Restoring Purpose in Behavior

Henry H. Yin

Abstract The dominant paradigm in the study of behavior today is the linear causation paradigm. This paradigm, inspired by classical physics, assumes that causes precede effects, that the behavior of organisms is caused by antecedent events inside and outside the organism, and that future states such as goals and purposes cannot possibly cause behavior. It is the basis of the general linear model in psychology and the input/output analysis in neuroscience. But linear causation does not apply to any control system with negative feedback. Here I shall argue that organisms are collections of such negative feedback systems that control their perceptual inputs. The chief difference between the behavior of living organisms and that of non-living things is the presence of control. Rather than the effect of some prior cause, behavior is the observable manifestation of control in teleological systems that act on the environment to make inputs match their internal reference values. The previous rejection of control theory in the behavioral sciences was largely based on a misunderstanding of the principles of negative feedback control. I discuss the types of behavioral control enabled by the hierarchical organization and the experimental method for testing the controlled variable.

1 Introduction

The fifth way begins from the guidedness of things. For we observe that some things which lack knowledge, such as natural bodies, work towards an end. This is apparent from the fact that they always or most usually work in the same way and move towards what is best. From which it is clear that they reach their end not by chance or by intention. For those things which do not have knowledge do not tend to an

---

H.H. Yin (✉)
Centre for Cognitive Neuroscience, Department of Psychology and Neuroscience,
Department of Neurobiology, Duke University, Box 91050, Durham, NC 27710, USA
e-mail: hy43@duke.edu

G. Baldassarre and M. Mirolli (eds.), *Computational and Robotic Models of the Hierarchical Organization of Behavior*, DOI 10.1007/978-3-642-39875-9_14,

end, except under the direction of someone who knows and understands: the arrow, e.g., is shot by the archer. There is therefore an intelligent personal being by whom everything in nature is ordered to this end, and this we call God.

In the above passage, the fifth argument for the existence of God, Aquinas argues that the "guidedness of things" is evidence for some omnipotent agent which directs everything in nature to the appropriate end. For Aquinas, knowledge of ends is needed for movement. Since "natural bodies" do not possess such knowledge, a human-like agent is needed to move them, as the archer shoots the arrow. Galileo and his followers, however, succeeded in explaining the movement of natural bodies, from balls on inclined planes to planetary motion, and the laws of physics have since demonstrated that neither a concept like purpose nor an all-knowing agent like God, is necessary to explain how things move. But a gap remains. A pigeon released into the air does not fall like a rock of the same mass, as predicted by physical law. Its flying, and indeed the behavior of any organism, remains a puzzle.

We now know that, at the level of atoms and molecules, there is no fundamental distinction between living and non-living things. Many therefore assume that the physical laws must also suffice to explain the behavior of organisms, which are just collection of atoms like everything else in the universe. The rock falling can be explained by gravitational force; it has no knowledge of the future state of lying on the ground. It is manifestly a violation of the law of causation that a future state can cause anything to happen in the present. There is no place in physics for purpose.

Although Aquinas and modern scientists differ on the question of what causes movement, they all assume that there can only be one kind of "guidedness," for physical things and for living things. For Aquinas it is purpose. For modern scientists it is physical law. Following the example of physics, modern psychology attempted to purge all teleological concepts from its theories. Since Descartes introduced the paradigm of physics into psychology, different schools have looked for the cause of behavior in different places, either external to the organism, in some stimulus, or internal to the organism, in some executive homunculus issuing a command. Over the centuries, these approaches have evolved largely by changing their labels, but little has changed in the underlying assumption that some antecedent event prior to behavior is the cause of behavior. For example, observing someone walking, we may say that the behavior of "walking" is caused by contraction of his leg muscles, which is caused by the release of acetylcholine at the neuromuscular junction, which is caused by firing of alpha motor neurons, and so on, tracing the causal chain backwards until we reach the cause of walking. This view is deemed so self-evident that hardly anyone has ever questioned it.

The dominant approach in the study of behavior is the "input/output" approach. The experimenter manipulates input to the organism, in the form of stimuli, and measures output, in the form of behavior. The input is the independent variable, and the output the dependent variable. The goal is to identify a function that relates the input to the output. So far, however, all attempts to find such a function have failed, despite considerable progress in our understanding of the structure and physiology of the brain. It is clear to the engineer attempting to build a behaving robot that the diagrams found in neuroscience textbooks do not actually produce viable behavior, even the behavior of a worm or of a fly.

Students of behavior have resorted to two major strategies. In neuroscience, researchers have often used "preparations" that yield high correlations between stimulus and response. For example, when Sherrington, a pioneer of the input/output analysis in neuroscience, used awake and intact animals as his experimental subjects, the behavior he observed was simply too variable and the relationship between brain stimulation and behavioral output too uncertain (Leyton and Sherrington 1917). Instead he often relied on the decerebrate preparation, which removes the influence of the brain to yield behavioral data that more closely approximate the predictions of the linear causation hypothesis (Sherrington 1906). Scientists like Sherrington insist that the organism yield a high correlation between the stimulus and the response, even if it involves removing the brain, keeping the animal anesthetized, or whatever it takes to arrive at a "preparation" that will satisfy the rigor of the physical sciences. In experimental psychology, on the other hand, the strategy is to move beyond the individual as the basic unit of analysis, by averaging across multiple individuals. The statistically average animal appears to show the stimulus-response correlation that satisfies the experimenter, even if the individual animal does not. Whereas neuroscience has focused on the "partial animal," so that whatever is left can "behave" according to the experimenter's a priori assumptions about what behavior should be, psychology studies the "average animal," or the "Gaussian animal," a creature not found in the woods.

Consequently, students of behavior today often share a set of implicit assumptions about the study of behavior. Above all, the experimenter must be able to manipulate input and measure output, so that some function can be found to describe the output in terms of the input. When behavior is variable, the low correlation between stimulus and response is attributed to noise, or changes in "conditions," "contexts," or whatever the experimenter cannot explain. In addition, because more subjects in a study means more statistical power, a study with many subjects is better than a study with a few subjects, single-subject designs or self-experimentation having been largely excluded from acceptable scientific practice. These assumptions all stem from the most basic assumption: that behavior is the effect of some antecedent cause. Something must act on the organism, or more exactly its nervous system, resulting in a chain of transformations within that ultimately leads to behavioral output. This something can come from within, in the form of reasons or homunculus, or it may come from without, in the form of stimuli and contexts.

2 The Calculation Problem

In 1935, Nicholai Bernstein pointed out a problem with the assumption that the output of the nervous system causes movement. Clearly some neural output leads to behavior. That much is seen. What is not seen, however, is additional forces acting on the muscles to generate what we observe as behavior. For example, in opening a door, the stiffness of the door handle, the wind outside, the weight of the door, the oil in the hinge all these variables are independent of the neural activity, yet they contribute to the observed behavior. Bernstein realized that, for natural movements to be possible, there cannot be an unequivocal relationship between neural output and behavior, because the environmental disturbances are always present and unpredictable, and together with neural output they jointly determine the behavior observed (Bernstein 1967). Just as the movement of a car is not determined by the engine output alone, so the movement of the pedestrian is not a result of his neural activity alone. It does matter whether there is a hill or how strong the wind is, for the car or the pedestrian. Because disturbances can accumulate quickly, the more complex the movements, the less correspondence there is between neural output and actual behavior, as more disturbances are added to the neural outputs, at each turn, to produce behavior as observed. Consequently, the sequence of neural activity recorded while you walk home, if repeated, will not take you home a second time. As Bernstein explained: “A trained athletes consecutive running steps are as identical as coins of the same value, but this identity results, not from the brains ability to send absolutely identical motor impulses to the muscles, but only from the faultless work of sensory corrections. . . . even if the muscles received ten absolutely identical motor impulses in a row, there would be, in the very best case, 10 ugly steps, each one different from the others and with a result quite different from running” (Bernstein 1967 p. 180).

If behavior as observed is determined by both neural output and environmental disturbances, then how can it be reproduced when only one of these is generated by the brain? How do organisms generate output that is specifically calibrated to oppose unknown and unpredictable disturbances? It is easy to overlook this “calculation problem” because we are accustomed to observing living organisms “behave.” What we observe is the achievement, in the Brunswikian sense (Hammond and Stewart 2001), but we do not see the hidden disturbances that also influence the behavior as achieved.

3 The Nature of Control

The calculation problem, that of calculating inverse kinematics and dynamics of behavior, is widely acknowledged, but there is no consensus on how it can be solved. One possible solution, popular today, is to predict exactly what the disturbances will be and to calculate exactly what is needed to overcome such disturbances (Franklin and Wolpert 2011; Shadmehr et al. 2010). Some believe that, given unlimited computational power, which is often attributed to the brain, the organism can predict exactly what is needed to overcome the disturbances. But these models are based on unwarranted assumptions and a misunderstanding of negative feedback systems. Consequently, much effort is devoted to the solution of the calculation problem using advanced mathematics. The “feedforward” solution to the calculation problem may appear simple, but in practice even the most powerful computers have yet to succeed using this approach to solve the calculation problem. Calculating output requires an understanding of the basic laws of physics, which are not known to the organisms. To compute the right outputs, the brain needs information about the masses and moments of inertia of the arm segments, the properties of muscle contraction, the properties of nonlinear muscle springs, the variations in mechanical advantage as the joint angles change, the physics of dynamic movements, the trigonometry of spatial relationships, and the initial state of the effectors, among other things. And all these elaborate computations must be performed in real time (Powers 1978).

There is, however, a far simpler alternative that does not require the calculation of inverse kinematics and dynamics (Powers 1973b). For the calculation problem is not unique to the field of behavior. Engineers face it every time they try to design a machine that can resist unexpected disturbances. Although they repeatedly stumbled upon the solution, only in the twentieth century did they finally understand the properties of a negative feedback control system (Black 1934). In their terminology, the phenomenon of resistance to disturbances is called control, and the principles that achieve control without calculating inverse kinematics are known as classical control theory.

The key principle in classical control theory is closed loop negative feedback. Although for engineers negative feedback is simply a convenient way of solving practical problems, Rosenblueth and Wiener, the founders of cybernetics, realized that it has important implications for psychology and the life sciences. Unfortunately, because they were not familiar with how control systems actually operate, the cybernetic models were fatally flawed (Rosenblueth et al. 1943; Wiener 1948). They introduced the vocabulary of control theory, without explaining correctly how it could be applied to the study of behavior. Partly for these reasons, classical control theory was abandoned by students of behavior long before they even understood it. Had they studied it carefully, they would have discovered that the phenomenon of control creates insurmountable challenges for the standard paradigm in neuroscience and psychology.

The major implication of negative feedback control is that linear causation is always wrong in explaining the behavior of control systems. For example, the thermostat is a simple negative feedback control system. If the room temperature deviates from the set temperature, the thermostat will act to oppose the change. If we apply the standard input/output analysis to the thermostat, we will obtain data similar to what is shown in Fig. 1. The room temperature, the input to the thermostat, is relatively stable, whereas the output, measured by the utility bill, varies a great deal. It is impossible to find a function to describe the relationship between these two sets of values. In fact, the major theories in the history of psychology all fail when applied to the analysis of the humble thermostat, precisely because all such theories attempt to find the function that relates the input to the output. The major implication of classical control theory is that all such attempts are necessarily in vain.

The thermostat senses input in the form of room temperature; it compares this input with a reference signal specifying what the temperature should be; and it

Month	Input (temperature in Celsius)	Output (Utility bill in dollars)
Jan	20	150
Feb	20	120
Mar	20	110
Apr	20	70
May	21	60
Jun	22	200
Jul	25	170
Aug	25	140
Sep	25	130
Oct	21	80
Nov	19	100
Dec	19	150

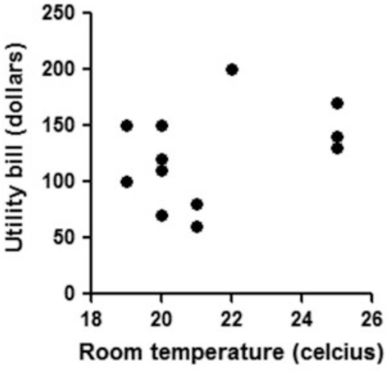

Fig. 1 Hypothetical data set from the input and output measurements of a thermostat

generates some error signal, the difference between current temperature and desired temperature, which is then converted into the output. The controlled variable is the room temperature, and a working thermostat will keep this variable as close to the set temperature as possible, despite fluctuations in the temperature outside. When the current temperature is the same as the desired temperature, there is no error, and the output is turned off. Because the output acts to keep the input close to the set temperature, there is no consistent relationship between the input and the output. Rather the output is correlated with the disturbances, i.e. outside fluctuations in temperature. The thermostat does not sense the disturbances directly. It only senses one input, current temperature, which is determined jointly by its own output and a variety of disturbances.

The thermostat could have used a “feedforward” approach popular in the motor control literature to predict disturbances to the temperature, calculating outputs needed to maintain the desired temperature, based on an inspection of the visual scene outside, the color of the leaves, the time of the year, and so on, and calculate the exact amount of output needed at any given moment. But even if it was equipped with a very powerful computer, it would not have been able to control room temperature.

Closed loop negative feedback control system has special properties not found in open loop systems. An example of an open loop system is the electric fan, which controls its output rather than its input. The thermostat, by contrast, controls its input—the sensed temperature in the room. It does not need to know anything about the disturbance to the controlled variable. The fluctuations in temperature could be caused by the fire in the bedroom or the snow on the roof. Whatever the sources, the thermostat only has to sense the variable being controlled, namely the current

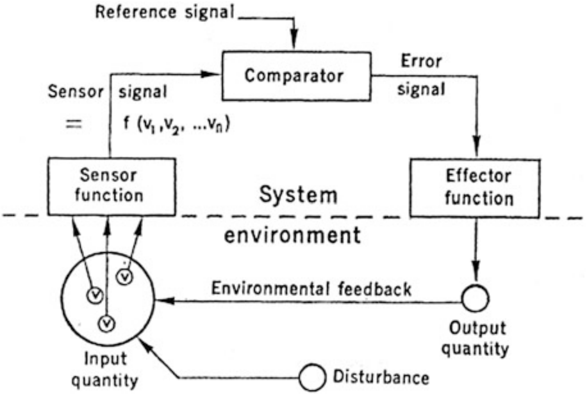

Fig. 2 The basic organization of a closed loop control system (from Powers 1973)

temperature at its sensor, to bring the input variable close to the value specified by the reference signal.

Although teleology has become a dirty word, the thermostat is a teleological machine. If you set it at 25 degrees Celsius when it is currently 10 degrees, the room temperature will eventually reach 25 degrees. The set temperature is a representation of a future state, but it contains no information about how the system should behave. We cannot predict the thermostat's behavior if we only know the reference signal. Nor can we predict its behavior if we only know the room temperature. To predict its behavior we need to know both the current temperature and the reference temperature, because the difference between the two generates the output.

3.1 The Organization of Control

A negative feedback control system, then, solves the calculation problem without performing calculations of the actual disturbances. It merely adjusts its output according to the difference between the present value of the controlled variable and the desired value. Control is an emergent property of the negative feedback loop. The linear cause and effect model breaks down in the analysis of such a system, because the output is acting on the input at the same time that the input is acting on the system. The closed loop is a loop of circular causation (Powers et al. 1960).

As illustrated in Fig. 2,a closed-loop control system has three main components: input function, comparator, and output function. The input function represents some variable to be controlled; the comparator computes the difference between a reference signal that represents the desired value of the variable and the input signal; the output function converts the error signal into an output that acts on the environment. A successful control system will keep the value of the controlled variable close to the value of the reference signal, by varying the output to counter the effects the effects of disturbances on the controlled variable. In traditional open loop models of behavior, the input is not controlled; outputs are generated either by transforming inputs, or spontaneously. In a negative feedback control system, only the input is controlled.

The comparator receives two types of signals, reference signals and input signals. Either could be positive or negative, but they should be opposite in sign. Thus, if the reference is positive, then increasing the reference signal also increases error if the input remains the same. The output will be generated continuously until that error is reduced.

Reference signals are internal to the organism. This is the chief difference between biological organisms and existing man-made control systems. In the thermostat, the reference signal (desired temperature) is set by the user. This is not possible in an organism, because its reference signals are its own. We can only influence reference signals indirectly, e.g. by depriving the animal of what is essential to it, or by manipulating their nervous system if we understand how the comparator functions are implemented.

The concept of set point (Sollwert) in physiology comes closest to the reference signal. The maintenance of a stable internal environment (milieu interieur) in various physiological systems was first noted by Claude Bernard. Indeed, homeostasis, the popular term coined by Cannon, has become a core concept of physiology. Through Cannon's student Rosenblueth, it inspired the field of cybernetics. Surprisingly, none of these concepts had a major impact on the study of behavior, with the possible exception of the work on feeding systems (Staddon 1983). Homeostasis refers to a control system with a relatively fixed reference signal. A common experiment in nineteenth century physiology was to study the wiping reflex of frogs by leaving acid on the skin of frogs with transected spinal cords. The wiping is easily elicited, presumably because the tolerance for acid (reference signal) is close to zero. Whenever the reference signal is negative and close to zero in magnitude, a small perceptual signal can generate an error signal, thus generating an output. Influenced by the work on reflexes, many physiologists implicitly assume a constant reference signal of zero. Freud, for example, often assumed that the primary purpose of behavior is to remove sensory stimulation altogether (Freud 1915). Such systems have the appearance of a simple input causing a simple output. But the appropriate stimulus is simply a disturbance to a controlled variable, creating an error signal from the constant reference signal. Thus homeostasis is a particular type of control phenomenon, found in control systems with relatively constant reference signals, e.g. body temperature in mammals. In many behaviors, the reference signal can change quickly, though a similar negative feedback organization is still found.

The comparator sends an error signal to the output function, which transforms it into some output. There is not necessarily a one-to-one correspondence between error signal and output. Even a simple operation like integration can make the output appear very different from the error signal (Robinson 1989). The output of a control system, however, should not be equated with the behavior per se. Most control systems do not generate outputs that correspond to what we observe as the behavior of the organism. Their output functions are located inside the organism, or more exactly in the nervous system, and do not affect motor neurons directly. In fact, as we shall discuss below, they can act on other control systems in a hierarchical arrangement.

The input function, output function, and comparator are all located inside the control system, but the feedback, which describes the input as a function of the output, is located outside, in the environment. The environment is everything that is outside the control system. With feedback the loop is closed, so that the output will ultimately reduce the error signal—the definition of negative feedback. A positive feedback system, by contrast, will increase the error and moves input away from its desired state. Such a system is not desirable for controlling any variable, as it will only amplify the effects of disturbance rather than resist them (Powers 1978).

3.2 *Wiener's Error*

Although words like feedback and control have become a part of the everyday vocabulary, the principles outlined above are misunderstood. For by the time engineers discovered the principles of control, the life sciences were already entrenched in the paradigm of linear causation, imported by Descartes from physics. Descartes himself understood the logical consequence of his position. On the one hand, the explanation of the movement of natural bodies offered by physics can be adopted to explain animal behavior, using the model of the reflex arc. On the other, subjective experience convinced him that he was able to act voluntarily, independent of environmental stimuli. These two possibilities being incompatible, Descartes was forced into the dualism for which he is known. Where a clear antecedent is found, the reflex arc is proposed to transform external energy into motor output; where no stimulus can be identified as the cause of behavior, it is assumed that, somewhere within the brain, a homunculus issues commands which are ultimately sent to the muscles. The homunculus, on this account, cannot possibly be material, for that would place it in the chain of linear causation. Behind dualism is the assumption of linear causation. Descartes was therefore quite consistent. Only those who came after him became confused. Instead of accepting his dualism, most chose either the stimulus-driven reflex arc model or the command-driven model. These two positions were eventually transformed into pairs of opposing schools: rationalism and empiricism, vitalism and mechanism, associationism and idealism, and most recently, behaviorism and cognitivism. As is common in intellectual history, the opponents typically share key assumptions, in this case the linear causation model.

Today, many students of behavior would be surprised to hear that they are following Descartes. It is therefore important to understand how they came to be Cartesians without knowing it, by exposing their unquestioned assumptions. Both the stimulus-driven and the command-driven models assume that behavior is open loop. Even the founders of cybernetics, who introduced control theory to the general public, could not free themselves from the grip of the linear causation paradigm (Rosenblueth et al. 1943; Wiener 1948).

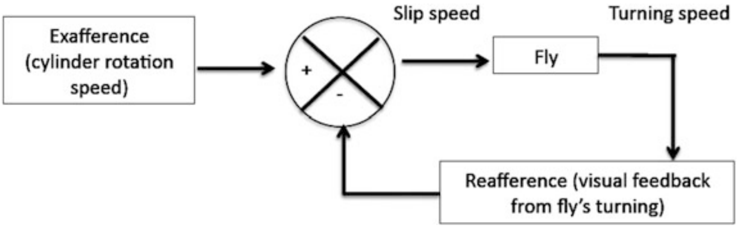

Fig. 3 The wrong model of closed loop organization influenced by Wiener (after Camhi 1984). In this model, the comparator function is located in the environment rather than in the organism

Wiener, in particular, was responsible for the widespread misunderstanding of negative feedback in the life sciences. In Wiener's model, the organism is not a control system, but merely the output function of a control system (Wiener 1948). He placed the comparator outside of the organism, in the environment (Fig. 3). The organism simply receives the error signal and transforms it into output—a stimulus-response device. Wiener's error is not easy to discern, for his model contains exactly the same components as a negative feedback control system. His mistake is in the assignment of the components of the control system to the organism–environment relationship. Instead of placing the collection of control systems inside the organism, Wiener simply turned the organism into an output function (Fig. 3). In itself, of course, the output function is not a control system, because control is an emergent property of a collection of components in a particular arrangement.

Wiener was probably misled by man-made control systems, in which the reference signal is specified by the human user. The engineers have made a control system seem to act like any other input–output system. Although it has the negative feedback control loop, it does not have its own reference signals. That is why one can adjust the setting on thermostat, but the thermostat itself does not have a preferred temperature. The preference belongs to the user. The reference signal in the man-made control system is normally designed to be accessible to the human user. Biological organisms do not serve any user. They serve their own essential variables, and their reference signals belong to them alone. Environmental influences act as disturbances to the controlled variable.

Unfortunately, Wiener's error became his major legacy, resulting in nearly universal confusion regarding the use of negative feedback in explaining behavior. A good example is standard textbook explanation of the optomotor turning response (Camhi 1984). In the typical experiment, a fly is placed on a platform inside a cylinder with a striped pattern. If the cylinder is rotated, the fly will turn in the same direction as the cylinder. The difference between the fly turning and the cylinder turning is called "slip speed." According to Camhi, the slip speed serves as the error signal in this closed loop control system; it is needed for the fly to generate the output of turning. In other words, slip speed is used as a stimulus which "causes" the response.

Moreover, Camhi assumes there are two inputs, one that is independent of self-motion (exafference) and another that is self-induced (reafference). The turning of the cylinder itself is the exafference, which is then compared with the reafference, visual motion as a result of turning (Von Holst and Mittelstaedt 1950). But Camhi fails to see that the slip speed, the difference between the cylinder rotation and the self rotation, is the perceptual input—the actual visual motion sweeping across the retina. The slip speed is not an error signal, but the input to the optomotor control system, and the variable being controlled. What Camhi calls the input, the cylinder rotation, is really the input from the perspective of the experimenter, not the actual image velocity on the retina of the fly. The comparator is located somewhere in the fly brain; and it compares the perceived visual motion with the reference signal, which specifies how much "slipping" of the cylinder is permitted. If the latter is zero, then any perceived slip speed will create an error signal (i.e. cylinder is moving too quickly). The output ultimately generated is turning in the same direction as the cylinder, which reduces the value of the controlled variable.

Camhi writes: "one surprising feature of the optomotor feedback loop is that it is actually impossible for the fly to keep up perfectly with the cylinder's rotation; for to do so would ... create a slip speed equal to zero. But a slip speed of zero would produce zero behavior, so that fly's turning would, of necessity, stop. Given this circumstance, the best that the fly can do is to keep its turning speed very close, but not equal, to the cylinder's rotation speed. This it does, especially for slow cylinder rotations of up to about 20 per second." (Camhi 1984). Camhi assumes that closed loop feedback control requires the animal to generate sufficient error in order to generate behavior. This he considers a flaw in the system. But if we assume that the reference signal is zero (no slipping allowed), then error and output will be proportional to the input. Turning continues because additional error is created by the disturbance introduced by the experimenter (i.e., cylinder rotation), not by the fly. So long as the cylinder is rotating, the error is continuously generated.

4 External Causation and the Reflex Arc

The stimulus-response model is the modern day version of classical empiricism: "sensory vibrations which are excited in the external organs, and ascend towards the brain by agitating the small particles of the muscular fibers . . . excites them to contraction" (Hartley 1749). This model still dominates psychology and neuroscience. Even many models in cognitive psychology (e.g., connectionist neural network models) are just stimulus-response models with intervening variables or hidden layers (Hull 1943; Millerm and Cohen 2001). But few have asked whether the stimulus-response model can explain any behavior in any organism.

By universal consensus, the reflex arc is the best example of a stimulus-response mechanism in its purest and simplest form. According to Sherrington, the reflex arc, comprising a receptor, a conductor, and an effector, is the fundamental unit of integrative activity in the nervous system (Sherrington 1906). Such a model appears similar to the control system, the receptor resembling the input function and the effector the output function. But there is no feedback in the reflex arc, and the conductor does not receive reference signals for desired value of the input. It assumes that nothing happens to the value of the stimulus, once the response is initiated.

Sherrington did not hesitate to title his book “The Integrative Actions of the Nervous System” even though it only describes his work on spinal reflexes using the decerebrate preparation. In this preparation, the spinal cord is cut to remove the influence of the brain. Sherrington focused on the scratch reflex in dogs. Since the dog can no longer behave, it will not respond to the stimulation as it would normally do (which might include attacking the experimenter). Having removed the brain, Sherrington could better manipulate the input (e.g., drops of acid or electrical stimulation) and measure output (a scratching of the hindleg). But even in his reduced preparation, the relationship between input and output was far from clear. The observed irregularities (e.g., variable latency and after-discharge) was attributed to synaptic transmission. As a leading proponent of the neuron doctrine, Sherrington was convinced of the existence of junctions between neurons. He therefore concluded that his data on reflexes not only supports the reflex arc model but also suggests the existence of synapses (Sherrington 1906).

4.1 The Stimulus-Response Illusion

What Sherrington neglected is the feedback acting on the controlled variable. The “stimulus specificity” of reflexes provides a clue about the controlled variable, which is always disturbed by the “effective” stimulus. For the scratch reflex, it is the amount of irritation on the skin; for the pupillary light reflex, the amount of light on the retina. The reference signal is negative for protective reflexes, determining what is “excessive” activation of the relevant receptors. In the case of protective reflexes, the output usually reduces the value of the controlled variable.

Sherrington’s input/output analysis only works for open loop systems. When applied to control systems, it can produce a powerful illusion. The illusion of a stimulus-response system emerges when the observer can identify a clear antecedent stimulus event and a behavior that follows (Powers 1978). The antecedent event appears to be the cause of behavior. The fire appears to cause withdrawal of the hand; the acid on the skin appears to cause scratching; and the nervous system appears to transform the sensory input into motor output. If there is any correlation between the stimulus and the response, then the function defining this relationship appears to be a function of the organism.

Appearance, however, can be deceiving. For stimuli to be effective, they must act as disturbances to the input variable being controlled. Since the error signal from the comparator can only be produced by variations in the input and the reference signal, when the reference signal is constant and low (e.g., zero), the error signal is largely determined by the disturbance. Thus with a fairly constant reference signal, we can observe a regular relationship between stimulus and response, thus creating the illusion that stimuli cause responses.

Since the definition of control is systematic resistance to disturbance, the disturbance and the output are highly correlated. Yet this correlation is not a function of the organism. It is a result of disturbance and output simultaneously acting on the same input variable, but in opposite directions, one pushing and the other pulling. For example, if I am trying to hold my umbrella in the rain, the wind is a disturbance to the position of my umbrella. My arm counters this disturbance to keep my umbrella straight. Recording activity from my arm muscles and the wind simultaneously, we can find a high correlation. The wind is transformed by my nervous system into muscle contraction in my arm—so it appears. But there is no such transfer function inside of me. For instead of a "wind to muscle" sensorimotor transformation, there is the connection between the wind and the umbrella, which is a feature of the environment, and there is the connection between the arm and the umbrella, which is another feature of the environment. These two functions are correlated, so that the latter cancels the effects of the former. Neither connection belongs to the organism, and neither can be found inside the brain. Both the wind disturbance and the neural output act on the same variable—umbrella position—which is perceived by my nervous system. So long as the umbrella position is being controlled, the wind input will be correlated with the arm activity, creating the illusion of a stimulus-response system. But as soon as I stop controlling the umbrella, the correlation disappears. This illusion is the chief source of spurious sensorimotor transformations in many studies.

4.2 Fixed Action Patterns, Central Pattern Generators, and the Cognitive Homunculus

Some might argue that more enlightened "cognitive" theories of behavior have long ago replaced the simple reflex arc models. Most cognitive models are variants of the command-driven model, which is on the other side of Cartesian dualism. On this view, the cause of behavior is not found outside the organism, in some environmental stimulus, but within the organism itself. Independence from sensory inputs is considered the basis for cognition. Voluntary actions are produced by some internal, spontaneously active process, a homunculus issuing commands to output functions. This view has many branches, such as rationalism and its modern variants in linguistics and cognitive science (Chomsky 1965; Miller et al. 1960).

In neuroscience, the command-driven model is based on the concept of central pattern generators, which originated from the work of Graham Brown, one of Sherrington's students (Graham Brown 1911). Using the same spinal preparation, Graham Brown showed that the removal of sensory afferents to the spinal cord did not eliminate locomotion in the cat, suggesting the existence of intrinsic mechanisms capable of generating rhythm independently of sensory afferents (Lashley 1951). That the nervous system, like the heart, can generate spontaneous behavior independent of inputs was not a novel idea. For the scratch reflex, there is no one-to-one correspondence between stimulus and response; the rhythm of scratching, at about four times per second, is independent of the rhythm of stimulation (Sherrington 1906). Sherrington himself compared the fixed rhythm of the scratch reflex to the beating of the heart. But Graham Brown argued that the independence of such intrinsic rhythm from sensory afferents falsifies the reflex arc theory. We now know of mechanisms which allow neurons themselves to act as pacemakers, even in the absence of synaptic inputs (Llinas 1988), which appears to support the idea that intrinsically generated patterns can generate behavior independent of sensory inputs.

The command-driven model is intended to replace the reflex arc model. But like Descartes' homunculus, it shares the same underlying assumption of linear causation. The cause is internal, and the effect or observed behavior is external—an inside-out version of the stimulus-response reflex arc paradigm. And to explain the unpredictability of behavior, what occurs inside the command-issuing center is often assumed to be some random process (Glimcher 2005; Neuringer 2002; Neuringer and Jensen 2010). But the conceptual confusion of linear causation remains. If the behavioral patterns control specific variables, then the intrinsic rhythm may simply reflect properties of the output function, which must still vary as a function of the relevant error signal. If the spontaneously generated outputs do not control any input variable, then the calculation problem cannot be solved at all.

5 Hierarchy of Control

As I write, my head position, body temperature, blood sugar, and direction of my gaze are all being controlled. A single control system cannot possibly control all these variables. The control system described above is therefore only the basic building block out of which the organism is constructed.

The comparator in any control system must receive input signal from the input function and reference signal from somewhere else, and its output must also go somewhere else. Where does the reference signal come from? And where does the output signal go? According to Powers, the level that supplies the reference signal is another control system (Powers 1973b). As shown in Fig. 4, a control system can send a reference signal to the comparator of another control system. The system that sends the reference signal is hierarchically higher. In other words, the higher level supplies the purpose for the lower level.

In traditional hierarchical models, "commands" are elaborated by lower levels into more detailed commands (Fuster 1995). Each level is open loop (Millerm and Cohen 2001). By contrast, descending signals in a control hierarchy do not specify the output of the lower system (Marken 1993). Instead of telling lower systems what to do, the higher systems tell lower systems what to sense. The lower output is proportional to the difference between the top down reference signal and the input

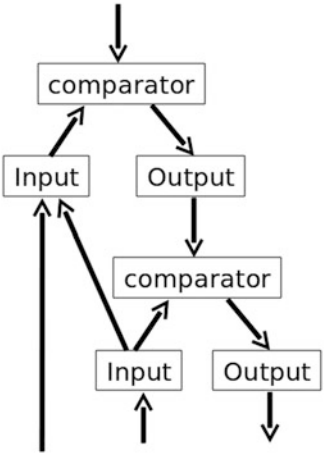

Fig. 4 Hierarchical organization of control systems. The key feature of the hierarchy is that a higher control system may send outputs to the comparator of a lower control system, thus changing its reference signal

to the lower control system. Consequently there need not be any correlation between the top-down reference signal and the output of the lower systems.

It is not possible to specify the output of lower systems by directly activating their output functions. If the reference of the lower control system remains the same, the direct command to the output function will change the input to the lower system, which will produce an error signal, and an output from the lower system that cancels the effect of the command. To use the lower system, the higher system must change the reference signal, to produce the needed output to reduce its own error signal. The lower system knows nothing about the variable being controlled by the higher systems. But it will counter any disturbance to its controlled variable, the value of which is determined by the top down reference signal. It just matches its input with the reference signal it receives. The higher system does not sense the effect of such disturbances to the lower controlled variable, if the inputs at the lower level are already controlled.

5.1 *Essential Variables and Homeostatic Control*

As we ascend the control hierarchy, we ask why a certain output is produced, which is a question about the reference signal and the variable being controlled (Powers 1973a). Conversely, when we descend the hierarchy we ask how a variable is controlled. If we keep asking the “why” question as we ascend the hierarchy, i.e. look for the source of the reference signal for the present level, there must be a level at which we cannot go any higher. This would be the top level, for the sake of which other control systems operate. Ultimately, why do living organisms do what they do? The concept of “survival” comes to mind. But survival is merely an abstraction, comprising a collection of variables, which Ashby called “essential variables.” They are essential, because death results if these variables are not controlled (Ashby 1960). The traditional examples of homeostatic control, e.g. blood glucose, body temperature, are all essential variables (Cannon 1932).

For any essential variable, there is some innately organized control system that opposes the effects of disturbances. For example, blood pressure is sensed by baroreceptors located in the carotid sinus. When blood pressure increases, activating the baroreceptors, the signal is sent to the brainstem, where specialized neurons with a built-in discharge rate serving as the reference signal are located. The difference between the input and the desired rate is the error signal, which can change the heart rate. Such autonomic functions are performed by many control systems specialized for the control of specific variables. At any given moment, all of these specialized control systems must function for the organism to survive. But here is a paradox. If the essential variables are controlled by first order control systems, each associated directly with the appropriate output functions that can act to control the value of the variable, then why is a hierarchy needed at all?

Although all control systems within the organism must somehow serve the essential variables, this organization is more complex than it appears at first glance. Merely having the first order homeostatic control systems is not enough. The new born infant is already equipped with functional control systems for essential variables, yet it is quite helpless, its ability to resist perturbations to these variables limited. For example, it can sweat when the body temperature increases, but it cannot escape from a fire, or to extinguish the fire with water. It is instructive to understand the limitations of first order control systems.

In simple negative feedback systems, the sensors sense the state of the essential variable directly. For example, a body temperature control system does not get information about the disturbance to the temperature until the disturbance has already started to change the value of the controlled variable. Should a fire be started in the room, it would know nothing about the fire. As far as the first order control system is concerned, there is no such thing as fire. There is only the sensed temperature. No other input is available to this control system—not the smoke or the flame. The simple control system is therefore limited by the sensory variables it can detect. To detect other variables, signals from various sensory receptors must be combined. A hierarchical organization is required for the progressive transformation and recombination of “primitive” sensory inputs (Konorski 1967). At the same time, the outputs of the control systems for essential variables are also limited. The autonomic outputs (e.g., sweating) are not general purpose output functions. One cannot use them to walk. As we shall see, such limitations in the output function also requires hierarchical organization and the development of a general purpose skeletomotor system.

5.2 Anticipatory Control

Facing a fire, most organisms will not wait for the flame to engulf them. They will act as soon as they detect the fire. To many, the fact that organisms can anticipate disturbances suggests that the nervous system makes predictions about the future (Wolpert et al. 2011). But before jumping to this conclusion, we must first understand exactly what "prediction" means, and how the anticipatory behavior is produced.

Starting with Pavlov, anticipatory control has been studied extensively. In Pavlovian conditioning, the unconditional stimulus (US) can be a drop of acid, and the UR is salivation, which dilutes the acid. The controlled variable is irritation of the mouth cavity. The US is merely a disturbance that changes the value of some essential variable, and the action of the control system is the unconditional responses (UR). The conditional stimulus (CS) is a neutral stimulus, such as the sound of a metronome, which does not evoke any response. But after pairing with US, it will reliably evoke a "conditional response" or (CR). The CR, then, is the anticipatory action. In Pavlov's original studies, the same output—salivation—simply occurs earlier in time. As a result of conditioning, the body will react before the actual disturbance or the US occurs. The CR is an anticipatory version of the UR and prepares the animal for the US. Pavlov called this process stimulus substitution, since the CS becomes a substitute for the US in its ability to elicit the same response (Pavlov 1927).

If the US is a disturbance to some controlled variable, then the CS, as the predictor of the disturbance, can act like the disturbance itself. For example, the introduction of glucose into a rat's stomach will cause a rise in blood sugar. This increased blood glucose is not the UR. It is merely a consequence of the glucose absorption. Repeated introduction of glucose following a CS will result in the development of a hypoglycemic response to the CS. The reduction in blood sugar will compensate for the glucose disturbance. The CR attenuates the impact of the US (Dworkin 1993). This can be accomplished by redefining the input function, so that US input function is now responsive to the CS input. The weight of the existing connection may be negligible, which explains why the CS is originally "neutral", i.e. does not evoke any response. Pavlovian conditioning can therefore strengthen this connection, adding a new source of signal to the input function of the UR control system. This type of mechanism can explain some of the constraints on learning. For example, the closer together the two input functions are, the better will be the consequent conditioning. Obviously, if the CS is in the same sensory modality as the US, the conditioning is easier because there may already be existing weak connections between the CS input function and the motor neurons producing the UR.

It is beyond the scope of this essay to discuss the mechanisms of learning in any detail, but it is important to point out the critical role of the homeostatic control systems in driving the reorganization. The error signals from the essential variables can produce widespread random fluctuations in the parameters of other control systems (Powers 1973b). With the reflexive withdrawal of a hand from the approaching fire, the controlled variable is sensed temperature, and the UR reduces that temperature. When there is a burn, the error signal can randomly change the weights of the system. And a reduction in the error signal can stop this reorganization process. For example, if there is a CS that predicts the fire, the weight connecting this signal to the same motor neurons that produce the UR will sometimes be larger due to the random fluctuations, producing earlier withdrawal from the heat source and a reduction of error, which in turn slows down the reorganization process, retaining the increased weight.

But pure stimulus substitution of this type cannot be common for a number of reasons. A predictor of the disturbance can be detected by sensors from many modalities, which are not located in the local spinal circuitry critical for the withdrawal UR. The projections from these detectors to the detector of the US are not usually direct. Consider the noise of the crackling fire, the smell of the smoke, or the sight of the fire—in all these cases the sensory signal cannot go directly from the perceptual channels to the spinal motor neuron innervating the flexors. The expansion of the input function in stimulus substitution is not sufficient. The input function must now associate a high level perception, e.g. sight of fire, to the low level perception of the US disturbance.

One possibility is that the output function of the higher-order system that receives the CS can be modified, so that a projection is sent to the lower system to alter its reference signal. This is supported by experimental evidence. In eyeblink conditioning, an air puff to the eye elicits a reliable blink, but if the air puff is predicted by an auditory stimulus, the blink will occur earlier in time, following the CS. Although the CR and the UR are similar, pharmacologically blocking the red nucleus prevents expression of the CR but not the UR (Thompson et al. 1998). The CR circuitry can thus be dissociated from the UR circuitry, though they share the final common path, using the motor neurons in the cranial motor nuclei. The CR can be generated by a hierarchically higher level sending a reference signal to the motor neuron generating the UR. The higher system, which involves the interpositus nucleus of the cerebellum, generates an output that, via the red nucleus, changes the reference signal for the motor neuron, which serves as a comparator. A change in the reference signal can also produce the blink, in the same way that a change in the perceptual signal does. The reference signal “simulates” actual disturbance. The comparator takes the difference between the actual disturbance to the controlled variable (proportional to US magnitude) and the reference signal from the higher system, generating an error signal that produces the CR. The CR reduces disturbance, and as a consequence there will be less error from the output function when the US arrives. The UR is reduced, since the effect of US is already resisted by the CR. This phenomenon is commonly called conditional diminution of the UR, an example of US processing (Domjan 2005), but it is just due to the disturbance resisting effect of the CR.

Traditionally, the term Pavlovian conditioning is simply used to describe whatever happens when some neutral stimulus is paired with some disturbance to an essential variable. But this classification is not based on any common underlying mechanism. In some cases, whether by redefining the input function or by changing the reference signal from the output of a higher level, the CR and the UR share the final common path. This is true for salivation and eyeblink conditioning. But in other cases, the CR and the UR are different. For example, in fear conditioning, an auditory tone predicts electrical shock to the feet. The URs to shock include jumping up and down. The CR is freezing, a stereotypical defense reaction to danger in rodents. When the predator approaches and there is nowhere to escape, the animal freezes. The sight of predator is a danger signal, not a pain signal, and the behavior in response to danger and to pain can be quite different. Of course, when the predator is attacking the prey, the prey does not freeze—it struggles to escape, much as the shock elicits avoidance behavior. If the sight of the predator actually triggers the unconditional response to pain, as in stimulus substitution, then the animal will simply alert the predator.

Thus, it is not enough just to produce same output earlier in time, in response to a predictor of the disturbance. A different type of output is required to affect the disturbance in question. Instead of resisting the same variable as the UR does, the CR controls a different perceptual variable. By controlling a different aspect of the disturbance, it is able to reduce or prevent the error signal in the essential variables. It is by not alerting the predator that the rat can avoid injury and death altogether.

As we have seen, anticipatory control, as revealed by Pavlovian conditioning, can involve three possible mechanisms. First, conditioning can redefine the input function within the same control system that generates the UR. The US is always a disturbance to some essential variable under homeostatic control, and the additional input from the CS detector to the UR motor neuron can become stronger. This mechanism requires the CS input function to be at the same level as that of the US. No hierarchical organization is needed, since the same system that controls for the effect of the US can be modified to produce the anticipatory output. Second, a higher system detecting the CS sends an output to the lower system which controls the variable disturbed by the US. The CR and the UR share the final common path, but the CR is generated by changing the reference of the US control system. To the motor neuron, such a change in the reference signal is similar to a change in perceptual input caused by the disturbance. The only difference is in the timing. The CR is produced earlier and consequently reduces the impact of the US disturbance when it actually arrives, thereby also reducing the UR. Finally, it is possible for higher control systems to alter the reference of alternative control systems that can control for different perceptual variables and generate different behavioral outputs. As a result of controlling these other perceptual variables, the error signal in the US control system is reduced or even prevented.

5.3 Instrumental Control

In Pavlovian conditioning, the animal is usually prevented from affecting the CS, which is under the control of the experimenter. If it is free to behave as it normally would, without the interference of the experimenter, the animal will attempt to control the CS. Imagine what would happen to a dog if acid is introduced into its mouth. It is only under very artificial conditions that we can measure the anticipatory salivation and repeat the acid treatment. The restrained animal simply cannot behave freely. It would be a very unusual dog indeed if it simply accepts this state of affairs. Had it not been restrained, it might not be so kind to the experimenter who is injecting acid into its mouth. The way to avoid acid is not to dilute it with saliva, but to attack the man in the lab coat. This type of behavior, acting on the source of the disturbance and exerting control over it directly, is familiar to us. More than the reflexes and simple anticipatory behaviors, it is the most common behavior in humans. Such goal-directed, voluntary behavior requires yet another type of control, which I shall call "instrumental control," often studied in instrumental or operant conditioning experiments.

To understand the distinction between instrumental control and anticipatory control discussed above, consider the example of approach behavior. Animals will often approach predictors of reward. In the well-known experiments on pigeons, a key light was lit before food presentation—a basic Pavlovian experiment with key light as the CS and food as the US (Brown and Jenkins 1968). Pigeons learned to approach and peck the key light that predicted food delivery. The US is not contingent upon the pecking: there is no feedback function between the anticipatory behavioral output and the food delivery. Then the experimenter reversed the contingency between the pecking and the food delivery. Now food was delivered following the CS as before, but any pecking of the key light would cancel it. Surprisingly, when this "omission" contingency was imposed, the pigeons did not immediately reduce their pecking behavior. The CR persisted, even though it canceled the reward.

Before calling the persistence of the CR maladaptive, it is worth noting that, in the natural environment outside the laboratory, sign tracking is very effective. Approaching the sound of the apple falling from the tree, e.g., usually gets you closer to the apple. There is a stable relationship between the sign and the goal, so that proximity to the sign means proximity to the goal itself. Of course, it is possible to imagine a world with a different arrangement, but evolution does not operate on imaginary worlds.

Once the animal has acquired the anticipatory tracking behavior, it cannot withhold it even when doing so increases the error signal. This is trivially true for first order control systems. One cannot refrain from sweating in order to cool the body. If some mischievous experimenter arranged the contingency so that sweating will increase room temperature, then the body will still continue to sweat, even though it is increasing the error. The error signal that represents "too much heat" can only be translated into sweating. It cannot be used by the same control system to stop sweating. A first order control system cannot adapt to a reversal in the polarity of the feedback function, though of course the error signal can engage different control systems in a hierarchy that allows one to leave the heated room. More importantly, the control hierarchy that produces anticipatory control also cannot cope with this type of reversal, as revealed by the omission test. The ability to alter behavior following a reversal in feedback function is a distinctive feature of instrumental control. A typical instrumental contingency could be “push the door to open it,” or “pull the door to open it.” The same goal is reached with different actions. The same error signal can be used to perform two opposite actions. Distinct lower level control systems can be selected depending on the feedback.

What is the difference between instrumental control and anticipatory control? If we hear the sound of an apple falling, we may approach the location of the sound. The fact that we are closer to the apple when we approach the sound is due to the environmental contingency. Apples normally stay close to where they fall—a feature of the environment independent of any organism, though the efficacy of the sign tracking behavior depends upon it. As a stable feature of the environment, it is probably responsible for the evolution of the sign tracking behavior. Note, however, that neither salivation nor approach changes the rate of the apple falling. These behaviors control for variables related to the apple—proximity or ease of digestion, but they do not cause the falling of the apple. They do not operate on the distal environmental variable that is responsible for the apple falling. On the other hand, instead of waiting for the apple to fall from a tree, we may shake the tree to make more apples fall, or we can climb the tree to get the apples. We can even grow apple trees. These are all examples of instrumental actions.

A key feature of instrumental control is the diversity of the means to achieve the same end. Wittgenstein once observed that the category of games have nothing in common, in the sense that there is no Platonic essence in which all games partake. He described the relationship among members of this category as “family resemblance” (Wittgenstein 1953). But Wittgenstein failed to consider the function or the goal of the game. Similarly, Skinner has defined the “operant” as a family of movements whose only common feature is that they effectively earn the reward. For example, there are many ways to move the lever in order to produce food delivery. One cannot even define a priori what these movements will be, just as one cannot predict ahead of time all the members of the category of game. If we carefully measure the physical attributes of each lever press in terms of the kinematics, or the pattern of neural signals sent to the muscles, we will not find some essence that all of them have in common, except their goal. But the goal is not a property to be abstracted from the physical properties of the movements, just as the reference signal is not something you can extract from the behavior of the thermostat.

Anticipatory behavior is based on a stable environmental contingency—approach the location of the sound, you approach the apple itself. The CR is generated by a control system controlling for a particular perceptual variable with a fixed relationship with the essential variable controlled by the UR system. The only way to reduce the CR is to change the relationship between the CS and the US. For example, the organism does not affect the reliability of the key light in predicting food presentation. One can degrade this relationship and reduce the CR (Schwartz and Gamzu 1977). By contrast, the instrumental action is based on the contingency between the action and the goal. The rate of the action can itself be a controlled variable. The system that controls the rate of reward, e.g., can send a reference signal to the lower system that controls the rate of action, which in turn tells other lower systems what to sense in completing the action.

During instrumental conditioning, by trial and error the animal identifies the mechanism underlying this environmental contingency (that apple falls from the tree), and acts on it by shaking the tree or climbing it. There is therefore learning about the action-outcome or instrumental contingency (Balleine and Dickinson 1998). The goal is related to the error signal of some essential variable. A reduction in blood sugar, e.g., may produce an error signal in the homeostatic control system, triggering a number of autonomic responses, but at the same time such an error signal may be transformed into higher level perceptual signals, related to hunger and appetite, or to specific representations of goals like apples. Coupled to knowledge of the instrumental contingency between actions and outcomes, the higher levels may choose from multiple action systems, which are themselves hierarchically organized, and use these as output functions to control a variable like the rate of apples obtained.

5.4 Perception and Action

Hierarchically higher control systems can send reference signals to the homeostatic control systems. Consider the behavior of withdrawing your hand from excessive heat. Let us assume that the reference signal has a net inhibitory effect on the motor neurons innervating the flexors. In the lower control system, the pain signal, which exceeds the value allowed by the reference signal, can produce the withdrawal reflex. The reference signal can be translated as “do not let the input exceed this value” and the output reduces the input, by moving the hand away from the source of pain. But we can hold a hot cup without dropping it. Thus changing the reference signal can increase pain tolerance of the flexion withdrawal system. The reference signal of the higher level might be “do not drop this cup”; and of a still higher one, “do not embarrass yourself.”

When a higher control system sends a reference signal to a lower one, the lower one essentially serves as an extension of the output function of the higher level. Yet the higher systems can only have limited influence on the reference signals, because the value of the essential variables must stay within a range before survival is threatened (Mrosovsky 1990). Although the top down reference signals to the homeostatic control system can change their outputs, the autonomic outputs are not sufficient to serve as an extension of their output function. This may explain the difficulty with operant conditioning of the autonomic outputs (Dworkin 1993). Moreover, the homeostatic control systems can only have very limited effects on the environment. To act on the environment, to extinguish the fire or to climb the tree for apples, the skeletomotor system is needed, which has a parallel role in extending the output function of higher control systems.

The skeletomotor system is itself a hierarchy of control systems. The last output function of this hierarchy is the final common path, comprising the alpha motor neurons and the muscles they innervate. The force of the muscle contraction, a function of the motor neuron output, is sensed by the Golgi tendon organs. The reference signal for force control can come from multiple sources, such as the Ia afferent and descending projections to the alpha motor neuron. The output produced by the alpha motor neuron is sensed by the Golgi tendon organs, and the sensed signal representing load in turn is compared to the sum of the force reference signals. The alpha motor neuron can receive reference signals from multiple higher sources. The brain cannot send a command to specify the tension of the muscle. It can only specify how much tension is to be sensed. This does not determine how tense the muscle will be, for if there is a disturbance, the muscle will adjust its degree of contraction to compensate for the disturbance. If the sensors are activated to the desired value artificially, e.g. by pulling the tendon organs, the muscle tension will not increase because the reference signals merely specify the requisite level of sensed tension.

Movement per se involves multiple levels controlling for multiple variables such as muscle load, length, joint angle, and sequential activation of agonist and antagonist. But distal senses like olfaction, audition, and vision are not necessary for these types of control. The modality of the perceptual input is primarily proprioceptive, which come from specialized sensors that report the current state of the muscles, tendons, and joints. The distal senses play an auxiliary role in movement. They are critical for goal-directed behavior, particularly for instrumental control and anticipatory control, but they are not necessary for the very character of movements themselves. The spinal cord, brainstem, and cerebellum are sufficient for the control of the proprioceptive inputs, producing the needed output from motor neurons. Thus, it is not surprising that anticipatory control and instrumental control generally require the cerebral hemispheres. Decerebrate animals can still move, but they can no longer control abstract variables such as proximity to predator or the rate of food reward. Despite the enormous complexity in the skeletomotor system, it can be viewed as an extension of the output function of higher order control systems, which use the skeletomotor system to reach specific goals.

Likewise, the perceptual hierarchy for the distal senses like vision and audition are also extensions of the input function for the essential variables. They permit the perception of distal aspects of disturbances and predictors of disturbances. Whereas first order control systems (simple reflexes) receive the first order inputs directly from receptors in different sensory modalities, hierarchical control requires the construction of new perceptual variables. From homeostatic to instrumental control, there is a continuous spectrum of perceptual inputs. They permit the representation of abstract variables like the rate of reward by combining lower level inputs. New variables continue to be formed through experience, becoming capable of being controlled by the skeletomotor outputs through the diverse and ever changing feedback functions in the environment. Note that neither anticipatory nor instrumental control is accomplished by a feedforward model making detailed calculations about the outputs needed ahead of time. In spite of the tremendous complexity in the hierarchical organization, each control system still operates by negative feedback and control of input.

6 Test for the Controlled Variable

As mentioned above, traditionally the input is introduced and manipulated precisely, while output is measured. No question is asked about what the impact of the output is on the controlled variable, because the existence of negative feedback is not recognized. Based on the linear causation paradigm, many experiments were designed to turn organisms into open loop systems. For example, the cuttlefish attacks prey by ejecting its tentacles. If the target is pulled away just after the cuttlefish has begun to eject its tentacles, the direction of the strike is not adjusted, and the tentacles miss their target (Camhi 1984). This type of manipulation is too abrupt and tells us little about the behavior in question. Any control system can fail, especially when the environment presents disturbances that exceed the capacity for control. During a hurricane, the behavior of a pigeon may be indistinguishable from that of a rock of similar mass; but that does not mean that the pigeon is an open loop system. It could be an example of open loop control, but it could also be an example of failed closed loop control. Manipulations that do not take into account the timescale and the capacity for control just force the control systems to fail. It is like killing an animal to show that it cannot control its body temperature.

The principles of negative feedback control, however, suggest a very different way of studying behavior. To study any behavior, it is necessary first to identify the controlled variables. This has been called the test for the controlled variable. Rather than asking how an organism responds to a stimulus or generates intrinsic activity, we ask which variable it is trying to control, i.e. what the purpose is.

Fatigue, e.g., can be a consequence of strenuous activity, but it is not usually the purpose. If we hypothesize that fatigue is the purpose of someone moving a sofa from a truck to his living room, we can introduce disturbances systematically and measure resistance to the disturbances to the controlled variable. We can move the sofa for him, thus preventing fatigue unless he insists on moving it himself. If the controlled variable is fatigue, he will resist the attempt to reduce fatigue for him. On the other hand, if the controlled variable is not fatigue but the position of the sofa, he would gladly accept the offer. Introducing disturbance to the "sofa position" variable would then encounter predictable resistance.

To test for the controlled variable, then, we can manipulate the disturbance to a hypothetical variable or the feedback function (Marken 2001). Systematic resistance to such manipulations is evidence for control. In the scratch reflex, to terminate the irritation on the skin, it is necessary for the scratching movement to target specifically the area being stimulated. To discover the controlled variable, we can manipulate the relationship between the scratching and the sensory input. For example, a frog with its brain disconnected from its spinal cord can still remove a chemical stimulus from the skin of the forelimb by wiping with the ipsilateral hindlimb (Fukson et al. 1980). When the experimenter moves the forelimb, the coordinates of the stimulation site in relation to the hindlimb are altered. Since the same receptors are stimulated, the reflex arc model predicts that the same output will

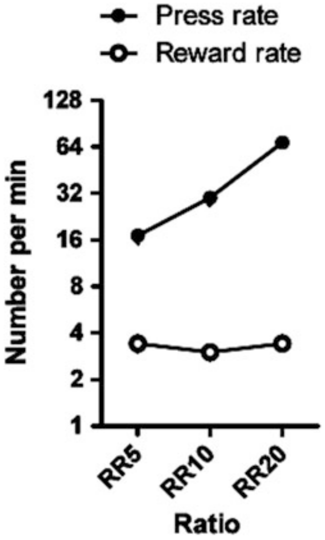

Fig. 5 Control of reward rate in operant conditioning. We used a random ratio (RR) schedule in which the ratio varies randomly around a mean specified ratio value. As we manipulated the average ratio (RR5, RR10, and RR20), the rate lever pressing changed accordingly, while the rate of pellet delivery remained relatively constant. Shown are the steady state values of the rate of lever pressing and the rate of pellet delivery after at least two sessions of training on each ratio ($n = 6$, error bars represent standard error of the mean). The rate of reward appears to be the controlled variable

be elicited, though of course repetition of the same scratching movement would have missed the target region. But in fact the spinal frog successfully moves the hindlimb to the correct target region, the output compensating for the effect of the disturbance (change in stimulation location). For example, wiping at the same site occurs not only when the forelimb is parallel to the body but also when it is perpendicular to the body. The outputs necessary for targeting these two locations are quite different, but the end result is the same. In this study, therefore, the researchers simply manipulated the feedback function, so that a very different output was required to reach the same physical location.

Another illustration comes from our own work on operant conditioning. We trained six hungry rats to press a lever for food pellets. We manipulated the schedule of reinforcement, the feedback function that links the behavioral output to food input. The ratio of a schedule defines how many lever presses must be produced before a food pellet is delivered. As we increased the number of presses required to earn a food pellet from 5 to 20, the rats increased their rate of pressing accordingly, but the overall rate of reward remains relatively constant (Fig. 5). Thus by varying the feedback function, we can show that the rate of reward is the controlled variable. Of course, if the ratio is too high, then the rats will not be able to press rapidly enough to defend their preferred rate of reward. The cost of lever pressing is assumed to be negligible so long as the ratio is not too high, which appears to be the case when the average number of presses required per reward (45 mg pellet) is less than 20. The controlled variable in this case is therefore the rate of pellet delivery, and the output (rate of lever pressing) simply varies in order to defend the preferred rate of reward. Such results cannot be explained by traditional models based on the concept of reinforcement (Sutton and Barto 1998). These models inevitably predict that decreasing the amount of reinforcement will also reduce the behavior that is reinforced, whereas we observed the opposite. Nor do reinforcement-based models predict a constant rate of reward—the controlled variable being defended as a result of variable behavioral output.

7 Conclusion

Behavior is not usually recognized as the manifestation of control. Although we can easily observe the movements of animals and inanimate objects, we cannot just as easily observe purpose. Folk psychology, though much maligned, is not so absurd as to deny the existence of purpose.

Words like purpose and goal are of course used often, but the verbal acknowledgment of such ideas is deceptive. Underneath these models we still find the same linear causation paradigm. Many, e.g., argue that internal representations of the goals can cause behavior. This type of "cognitive" or "purposive" model is meant to replace the S-R model. But as we have seen, this is an inside-out S-R account. It is incorrect to say that the purpose or representation of the goal causes behavior. In the analysis of negative feedback control systems, there are two simultaneous equations to be solved, because the output acts on the input at the same time that the input is acting on the output. To account for purposive behavior, it is necessary to take into account both equations.

In neuroscience, on the other hand, the function of the brain is often described as a "sensorimotor transformation." As the information from the peripheral receptors travels to the brain, at each step some transformation occurs, ultimately resulting in the commands from the brain to generate the behavior (Sherrington 1906). Regardless of how complex the transformation is in the brain, this is also a linear causation model.

Thus whether or not words like "purpose" or "goal" or "prediction" are used is quite irrelevant. We can always find out exactly what the underlying assumptions are by examining the experimental methodology. Does the experimenter attempt to identify the controlled variable (purpose)? Or does he simply manipulate sensory input while recording behavioral output?

Purpose must also be distinguished from equilibrium or consequence. The equilibrium point is not a controlled variable. In a pendulum, the restoring force can be predicted precisely given the disturbance. The displacement itself is solely responsible for the corrective oscillation of the pendulum. In a control system, the resistance to disturbance cannot be predicted from the disturbance, because it cancels the disturbance in relation to an internal reference condition. How much resistance is offered depends on the reference signal as well as the disturbance. All the energy used to correct a deviation of a variable from the equilibrium point comes from the disturbance that caused the deviation. The loop gain is very low. In a control system, the energy used to resist the disturbance does not come from the disturbance itself. In animals, it comes from food. The closed loop system has a high loop gain, which offers a lot of resistance in response to a little disturbance.

Descartes, impressed by the machines of his day, tried to apply the mechanistic reasoning to explain the behavior of organisms. Unfortunately, he only knew about open loop machines, and it was this narrow concept of mechanism, based on linear causation, that has become synonymous with science. In physics, Galileo was right to reject purpose—since natural bodies are not control systems. The success of physics, however, has led to the erroneous conclusion that purpose does not exist.

Until the formal analysis of negative feedback control was properly introduced into the study of behavior, no one understood exactly what a purpose is (James 1890; Tolman 1932). Most conceptual confusions are created by the campaign to remove teleological explanations of any biological phenomenon. After all, modern physics owed its beginning to a victory—no easy victory—over teleological thinking. Yet purpose is not a metaphysical fiction, though the reflex-arc increasingly appears to be one. Nor is it an illusion created to satisfy our vanity, though the illusion of a stimulus-response transformation can be demonstrated easily. And ironically even the argument that the illusion of purpose is created in order to satisfy our vanity is a teleological explanation. Not only is purpose an operationally defined and experimentally testable concept, it is the only one that can make exact experiments on behavior possible for the first time. A restoration of purpose is therefore a new beginning for the sciences of behavior.

References

Ashby, W. (1960). *Design for a brain*, 2nd edn. New York: Wiley.

Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, *37*(4–5), 407–419.

Bernstein, N. (1967). *The coordination and regulation of movements*. Oxford: Pergamon.

Black, H. S. (1934). Stabilized feedback amplifiers. *Electrical Engineering*, *53*, 114–120.

Brown, P. L., & Jenkins, H. M. (1968). Auto-shaping the pigeon's key peck. *Journal of the Experimental Analysis of Behavior*, *11*(1), 1–8.

Camhi, J. M. (1984). *Neuroethology*. New York: Sinauer.

Cannon, W. (1932). *The wisdom of the body*. New York: W. W. Norton.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge: MIT.

Domjan, M. (2005). Pavlovian conditioning: a functional perspective. *Annual Review of Psychology*, *56*, 179–206.

Dworkin, B. (1993). *Learning and physiological regulation*. Chicago: University of Chicago Press.

Franklin, D. W., & Wolpert, D. M. (2011). Computational mechanisms of sensorimotor control. *Neuron*, *72*(3), 425–442.

Freud, S. (1915). Instincts and their vicissitudes. In *Collected papers*. New York: Basic Books.

Fukson, O., Berkinblit, M. B., Feldman, A. G. (1980). The spinal frog takes into account the scheme of its body during the wiping reflex. *Science*, *209*(4462), 1261–1263.

Fuster, J. M. (1995). *Memory in the cerebral cortex*. Cambridge: MIT.

Glimcher, P. W. (2005). Indeterminacy in brain and behavior. *Annual Review of Psychology*, *56*, 25–56.

Graham Brown, T. (1911). The intrinsic factors in the act of progression in the mammal. *Proceedings of the Royal Society of London*, *84*, 308–319.

Hammond, K. R., & Stewart, T. R. (2001). *The essential Brunswik*. New York: Oxford University Press.

Hartley, D. (1749). *Observations on man*. Bath: Leake and Frederick.

Hull, C. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.

James, W. (1890). *The principles of psychology*. New York: Henry Holt.

Konorski, J. (1967). *Integrative activity of the brain*. Chicago: University of Chicago Press.

Lashley, K. S. (1951). The problem of serial order in behavior. *Cerebral mechanisms in behavior: The Hixon symposium* (pp. 112–146). New York: Wiley.

Leyton, A. S. F., & Sherrington, C. (1917). Observations on the excitable cortex of the chimpanzee, orangutan, and gorilla. *Experimental Physiology*, *11*(2):135–222.

Llinas, R. (1988). The intrinsic electrophysiological properties of mammalian neurons: insights into central nervous system function. *Science*, *242*, 1654–1664.

Marken, R. (1993). The hierarchical behavior of perception. *Closed Loop*, *3*, 33–54.

Marken, R. (2001). Controlled variables: psychology as the center fielder views it. *American Journal of Psychology*, *114*(2), 259–281.

Miller, G. A., Galanter, E., Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt.

Millerm, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neurosciences*, *24*, 167–202.

Mrosovsky, N. (1990). *Rheostasis: the physiology of change*. Oxford: Oxford University Press.

Neuringer, A. (2002). Operant variability: evidence, functions, and theory. *Psychonomic Bulletin and Review*, *9*(4), 672–705.

Neuringer, A., & Jensen, G. (2010). Operant variability and voluntary action. *Psychological Review*, *117*, 972–993.

Pavlov, I. (1927). *Conditioned reflexes*. Oxford: Oxford University Press.

Powers, W. (1978). Quantitative analysis of purpose systems. *Psychological Review*, *85*, 417–435.

Powers, W. T. (1973a). Feedback: beyond behaviorism. *Science*, *179*, 351–356.

Powers, W. T. (1973b). *Behavior: control of perception*. New Canaan: Benchmark Publications.

Powers, W. T., Clark, R. K., McFarland, R. L. (1960). A general feedback theory of human behavior. *Perceptual and Motor Skills*, *11*, 71–88.

Robinson, D. A. (1989). Integrating with neurons. *Annual Review of Neurosciences*, *12*, 33–45.

Rosenblueth, A., Wiener, N., Bigelow, J. (1943). Behavior, purpose, and teleology. *Philosophy of science*, *10*, 18–24.

Schwartz, B., & Gamzu, E. (1977). Pavlovian control of operant behavior. In W. Honig, & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 53–97). Old Tappan: Prentice Hall.

Shadmehr, R., Smith, M. A., Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neurosciences*, *33*, 89–108.

Sherrington, C. S. (1906). *The integrative action of the nervous system*. New Haven: Yale University Press.

Staddon, J. E. R. (1983). *Adaptive behavior and learning*. Cambridge: Cambridge University Press.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. Cambridge: MIT.

Thompson, R. F., Thompson, J. K., Kim, J. J., Krupa, D. J., Shinkman, P. G. (1998). The nature of reinforcement in cerebellar learning. *Neurobiology of Learning and Memory*, *70*(1–2), 150–176.

Tolman, E. C. (1932). *Purposive behavior in animals and man*. New York: Macmillan.

Von Holst, E., & Mittelstaedt, H. (1950). The reafference principle. In *The collected papers of Erich von Holst*. Coral Gables, FL: University of Miami Press.
Wiener, N. (1948). *Cybernetics*. Paris: Hermann and Cie Editeurs.
Wittgenstein, L. (1953). *Philosophical investigations*. London: Blackwell.
Wolpert, D. M., Diedrichsen, J., Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, *12*, 739–751.