can conceive of the stimulus as a disturbance to a reference level for oral acidity and the response of salivation as an action designed to restore oral acidity to its normal level. These and other unconditioned responses (for example, eye-blinking to a puff of air, the startle response to a sudden loud noise) can be seen as very basic, inherited control systems having obvious survival value that evolved through natural selection to protect the organism from harmful environmental disturbances.

Now, if a certain neutral stimulus, such as the sounding of a bell, regularly precedes the unconditional stimulus, it can be used by the organism as a signal that a disturbance is about to occur. So by producing the response after the neutral stimulus (now a conditional stimulus) but before the unconditional stimulus, the organism can prevent or at least lessen the disturbing effect of the unconditional stimulus. By salivating at the sound of the bell that precedes the introduction of acid into its mouth, the dog can buffer its mouth against the effect of the acid.

Considering now the operant conditioning of Thorndike and Skinner, we have seen how stimulus-response conceptualizations of learning cannot account for the purposeful nature of behavior as noted by William James in 1890--the ability to achieve fixed ends by varied means. Construing learning as the acquisition of fixed patterns of behavior cannot explain how organisms can be successful in achieving important goals, such as finding food, mates, and shelter, in the face of unpredictable disturbances.

Another problem with operant conditioning theory is that it provides no explanation for why certain events reinforce the organism's behavior and others do not.[19] Why is it that a hungry but well-watered rat will work a lever to obtain food but not water, while a thirsty but well-fed one will do the opposite? Perceptual control theory answers this question by seeing the reward as a controlled variable, that is, a variable that is controlled by the organism by varying its behavior. If a hungry rat pushes a lever to obtain food, it is only to bring his perceived rate of food intake close to its reference level, which has been chosen through natural selection during the evolution of the rat as a species.

Finally, perceptual control theory explains an intriguing pattern of behavior observed by Skinner that contradicts the basic notion that reinforcement increases the probability of the response preceding the reinforcing stimulus. Skinner found that he could obtain very high rates of operant conditioned behavior (such as a hungry pigeon pecking at a key to obtain food) by gradually decreasing the rate of reinforcement. Very high rates of behavior could be shaped by starting out with an easy reinforcement schedule that provided a speck of food for each key peck, and gradually moving toward more and more demanding schedules requiring more and more pecks (2, 5, 10, 50, 100) for each reward. Skinner was thereby "able to get the animals to peck thousands of times for each food pellet, over long enough periods to wear their beaks down to stubs. They would do this even though they were getting only a small fraction of the reinforcements initially obtained."[20]

But if, according to Thorndike's law of effect and Skinner's theory of operant conditioning, more reinforcement is supposed to cause more of the type of behavior that resulted in the reinforcement,[21] how could it also be that less reinforcement could also cause more of the behavior? This problem is effectively solved when we see reinforcement not as an environmental event that increases the probability of the specific behavior that preceded it, but rather as a means by which the organism can achieve a goal by controlling a perception. If the circumstances are arranged so that the hungry rat must perform more bar presses to be fed, and it has no other way to obtain food, the rat will adapt by increasing its rate of pressing to obtain its desired amount of food. And if the rate of reinforcement is increased to the point at which the rat can maintain its normal body weight, a control system model of behavior would predict that further increases in reinforcement should lead to decreases in the rate of behavior. Indeed, this is exactly what happens.[22]

It should now be obvious that a perceptual control theory interpretation of adapted behavior is radically different from a behaviorist view of operant conditioning. Whereas behaviorism sees the environment in control of the behavior of the organism, perceptual control theory sees the organism in control of its environment by means of varying its behavior. In other words, to behaviorists, behavior is controlled by the environment; to perceptual control theorists, behavior controls the environment. This is not to say that the environment has no influence on behavior. Rather, behavior can be adapted only if it is part of a larger control process that varies behavior to produce the perceptions specified by internal reference signals leading to the accomplishment of goals important for survival and reproduction.

Perceptual control theory also makes an important distinction between changes in performance and learning that is not made by behaviorist theories. According to the behaviorist theory of operant conditioning, the rat's increase in rate of bar pressing in response to a decrease in reinforcement is an example of learning--the animal has somehow learned that more presses are required to obtain a pellet of food and will consequently increase its rate of behavior in response to this new demand. But according to perceptual control theory, no learning has taken place, as this is just the normal functioning of a control system. Recall from our earlier description of an automobile cruise control system that control systems vary their output to control their input to compensate for environmental disturbances. No rewiring or other modification of the control system is necessary for this to happen. Consequently, the rat's increase in the rate of bar pressing is a change of performance that can be explained as the functioning of an existing control system, in the same way that the cruise control system will deliver more and more fuel to the car's engine as external disturbances (hill, headwind, the added weight of accumulating ice and snow) act to slow it down.

However, instead of decreasing the rate of reinforcement for the rat, let us imagine that the situation is changed so that food is delivered not when the rat presses the bar down but when it pushes the bar up. Now, the rat's existing control system will no longer prove effective in controlling the amount of food obtained. This is analogous to reversing certain electrical connections in the cruise control system so that more fuel (instead of less, as before) is delivered to the motor when the speed rises above the reference-level speed, changing a negative-feedback closed loop to a positive-feedback one. A striking difference emerges between the rat as a complex living control system and the cruise control as a much simpler artificial one. In the case of the latter, the car will accelerate until the throttle is wide open and its maximum obtainable speed is reached, a maximum speed that will not be controlled but will vary as a function of the external disturbances. In other words, the cruise control system will simply fail to do what it is supposed to do, and the car's speed will be literally out of control. But the rat will act differently. Although immediately after the change it will be unsuccessful in obtaining food, it will start to reorganize its pattern of behavior so that after a while it will be busily pushing up the same lever that it was busily pressing down just a short while ago.

But how does this reorganization take place? How is the rat to know that pressing the lever down is no longer going to do any good, and that it must push the bar up to be fed? Of course it cannot know this in advance. But the persistent, increased error caused by this change in the environment initiates blind, random changes in the control systems. In this case the reorganization involves a rather simple change in the direction of force applied to the lever. But more complicated situations can be imagined in which changes in the environment require more elaborate reorganizations involving the perception and control of new variables and the resetting of reference levels. As described by Powers:

> Reorganization is a process akin to rewiring or microprogramming a computer so that those operations it can perform are changed. Reorganization alters behavior, but does not produce *specific behaviors*. It changes the parameters of behavior, not the content. Reorganization of a perceptual function results in a perceptual signal altering its *meaning*, owing to a change in the way it is derived from lower-order signals. Reorganization of an output function results in a different choice of means, a new distribution of lower-order reference signals as a result of a given error signal. Reorganization is an operation *on* a system, not *by* a system.[23]

Reorganization, then, is also the functioning of a control system, but it is different in one crucially important respect from the functioning behavioral control systems we have considered up to now. The latter are organized so that response to error tends to remove the error. In a certain sense, the cruise control system knows to open the throttle if the speed of the car drops below the reference level, and the rat knows it must produce more bar presses if more are required for each reward. But the special type of control system designed to monitor and reorganize other working control systems cannot know what to do when it begins to show chronic uncorrected error where there was little error before. All it can do in this case is to modify the control system blindly in some way. If the change results in a reduction of error, any further modifications will be delayed. But if the change has no effect on the error or actually increases it, the next modification will come quickly.

In this respect, the reorganization system that rewires control systems to eliminate the error that the existing systems cannot reduce must act very much like *E. coli*. This common microorganism can either swim in a more or less straight line or tumble blindly. If it senses that it is getting closer to food it will continue along its merry way. But if it senses that it is not doing so, it will stop in its tracks, tumble a while, and then head off in a new, randomly chosen direction. If this new heading is perceived as better than the previous one, it will continue moving in this direction, and the next tumbling act will be put off for a while longer. If, however, the error has not been reduced, it will soon tumble again. Although this method of locomotion may initially appear quite crude, it turns out to be a remarkably useful and virtually foolproof way for the bacterium to move around its environment.[24]

Accordingly, the reorganization of control systems is hypothesized to be an evolutionary process dependent on cumulative blind variation and selection. If existing control systems are ineffective in controlling important perceptions, they will be randomly modified until the error is reduced. Learning to play the piano, ice skate, or speak a foreign language requires reorganization of control systems that is achieved by cumulatively modifying the systems that produce error and selectively retaining those modifications that produce less error. And although the environment certainly plays an important role in influencing the state of the control systems that will be retained, it does not determine the organism's behavior. Instead, the organism is actively involved in the selection process, with the chosen control system parameters and controlled variables depending on the higher-level goals that were selected as being important, and the effectiveness of these parameters and reference levels in the organism's current environment. That higher thought processes play a determining role in what appear at first to be simple Pavlovian and operant conditioning in humans was mentioned in the discussion of Brewer's review in the previous chapter.[25]

Although first developed by Powers and his colleagues over 30 years ago[26] and given a detailed description by Powers in 1973, perceptual control theory is only now becoming more widely known, appreciated, and applied as a theory and research tool in the behavioral and social sciences. The insight that behavior is adapted to its environment only to the extent that it allows an organism to control crucial aspects of its environment has far-reaching implications for all aspects of the life sciences. In addition to a rapidly growing body of experimental psychological research,[27] it is now being applied to clinical psychology,[28] sociology,[29] law,[30] ethology,[31] business administration,[32] and philosophical and educational issues.[33] Such work is still in its infancy, but the power of perceptual control theory for understanding human behavior has been demonstrated in the construction of generative models of behavior that account for over 95% of the variance in a variety of tasks.[34]

But of most importance for our purposes, the theory provides a plausible explanation for how behavior can become *and remain* adaptedly complex. It is clearly not the case, as believed by Darwin and Lorenz, that organisms with useful fixed behaviors are selected during the course of evolution, resulting in innate, fixed patterns of behavior known as instincts. And it is also not the case that specific behaviors are selected by the environment by contingencies of reward during the life of the organism, as believed by Thorndike and Skinner. It is the selection of organisms with useful, adapted perceptual control systems over the course of evolution, coupled with the organism's cumulative variation and selection of its own perceptual control systems during its relatively brief life, that accounts for the adapted nature of behavior. There is no instruction by the environment, no stamping in of stimulus-response connections within the nervous system. Rather, we find a very Darwinian process of selection, not of behaviors, but of closed, negative-feedback loops encompassing perception, comparison with the reference level, and action, which allow patterns of behavior to remain functional, not only from one occasion to the next, but also within the continually changing environment of the behavior itself.

What may seem mysteriously ironic in all this is to realize that the purposeless process of natural selection has led to the evolution of purposeful organisms. But the irony fades when one considers the great survival and reproductive advantages of organisms that are able consistently to achieve goals essential to their survival and reproduction despite an unpredictable, uncaring, and often hostile environment.

[1]Dewey (1896, p. 363).

[2]ames (1890, p. 7).

[3]James (1890, p. 7).

[4]Skinner (1974, p. 224).

[5]Dewey (1896, p. 363).

[6] Described in Tolman (1932, pp. 79-80) and in Boakes (1984, p. 232).

[7] Tolman (1959, p. 100).

[8] Tolman (1959, p. 103).

[9] Peckham (1905, p. 123).

[10] This description of a cruise control as an example of a control system was inspired by McClelland (1991).

[11] Ashby (1952, 1956).

[12] See Richardson (1991) for a historical account of the application of cybernetic and control system concepts to behavior, which includes Wiener, Ashby, and Powers, among others.

[13] Powers (1973).

[14] To show how control systems actually behave, Powers has created and made available two programs for IBM-compatible personal computers (*demo1* and *demo2*) to demonstrate both the phenomenon of control and the use of control systems as generative models of behavior. These programs are available on the Internet at http://www.ed.uiuc.edu/csg/csg.html and at gopher://gopher.ed.uiuc.edu following the path Higher Education Resources/ Professional societies and journals /Control Systems Group.

[15] See Robertson & Powers (1990, pp. 67-80).

[16] McPhail, Powers, & Tucker (1992).

[17] Powers (1991b).

[18] See Hershberger (1990).

[19] See Powers (1973, p. 189).

[20] Powers (1991a, p. 9).

[21] Note that this describes a positive feedback loop.

[22] See Staddon (1983, p. 241, figure 7.18).

[23] Powers (1973, p. 179).

[24] See Koshland (1980, pp. 14-15).

[25] Brewer (1974).

[26] Powers, Clark, & McFarland (1960).

[27] See for example Pavlovski et al. (1990) and Bourbon (1990).

[28]Ford, (1989); Goldstein (1989, 1990).

[29]McClelland (1991, 1994); McPhail (1991); McPhail & Tucker (1990).

[30]Gibbons (1990).

[31]Plooij (1984).

[32]Forssell (1993, 1994).

[33]Petrie (1981); Cziko (1992).

[34]That is, generative control system models have simulated behavior that typically correlates with individual human performance at values between .97 and .99, a level of precision that surpasses by far the explanatory and predictive power of the more traditional and established approaches to social and behavioral science (e.g., Bourbon, 1990; Marken, 1986, 1989, 1991).