PART 2

Cognition in Engineered Systems

CHAPTER

1 The Closed-Loop Dynamics of Cognitive Work

John M. Flach, Kevin Bennett, Richard J. Jagacinski, Max Mulder, *and* Rene van Paassen

Abstract

This chapter provides a tutorial introduction to the logic of closed-loop systems. A series of examples is presented to illustrate the dynamics of closed-loop systems and to contrast the behavior of these systems with expectations suggested by the logic of open-loop causal systems. In particular, the examples show that "feedback" per se is insufficient to guarantee convergence on a target. The chapter explores some of the issues that must be addressed in order to determine whether closed-loop systems will be stable. Parallels to the phenomena of cognition in the wild are highlighted, and the case is made that the logic of closed-loop systems is an essential foundation for understanding the behavior of cognitive systems.

Key Words: adaptive control, abduction, closed-loop control, observer, self-organization

Introduction

In 1942, at a conference sponsored by the Macey Foundation to promote interdisciplinary discussions about neuroscience, Arturo Rosenblueth introduced a radical challenge to classical ideas of cause and effect. He introduced the construct of "circular causality." This new construct was motivated by collaborations with Norbert Wiener and Julian Bigelow to understand stability in feedback control systems (Rosenblueth, Wiener, & Bigelow, 1943; Wiener, 1948/1961). The neuro- and social scientists at the meeting (including Warren McCulloch, Gregory Bateson, and Margaret Meade) immediately identified this idea as relevant to their work—easily identifying concrete examples of circular coupling in the biological and social systems that they were studying (Conway & Siegelman, 2005).

There is no doubt that the "cybernetic hypothesis" that emerged from the work of Norbert Wiener and his colleagues had an enormous impact on the trajectory of cognitive science. For example, feedback was a central theme in Miller, Galanter, and Pribram's (1960) influential book *Plans and the Structure of Behavior*.

> The general pattern of reflex action, therefore, is to test the input energies against some criteria established in the organism, to respond if the result of the test is to show an incongruity, and to continue to respond until the incongruity vanishes, at which time the reflex is terminated. Thus, there is "feedback" from the result of the action to the testing phase, and we are confronted by a **recursive loop**. (p. 26, emphasis added)

Miller et al. introduced the TOTE unit (Test-Operate-Test-Exit) as a simple process to illustrate the concept of feedback using the example of hammering a nail. They conclude:

> It may seem slightly absurd to analyze the motions involved in hammering a nail in this explicit way, but it is better to amuse a reader than to confuse him. It is merely an illustration of how several simple TOTE units, each with its own test-operate-test **loop**, can be embedded in the operational phase of a larger unit

> with its particular test-operate-test **loop**. Without such an explicit illustration it might not have been immediately obvious how these **circles within circles** could yield hierarchical trees. (p. 37, emphasis added)

Miller et al. used the "feedback" concept as a means to motivate their exploration of plan hierarchies, and in turn they helped to motivate a 50-year program of research to explore human information processing. However, with some notable exceptions, the cybernetic hypothesis has generally been trivialized, and the field of cognitive science as a whole has failed to appreciate the dynamics of circles and loops. Most in the field have persisted in viewing the phenomenon of human information processing through the lens of the logic of simple open-loop causal systems. Most have failed to understand the full implications of the concept of circular causality.

The classical image of an information processing system is an open-loop series of "stages." These stages form a chain from the stimulus to the response, which is treated as a string of dominoes with the "response" from one stage being the "stimulus" for the next in a sequence of discrete acts. In this framework, there is an implicit arrow of time, such that some stages are seen as logically prior (as input or cause) to other stages. In fact, whole research programs are designed around single stages—as if each stage can be understood independently from the other stages.

However, if the dynamic of cognition is closed-loop, this image will not work. A better image is to think of the stages as being linked within a web—so that an impact anywhere within the web reverberates through the whole web. In a circular dynamic, each stage is simultaneously providing input to other stages and receiving input from those stages. There is no implicit arrow of time. Therefore, relative to the overall recursive, circular flow of influence, no stage can be simply categorized as either "cause" or "effect" relative to another stage. Also, note that the loop is closed through the ecology, so that the ecology is an intrinsic component of the web, simultaneously contributing as both cause and effect (consequence).

The goal of this chapter is to reintroduce the idea of circular causality—to provide a pedagogically sound but analytically simple introduction to the logic of feedback control. We hope to correct some common misconceptions about the nature of feedback and to help those interested in cognition in the wild to appreciate the dynamics of circles and loops and the implications for studying situated cognition.

The Stability Problem

One of the most pervasive misconceptions about feedback is that the presence of feedback is sufficient to explain stable progression toward a goal. In fact, the central problem of control theory as a field of study is to identify the special conditions under which feedback will result in stable control. Stable progression to a goal is NOT guaranteed by the presence of feedback. In Miller et al.'s example of hitting the nail, there is an assumption that each swing is executed exactly right so that the only relevant feedback is whether the nail is flush or not. However, what if the nail position shifts from vertical or the nail bends? How will the swing need to be adjusted to achieve the goal of the carpenter? Or should the goal be shifted to removing the nail and starting over? Hammering a nail so that it satisfies the goals of the carpenter is not as trivial a problem as Miller et al.'s description suggests. They prematurely dismiss the "circles within circles" in order to get to hierarchies, leaving the dynamics of feedback systems unexplained; for the most part, cognitive science has persisted in ignoring the dynamic of circles.

To illustrate the stability problem, consider the simple controller illustrated in Figure 1.1. This system consists of three components: a time delay that reflects the time to process error feedback, a gain that determines the magnitude of correction as a result of the error at each moment, and an integrator that effectively sums or concatenates the corrections across moments to determine the next response. Let's consider the time delay to be a fixed property of the system (e.g., representing the reaction time from seeing an error to initiating a corrective action). Thus, the lone free parameter is the gain, which can be treated as a simple multiplier of the error signal. As input to the system, we will use a step function. Tracking a step can be considered analogous to a target acquisition task. When the signal changes from one level (position) to another, this is analogous to the target appearing—requiring the person to move a cursor from one fixed (home) position to another fixed position (the target).

Can you guess what value of the gain will insure that the output of the system illustrated in Figure 1.1 will follow the target (i.e., so that the output will shift from the home position to reliably converge on the new target position commanded by the step input)? You might guess that a gain

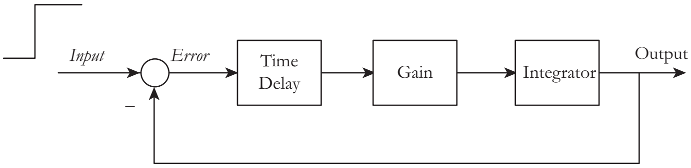

Figure 1.1 A simple feedback system with a step input. How will the output response vary as a function of the gain parameter? Under what conditions will the output converge to match the input?

of 1 would be ideal, insuring a one-to-one match between the input and the output. This is not a bad guess, and depending on the magnitude of the time delay, this may well produce a satisfactory solution. However, it probably will not be optimal. What would happen if the gain were reduced to less than 1 or increased to more than 1? How would this affect the output? Would it cause an offset between the target and the output, such that the size of the output step would scale proportionally to the magnitude of the gain? This is what would happen if this were an open-loop system. However, in this closed-loop system the gain does not determine the size of the output signal, but rather the speed at which the output will converge toward the input. Very low gains will result in a "sluggish" response to the input. The output will eventually reach the target position, but it will take a long time to get there (Figure 1.2). As gain is increased (i.e., sensitivity to error is increased), the speed of correction will increase. However, at some point the higher gains will cause the system output to overshoot the target, before gradually converging back to the target. At even higher gains the output will oscillate around the target before settling down, and at still higher gains the oscillations will actually grow over time—so that the output diverges, never settling down on the target (Figure 1.2).

In general, the range of stable gains will depend on the size of the time delay. As time delays become larger, the range of gains that will produce stable control will get smaller. That is, with longer time delays the system will need to be less sensitive (i.e., lower gains—more conservative or more cautious) in responding to errors in order to avoid oscillatory or unstable responses (e.g., Jagacinski, 1977). In designing automatic control systems, much of the attention is given to eliminating unnecessary time delays and then determining an appropriate gain for the forward loop—in order to insure a fast but stable response to errors. In cases where long time delays are unavoidable (e.g., tele-operation over large distances), there may be no satisfactory gain for proportional control. In these cases, a discrete style of control (small adjustment, wait, small adjustment, wait, . . .) will generally prove to be a more satisfactory means of control.

Figure 1.2 illustrates the range of output behaviors as a function of the gain parameter. It should be clear from this illustration that feedback alone does not guarantee that the output will converge to the target. This is one of the fundamental issues of control theory—to identify those special conditions that lead to stable control, or more generally to study those factors that determine the boundaries of stable control. In fact, it was the parallels between the behavior

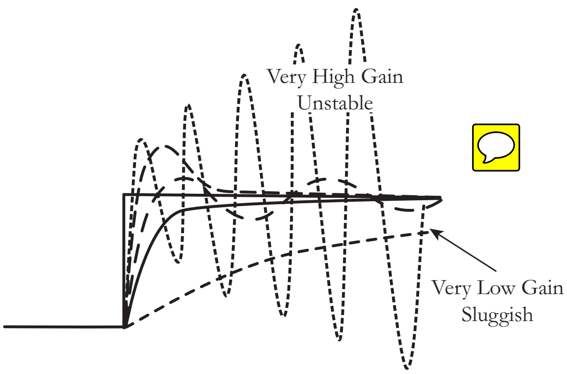

Figure 1.2 This diagram illustrates a range of outputs for the control system illustrated in Figure 1.1 as a function of the gain parameter. For very low gains, the response will "sluggishly" converge to the target without overshoot. For very high gains, the response will become unstable with oscillations that diverge from the target. At intermediate gains, the response will range from relatively rapid convergence to the target without overshoot to converging oscillations as gain is increased within this range.

of unstable engineered control systems and behavior associated with some motor ataxias (purpose tremor) that stimulated the interdisciplinary collaboration between Wiener and Rosenblueth. These parallels suggested that the logic of circles was not simply an engineering problem, but a general problem of any system that adjusts its behavior based on feedback.

In natural closed-loop systems (with inevitable time delays) there will always be a speed-accuracy trade-off. That is, there will always be a limit to how fast a goal can be reliably approached (i.e., a limit to the gain). In human-machine systems, the gain is typically a joint function of the plant (e.g., the gain setting for the mouse, the gain of the control stick or steering wheel, and/or the vehicle dynamics) and the human (e.g., the scaling of the manual response to observed error). You may notice this when you use a computer system that has the mouse gain set differently than on your own computer. If the gain is higher, you may find yourself initially overcorrecting and oscillating around the target before finally capturing it. If the gain is lower, then the new system will feel "sluggish." However, it will typically not take you long to recalibrate to the system gain and to make adjustments so that the combined gain (human + plant) results in reliable target acquisition.

Controls on high-performance aircraft typically have very high gains, so that pilots must compensate by making only small adjustments (lower gain) in response to errors so that the net gain (human + plant) will meet the constraints for stable control. It is not unusual to see novices lose control (i.e., generate pilot-induced oscillations) when they first try to fly a simulation of a high-performance aircraft. That is, their initial gain is too high! Research on human tracking shows that humans can learn to adjust their gain so that the human-vehicle system behavior is stable. Effectively, McRuer's classic crossover model predicts that the effective control dynamics of the human-machine system will (with practice) behave much like the controller illustrated in Figure 1.1 (McRuer & Jex, 1967). That is, the system dynamics will approximate the behavior of a system with a time delay, gain, and integrator in the forward loop. This will be true for a fairly broad range of different plant (or vehicle) dynamics. In fact, the optimal control model predicts that for a range of vehicle dynamics humans can learn to conform fairly closely to the ideals prescribed by normative linear models for optimal tracking of randomly appearing input signals (Kleinman, Baron, & Levison, 1971). That is, the human will choose gains that minimize a weighted combination of mean squared control speed (effort) and mean squared error (accuracy). Similar adaptations regarding plant dynamics, speed, and accuracy can be seen in the Fitts' Law paradigm (e.g., Jagacinski & Flach, 2003).

There are several key points that we hope to have made in this section:

1. The intuitions derived from simple causal assumptions (based on open-loop dynamics) can be misleading when applied to closed-loop systems. That is, increasing gain determines the speed of convergence to the command input in a closed-loop system, NOT the relative magnitude or scale of the response as would be true in the case of an open-loop system.
2. All natural closed-loop systems will exhibit a speed-accuracy trade-off. That is, there will generally be a limited range of gains that result in satisfactory stable performance. When gain is too low, the response will be sluggish (too slow). When gain is too high, the response will become oscillatory, and with very high gain the oscillations can become unstable (increasingly diverging from the target).
3. The main point of this section is that *the presence of feedback alone does not insure stable convergence on a target or goal.* Thus, the presence of feedback is not the answer that explains behaviors guided by a goal. Rather, the logic of circles provides the context for asking the appropriate questions and for generating viable hypotheses about the conditions that lead to stable, goal-directed behaviors.

The Regulator Paradox: Control and Observation

The challenge that intrigued Wiener and Bigelow as they struggled with the design of systems for guiding artillery during World War II was not the design of simple servomechanisms, as is part of social science lore. Rather, the key problem was to predict the future positions of aircraft. That is, because of the speed and altitude of aircraft, you would not hit an aircraft if you fired at its current position. You had to fire at a point that corresponded to a future position of the aircraft (where the plane would be when the missile arrived). This problem was complicated by the fact that pilots would not fly in simple paths. They would maneuver evasively with the goal to make it difficult to predict where they would be by the time a missile arrived. This problem can best be conceptualized as an observation and prediction problem. The challenge is to predict the future based on noisy samples of past behavior. This is a problem of distinguishing signal from noise and extrapolating to predict where the aircraft will be when the missile arrives. This is a problem of anticipating or predicting the future. Whereas the goal of the servomechanism (control system) is to minimize "error" (i.e., the deviation between an input goal and behavior output), the goal of the observer is to minimize "surprise" (i.e., the deviation between estimates of target position and velocity based on noisy observations and the actual position and velocity of the target), and the goal of the predictor is to extrapolate that estimate forward in time.

Figure 1.3 illustrates a simple observer system. Note that this is essentially the same dynamical system as in Figure 1.1. The components include a time delay reflecting the time to register the surprise, a gain reflecting the magnitude of adjustment based on the current surprise, and an integrator that effectively sums or concatenates past corrections to determine the next estimation. As in the previous section, let's assume a fixed time delay that reflects hard constraints on the system. This again leaves the gain as the lone parameter to consider. Let's consider the input to this system to be a noisy stream of data to which we want to estimate the mean as a reasonable prediction of future samples. For this example let's consider this a stream of sampled data about the height of people in a particular population as illustrated in Figure 1.4. Further, let's consider the possibility that in addition to the variability from sample to sample, there is a step change in the mean height of the population (e.g., due to a social change improving diet and access to health care) that occurs during the period of sampling.

In designing an observer for this problem, the goal is to design a system that will not be fooled by the noise (i.e., that will filter out the changes due to noise) but that will be sensitive to real changes in the underlying population (i.e., that will detect the step change when it occurs). The lone parameter you have to adjust the system is the gain. What value for the gain would be best? Certainly, we know from the discussion in the previous section that if the gain is set too high, the system will become unstable and the estimates will oscillate wildly and will not converge to the input signal. So, we know that there is an upper limit to the gain, but how low should we go? At the setting that is ideal for solving the control problem as it was described in the previous section, the output will follow the input signal very closely. That is, it will respond to every change in input, whether resulting from noise (sampling variability) or from a real change in the underlying distribution, because this distinction was not previously considered in characterizing the input. If the goal is to track the "true signal" based on noisy observations, then determining a satisfactory gain requires simultaneous consideration of both the constraints on control and the constraints on observation. So, it is likely that to reduce the impact of noise, the ideal gain for solving the tracking problem would be lower than if there were no observation noise. In this simplified context, the tracking system in Figure 1.1 has the same structure as the observer in Figure 1.3, so the structure has multiple interpretations.

As in the previous section, there is a fundamental trade-off associated with the gain. Figure 1.4 shows a range of responses for the observer problem as a function of changes in gain. As the gain is lowered, the response to the noise signal will be "smoothed" so that the impact of sampling variability on the estimation will be reduced. However, the price of this "smoothing" is that the observer will be slower to detect real changes when they occur. Thus, in the design of an observer, the engineer must weigh the cost of following the noise against the costs of a slower detection of real changes when they occur. The observer can be contrasted with the Kalman filter, which is an optimal solution to the observer

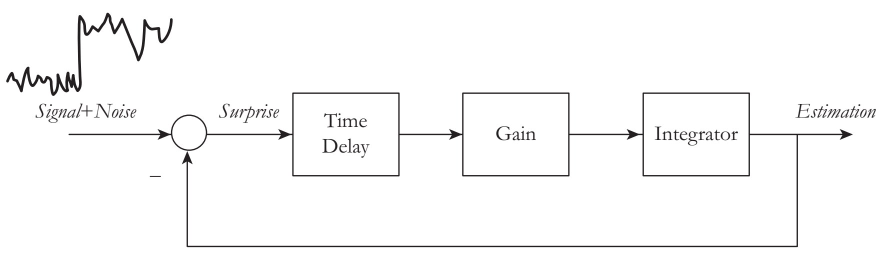

Figure 1.3 This illustrates a simple observer system. The input reflects a noisy signal and the output is an estimate of the signal. In essence, the observer is designed to "filter" out the noise.

problem in a stationary environment. The Kalman filter employs a time-varying gain that minimizes the mean squared error in the estimate. However, if the period of observation is long prior to a change in the environment (Figure 1.4), the Kalman filter will be slower to respond to this change than a well-designed observer.

In essence, this is analogous to a signal detection problem, and the gain parameter will determine the ultimate balance between false alarms (following noise) and hits (detecting real changes). That is, the gain parameter determines the relative sensitivity to signal and noise, as might be captured by the beta parameter in the signal detection model. However, there are important distinctions between the problem as presented here and conventional treatments of signal detection in the social sciences. First, the conventional treatments typically don't consider the factor of time—integration of information from one sample to the next and changes over time in the underlying signal distributions. In essence, conventional treatments treat the signal detection problem as open-loop.

A second important point is that conventional open-loop information processing models and linear closed-loop control models such as the optimal control model (e.g., Pew & Baron, 1978) partition the observation and control processes into separate stages. In a closed-loop model, this elaboration allows different gain settings for the observer and controller. The gain of the observer is chosen to separate signal from noise. The gain of the controller is chosen to satisfy some effort-accuracy trade-off. This partitioning is advantageous for overall performance. However, it does not capture the true intimacy between perception and action that often exists in closed-loop systems. As Wiener and Bigelow discovered in the process of solving practical control problems like targeting evasive aircraft, there is often an intimate coupling between the observer problem (separating signal and noise to estimate the present position and velocity of an airplane), the prediction problem of extrapolating the path of the aircraft forward in time, and the control problem (aiming projectiles relative to the path of the aircraft). For example, a series of projectiles could be fired either to influence or to constrain the path of the evasive aircraft (i.e., to simplify the observation and prediction problems) as well as to actually hit the aircraft.

In the context of discussing adaptivity, Weinberg and Weinberg (1979) summarize the close coupling between observing and controlling in their Fundamental Regulator Paradox:

> The task of a regulator is to eliminate variation, but this variation is the ultimate source of information about the quality of its work. Therefore, the better the job a regulator does, the less information it gets about how to improve.

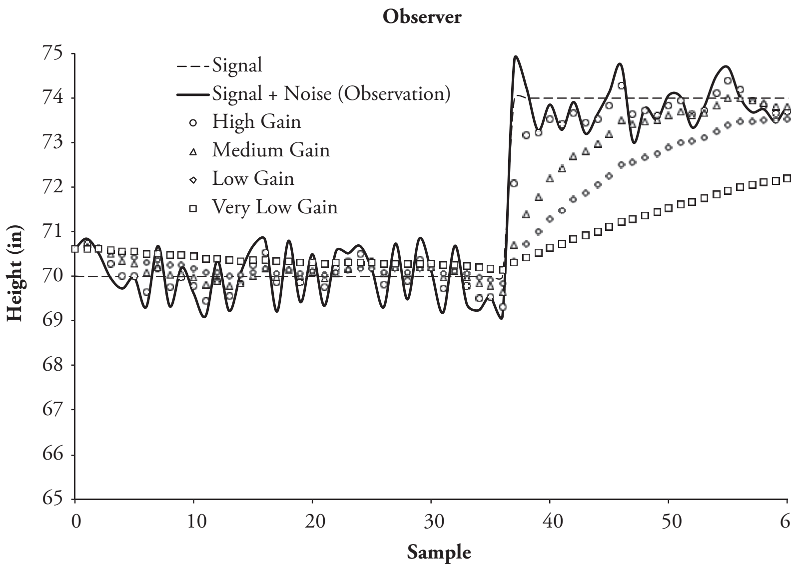

Figure 1.4 This graph illustrates the response of a simple observer as a function of the gain. When the gain is very low, the system will effectively filter out the noise but will be sluggish in responding to changes in the signal. As gain is increased, more of the noise is reflected in the response, but the system is also more responsive to real changes in the signal. At the high gain the output responds to both signal and noise.

or

Better regulation today risks worse regulation tomorrow. (p. 250)

The point is that often cognitive systems and biological control systems must solve the observer, predictor, and control problems simultaneously. That is, they must discover the goals (find the signal in the noise) in real time while simultaneously extrapolating forward in time and managing stable control with respect to those goals. The Weinbergs used the example of controlling a car during the winter to illustrate this coupling. On the one hand, the driver must function as a controller, keeping the car safely on the road. On the other hand, the driver must function as a higher-order observer of system performance—in order to discover potential changes in the control dynamics (e.g., resulting from changes of the road surface such as black ice). In balancing the demands of these two tasks, good drivers learn to act to both steer the car and to test for changes in road conditions (e.g., put in test signals such as jiggling the steering wheel or tapping the brakes). Note that the test signals produce "error" with regards to the control problem, but they provide information with regards to the changing dynamic context (e.g., detecting control-relevant changes in the situation). This information may be critical for maintaining stable control when road conditions change. Gibson (1979/1986) used the terms "performatory action" and "exploratory action" to distinguish between those actions motivated by the demands of control and those actions motivated by the demands for information.

To reiterate, in designing automatic control systems, engineers typically "solve" the control problem and the observer problem offline and then implement the solutions as an automatic control system (with gains tuned to reflect both the signal-to-noise and the stability constraints of the problem). This works well if the environments are relatively stable (i.e., stationary). However, cognitive and biological systems have to design themselves while simultaneously interacting with the problem situation. For example, the new driver or new pilot must discover the appropriate gains with respect to 1) signal-and-noise (i.e., determine what aspects of variability to filter out) and 2) good control (i.e., efficient, stable correction of errors) while engaged in the control task. These systems must learn by doing. These systems must self-organize. Further, to survive, the cognitive and biological systems must be prepared to adapt to changing contingencies. That is, they cannot assume a stable environment. For example, the operators of the anti-aircraft artillery cannot assume that the enemy pilots will not change their evasive tactics, and the operators may even act to influence those tactics.

Today, control theory and information theory are typically treated as distinct fields of research. This is in part due to the different analytic tools (differential equations for control theory; probability statistics for information theory). However, both fields emerged as a result of the work of Wiener and Bigelow on the problem of predicting aircraft movements in order to target them. Note that the subtitle of Wiener's (1948/1968) classical work *Cybernetics* was *Control and Communication in the Animal and the Machine*. In closed-loop systems, observation and control are two sides of a single coin!

The main points for this section were:

1. To show the feedback dynamic from the perspective of information processing or perception. In this case, the gain functions to tune the filtering properties of the dynamic—in order to distinguish signal from noise based on observations over time.
2. To suggest that in most natural systems the control and observer problems are intimately linked. There could be a single gain parameter to meet both the constraints on observation and control in a simplified context, or there could be two or more independent gain parameters for observer and controller that together determine the overall responsiveness and stability of the closed-loop system. In addition, the actions of the system may have multiple goals of influencing the observation/prediction process and well as tracking the input signal.
3. Finally, it is important to appreciate that research into closed-loop systems must be guided by the intuitions of both control and information theory.

The Comparator Problem

In a control or observer system, the comparator is the point at which the output is fed back and "compared" with the input to compute an error (or surprise) signal. In typical engineered control systems such as those illustrated in Figures 1.1 and 1.3, the system is designed so that the comparator simply involves subtraction of one signal from another to get a third signal. That is, all signals are in a comparable currency that allows subtraction of one from another to get the third. However, consider

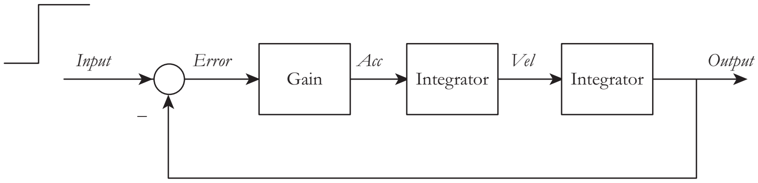

Figure 1.5 A second-order control system in which the gain determines the acceleration of the response (output).

the novice pilot learning to land. The goal might be a particular approach path or simply a soft contact at a particular region of the runway. The feedback would be in the form of optical flow through the windows and/or information presented on the cockpit instruments. These two very different types of signals must somehow be compared in order to specify appropriate movements of the controls (stick, throttle, rudders).

It should be clear that for the pilot, and more generally for most cognitive or biological control systems, the signals involved in the comparator process may be in diverse forms or currencies. Thus, comparing feedback to intentions in order to specify actions is not a trivial process. In fact, this is a central issue for control theory—to determine the dimensionality of the state space or, in other words, to identify what variables must be fed back in order to guide action in a particular situation (e.g., as a function of different vehicle dynamics). This is also probably the central issue for skill development—attuning to the feedback that specifies the appropriate actions with respect to the opportunities and consequences (e.g., E. J. Gibson, 1969; Ericsson & Charness, 1994). In Gibsonian (1979/1986) terms, this is the problem of specification of affordances.

Figure 1.5 shows a simple feedback system to illustrate the comparator problem. This system has a gain and two integrators in the forward loop. As with the other systems, the gain determines the sensitivity to error. Because of the two integrations in the forward loop, the output from the gain element determines the acceleration of the output. This is a dynamic that is consistent with most movement tasks (e.g., vehicle control or body movement) in a world governed by inertia. For example, the initial response of deflection of the accelerator or brake is an initial change in the acceleration or deceleration for your car.

What do you suppose the response of this system would be to a step input, and how might this response change as a function of the lone parameter in the forward loop—the gain? Is there any gain value that will result in an asymptotic approach to the step target? Somewhat surprisingly in the context of naive discussions of feedback systems, the answer is "No." Here again is a situation that illustrates that feedback is not sufficient to insure convergence with the input. In fact, the response of this system to a step input is a sine wave output. The speed of oscillation of this sine wave (i.e., its frequency) is determined by the value of the gain parameter. A higher gain produces a higher-frequency response. There is NO value of gain that will lead to convergence of output with the input target!

Figure 1.6 shows an alternative system that includes feedback of both the output position and the output velocity. This system will result in an output that will converge to the input target, if the feedback of position and velocity are combined with the appropriate weights. Heavy relative weight on the velocity component will lead to "sluggish" or conservative approaches to the target. Less relative weight on velocity will lead to more aggressive approaches—with a damped oscillation at very low

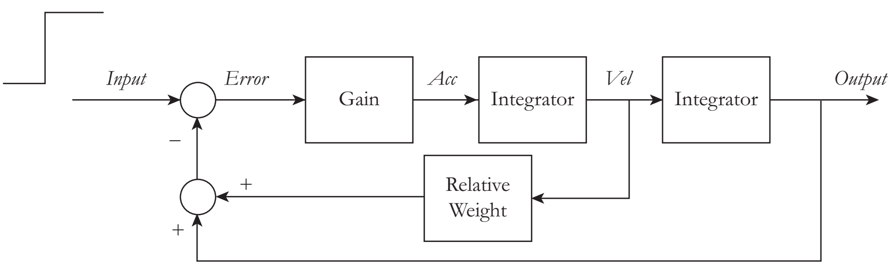

Figure 1.6 Stable goal following can be achieved in a second-order system if feedback includes both position and velocity of the output.

weights and, as described above, a non-converging oscillation at zero weight on velocity. Depending on the relative weights on position and velocity, this system will achieve a similar range of responses (from sluggish to oscillatory) as illustrated in Figure 1.2.

A general implication of the behavior of the systems illustrated in Figures 1.5 and 1.6 is that for control of an inertial system (e.g., a car), position feedback alone is insufficient. The system must have feedback about both position and velocity in order to control the vehicle. For example, in order to stop the car in front of an obstacle on the road (e.g., stopped line of traffic), a driver must take into account both distance to the obstacle and the speed of approach. Current research suggests that for visual control of locomotion, this information is specified in terms of the angular extent (e.g., visual angle of the taillights of the preceding vehicle) and the angular velocity (e.g., rate of expansion of the taillights) of the object in the visual flow field (e.g., Lee, 1976; Smith, Flach, Dittman, & Stanard, 2001; Flach, Smith, Stanard, & Dittman, 2004).

In most natural situations, cognitive systems must deal with many different potentially useful sources of feedback. Figure 1.7 provides a simplified illustration of the multiple variables that need to be considered in controlling the lateral position of an aircraft when trying to track a specific approach path. The top portion of Figure 1.7 illustrates the multiple "state variables" that a controller must be aware of in order to control effectively. The input from the pilot directly affects the position of the aileron, which is integrated (first-order lag) to determine the roll rate, which in turn is integrated to determine the roll angle, which is integrated to determine the heading angle or turn rate, which ultimately is integrated to determine the lateral position of the aircraft. The output of each integration represents a "state variable" of the aircraft that must be fed back in order to achieve reliable control. Note that the lateral control problem is only a subset of the variables that need to be considered in landing.

The lower portion of Figure 1.7 illustrates how each of these state variables might be fed back and combined to determine the appropriate control adjustments. Each loop has a specific gain that in effect reflects the relative weighting of each of the state variables in determining the next adjustment of aileron position. In designing an effective control system, the engineer would set the gains in each loop to reflect the dynamics of a particular aircraft. Similarly, a pilot who is learning to fly the aircraft must discover the appropriate mappings from the perceptual information associated with each variable (e.g., properties of optical flow or instrument

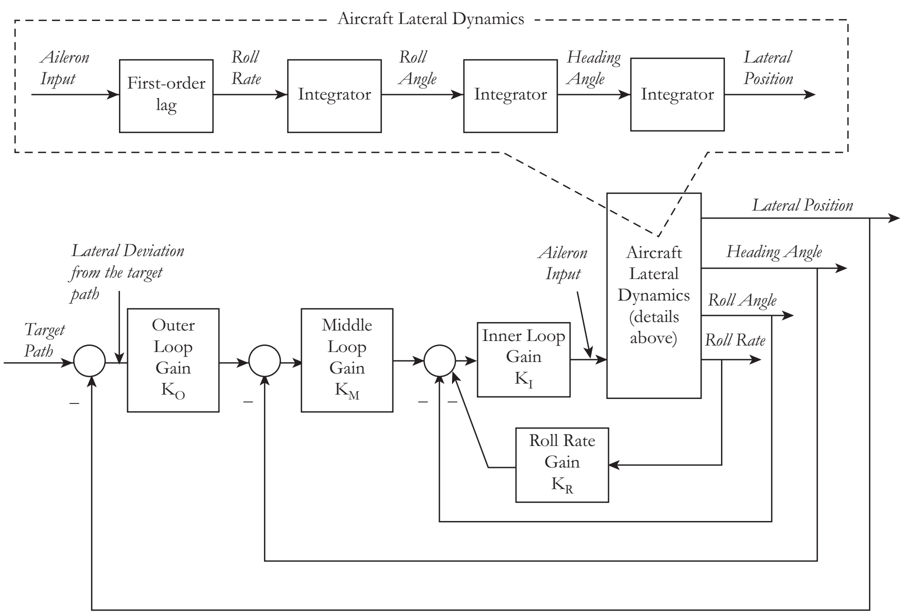

Figure 1.7 The top diagram provides a highly simplified representation of the aircraft lateral dynamics and the state variables involved in controlling lateral position of an aircraft (e.g., in trying to track the target approach path to an airport). The bottom diagram illustrates the multivariable control problem. Each of the state variables are perceived, weighted, and combined to determine the appropriate control actions.

readings) and the adjustments of his or her control interface (e.g., control stick).

In addition to considering the multiple state variables associated with the lateral and vertical positions of the aircraft, pilots must consider potentially conflicting goals (e.g., aborting the approach when a taxiing aircraft encroaches on the runway), and many potential actions or means for moving toward those goals (e.g., whether to make stick or throttle adjustments to correct for excessive air speed). Thus, the comparator problem involves attuning to many potentially useful sources of feedback; it involves setting priorities to appropriately trade off the potential consequences of multiple competing goals; and it involves choosing among many potential means for reducing the error between output and intention. These multiple degrees of freedom in terms of multiple consequences, multiple sources of information, and multiple potential control actions emphasize the intimacy between observation and control introduced in the preceding section.

Also, it is important to understand that the setting of the gains in Figure 1.7 that result in satisfactory control depend on the particular aircraft dynamic. However, the aircraft dynamic itself may change as a function of context. For example, the dynamics can change as a function of both speed and altitude. For example, gains that lead to stable control at low altitudes may lead to instabilities at higher altitudes. In the design of automatic control systems, engineers address this by designing adaptive control systems. As illustrated in Figure 1.8, an adaptive control system is capable of adjusting the gains on the inner loop dynamics as a result of monitoring the context or the quality of inner loop responses.

In Figure 1.8, the thin (wire) arrows represent the flow of signals or information that is then acted on or processed according to the workings (i.e., the transfer function—e.g., control gains) of each box. However, the fat arrows that close the outer loops are signals that change the properties of the boxes (e.g., change the control gains or the expectations within the boxes). Note that the inner loop includes three modes of action: performatory actions are intended to reduce error, exploratory actions are intended to test hypotheses, and anticipative actions reflect direct action to achieve a goal. The anticipative path reflects direct response to the reference (i.e., not dependent on error feedback). This is an open-loop path from the reference or goal and could reflect actions that are shaped by previous experience with the plant dynamics (e.g., what has often been called a mental model or schema). That is, these are responses in anticipation of a consequence, rather than in response to error feedback. Actions may fulfill any or all of these roles at any moment.

Figure 1.8 illustrates three different ways that an engineer might close the outer loop to achieve stable adaptive control. First, the adaptation might be a direct function of the changing context, as illustrated in the outermost loop. For example, the engineer might compute the appropriate gains for different altitudes and preprogram these different gains into the automated control system. The gains would be changed as a function of a direct measure of the appropriate context variable (e.g., the altitude). This is typically called "gain scheduling." In human performance, this path may be representative of the phenomenon of context sensitivity. That is, the strategy for controlling action or the expectations of the human agent may change as a function of the situation. For example, data suggests that a teenage driver with a parent in the car is one of the safest drivers, while a teenage driver in a car with other teenagers is one of the most dangerous drivers.

The next outer loop (hypothesis) in Figure 1.8 reflects conscious exploration of the dynamics through exploratory actions. In engineered systems this might involve a low-amplitude test signal that is constantly input to the plant (i.e., dithering). This input is designed to have minimum consequences relative to the performance objective (i.e., to produce minimal error). However, the changes in the properties of the output from this signal can be information relevant to detecting changes in the plant dynamics. Deviation in the output related to the dithering can be fed back and used to adapt the control gains (and the expectations). Skilled human drivers use a similar strategy to test for possible changes in the driving dynamics due to changing road conditions. They may dither the steering wheel to test for changing traction (Weinberg & Weinberg, 1979). The reference signals for this loop are labeled "local expectations" to reflect that this loop reflects explicit tests of local hypotheses. In this loop the human agent is acting as a test signal generator and observer—in order to detect control-relevant changes.

The innermost of the outer loops (surprise) represents an approach to adaptive control that engineers call "model reference" control. With this style of adaptive control, a normative model of the plant dynamics can be simulated in parallel with the actual performance. For example, this model might be a simulation of aircraft performance at the typical

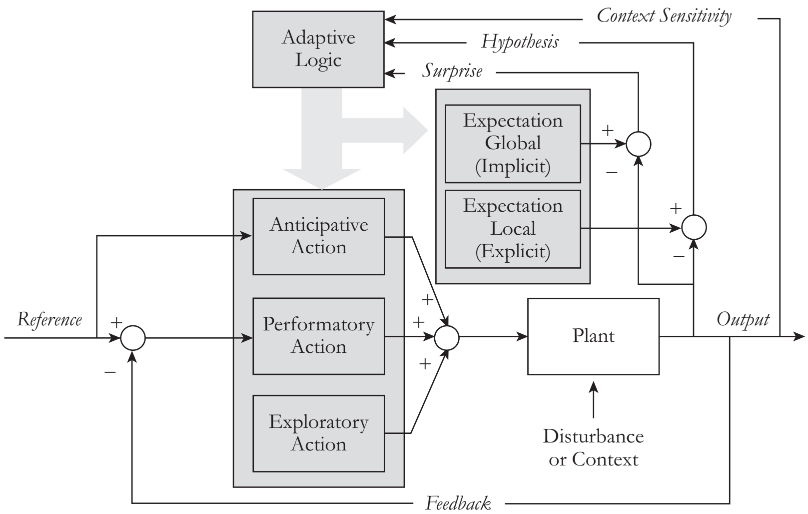

Figure 1.8 This diagram combines an inner (control) loop with outer (adaptive or metacognitive) loops to illustrate three styles that engineers use in the design of adaptive control systems.

altitude. Expectations based on the normative simulation can then be compared with the actual behavior of the vehicle, and deviations from expectation can be fed back to adapt the control strategy (and the expectations). Again, the engineer's "simulation" may be somewhat analogous to what cognitive scientists refer to as knowledge or a mental model. This "model" reflects integrated experiences from the past that provide a backdrop that experts can use to assess situations. In many cases, this "internal model" operates implicitly—so that experts may not become aware of the expectations until there is a mismatch. And even when the mismatch is noticed, the feedback may be experienced only as a vague sense that something is not normal. Thus, this might be the basis for intuitive aspects of expertise (Klein, 2004).

Overall, the outer loops in Figure 1.8 may provide a constructive way to think about the general phenomena of "metacognition," that is, self-awareness or our ability to monitor and critique our own performance. This is another layer of closed-loop, iterative processing in which the processes are simultaneously shaping their responses and being shaped by those responses. Again, this reflects the self-organizing aspect of cognitive systems. These are systems that are capable of learning from their mistakes.

Adaptive control systems where outer loops change parameters of inner loops are inherently nonlinear. The stability of these systems becomes a much more difficult analytic problem for control theory, particularly considering that all three outer loops may be acting simultaneously. These systems can be trapped in local minima (e.g., converge on a degenerate model of the plant; superstitious behavior), and they are vulnerable to butterfly effects (e.g., a small change can cascade, having dramatic effects with regards to stability).

Finally, it should be apparent that as we increase the number of variables that must be fed back and as we add outer loops that change inner loops, the complexity of the computations involved expands rapidly, thus raising questions with respect to the capacity of the human agent to manage this complexity. This issue will be addressed in the next section.

The main points of this section are:

1. To increase appreciation of the natural complexity of the comparator process whereby feedback is compared with intentions in order to specify corrective actions. The comparator process is typically represented as simple subtraction. This formalism greatly trivializes the problem faced by cognitive systems.
2. To reinforce the intimacy of observation and control in closed-loop systems. The closed-loop dynamic demands that questions of intention, perception, and action be framed in relation to each other. A program that studies these as independent stages in a linear causal stream will result in trivializing the dynamics of cognition.
3. To reinforce the fact that despite the common misconception, feedback does not insure stable convergence of output to the target input due to the complexity of multi-loop dynamics and nonstationary (i.e., changing over time) parameters.

Smart Mechanisms and Situation Awareness

There are two important consequences of having multiple degrees of freedom in terms of intentions, feedback information, and possible actions. On the positive side, multiple degrees of freedom mean that there are many ways to "skin a cat." That is, there will typically be multiple ways to satisfy an intention. This reflects a range of goal trade-offs that might be acceptable, multiple potentially redundant sources of information relative to guiding action, and multiple combinations of actions that move the system toward satisfactory goals. On the negative side, multiple degrees of freedom can greatly increase the computational demands on the system. That is, the computational load associated with the comparator problem will be a rapidly increasing function of the number of variables involved—due to the potential for multilevel interactions among these variables and across inner and outer control loops.

It is clear that the human is a limited bandwidth information processor. That is, there is a limit to how many independent things the system can attend to or be aware of at the same time (e.g., Broadbent, 1971; Miller, 1956). Thus, there will be limits to the complexity of the comparator problem that humans (and other biological systems) are able to solve. In general, for "good" or "smart" control, the system should prefer solutions that both satisfy the intentions and minimize the computational demands (i.e., demands on awareness).

Bernstein (1967) was one of the first to draw attention to the awareness or computational constraints on feedback control in the context of motor skills. He noted that from the many possible means for achieving a goal, skilled athletes tended to choose solutions that reduced the degrees of freedom that needed to be monitored (or controlled) in real time. This was typically accomplished by locking out or constraining other potential degrees of freedom. This is easily illustrated by considering the swing of a skilled golfer when driving a golf ball. The way that the golfer holds the club, addresses the golf ball, and moves during the swing is explicitly designed to lock out many degrees of freedom (e.g., keeping head fixed with eyes on the ball and keeping the elbow of the left arm straight for a right-handed golfer). The result of locking out or constraining many potential degrees of freedom is that the complexity of the comparator problem is greatly reduced. It reduces the number of variables that need to be monitored/controlled in real time (i.e., attended or adjusted) in order to control convergence to a satisfactory solution. For example, if the elbow position of the left arm were also changing during a swing, its interactions with all the other variables (e.g., wrist angle) would have to be monitored and compensated for in order to reliably make contact with the ball. The term "coordinative structures" has been suggested for the patterns of constraint seen in skilled motor activities (Turvey, 1977).

Runeson (1977) contrasted the solutions such as coordinative structures with some of the early conventional approaches to engineering and general problem-solving processes. The conventional approaches tended to start with a fixed general coordinate system (e.g., orthogonal three-dimensional spatial coordinates), and all control problems were then organized with respect to those dimensions (e.g., describing the motion of all arm components with respect to the positions and velocities in this orthogonal space). The result of using this fixed coordinate space is that many common natural motions required relatively complicated descriptions; thus, very complex computations were implied. To simplify computations, engineers typically choose solutions that are more compatible with the fixed coordinate system. This results in very unnatural, "robotic" motion. Runeson called the systems resulting from reliance on a single, fixed coordinate system "rote mechanisms," suggesting that these reflected brute, one-size-fits-all solutions.

Alternatively, Runeson suggested that biological systems change the coordinate system to reflect the intrinsic constraints on the degrees of freedom (i.e., the coordinative structures). In these cases, a coordinate system was chosen to make the description of the solutions simpler—in effect reducing the computational demands. Runeson called systems that organize the comparator processes around "intrinsic" rather than "extrinsic" coordinates "smart mechanisms."

The constructs of "coordinative structure" and "smart mechanism" have important implications for understanding the dynamics of cognitive systems. Much of the research on perception defines "space" with respect to a fixed coordinate system, where human performance is evaluated with respect to fixed rulers (e.g., length measured in cm) and fixed coordinate systems (e.g., height, length, depth perception). This classic approach is designed around open-loop tasks where a stimulus is presented (e.g., an object of a particular size and position in the world), and subjects are asked to make passive judgments (e.g., how big is the object? or how far away is it?). Great care is typically taken to insure that the subjects' responses do not change the stimulus (e.g., stabilizing the eyes using chin rests or bite bars).

The constructs of coordinative structure and smart mechanism suggest that the closed-loop system may self-organize around the intrinsic geometry of the problem (e.g., vehicle dynamics). This suggests that intrinsic aspects of the situation (e.g., imminence of collision relative to the braking dynamics) may be more "basic" to perception than the extrinsic coordinates used by physicists to describe space (e.g., depth perception). In other words, this suggests that "space" in the classical extrinsic sense is not the dominant coordinate frame for perception. That is, the perceptual system is not designed to "see" space, but rather to detect those properties most directly relevant to guiding action (e.g., the imminence of collision). It is designed to solve the comparator problem. Gibson (1979/1986) has argued for this hypothesis in terms of his construct of "direct perception of affordances." That is, a smart mechanism organizes itself so that the functionally relevant properties (i.e., the affordances) are specified in a simple or direct way to reduce the computational demands. This idea, which once seemed quite radical, becomes quite plausible when considered in the context of feedback control.

In the spirit of Gibson's approach, the challenge is not to understand "space perception" per se. Rather, the challenge is to understand the ability for experts to discover those coordinates that provide the most direct mapping across the dimensions of intention, feedback, and action. In other words, the challenge is to find the smartest solutions to the comparator problem (e.g., the simplest solution that satisfies the demands of the situation). While in some cases the solutions might be based on global constraints (suggested by the search for optical invariants), recent research suggests that even in relatively simple perceptual judgments, people are likely to utilize more local constraints of specific task environments when they are locally reliable (e.g., kinematic correlates to the specifying mass ratios for predicting the results of collisions) (Jacobs, Runeson, & Michaels, 2001). This suggests that even solutions to basic perception tasks may be more heuristic in nature than had previously been thought by those following a Gibsonian approach. That is, people will often organize behavior around local constraints, rather than always utilizing global invariants. However, Gibson's main point remains—the key to understanding the coupling between perception and action is to discover the intrinsic coordinates that allow direct mapping between intentions, consequences, and action.

The key point of this section is to reframe the challenge of cognition. In classical approaches the challenge of cognition is framed as the ability to construct a valid internal model of the world (that then becomes the basis for motor control and decision making). This "internal model" is typically judged relative to extrinsic physical models of space (e.g., standard rulers and coordinate systems). We suggest that the problem of cognition is simply to guide or direct successful action. That is, the function of cognition is to solve the comparator problem. This requires that our understanding of "awareness" (e.g., internal models) be grounded in an understanding of "situations" (e.g., Flach & Rasmussen, 2000; Flach, Mulder, & van Paassen, 2004). The "test" of a belief is not based on classical induction or deduction, but rather it is based on the consequences of actions based on that belief.

It is important to understand that this mapping between perception and action is not a trivial problem (e.g., consistently hitting a nail is not as trivial as suggested by Miller et al.'s TOTE example). It is very likely that as a result of searching for and discovering smart control mechanisms for the comparator problem, we build internal representations of the world that might generalize to more abstract models that guide scientific exploration and discovery and that create new ways for interaction. However, from the circular systems perspective, the abstract models of science are a product of the coupling of perception and action, not a necessary prerequisite for the coupling. We discover the world through acting in it.

Generalizing to Problem Solving and Decision Making

The relevance of the dynamics of closed-loop systems is easy to discover in the context of programs of research on perceptual-motor coordination (e.g., manual control). In that research context, there is an obvious mapping between the normative models (e.g., optimal control models and automated control systems) and the cognitive phenomena (e.g., piloting an aircraft), and there is a history of work utilizing the analytic tools of control theory (e.g., frequency domain analysis) to describe human performance (e.g., Sheridan & Ferrell, 1974; Jagacinski & Flach, 2003). However, do these principles have implications more generally for cognitive systems? Is the logic of circles also relevant to phenomena associated with decision making and problem solving?

We believe that the answer to this is obviously "Yes." This belief is supported by observations of naturalistic decision making that discovered that in many decision contexts there is an intimate link between "recognition" (i.e., perception) and "choice" (i.e., action) (Klein, 1993). For example, choices (decisions/actions) of a fire ground commander (e.g., where to direct the hoses) are primed by recognition processes or feedback (e.g., observations of the fire relative to the goal priorities—saving lives and property). It is also supported by work on problem solving in the wild (Hutchins, 1985). For example, Hutchins (1985) describes how the solutions to navigation problems are shaped by the available tools and representations (e.g., maps). Again, we believe that these analyses of "cognition in the wild" are unpacking the comparator problem. That is, these analyses are discussing the relations between input (e.g., intentions/goals), output (e.g., consequences, feedback), and action. Different strategies can then be evaluated relative to whether they lead to satisfaction of the goals (i.e., whether the system converges to stable solutions) and relative to the computational load (i.e., how smart are the computational mechanisms?). Finally, a big part of the story is how experts in natural work domains leverage intrinsic problem constraints to improve efficiency (i.e., coordinative structures).

Cognitive theory and research continues to be dominated by the logic of open-loop causal reasoning, despite the growing evidence that cognitive phenomena in nature are typically closed-loop. For example, human decision making continues to be evaluated against normative models that are based on open-loop logic (induction, deduction). However, the logic of abduction that was first proposed by Charles Sanders Peirce (1978) provides an alternative basis for rationality that is more consistent with the logic of circles. For Peirce, the test of a hypothesis was not the form of the argument (as suggested by classical logic), but rather, the test of the hypothesis was the pragmatic consequences of actions directed by that hypothesis.

The abductive process described by Peirce is directly analogous to the observer illustrated in Figure 1.3. That is, an abductive system is a system designed to eliminate surprise. In an abductive system, as in an observer, beliefs that lead to accurate perceptions and predictions about the world (i.e., converge in ways that reduce surprise) are strengthened. However, remember that this observation process is not independent from the process of control. Thus, the real test of our beliefs in an abductive system is whether the actions guided by these beliefs lead to satisfactory consequences (i.e., reduce error with respect to intentions). In essence, the targeting problem that challenged Wiener and Bigelow was a problem of abduction. The problem was to anticipate the actions of an aircraft in order to successfully target it. The ultimate test of any belief/guess about the behavior of the aircraft was the consequences of an action (either firing projectiles or simply comparing predicted and actual behavior) that were then fed back, compared, and integrated to both change the belief (prediction) and guide the next prediction and/or action.

Conclusion

The primary motive for this chapter is that we believe that cognition is a closed-loop phenomenon. Thus, the logic of circles is fundamental. We believe that many of the conundrums in cognitive science and the failures of that science to inform design are the result of framing the questions in the context of simple open-loop causal systems. Although Wiener's cybernetic hypothesis has helped to inspire the cognitive revolution, and although feedback loops are often included in the images we create to represent cognitive systems, we believe that for the most part the logic of circular systems has not been fully appreciated or applied. For the most part, cognitive theory and research continues to be framed in an open-loop context.

This is not simply a problem for cognitive science, but it seems to be integral to the more general logic of experimental science, where experiments are explicitly designed to explicate open-loop cause-effect relations. Note that in considering a circle as a whole, there is no unambiguous direction for distinguishing cause from effect. In a recursive circular dynamic, what we see determines what we do, while simultaneously what we do determines what we see. Observation and control are intimately coupled. There is no basis for one (either perception or action) to have causal precedence over the other.

There is a tendency for experimental science to organize around general, extrinsic coordinate systems that are at least implicitly accepted as ground truth. These truths become the ruler against which behavior is measured. For example, perception of size is judged against a standard measure in centimeters. Or choices are judged against the prescriptions of inductive or deductive logic. However, circular systems are self-organizing. This means that understanding will generally depend on the ability to discover the intrinsic coordinates (or standards) that are structuring that organization. Thus, the question is not "how big?" or "how far?" in any absolute sense (e.g., *X* cm), but rather whether it is graspable, or pass-through-able, or whether collision is imminent relative to the capacity to brake. In a circular system, the value of any variable (intention, action, or perception) must be "measured" relative to the other variables. The meaning of "too close" is contingent on the action dynamics (i.e., the capacity for braking or evasive maneuvering). The question is not whether the reasoning is "sound" in an absolute sense (e.g., valid relative to normative prescriptions), but rather does the reasoning lead to successful action?

In the classical approach, performance is gauged with respect to ideal norms derived from abstract mathematical models that were intended to generalize across a myriad of different situations. We believe such abstract norms can be very important for bounding the space of possible behaviors, but that they are often not the appropriate gauges for understanding the dynamics of actual behavior. Actual behavior is grounded in the pragmatics of situations, and while an artful application of mathematics may be essential for describing the resulting patterns and constraints, it needs to be particularized to distinct situations. Actual behavior of circular systems is situated. Closed-loop systems are adapting to their environments, and they are simultaneously adapting (i.e., changing) their environments. This coupling between environment and organism is reflected in the dynamics of prey-predator systems, where the size of each population is shaping and being shaped by the size of the other population. This coupling and the resulting self-organization are also nicely illustrated by Kugler and Turvey's (1987) metaphor of insect nest building. In this system, insect behavior creates the pheromone landscape at the same time the pheromone landscape is shaping insect behavior. The resulting structure of the insect nest is a product of this coupling.

In this chapter, our goal was to use a few very simple examples to help people to appreciate the dynamics of closed-loop systems. These examples were chosen to highlight differences between the expectations developed from an open-loop theory of behavior and to increase appreciation for the complexity of closed-loop systems. We intentionally chose very simple closed-loop systems so that the examples would be understandable and to minimize any need for complex analysis. However, it is important to realize that we are not suggesting these examples as specific "models" or "metaphors" of cognition, although we do believe that intuitions developed through understanding these simple circular dynamics may be important in shaping a theory of cognition that captures the dynamics of particular situations.

We believe that the evidence for circularity in natural cognitive phenomena is pervasive. Although one can find local behaviors that are open-loop (e.g., ballistic braking), even in these cases the local actions will typically be components of a larger circular dynamic (e.g., safe driving). However, the point here is not to fully elucidate cognition, but to elucidate the logic of circular systems. Our hope is that a deeper appreciation and understanding of closed-loop systems will enrich theory, experimental practice, and application of cognitive science. We hope that a deeper appreciation of closed-loop systems will lead cognitive scientists to ask better questions.

Finally, we believe that posing a question correctly will take us much further toward finding satisfactory answers to that question. We are very optimistic that trends in cognitive science (e.g., associated with neural nets, dynamical systems, artificial life, etc.) and in cognitive engineering (e.g., associated with ecological interfaces and semantic computing) suggest that appreciation for the circularity inherent in cognitive phenomena is growing. However, unfortunately, this appreciation is hampered by pervasive temptations to trivialize the dynamics of circles in order to satisfy conventional assumptions about the open-loop nature of explanation and the open-loop nature of experimental inference. Although most programs in cognitive science require students to take experimental methods courses (e.g., to learn analysis of variance and regression), we are aware of no academic program in cognitive science that requires a course to introduce the dynamics of closed-loop systems.

Future Directions

We suggest three challenges for the future associated with theory, methodology, and practice. On the theoretical side, it is important to move beyond trivial control metaphors to more carefully consider the dynamics of closed-loop systems. There are two levels of theory to consider. At the meta-level, we need to give up the domino theory of open-loop causality to consider general models of complex, dynamical systems. At the base theoretical level, we need to attend to the situated (e.g., Suchman, 1987) or embodied (e.g., Clark, 1997) dynamics of cognition. We need to move beyond an exclusive focus on formal logic and other context-free models of rationality to consider the pragmatic, ecological

AQ: Provide reference forClark 1997.

constraints of perception and action (e.g., abduction) (e.g., Todd & Gigerenzer, 2003).

Methodologically, it is important to broaden the empirical base. This includes attending more to naturalistic observations of cognition in complex work domains (e.g., Hutchins, 1985). It also means considering more representative designs for controlled experimentation (e.g., Kirlik, 2006). The representative designs need to provide more degrees of freedom to the participants with respect to both the means and ends for performance. If the laboratory tasks are designed around the simple servomechanism metaphor, then of course the humans will adapt and behave like a simple servomechanism. If we want to explore the creative, adaptive capacity of human cognition, then we have to create laboratory situations that invite creativity and adaptation. Synthetic task environments provide a unique opportunity here (Flach, Schwartz, Courtice Behymer, & Shebilske, 2010). In these environments it is possible to empirically link performance at the micro-level (e.g., reaction time to specific events) to functional ends at a macro-level (e.g., success in domain terms, such as successfully completing a mission). The capacity to empirically link variations at the micro-performance level with more global functional consequences will be critical to modeling self-organizing dynamics.

Finally, at a practical level, it becomes necessary for enhanced collaborations between those who focus on awareness (e.g., cognitive scientists) and those who focus on situations (e.g., engineers, domain experts). Without cooperation and mutual respect across disciplines, it will never be possible to achieve a deep understanding of situation awareness. This will require that we escape from the view where basic and applied sciences are seen as competitors in a zero sum game. We need to embrace the spirit of Pasteur's quadrant, where the goals of theory and practical utility are respected and valued equally as complementary components of a mature science (Stokes, 1997).

References

Bernstein, N. A. (1967). *The control and regulation of movements.* London, England: Pergamon Press.

Broadbent, D. A. (1971). *Decision and stress.* New York, NY: Academic Press.

Conway, F., & Siegelman, J. (2005). *Dark hero of the information age: In search of Norbert Wiener, the father of cybernetics.* New York, NY: Basic Books.

Ericsson, K. A., & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist*, *49*(8), 725–747.

Flach, J., Mulder, M., & van Paassen, M. M. (2004). The concept of the "situation" in psychology. In S. Banbury and S. Tremblay (Eds.), *A cognitive approach to situation awareness: Theory, measurement, and application* (pp. 42–60). Aldershot, England: Ashgate.

Flach, J. M., & Rasmussen, J. (2000). Cognitive engineering: Designing for situation awareness. In N. Sarter & R. Amalberti (Eds.), *Cognitive engineering in the aviation domain* (pp. 153–179). Mahwah, NJ: Erlbaum.

Flach, J. M., Schwartz, D., Courtice A., Behymer, K., & Shebilske, W. (2010). Synthetic task environments: Measuring macrocognition. In E. Patterson & J. Miller (Eds.), *Macrocognition metrics and scenarios: Design and evaluation for real-world teams* (pp. 201–284). Aldershot, England: Ashgate.

Flach, J. M., Smith, M. R. H., Stanard, T., & Dittman, S. M. (2004). Collisions: Getting them under control. In H. Hecht & G. J. P. Savelsbergh (Eds.), *Time to contact, Advances in Psychology* (Vol. 135, pp. 67–91). North-Holland, Amsterdam: Elsevier.

Gibson, E. J. (1969). *Principles of perceptual learning and development.* New York, NY: Appleton-Croft.

Gibson, J. J. (1979/1986). *The ecological approach to visual perception.* Hillsdale, NJ: Erlbaum.

Hutchins, E. (1985). *Cognition in the wild.* Cambridge, MA: MIT Press.

Jacobs, D. M., Runeson, S., & Michaels, C. F. (2001). Learning to visually perceive the relative mass of colliding balls in globally and locally constrained task ecologies. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(5), 1019–1038.

Jagacinski, R. J. (1977). A qualitative look at control theory as a style of describing behavior. *Human Factors*, *19*, 331–347.

Jagacinski, R. J., & Flach, J. M. (2003). *Control theory for humans: Quantitative approaches to modeling performance.* Mahwah, NJ: Erlbaum.

Kirlik, A. (2006). *Adaptive perspectives on human-technology interaction.* New York, NY: Oxford University Press.

Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein , J. Orasanu , R. Calderwood, & C. E. Zsambok (Eds.), *Decision making in action: Models and methods* (pp. 138–147). Norwood, CT: Ablex.

Klein, G. A. (2004). *The power of intuition.* New York, NY: Random House.

Kleinman, D. L., Baron, S., & Levison, W. H. (1971). A control theoretic approach to manned-vehicle systems analysis. *IEEE Transaction on Automatic Control*, *AC-16*, 824–832.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement.* Hillsdale, NJ: Erlbaum.

Lee, D. N. (1976). A theory of visual control of braking based on information about time-to-collision. *Perception*, *5*, 437–459.

McRuer, D. T., & Jex, H. R. (1967). A review of quasi-linear pilot models. *IEEE Transactions on Human Factors in Electronics*, *HFE-8*, 231–249.

Miller, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.

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

Peirce, C. S. (1978, August). Deduction, induction, and hypothesis. *Popular Science Monthly*, *13*, 470–482.

Pew, R. W., & Baron, S. (1978). The components of an information processing theory of skilled performance based on the optimal control perspective. In G. E. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 71–78). New York, NY: Academic Press.

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

Runeson, S. (1977). On the possibility of "smart" perceptual mechanisms. *Scandinavian Journal of Psychology, 18*(1), 172–179.

Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems: Information, control, and decision models of human performance.* Cambridge, MA: MIT Press.

Smith, M. R. H., Flach, J. M., Dittman, S. M., & Stanard, T. W. (2001). Alternative optical bases for controlling collisions. *Journal of Experimental Psychology: Human Perception and Performance*, *27*, 295–410.

Stokes, D. E. (1997). *Pasteur's quadrant: Basic science and technological innovation.* Washington, DC: Brookings Institutions Press.

Suchman, L. A. (1987). *Plans and situated action.* Cambridge, England: Cambridge University Press.

Todd, P., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology, 24*, 143–165.

Turvey, M. T. (1977). Preliminaries to a theory of action with reference to vision. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology*. Hillsdale, NJ: Erlbaum.

Weinberg, G. M., & Weinberg, D. (1979). *On the design of stable systems.* New York, NY: Wiley.

Wiener, N. (1948/1961). *Cybernetics: Or control and communication in the animal and the machine.* Cambridge, MA: MIT Press.

AQ: Please provide page range for Turvey 1977.