3

DIRECT ORIENTATION AND FEEDBACK

Kinesic (indirect) orientation works by means of successive comparisons. Animals orient in this way either when the gradient to be detected is too shallow to permit simultaneous comparison, or when the source of stimulation cannot be sensed at a distance because the direction of local gradients is not perfectly related to the direction of the source. Visual and auditory stimuli allow for a more efficient strategy: The steepest local gradient always points to the source, and gradients are usually quite steep. Thus, simultaneous comparisons can provide immediate feedback about the proper direction of movement, allowing direct orientation with no wasted motion.

In this chapter I discuss simple mechanisms of direct orientation, first in a descriptive way, and then in terms of feedback mechanisms. The discussion of feedback serves several purposes: It shows how feedback is involved in orientation, and how a mathematical model can account for experimental results. It shows how apparently different orientation reactions can be considered as variants of the same process. It illustrates the difference between *static* and *dynamic* behavior theories. The remainder of the chapter discusses the type of explanation given by feedback accounts and shows how the concept of feedback control provides a mechanistic explanation for motive and purpose.

TAXES

Fraenkel and Gunn classify direct orientation (taxic) reactions into four main types. These do not, in fact, correspond to completely different mechanisms, but the mechanisms will be easier to understand after I describe the different types and the experiments used to tell one from the other

Klinotaxis

The first taxic reaction, *klinotaxis*, is really an intermediate case, since it involves both direct orientation and successive comparisons. Figure 3.1 shows an example of the relevant behavior: Mast tested each of four maggots (*Lucilia sericata*) four to six times for its orientation to a beam of horizontal light. Movement was always directly away from the light, although the tracks are not perfectly smooth. Substantial head movements (wiggles in the record) are apparent, particularly in the track of maggot A. These head movements provide a clue to the underlying process. Fraenkel and Gunn write: “During steady forward crawling the head is sometimes put down symmetrically in line with the axis of the body, but from time to time it comes down alternately to the right and to the left. When the maggot first begins to crawl, these lateral deviations of the head are usually considerable, and may even result in the body assuming a U-shape” (1940, p. 60). The animals look as if they are searching in some way for the direction of the light. Once they find it, they move in the opposite direction.

The animals find the direction of the light by successive comparison. This can be shown by tricking the animal: turning on an overhead light when the animal is turning in one direction and turning it off when it faces the other direction. Under these conditions a negatively phototactic animal moves in the direction associated with the low light level.

This process is similar to so-called trial-and-error learning, but is much simpler than the learning of higher organisms. No long-term memory is involved; the maggot does not remember where the light was or anything of the sort. The maggot has fewer modes of reaction than an animal such as a rat, and can, therefore, be led into maladaptive behavior by relatively simple means. For example, the animals may move away from a directional source even if it means moving into a region of higher, but non-directional (e.g., overhead) illumination. They can also be induced to move up (rather than down) a gradient of illumination by turning on a second light only when the animal is facing down the gradient. In the absence of a long-term memory, the animal has no way of discovering that its situation, while improving on a moment-by-moment basis, is getting worse overall.[1]

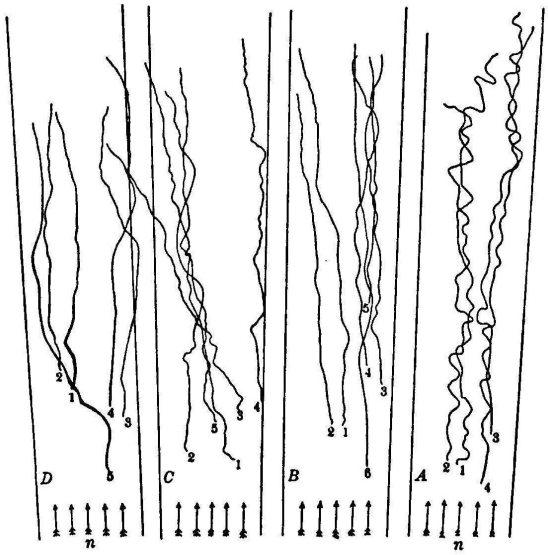

Figure 3.1. Tracks of four maggots (A to D) in a beam of horizontal light (arrows.) Each maggot was tested 4-6 times. (From Mast, 1911.)

Lest we begin to feel arrogant, I should point out that very similar tricks can be played on the human visual system. Figure 3.2 shows at the top a sawtooth gradient of illumination: from left to right, the illumination repeatedly grows slowly dimmer then rapidly brighter. The bottom part of the figure shows what this pattern looks like to the eye. Because of inhibitory, rate-sensitive mechanisms, the rapid increases in illumination have a greater perceptual effect than the slow decreases. Consequently what is seen is a staircase pattern of progressively increasing brightness. Edge-sensitivity is a very general property of perceptual systems — even honeybees show it (Davey, Srinivasan & Maddess, 1998).

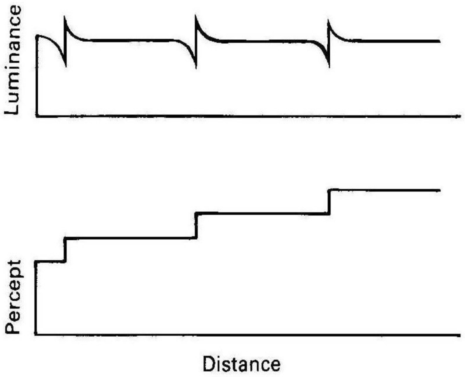

Figure 3.2. *Top*: gradient of luminance across a striped array. *Bottom*: perceptual effect of this luminance disturbance - a staircase pattern of progressively increasing brightness.

These effects arise because sensory systems, and the orientation mechanisms that depend on them, are usually more sensitive to the *rate of change* of the relevant stimulus than to its absolute value: The important variables are the *higher derivatives* if the stimulus value, with respect to space and time, rather than the stimulus value itself.[2] *Sensory adaptation* is one of many labels for this sensitivity rate of change, since it refers to the waning effect of a stimulus with time following its onset. This mechanism is enormously useful to animals, and represents in most cases a great improvement over dependence on absolute stimulus values. As we have seen, however, in the absence of other, longer-term memory processes it can sometimes lead to maladaptive behavior.

The details of the mechanism of klinotaxis, in *Lucilia* or any of the other species that show the pattern (e.g., the protozoan *Euglena*, earthworms, the ascidian *Amaroucium*), have not been fully worked out. It is likely that the behavior of head-swinging (which also guides the animal) occurs only when the rate of change of illumination of the receptor(s) exceeds some threshold value. In a stable, uniform light field, the animal moves little and swings its head from side to side in a relatively narrow arc; but if the light falling on the receptors increases greatly (which might be because the animal has entered a lighter region, or because the ambient illumination has changed) a large head excursion occurs, followed by a smaller return swing. If any of these subsequent swings again produces an increase in receptor illumination, further swings are produced. As time elapses, the effect of increases in receptor illumination diminishes because of adaptation, so that successive head swings are likely to be of smaller and smaller amplitude. If the source of illumination is fixed, the animal will, in the meantime, have become oriented away from it, so that tracks similar to those in Figure 3.1 result. If the source of illumination is moving, then obviously the animal's ability to orient away from it will depend critically on quantitative details involving its rate of movement, of head swinging, and of sensory adaptation. I return to quantitative questions in the discussion of feedback later in the chapter.

The remaining taxic orientation reactions, tropo- and telotaxis and the light-compass reaction, all involve simultaneous comparison of the stimulation of two or more bilaterally symmetrical receptors. One clue to the type of reaction is provided by the track of an animal orienting in a uniform gradient: The track is convoluted if the mechanism is a kinesis, wavy for klinotaxis, but straight for the taxes that use simultaneous comparison. Tropotaxis and the light-compass reaction differ chiefly in their quantitative properties.

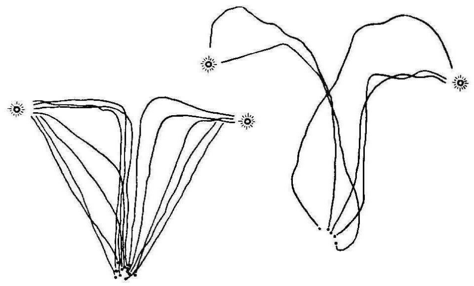

Figure 3.3. Tracks of photo-positive *Armadillium* toward two equal lights. (After Müller, 1925.)

Tropotaxis

The two-light experiment, illustrated in Figure 3.3, is critical to the identification of tropotaxis. The figure shows the tracks of several pill-bugs (*Armadillium* sp.), a terrestrial crustacean that lives under rocks and decaying wood, placed some distance from a pair of equal lights. *Armadillium* behaves in a positive phototactic fashion in this experiment: It approaches the two lights. But when placed equidistant between the two lights the animal will often follow a path in between them, rather than going straight to one or other of the two.

Tropotaxis is the outcome of a balancing process: The animal turns until the two eyes are stimulated equally and then proceeds forward. As the animal approaches a point directly in between the two lights, it will go directly to one or other of them, as the tracks show. At that point, a head swing that faces the animal directly toward one of the lights will also eliminate the influence of the other one.

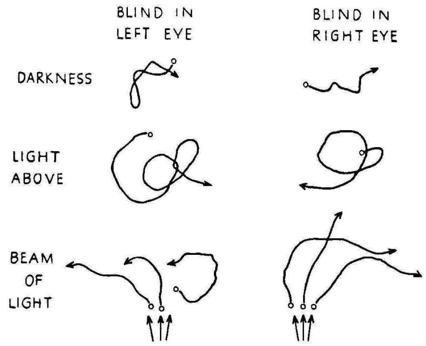

Figure 3.4. Tracks of photo-negative *Ephestia* larvae, blinded on one side, under three illumination conditions. (After Brandt, 1934.)

A second means of identifying tropotaxis is to eliminate the stimulation on one side of the animal, either by blinding it in one eye or painting over the eye if this is possible. The result is *circus movements*: In a uniform light field, such as an overhead light, elimination of one eye means that stimulation appears to come from only one side. If the animal is positively phototactic it then turns continuously toward the illuminated side; if it is negatively phototactic, it turns continuously toward the blinded side. In the normal animal with an eye-level light source, these reactions would lead eventually to equal stimulation of both eyes and leave the animal facing either away from, or toward, the light source. Figure 3.4 shows an example of circus movements produced in unilaterally blinded *Ephestia* larvae.

Light-compass reaction

The third taxis in Fraenkel and Gunn's basic list is the *light-compass reaction*. This pattern is very common in animal orientation. For example, bees use it in returning to the hive, and ants, although often able to rely on odor trails, can also orient with respect to the sun, finding their way back to the nest by keeping a fixed angle between their path and the direction of the sun. Moreover, many species are able to compensate for the movement of the sun by a change in this light-compass angle, so that a fixed direction is maintained (this is sometimes termed *sun-compass orientation*, and it appears to be one of the main components in the feats of navigation performed by migrating birds). This compensation depends on an "internal clock," tied to the 24-hour cycle (a *circadian rhythm*). I show in a moment that light-compass orientation can be derived from the same feedback mechanism as tropotaxis

Telotaxis

In tropotaxis, animals usually orient between two equal, symmetrically disposed lights, and only directly approach one of them when the angle subtended by the lights is large (see Figure 3.3). In telotaxis, the animals generally head straight for one of the two lights (although they may switch from one to the other as they approach) even if the lights are quite far away. Fraenkel and Gunn give a graphic, natural-history description of *telotaxis* as follows:

> Large numbers [of the little mysid crustacean, *Hemimysis lamornei*] are to be found in the aquarium tanks of the Marine Biological Stations at Plymouth and Naples. When a single light is placed at the side of a glass tank containing *Hemimysis*, the animals swim to and fro continually, always keeping in line with the beam of the light. They swim about 10cm towards the lamp, then turn sharply through 1800 and cover about 10 cm again before turning back towards the lamp, and soon. If an additional light is arranged so that the two beams cross at right angles, some of the mysids are quite unaffected in their behaviour while others switch over to this second light and behave as if the first one were nonexistent. The result is that the mysids form two streams which, so to speak, flow through one another, crossing at right angles and not interfering with one another. (1940, p.90)

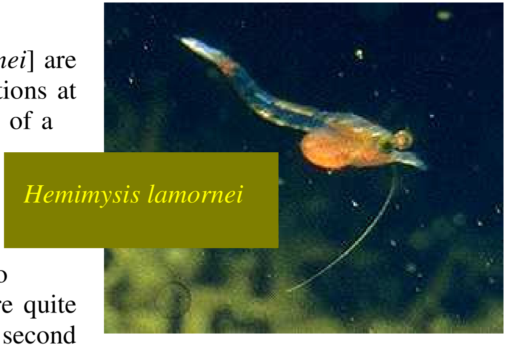

Hemimysis lamornei

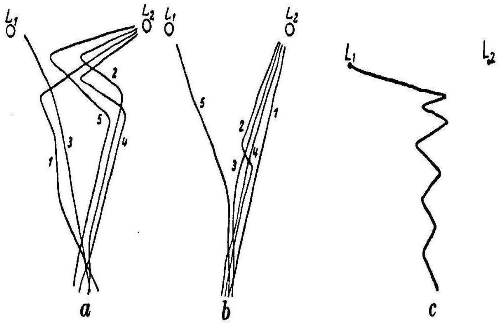

Figure 3.5. Tracks of hermit crabs (a) and (b) and an isopod (c) in a two-light experiment. Each part of the track is directed toward one light only. (a and b after Buddenbrock, 1922; c, Fraenkel, 1931.)

Examples of telotactic tracks of hermit crabs and an isopod (*Aega*) in a two-light experiment are shown in Figure 3.5

The mechanisms involved in telotaxis depend on the type of receptor the animal possesses. If the eye is capable of forming a rudimentary image, that is, of immediately identifying the bearing of a source of stimulation, then the animal has the necessary information available on the receptor surface. The animal can orient correctly by placing its walking apparatus under the control of one image exclusively. This is not a trivial problem, since the image-identification system must compensate for changes in retinal position caused by the animal's own movement. Both the lensed eyes of vertebrates and the compound eyes of insects and other arthropods provide the necessary directional information.

Many animals with image-forming eyes prefer to orient to a light using the balance mechanism of tropotaxis — shown both by orientation between two lights, and by circus movements when unilaterally blinded — even though the same animal may respond telotactically under other conditions. For example, at first a unilaterally blinded bee will show circus movements, but after a while they cease and the animal will head straight toward a light source — a telotactic response. The preference for tropotaxic orientation may be a reflection of the complex computations required for telotaxis.

Animals show many other simple orientation reactions. For example, so-called skototaxis (approach to dark areas or objects); a variety of reactions to temperature gradients (usually favoring a particular zone of temperature, either because it is appropriate for the organism's metabolism, or because it is a signal of the animal's prey, as when ticks prefer the skin temperature of their host); postural reactions to light and gravity — geotaxis and the dorsal light reaction; reactions to physical contact — many small and vulnerable animals seek out and conform to crevices and comers; reactions to fluid flow (rheotaxis) — fish generally face upstream, flying insects orient into the wind, some tidal crustaceans orient using cues from wave currents; reactions to chemical and humidity gradients such as those discussed in Chapter 2; orientation in sound fields (sound localization). Similar principles — simultaneous or successive comparison, feedback, time-dependent effects — are involved in all these reactions. Despite the relative simplicity of each reaction analyzed in isolation, in combination they can lead to quite complicated, "intelligent" behavior. More on this later.

FEEDBACK ANALYSIS

Tropotaxis depends upon a balance between stimulation of receptors on opposite sides of the body. The animal turns until stimulation of both receptors is equal and then approaches (positive phototaxis) or moves away from (negative photo-taxis) the source of illumination. The process can be represented as a simple feedback mechanism, as shown in Figure 3.6. The relations between organism and environment are here represented as a *feedback loop*, which shows the mutual relations between some aspect of the environment and some aspect of behavior that is linked to it.

The picture in Figure 3.6 applies to many things of interest to psychologists and ethologists. It will be helpful to have standard terms for its elements. Variable *x*, the aspect of the environment to which the system is sensitive (the *input*) corresponds to different things in different experimental situations. Here it is simply the receptor disparity (measured in illuminance units) associated with a given orientation of the animal in relation to the light source. The *output* or *response*, *y*, is the aspect of the organism's behavior that affects, and is affected by, feedback; here it is just the angle of the animal in relation to "north" (i.e., a standard direction), or in relation to the source, if the source is fixed. The two boxes labeled "organism" and "environment" in Figure 3.6 contain two functions, *O(x)* and *E(y)*, that describe how the input (*x*: receptor disparity) is related to the output (*y*: direction) by the animal (function *O(x)*), and how the output is related to the input by the environment (function *E(y)*). In most experiments *E(y)*, called a *feedback function*, is known, and the objective is to discover the *control function*, *O(x)*, imposed by the organism.

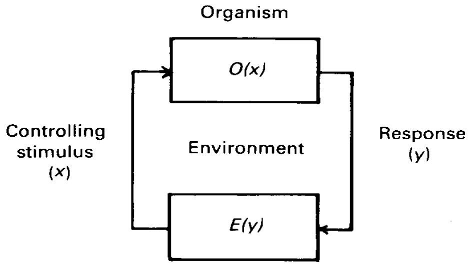

Figure 3.6. Organism-environment relations in adaptive behavior.

In tropotaxis, the feedback function, *E(y),* depends on the shape of the bilateral receptors, their sensitivity over their surface, and the geometry of the situation. In the simplest case, the receptors can be considered as points or spheres, and the stimulation falling on each receptor will, therefore, be determined solely by its distance from a source of illumination, according to the inverse square law. If the distance from the source is great, it is easy to show that the difference in stimulation of two bilateral receptors, y, is directly proportional to their separation, and to the sine of the angle between the midline of the animal and a line drawn from the source to the center of the animal (i.e., the animal's *heading* with respect to the source).[3]

Consider now how the control function, $O(x)$, might work from moment to moment. The animal is presented with a certain disparity of stimulation, x, which will, in turn, be a function of its heading with respect to the fixed light source (call that angle θ). Ignoring adaptation for the moment, the disparity must have an effect on the animal's tendency to turn. Perhaps the simplest possibility is just that the animal's *rate of turning* (measured in angular units per second) is proportional to the receptor disparity, with a sign such that the turns are toward the more illuminated side (for positive phototaxis). This assumption can be expressed formally by saying that the *first derivative* of the heading angle, $d\theta/dt$, is proportional to the receptor disparity. This is termed *integral control*, because the controlled variable, θ, is determined by the time integral of the controlling variable, x. We thus arrive at a simple formal model of this situation:

$$x = A\sin\theta, \tag{3.1}$$

giving the receptor disparity as a function of heading angle (the feedback function), and

$$y = d\theta/dt = -Bx, \tag{3.2}$$

giving the relation between rate of turning and disparity (the control function), where A and B are constants (A is proportional to receptor separation, and incorporates a scale factor; B represents the *gain* of the system: how fast the animal turns for how much disparity).

Two aspects of this little model are of interest. One is the *steady-state* solution: Assuming the light source is fixed, what will the animal's final heading be? A second is the dynamics of the situation: How does the animal's heading change with time?

The static problem is simple. When things have settled down, we know that the animal will have ceased turning; that is, $d\theta/dt$ will equal zero, which means that $x = 0$ (from Equation 3.2), which means that $\theta = 0$ (from Equation 3.1). Hence the animal will eventually point toward the light source. The dynamic solution is a little more difficult to arrive at, and requires that we know the *initial conditions* (i.e., the animal's initial heading) and, if the source is moving, the way in which it is moving as a function of time. I return to the dynamic problem in a moment

This model illustrates the distinction between closed- and open-loop control. The stability of the system is ensured under normal conditions by the immediate effect of the response, turning, on the feedback input, receptor disparity. But suppose receptor disparity were independent of turning (this is known as "opening the loop")? This could be done by blinding the animal unilaterally and leaving it in uniform light. Under these conditions, $x = C$ (constant). You can see at once from Equation 3.2 that $d\theta / dt$ must then equal BC, which is also a constant. Thus, the animal must turn constantly. This is, of course, exactly what is observed, in the form of the "circus movements" already described. In a slightly more elaborate experiment, fixed amounts of light might be presented to each eye, allowing $d\theta/dt$ (turning rate) to be measured as a function of C (receptor disparity). The slope of the resulting line then gives the value of parameter B, the gain of the system. Opening the loop is obviously a useful technique for getting at the properties of feedback mechanisms.

The virtue for the animal of a negative-feedback mechanism is that it minimizes the effect of changes in the control function, $O(x)$. For example, suppose that because of age or injury the animal is no longer able to turn as rapidly as before. In terms of our model this change might be represented by a decrease in the value of parameter B in Equation 3.2. The steady-state solution is unaffected by this change: Even with impaired movement, the animal must still eventually orient toward a fixed source of illumination. This is just what is observed when the locomotor apparatus of a tropotactic animal is surgically interfered with (e.g., by immobilizing or amputating legs):

The response may be slowed, but the animal continues to orient appropriately.

The light-compass reaction seems more complicated than the mechanisms discussed so far, but it need not be. Consider, for example, a simple modification of Equation 3.2, our illustrative model for tropotaxis: Instead of $d\theta/dt = -Bx$, where x represents receptor disparity and θ the heading angle, suppose we add a constant, so that

$$d\theta/dt = -Bx - C \qquad (3.3)$$

When $d\theta/dt = 0$, $x = -C/B$, rather than 0, the previous steady-state solution. Consequently, the resting value of θ, $\hat{\theta}$, from Equation 3.1, is $\hat{\theta} = \sin^{-1}(-C/AB)$, where A, B, and C are constants. Thus, the system represented by Equation 3.3 maintains a constant angle to a light source (light-compass orientation), rather than going directly toward it as in tropotaxis.

The dependence of light-compass angle on time (sun-compass orientation) can be incorporated simply by making C a function of time, so that the model becomes:

$$d\theta/dt = -Bx - C(t), \qquad (3.4)$$

where $C(t)$ is a function with a 24-hour period chosen so that light-compass angle shows the appropriate temporal shift.

Thus, feedback analysis shows that three apparently different orientation reactions can be considered just as variants on the same integral-control system.

Dynamic analysis

Equations 3.1 and 3.2 describe rate of turning as a function of time; they constitute a *dynamic model* for behavior. These equations illustrate two properties of dynamic models: First, such models generally express the effect of an independent variable on the *rate of change* in behavior, that is, as a differential or difference equation. In the tropotaxis example, equation 3.2 shows how a given interocular disparity changes the animal's rate of change of direction. But we are generally interested not so much in the change in behavior as in its actual value: We want to know where the animal is heading at a given time, not its rate of turning. The differential equations must be solved before we can make predictions about behavior as a function of time.

The second feature of dynamic models is that they have a steady-state solution that can often be deduced without having to solve the differential equations of the model. Much can be learned from the static analysis: Does the animal orient to the light? Does it show circus movements or not? Where does it come to rest? Most of the theories described in this book are static theories. Nevertheless, static theories omit important system properties. They give few clues to the *stability* of the system, for example. An organism that under frequently encountered conditions shows uncontrollable oscillations, such as the tremor of Parkinsonism or the continual turning of circus movements, is at a severe evolutionary disadvantage. Organisms must be stable in response to normal perturbations. Identification of the necessary and sufficient conditions for stability is one of the main objectives of the dynamic analysis of feedback systems.[4]

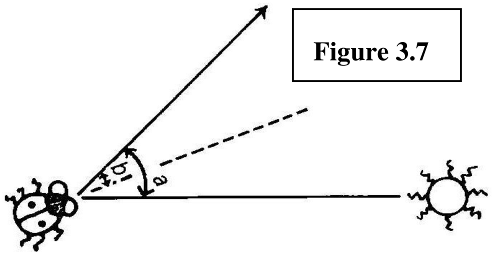

Figure 3.7

Solving Equations 3.1 and 3.2 is actually quite simple and can serve as a model for much more complicated analyses. The solution requires some familiarity with elementary calculus; readers lacking that background can skip the math and get on to the conclusions, which are all that is needed for later chapters.

First, it will be convenient to measure the bearing of the light source (angle a in Figure 3.7) and the heading of the animal (angle b) with respect to a fixed reference direction. Rewriting Equation 3.2 with these new designations yields:

$$d(a-b)/dt = -Bx, \qquad (3.5)$$

and Equation 3.1 becomes

$$x = A\sin(a - b). \qquad (3.6)$$

It will greatly simplify later discussion if we consider only small values of $a - b$ (i.e., headings close to the correct one) for which $\sin(a - b) = a - b$. Equation 3.6 then becomes

$$x = A(a - b). \qquad (3.7)$$

We wish to solve Equations 3.5 and 3.7 so as to obtain the change in the animal's heading angle, b, as a function of time, given that it starts from some initial heading b_0. If the source is stationary (a = constant), this is easily done by substituting for x in Equation 3.5 and integrating: $d(a - b)/dt = -AB(a - b)$. Since a is a constant, this reduces to $-db/dt = -AB(a - b)$, or

$$dt = db/AB(a - b). \qquad (3.8)$$

Integrating the left hand side of Equation 3.8 from 0 to the presenttime, t, and the right hand side from the initial heading, b_0 to the present heading, b, yields:

$$\int_o^t dt = (1/AB)\int_{b0}^{b} db/(a-b).$$

Evaluating the integrals yields:

$$ABt = {}_{b0}^{\;b}\left[-\ln(a-b)\right].$$

putting in the limits of integration and exponentiating yields:

$$\exp(-ABt) = (a - b)/(a - b_0), \qquad (3.9)$$

which is the desired function of time. Since a and b_0 are both constants, this reduces to an expression of the form:

$$a - b = \theta = \theta_0 \exp(-ABt), \qquad (3.10)$$

where $\theta = a - b_0$, which is the familiar exponential decay function illustrated in Figure 3.8.

Notice that the exponential function in Equation 3.10 has the properties it should: At $t = 0$ the bearing of the animal is θ_0; as $t \to \infty$, θ approaches 0; and the rate of change of θ with respect to time (the slope of the curve) is at first large, when θ is large, but slows down progressively as $\theta \to 0$. The rate at which the angle θ is reduced is determined by the quantity AB; the time that it takes for θ to be reduced to 0.368 (= $1/e$) of θ_0 is equal to $1/AB$, which is termed the *time constant* of the system: the smaller AB, the more slowly θ is reduced.

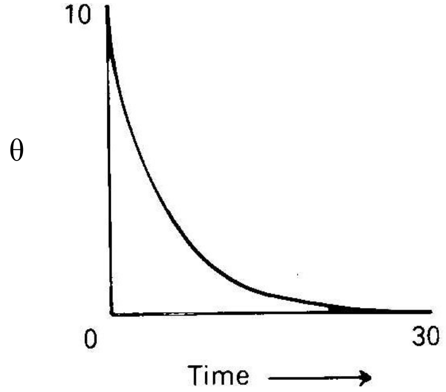

Figure 3.8. Output variable (θ: heading of tropotactic animal) of a simple integral-controlled servo-system following a "step" input (i.e., system starts with a fixed displacement between initial and desired orientation). Curve is a plot of Equation 3.10 in the text with $\theta_0 = 10°$ and $AB = 5$.

Frequency analysis

So far I have assumed that the light source is fixed. This may be realistic for light sources, but it certainly isn't true of prey items, for example, for the tracking of which predators possess comparable feedback mechanisms. What if the controlling (feedback) input varies with time?

Before looking at more theory, let's consider the kind of experimental results we want to explain. We will probably be interested in the limits of the system: Can the animal track any moving target, or just slowly moving targets? If so, how slowly must they move? If it fails to track, how does the behavior break down? Does the animal just stop tracking completely (remain motionless); does it track *out of phase* (i.e., look left when the light is on the right and vice versa); or does it begin to oscillate wildly from side to side? Does the tracking break down gradually or all at once?

To answer these questions we need to present the animal with a target that moves from side to side at different frequencies and in different time patterns. For each input time pattern ($x(t)$ in Figure 3.6) we will obtain a comparable output pattern of angular variation ($y(t)$). A useful explanation should summarize the results of all possible such experiments — give us $y(t)$ for any $x(t)$.

The theoretical problem can be handled in two main ways. One is simply an extension of the analysis already described: The constant term, a (the bearing of the light source), can be replaced by a time function, $a(t)$, and the resulting differential equation solved directly for particular forms of $a(t)$. This is tedious and may be difficult in particular cases. The second way is simpler, providing the system is linear, that is, describable by a linear differential equation (I explain the implications of this in a moment).

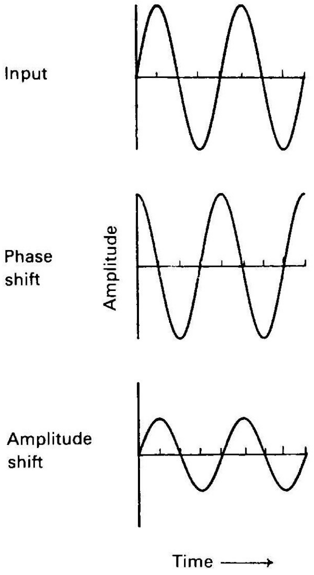

Figure 3.9. A sinusoidal input (top panel) shifted in phase only or amplitude only (bottom panel). In real control systems amplitude change is almost always accompanied by a phase shift.

For those with an interest in the mathematical basis for the ensuing arguments, I present a summary of linear-systems analysis in the Notes to this chapter.[5] The basic notions can be explained without mathematics in the following way. Problem: To describe the input-output behavior of a system to all possible time-varying inputs. How can this be done without having to list every possible input and its output? Suppose that the following three conditions are met:

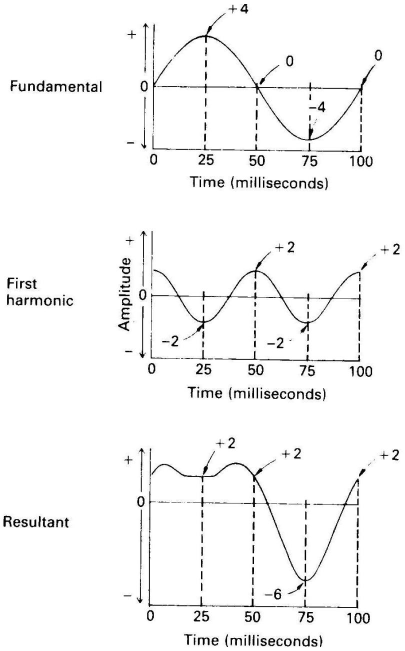

Figure 3.10. Graphic illustration of Fourier analysis of a period signal (bottom panel) decomposable into two sinusoidal components in harmonic relation to one another (top two panels: fundamental at top, first harmonic in the middle). At any point on the time axis, the amplitude of the compound wave is equal to the sum of the amplitude values of the two component waves.

1. The response of the system to a simple periodic input can be described by two quantities: the *phase* relation between input and output, that is, how much does the output lag behind the input; and the relative amplitude of the output and the input (system *gain*), that is, is the output smaller (or larger) than the input, and by how much?

2. Any periodic waveform (or any aperiodic waveform of finite duration) can be built up by adding together simple periodic waveforms.

3. When two or more simple waveforms are added together and presented as input, the resulting output is what would be expected if the two bad been presented separately and their separate outputs added together; that is, the effect of a simple waveform on the output is the same whether it is presented by itself, or summed with others. This property is known as *superposition*; together with the first property it means that the system is *linear*.

I have not defined what I mean by "simple periodic input," but for the purposes of this argument it is any kind of cyclic variation that passes through the system without distortion — the shape of the output is exactly the same as the shape of the input. Thus, phase and amplitude are the only two properties of the output that can be different from the input.

The point of these three conditions is that when they hold, we can summarize the properties of the system by two functions of *frequency*: a function showing how the ratio of input amplitude to output amplitude varies with frequency (this is termed a *gain plot*); and a function showing how phase lag depends on frequency (this is termed a *phase plot*). The two curves together are termed a *Bode plot*; I show an example in a moment. If we know the Bode plot, and the component waveforms in the input, we can predict the output.

It turns out that many real systems obey these three conditions well enough to be useful. The first condition is satisfied by the sine wave (the pattern traced out on a moving paper by a swinging pendulum): a sine wave passes through a linear system without distortion. Examples of sine waves, changed in phase and amplitude, are shown in Figure 3.9. The second condition was shown to be analytically true by the French physicist Fourier: Any periodic waveform can be built up out of sine waves of frequencies *N*, *2N*, *3N*, etc., where *N* is the *fundamental* frequency and the others are *harmonics*. Figure 3.10 shows how a complex waveform can be built up out of (or analyzed into) two sine-wave components. The third, superposition, condition is satisfied by many natural systems, providing the input amplitude is relatively small, since the independence of input components depends on their sum not being too large. If the input amplitude is too large, real systems often cannot follow, so that the output signal is *clipped*. This is one kind of *ceiling effect* (see Chapter 2).

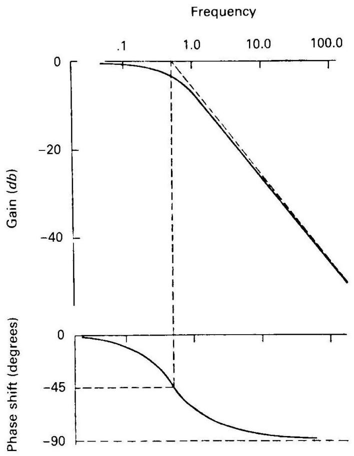

Figure 3.11. Bode plot for the exponential lag: $G(s) = a(s + a)$. Top panel: gain as a function of input frequency (both in log units). Bottom panel: phase as a function of input frequency.

Figure 3.11 shows an example of a Bode plot. The upper panel shows gain in decibels (dB), a logarithmic unit, so that equal *differences* represent equal *ratios* of output to input. The lower panel shows the phase lag, in degrees (360° represents one cycle at the fundamental frequency), as a function of frequency, also in logarithmic units. The behavior shown here is typical of many real systems: At low frequencies the output faithfully matches the input, gain is close to unity (zero on the log scale) and phase lag is small. As frequency increases, gain decreases and phase lag increases (indeed, in so-called minimum-phase systems, the two are interrelated).

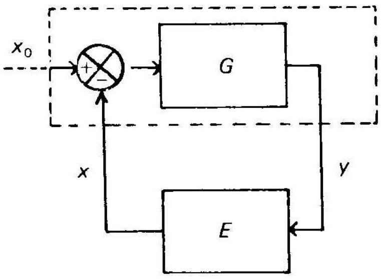

Figure 3.12. Conventional diagram of a control system: *G* is the transfer function of the organism, *E* the transfer function of the environment (feedback function), x_0 is the set point, which may be external, as in a servomechanism, or internal, as in a regulator. The environmental impact, *x*, is subtracted from the set point at the comparator, indicated by the circle and cross. The net input then goes to transfer function *G*.

The phase-amplitude-frequency information in a Bode plot is summarized mathematically by the *system transfer function*, which is described in more detail in note 5. In most respects, the transfer function can be treated as a simple multiplier. Its usefulness can be illustrated in the following way. The first step is to modify the generalized feedback diagram seen earlier in Figure 3.6 to incorporate the notion of a set point.

The modified diagram is shown in Figure 3.12. It has three elements: a transfer (control) function for the organism, *G*; a transfer (feedback) function for the environment, *E*; and a comparator, which subtracts the feedback input, *x*, from a *set point*, X_0. The difference between *x* and

x_0 is the input to G. The comparator and G are both inside the organism, as indicated by the dashed line; x_0 may be a purely internal reference, as in homeostatic mechanisms such as temperature regulation in which case the environment enters in as perturbations of input, x — or it may be an external input, as in orientation, where x_0 is given by the location of the (moving) object to be tracked. (In the first case the system is termed a *regulator*; in the second a *servomechanism.*) The input and output variables x and y are *frequency functions*, as explained more fully in note **5**. Because the feedback term, x, is subtracted from the set point, x_0, the feedback is negative and the function of the system is to diminish the difference between x and x_0.

G and E can be treated just as multipliers (*gain* functions). Thus, if the input to E is y, its output will be $x = Ey$; similarly the input for G is $x_0 - x$ and its output is $G(x_0 - x)$. The two simple, linear equations describing the relations around the loop are therefore

$$x = Ey \text{ (feedback function)}, \tag{3.11}$$

and

$$y = G(x_0 - x) \text{ (control function)}. \tag{3.12}$$

Eliminating y between the two equations and rearranging yields

$$x = x_0GE/(1 + GE), \tag{3.13}$$

which is the fundamental feedback equation.

There are three things to notice about Equation 3.13. First, there will always be some discrepancy between the set point, x_0 (zero retinal disparity in the tropotaxis example) and the actual input, x (actual retinal disparity will always be greater than zero for a moving target). The size of this error is inversely related to the product of the two gains, EG: the larger the loop gain, the smaller the error. Second, the gain terms (transfer functions) around the loop combine multiplicatively. Third, negative feedback reduces the net gain of the system, but at the same time ensures that the relation between x and x_0 is relatively insensitive to the absolute value of loop gain, providing it is fairly large: as $GE \rightarrow \infty, x \rightarrow x_0$. In any real control system, components such as muscles or motors are liable to deteriorate with time, so that an arrangement that protects performance from these effects is obviously of great value. Hence feedback systems are widespread in engineering and biology.

Equation 3.13 is a general description of any control system. Even if E and G are interpreted just as static multipliers (not transfer functions) this equation provides a passable model for feeding regulation that can explain otherwise puzzling properties of food intake and its relation to the taste of food and the effort necessary to obtain food (see Chapter 6). Many properties of reinforcement schedules can be elucidated by looking for the properties of reward and punishment that act as feedback to the animal (see Chapters 5 and 7).

Three simple feedback responses.

A linear system is completely defined by its Bode plot, which gives its response to all frequencies. This suggests two ways to discover the properties of a linear system: systematically present sine waves across the whole range of frequencies and measure the system output; or present a brief signal whose *spectrum* contains all frequencies (recall Fourier's theorem that any signal of finite length can be broken down into a series of sine-wave components). The first method is called *frequency analysis*, and the second (which is much more useful) is known as *transient analysis*. Signals particularly useful for transient analysis are the *unit impulse* (a very brief "spike" input) and the *step* input, since both these signals are made up of all frequencies, from the lowest (since the signal occurs only once its "fundamental" is at zero) to the highest (since both signals change level at an infinite rate, they must have components of infinite frequency). It turns out that a system's response to a step or an impulse is a unique "signature" that defines all its properties. This is why the quality of the "click" you used to get (in the days of vinyl) by touching the pickup stylus was a good indicator of the fidelity of a sound system.

There are three simple responses to a step input, and many natural systems correspond to one of these or to some combination. They are shown in Figure 3.13. The tropotactic system corresponds to either the exponential lag or the exponential lead, depending on how we define the dependent variable: If we are measuring θ, the angle between the animal and the target, then the system is an exponential lead — θ decreases with time after the initial target displacement. If we are measuring the angle between the animal's current and initial positions, then the system is an exponential lag — the angle begins at zero and increases in negatively accelerated fashion to a maximum. The third response, simple oscillation, would be produced by a step input if our tracker behaved like a pendulum, that is, no friction, a restoring force proportional to displacement, and (a key element) some inertia. Since a real animal will have some inertia we might expect some oscillation in response to rapidly moving targets, and a step response similar to the damped oscillation in the bottom panel of Figure 3.13.

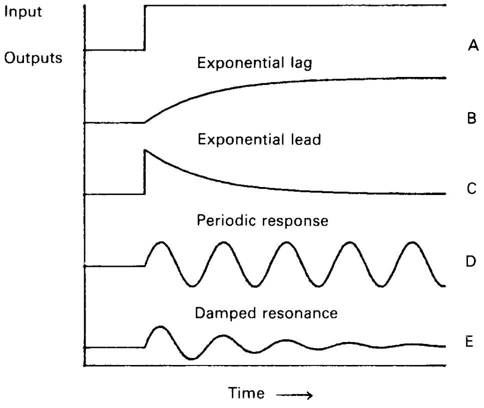

Figure 3.13. Simple responses to a step-function input. Panel A: input as a function of time. Panel B: exponential lag response. Panel C: exponential lead response. Panel D: oscillatory response. Panel E: combination exponential lead and oscillatory response (damped oscillation).

The exponential lag and lead both correspond to first-order systems (i.e., differential equations containing only first derivatives); these responses are derived formally in note 5. Oscillatory responses require at least a second-order system.

Numerous behavioral effects look like exponential lags. For example, many experimental situations show "warm-up" effects in which the organism takes some time at the beginning of each experimental session to come up to its stable level of performance. This effect is common in shock-postponement schedules, where an animal such as a rat is pressing a lever to avoid periodic, brief electric shocks. Motivational variables often show lagged effects: A sudden reduction in daily food intake shows up only gradually in a reduction of body weight. Adaptation and habituation are examples of exponential leads: A step change in stimulation produces an initial response that wanes with time. Pure oscillatory responses to stimulation are rare in the animal world, because they imply instability and a loss of contact between the organism and its environment; damped oscillation is commonplace, however, in systems where rapidly moving objects must be tracked and the physical mass of the animal's head and body enters into the feedback equations.

These three linear processes, plus a few nonlinear processes, such as delay (see the discussion of reflex latency in Chapter 2), account for most of the simple orientation mechanisms already discussed, for many physiological processes, such as aspects of temperature regulation and motor performance, for prey tracking in mantids and other animals, for simple reflexes such as pupillary contraction to light, for aspects of food and water regulation and even copulatory behavior, and for many aspects of operant behavior to be described in later chapters.

When is feedback useful?

Feedback mechanisms are found whenever the animal must cope with unpredictable variation, either in the external world, or in the properties of its own muscular system. Most motor skills involve feedback, because hand-eye coordination must change to accommodate growth and variation in muscular strength. On the other hand, feedback slows the response of any system (I show this formally in note 5, but it is pretty obvious intuitively). Consequently, activities that must be carried out rapidly usually do not involve feedback. For example, Mittelstaedt has shown that the praying mantis aligns itself with its prey using visual feedback but its strike is entirely ballistic and unguided. Thus, if the animal's "prewired" setting is wrong, the strike shows a constant error that never improves, there is no learning. The feedback that gets the "prewired" setting right is at the evolutionary level — phylogenetic, not ontogenetic. Evolution selects efficient strikers, not efficient strikes, as in the orb-weaving-spider example with which I began Chapter 2.

Even a vertebrate such as the frog shows persistent strike errors following surgical intervention. In a series of classic experiments, Roger Sperry[6] interchanged the optic nerves between the two eyes of immature frogs. When the frogs matured, they struck to the left when a fly was on the right, and vice versa, and these errors were not corrected with practice. Baby chicks with prism lenses on their eyes that shift the world over a few degrees peck to one side of the grain. Their pecks cluster more tightly as they mature, but the constant error remains (Figure 3.14).

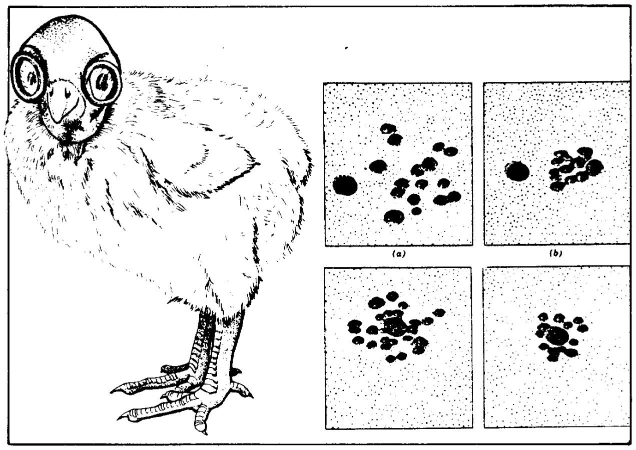

Figure 3.14. *Left panel*: chick with prism goggles. *Right panel*: clustering of pecks at a nail-head embedded in a soft surface as a function of age in normal chicks and chicks reared with prism goggles. (From Hess, 1956.)

Much human behavior can be either ballistic or feedback-guided, at the option of the individual. The Harvard physiological psychologist Karl Lashley raised the problem of feedback in connection with behavioral sequences, such as skilled typing or the arpeggios of the expert pianist, which occur too fast for tactile sensations from one keystroke to affect the next:[7] How are they organized, if not by feedback about the success or otherwise of each movement? What is being executed is not individual movements, one after the other, each guided by feedback, but a preprogrammed pattern of successive responses. There is feedback (the pianist will eventually notice if he has made a mistake), but not on a response-by-response basis: keystroke N might be affected not by the outcome of keystroke N-1, but by the effect of some earlier keystroke, say N-4. By allowing feedback to be delayed, the rapidity of a movement is limited only by muscle speed. On the other hand, feedback is not omitted entirely.

The difference between this system and the moment-by-moment one is in what is affected by feedback. In a moment-by-moment system, each response as it occurs is guided by feedback as in threading a needle, or when first learning to type, for example. In the arpeggio and prism-adaptation systems, the parameters of a ballistic motor program are adjusted on the basis of feedback from successive responses. This is a form of learning and allows the animal to anticipate future stimulus-response relations, because it is adjusting its "strike" for future action based on a history of past hits and misses — implicitly assuming that the world in the future will be essentially the same as the world in the past. A feedback system of this sort is sometimes termed a *mesh* system.

Thus, both mantids and people respond ballistically when speed of response is critical; but people (and some other mammals) can both alter the parameters of the ballistic system (as in their adaptation to prism lenses) and deal with the inevitable delays when feedback is allowed to lag behind action (as in fast typing). A ballistic system may itself be the result of learning: Ballistic typing only develops after a hunt-and-peck stage in which each keystroke is feedback-guided.

Feedback increases phase lag, which may, in turn, contribute to instability. Notice that in the Bode plot in Figure 3.11 phase lag increases with frequency. This is almost universal in real systems. The lag in Figure 3.11 does not go beyond 90º (a quarter of a cycle), but in many systems can go to 180º or more. At a lag of 180º, the output of the system is at its minimum when the input (to the set point) is at its maximum, and vice versa. Consequently, negative feedback, which is a stabilizing effect, has been turned into positive feedback, a destabilizing, "vicious circle" effect. It is not surprising, therefore, that many feedback systems get into difficulties at high frequencies. The ability to adjust system parameters, and to anticipate regularities in the input, are important protections against this kind of instability.

THE INTEGRATION OF BEHAVIOR

English writer H. G. Wells (*War of the Worlds*) once remarked that "The end of all intelligent analysis is to clear the way for synthesis." There is little point in breaking down the behavior of animals into simple units, be they reflexes, servomechanisms or "elementary information processes," if we cannot use these units to understand the behavior of the whole animal in its natural environment. The only way to tell if we have all the pieces of the puzzle is to put them together and see if we have reconstructed the natural behavior of the animal. I first discuss the role of variability in behavioral integration, then give examples of how a few simple mechanisms, acting together, can produce quite flexible and adaptive behavior.

The initial problem is to account for *variability* in behavior: If we have a well-defined set of elementary processes, it is natural to expect that they should lead to unique predictions about behavior — yet behavior is often unpredictable. What is the adaptive utility of behavioral variation, and on what does it depend?[8]

In cases like the simple kineses discussed in Chapter 2, the variability is intrinsic to the mechanism and serves the function of random sampling of an unknown environment. In other cases, a degree of variability even in an otherwise determinate mechanism can prevent the animal from getting trapped by particular circumstances. For example, a tropotactic animal that adjusts its orientation so as to balance the stimulation from bilateral receptors would pass completely between two equal lights and never reach either, if its mode of movement were completely regular. A little variability allows it to find the other balance point that involves going straight to an adjacent light even though it started out going between the two lights.

Variability is also helpful in getting around obstacles. Many insects are photo-positive in dim light and everyone has seen flies and other insects trapped in a room buzzing at the window. The fly does not search systematically for a way out, but if a gap exists the animal's random movement will eventually find it. As wasteful as it appears, the fly's random buzzing is much more successful than would be a precise homing mechanism that simply settled down at the lightest part of the windowpane.

Different individuals of the same species often behave differently under similar conditions. For example, photonegative *Ephestia* (meal-moth) larvae move generally away from a source of light, but different individuals show different amounts of error and the distribution of heading angles of a group of larvae is approximately normal (i.e., the bell-shape characteristic of a process determined by many independent causes). The nauplius larvae of the barnacle *Balanus* are a more interesting case: Some are photonegative and some photopositive, so that the distribution of heading angles to light is bimodal. Such bimodal distributions, an example of *behavioral polymorphism*, are quite common, and quite puzzling at first sight.

The functional basis for behavioral polymorphism seems to be a type of *frequency-*

dependent selection where the Darwinian fitness of one type of behavior is inversely related to its frequency relative to the other. For example, suppose that a particular prey species can find food in two different types of habitat that are not spatially contiguous; predators must decide to hunt in one or the other. Clearly it pays the predator to spend most of its time in the more popular habitat. Hence a prey individual that prefers the less popular habitat will be at an advantage. This system ensures that approximately equal numbers of prey individuals will be found in both types of habitat.[9]

Behavioral polymorphism can take two extreme forms: either the same individual can show two or more modes of action at different times, or there may be two or more *types* of individual, each showing one mode. Both kinds of variation are found, but the first kind is potentially more flexible, as it allows for the evolution of *systematic* variation: that is, selection by the animal of one mode or the other, depending upon circumstances. Systematic variation tends to be the strategy increasingly employed by mammals and other "higher" animals. In a simple form it is common even in invertebrates. For example, the protozoan *Euglena* is photo-positive in weak light and photo-negative in strong light. Consequently the animals congregate in intermediate levels of illumination that may represent an adaptive compromise between the bright light that provides maximum energy from their photosynthetic organ and the dim light that provides greatest safety from predators. Some individual cockroaches, *Blatta orientalis*, are hygro-negative (avoid high humidity), but after some time in dry air they lose water by evaporation and then become hygro-positive. In mammals and birds a change of this sort would be termed a *motivational* one, caused by an altered physiological water balance. Reactions also change as a function of time: habituation, warm-up effects, and circadian and other rhythms are examples that have already been discussed. A taxis may even reverse periodically, as in the *Mysis* crustaceans in the Naples aquaria.

The sense of a taxis may change as a function of some quite different variable. For example, under suitable conditions, *Paramecium* is geo-positive (it descends) in light, and geo-negative in darkness. Many marine invertebrates show a similar reaction, which may be designed to vary the animal's depth as a function of the time of day. The water flea, *Daphnia*, shows an interesting adaptive response to an increase in the carbon-dioxide concentration: It becomes photo-positive, a reaction that would normally take it to the water surface where the CO_2 concentration is likely to be lower and the oxygen tension higher. Exposure to dry air tends to make the woodlouse (sowbug), *Porcellio*, photo-positive rather than photo-negative. This case is more puzzling because it seems very unlikely that light is a reliable cue to humidity for this animal; indeed, in general, the opposite is more probable: Dark places such as crevices and holes are much more likely to be damp than bright, open areas. However, evidently the woodlouse, like a good politician, is a pragmatist: Ideology may say that dark means damp, but if one is in the dark, and dry, then maybe light is worth a try. This mechanism is a primitive, built-in, version of the "win stay, lose shift" strategy of trial-and-error learning in higher animals.

A simple set of mechanisms may, in combination, yield remarkably intelligent behavior. Consider the behavior of the sea-snail *Littorina neritoides*, as analyzed by Fraenkel:

> This animal is found several metres above high-water marks of European seas. It is usually geo-negative and never geo-positive. When out of water it is always photo-negative. In the water it is too, except when it is upside down, and then it is photo-positive. . These reactions may be expected to guide the animal from the bottom of the water to rocks (dark) and then up the rock face (geo-negative); if the light is very bright, the animal stops and settles down at the water surface. In crawling up under water, if it gets into a deep horizontal cleft, negative photo-taxis takes it inwards on the floor, negative geo-taxis upwards on the end wall, and positive photo-taxis outwards on the roof in the inverted position. Upward progress is therefore not barred by such a cleft. Above the water surface, the sign of photo-taxis does not reverse in this way, so the animal comes to rest in such a cleft. Under the influence of dryness and other unfavorable conditions the animal closes up; it may then fall back into the sea, but if conditions are not very unfavorable it may live for

> months out of water. This behavior provides examples of taxes in opposition (gravity and light) and of one stimulus (the presence of surrounding water) affecting the response to another (light). (Fraenkel & Gunn, 1940, pp.297-298)

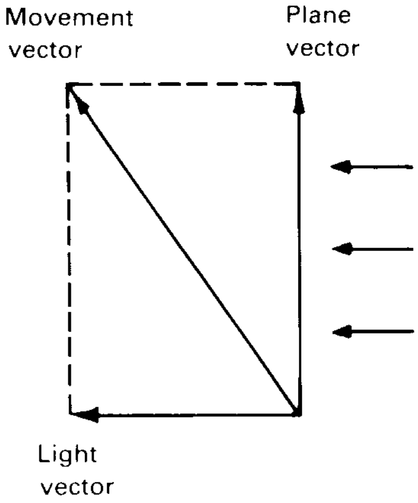

Figure 3.15. Vector model of the combined effect of negative geotaxis (tendency to move directly up the slope of a plane) and negative phototaxis (tendency to move directly away from a light). (Adapted from Crozier & Cole, 1929).

In these examples the animal's response to one variable, such as light, is modulated by another, *contextual*, variable, such as gravity or whether the animal is in or out of water. Often, however, two or more variables may each have an independent taxic effect; for example, both geotaxis and phototaxis may operate simultaneously. What happens when these variables act in opposition, if negative geotaxis dictates upward movement, say, and negative phototaxis urges downward movement, away from the light? Here, as in the case of opposed reflexes, there are only two possibilities — competition or cooperation. Most taxic mechanisms appear to interact cooperatively. For example, a slug placed on an inclined plane turns so that it travels straight up the plane, along the line of greatest slope (it appears to accomplish this by equalizing muscle tension on both sides of the body). Slugs are also negatively phototactic, and if a light is shone on the animal from one side as it moves up the plane, it will turn away and adopt an intermediate orientation; the stronger the light, the greater the turning away and the larger the angular deviation from a straight upward path. The animal's direction of movement can be modeled here as the resultant of two vectors, one representing the slope of the plane and the other the direction and intensity of the light, as illustrated in Figure 3.15. Under other circumstances one reaction may totally dominate the other, particularly if one is a protective reaction of some sort such as the withdrawal reaction shown by *Stentor* and other small, sessile invertebrates to a passing shadow.

These two kinds of combination rule, cooperation and competition, are not, of course, perfectly distinct, but merely extremes along a continuum. At one end the effects are additive (linear); this is cooperation. At the other end is total competition: The response with greatest strength occurs exclusively and the other is completely suppressed (nonlinear). As in reflex interaction, which of these possibilities actually occurs no doubt depends on the Darwinian fitness of intermediates: When a compromise is less adaptive than either alternative alone, as in the feeding versus flight choice for *Stentor*, one or other will occur exclusively.

There seems to be a trade-off across animal species between systematic and unsystematic variation. Systematic variation means that the rules the animal plays by depend upon the situation, on contextual variables such as time of day, condition of light, presence of another animal, and so on. In each situation, the animal's behavior may be very predictable, yet the animal may be sensitive to many different situations (not all of them, perhaps, known to an outside observer) so that its overall repertoire of behavior may be large. Unsystematic variation means that the animal always plays by the same set of rules, but these include a substantial unpredictable element — kineses are the obvious example. There seems to be a trade-off in the sense that the more predictable the animal in any particular situation, the richer the variety of situations to which it is sensitive. Human beings can carry out particular tasks with enormous precision, for example, yet there is no denying their unpredictability in general, or its dependence upon the individual's perception of situation.[10]

THE NATURE OF EXPLANATION

The emphasis of this book is on explanations for behavior in purely environmental terms: I ask, How does the animal's past history interact with rules built in to the animal by its evolutionary history and individual development to produce future behavior? I am not directly concerned with the animal's internal structure, or with intentional states we can relate to our own introspections (see Chapter 1). Since there is often confusion about the meaning of explanations at this level, I conclude the chapter with brief discussions, first of "black-box" explanations, and then of the relation between feedback theory and explanations of behavior in terms of motive and purpose.

The meaning of "black-box" analysis

Scientific explanations can be at many levels. The behavior of a moth circling a flame might be explained in a number of ways: (a) as an "instinct"; (b) as a form of taxic reaction (e.g., light-compass reaction); (c) as a particular kind of control system; or (d) as the response of a particular neural network connected to a set of receptors and effectors.

The instinct account says nothing about the process involved; it is really just a kind of classification. It places the moth's behavior in a group containing all those behaviors that seem not to depend on experience. Although this is a relatively primitive kind of explanation, it is not empty; it may in fact be essential to go through such a "natural history" stage at the beginning of any science.

Explaining the behavior as a taxis is also classificatory, but now in terms of the results of experiments — unilateral blinding, the two-light experiment — and begins to get much closer to the underlying process.

The control-system account depends on more experiments and more quantitative results and constitutes an explanation in terms of *process*. It gives the "rules of the game," those aspects of the behavior that are invariant, that is, independent of particular inputs and outputs (recall the constancy of the transfer function of a linear system). However, the control-system model is a black-box account: The equations of the model, and the boxes and connecting arrows that make them easier to understand, may have *no necessary relation* to with the elements of the fourth level of explanation: an account in terms of neural and other structures. Nevertheless, the control-system account says what the neural structures must do, and gives hints on how they might do it. But block diagrams cannot be interpreted literally as connections among neural circuits.

Black-box accounts provide a natural preliminary to accounts at the fourth level, in terms of neural structures, but cannot be directly interpreted in neural terms (see note 2, Chapter 2).

The lack of relation between block diagrams and the physical constitution of the system is especially obvious if we are modeling a real physical system. For example, the exponential lead describes the relation between current (output) and voltage (input) across an electrical capacitor (essentially a pair of parallel plates capable of storing a quantity of electrical charge): If the voltage across the capacitor is suddenly increased, there is an initial large current flow that subsequently decays to zero. Where is the "loop" here? There is no physical link corresponding to the feedback path in Figure 3.12, nor is there anything corresponding to a comparator or a set point. There is negative feedback, but it is in the form of the increased repulsion among charges as the number of charges on each plate increases. Many other physical and biological systems behave like exponential leads: a spring-loaded dashpot, a pendulum in molasses, and the inflow of organisms into a suddenly denuded habitat. The negative feedback (restoring force) in the first case is the spring; in the second, gravity; and in the third, competition among individuals for limited resources.

Each level of explanation has its uses. Explaining the moth's circling in terms of instinct is a useful starting point for the student of behavioral development and it may have implications for the evolution of behavior. The taxis account relates the behavior to other orientation reactions, and may be useful for understanding species differences. The control-system account is useful for making quantitative predictions, and as an element in accounts of the integration of behavior, that is, in theories which bring together a number of mechanisms so as to account for something approaching the organism's total behavioral repertoire.

Purpose, teleology, and mechanism

The concept of *purpose* has occupied a prominent place in the history of psychology. Disdained by behaviorists, it is covertly retained even by them in the form of so-called "control of behavior by its consequences." Alert readers may have noticed that the set point of a feedback mechanism (the *sollwert* or "should-be value" in German) corresponds pretty closely to the intuitive idea of a "goal" or "motive." This connection was pointed out some years ago in a classic philosophical paper by Rosenblueth, Wiener, and Bigelow (1943) in which they showed that the idea of feedback provides the conceptual link between mechanistic accounts of behavior, that is, explanations in terms of antecedent events (proximal causes), and teleological accounts, that is, explanations in terms of goals or motives (final causes). The superiority of a feedback account rests in its self-contained ability to account both for those cases where the goal is attained, and those where it is not. The former simply represent the domain of stability and negligible phase lag of the system, the latter its region of large phase lag and instability.

Commonsense explanation in terms of purpose does not easily account for failures to achieve a goal. For example, a student may do well on an organic chemistry examination, and this might be attributed to his goal of getting into medical school (say). But suppose he fails the exam, despite strong motivation to succeed? Explanation in terms of goals or motives must then postulate some competing motivation — he may have spent too much time with his girl friend, perhaps — or else say something about the student's capabilities and his methods of study, that is, resort to a mechanistic account. In short, conventional motivational accounts offer only two explanations for failures to achieve an objective: either competing motives, or an unsuspected structural limitation.

Both these escape routes have an ad hoc look to them, yet they roughly correspond to respectable explanatory alternatives: an optimality account, that provides the laws by which motives compete, and a mechanistic account, that dispenses with motives entirely and explains behavior by antecedent conditions and structural properties of the subject (see Chapter 1).

Different fields have inclined more to one or the other type of explanation. Classical economics, for example, explains people's allocation of time and money between different commodities, or between work and leisure, in terms of a balance of motives, a teleological account. The individual is assumed to optimize his total utility by allocating his resources so that the marginal gain from switching from one thing to any other is constant; that is, he spends so that his last nickel will buy him the same benefit, no matter where he spends it (equal *marginal utilities*). In this way, total benefit will usually be maximized. Optimality analyses (about which I shall have much more to say in later chapters) are the ultimate form of teleological account. On the other hand, the field of ecology, which has much in common with economics as the etymology implies, has generally favored mechanistic accounts: The distribution of species within and between habitats is usually explained in terms of their relative efficiencies and reproductive rates.

Psychology and behavioral biology have at different times favored both approaches. Reinforcement theory is basically teleological in the sense that behavior is assumed to be guided by access to a key event, the reinforcer, which in effect functions as a goal: The animal works for food or to avoid electric shock. Purposive terminology is studiously avoided, but the lack of any generally accepted mechanism to account for the effectiveness of reinforcers means that the term *reinforcement* is actually used as if it meant "purpose." For example, when several experiments showed a few years ago that hungry rats would press a lever for food even if free food were available, a popular response was that the lever-pressing behavior was "self-reinforcing." The difference between this and the naive explanation that the animal "likes" to press the lever may be too subtle for the naive reader to detect. As we'll see in later chapters the essence of reinforcement, both as an effective procedure and an explanation of behavior, is restriction of access, so that the notion of self-reinforcement (as a sufficient explanation for behavior) is a contradiction in terms.

Other areas of psychology have looked for mechanistic accounts. For example, Hullian theories of learning, though now of largely historical interest, were a praiseworthy attempt to specify the causal links between the stimulus, as cause, and the response, as ultimate effect.[11] More recent mathematical and computational theories of classical conditioning and the effects of food and other hedonic stimuli on general activity (arousal) are strictly mechanistic, looking to antecedent, rather than consequential, events for the explanation of behavior. Theories that explain choice and the distribution of behavior in terms of competition between the tendencies to engage in different activities are also, like their ecological counterparts, mechanistic in spirit.

Both mechanistic and teleological accounts have their uses, although mechanistic theories are obviously desirable where they are possible. But when the means an organism, or an economy, can use to attain some end state are many and poorly understood, but the goal is relatively clear, then a teleological theory may be the best that is available (see Chapter 1). An adequate teleological explanation (like the empirical principle of reinforcement) is certainly better than a premature mechanistic theory, as Hull and his followers found to their cost.

SUMMARY

The mechanisms discussed in this chapter and the preceding one represent the best that animals can do when limited to what might be termed *local memory,* that is, when their behavior is affected only by present events and events in the immediate past in an essentially context-free way. Although simple, and relatively easy to analyze, these mechanisms can nevertheless produce remarkably flexible behavior. In the rest of the book, the emphasis is on learned behavior, that is, behavior that is context-sensitive and depends on more remote past history.

Learned behavior can be studied in two ways: as learning, that is, as a *change* in behavior with time and experience; or as *habit,* that is, as a more or less fixed pattern, built up after much experience. The study of habit has the advantage that like the mesh feedback systems discussed earlier, once set up, habits may show rather little dependence on past history. They can be studied like orienting mechanisms, in individual animals, using reversible procedures. The next several chapters look at learning from this point of view. Later chapters deal with learning as change.

NOTES

1. This limitation is a general feature of all hill-climbing mechanisms: They find their way to the top of *a* hill, but not necessarily the *highest* hill. They are *local* rather than *global* optimizers.

2. ***Contrast mechanisms***. The initial advance in understanding these rate-sensitive effects of sensory mechanisms was made by the Austrian physicist and philosopher Ernst Mach (1838-1916). Floyd Ratliff, in his book *Mach Bands* (1965), has provided a fascinating account of Mach's sensory work and translations of relevant papers, as well as a summary of more recent work on the problem. See also Arend, Buehler, & Lockhead (1971); von Békésy (1967); and Ratliff (1974).

The intimate relation between the perceptual effect of a luminance gradient and its higher derivatives can be seen from Figure 3.16. The top curve is a gradient of luminance as a function of distance, $V(s)$, such as would be obtained from a set of progressively lighter bars (*ab*, *cd*, *ef*, etc.) separated by regions of gradually increasing lightness (*bc*, *de*, etc.). The next curve, labeled "response" shows the appearance of this luminance gradient. There are two features that are of special interest: First, the progressive change in lightness in regions *bc*, *de*, etc. is barely perceptible, and the successive bars *ab*, *cd*, etc. increase but little in perceived brightness. Second, narrow dark and light bars (*Mach bands*) are seen at the points of inflection *b*, *c*, *d*, etc. The third curve shows the first derivative of the luminance curve and the bottom curve is the second derivative. Mach saw at once that the light and dark bands are derivable (to an approximation) from the negative of the second derivative of luminance, and that the perception as a whole can be represented as the weighted sum of the original luminance distribution (transformed via a compressive function as described in Chapter 2) and the second derivative of that distribution:

$$R(s) = A\log\left[V(s) - (B/V)d^2V/ds^2\right] + C, \qquad \text{(N3.1)}$$

where *A*, *B*, and *C* are constants, and the second derivative is multiplied by 1/*V* because the effect of an inflection is inversely related to the brightness level.

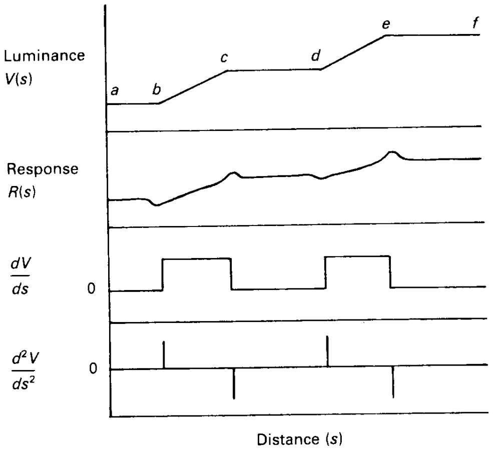

Figure 3.16. A luminance gradient (top curve), the perceived brightness gradient (second curve), and the first and second derivatives of the luminance gradients (third and fourth curves).

Figure 3.17 is redrawn from Figure 3.2 in the text, with the addition of the first and second derivatives. It is clear that it poses difficulties for Mach's scheme (which is itself equivalent to several more recent models, as Ratliff shows), since no simple combination of the original luminance distribution with the second derivative will yield the decreasing staircase that represents the perception.

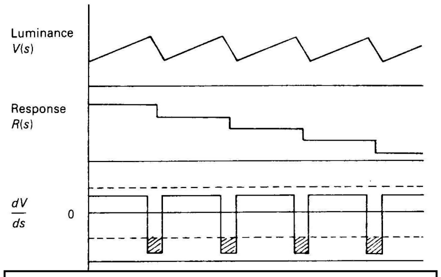

Figure 3.17. A sawtooth luminance gradient (in logarithmic units, top curve), its perception, and the first and second derivatives of the luminance gradient. Suprathreshold regions of the first derivative curve are crosshatched.

If for simplicity we ignore the compressive transformation (or, alternatively, simply assume that *V(s)* is measured in logarithmic units), then Equation N3.I becomes

$$R(s) = aV(s) + bd^2V/ds^2 + c. \qquad \text{(N3.2)}$$

It is obvious that *R(s)* cannot be derived from this equation for the luminance distribution shown in Figure 3.17. However, Equation N3.2 can be rewritten in a way that immediately suggests a modification to accommodate Figure N3.2 and related effects. First, we rewrite Equation N3.2 in terms of first and second derivatives only:

$$R(s) = a\int_0^s (dV/ds)ds - b\left(d^2V/ds^2\right) + c'. \qquad \text{(N3.3)}$$

(The integral term reduces to *V(s) - V(0)*, so that *V(0)* is incorporated in a new value of the additive constant, *c'*.) Equation N3.3 is equivalent to Equation N3.2 save for the assumption that absolute values of *V*(s) are not sensed directly, but rather affect sensation via the integration of the first derivative. Whatever the stimulus aspect that is sensed by a sensory system, there will always be a threshold. If rate of change of luminance, rather than absolute luminance, is important, then there will be a rate of change so slow as to have no effect. The value of this threshold will presumably be related to the difference limen (differential threshold) for brightness. Thus, the term *dV/ds* in Equation N3 .3 refers only to suprathreshold values of *dV/ds.*

The implications of this change can be seen in Figure 3.17 where the threshold is indicated by the two dashed lines (two because the first derivative can take on negative as well as positive values): The small positive values of *dV/ds* associated with the gradual part of the sawtooth are below threshold and therefore do not contribute to the integral, whereas the large, but briefer, negative values associated with the steep part of the sawtooth add to the integral yielding a fixed increment in the sensation at each cycle. The positive and negative parts of the second derivative are very close together because of the steepness of the falling part of the sawtooth, and so interfere with one another; since their average is zero (and since the contribution of the second derivative to perception is probably much less than the contribution of the first derivative), they have little net effect on perception. Thus, the perceptual effects of luminance gradients can be derived, to a first approximation, from the simple hypothesis that the visual system (and other sensory systems as well) is sensitive only to rates of change.

This analysis applies most easily to temporal variations because here the direction of integration is determined unambiguously by the flow of time: The system has no alternative but to integrate past stimulation up until the present. Spatial interactions are more difficult to handle: Should the system integrate from left to right, from top to bottom, or what? One possibility is that spatial changes are converted at once to temporal ones via eye movements, which occur all the time. Land and McCann (1971) have suggested other possibilities, based upon integration of differences along interconnecting paths.

As we have seen in previous chapters, habituation and adaptation are names for the processes that permit special sensitivity to changes in stimulation, and the curves in Figure 3.17, particularly, are very similar to those encountered earlier. For example, in Koshland's experiments on bacterial orientation a "pulse" of nutrient produced a transient decrease in tumbling, very like the transient increase in sensation level produced by a sudden stimulus change in Figure 3.17.

The transience of the stimulus effect in the bacteria was attributed to adaptation, which can, therefore, be seen simply as another name for the rate-sensitivity of many biological systems.

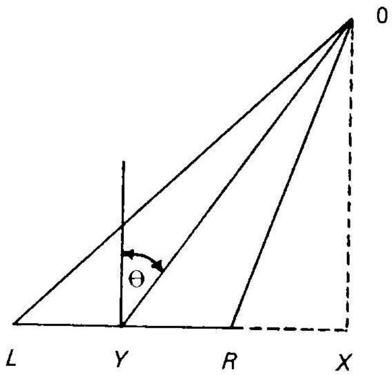

Figure 3.18. Geometric relations between a light source, *0*, and two bilateral, point receptors, *L* and *R*.

3. This result can be obtained as follows: Let A_0 be the intensity of the light source at *0* in Figure 3.18, and *L* and *R* be the two bilateral point receptors. *Y* is a point on the midline of the animal on the line joining the receptors and θ is the heading of the animal with respect to the source. Receptor disparity is the difference in light intensity falling on the two receptors; by the inverse square law this is:

$$\Delta I = A_0/0L^2 - A_0/0R^2 \quad \text{(N3.4)}$$
$$= A_0(0R^2 - 0L^2)/0L^2 0R^2.$$

Considering each term in the numerator separately:

$$0R^2 = 0X^2 + RX^2 \quad \text{(N3.5)}$$
$$0L^2 = LX^2 + 0X^2 = (LR + RX)^2 + 0X^2$$
$$= LR^2 + LR \cdot RX + RX^2 + 0X^2.$$

Substituting in Equation N3.4 yields: $\Delta I = -A_0\, LR(LR + 2RX)/0L^2 0R^2$. But by elementary trigonometry, $\sin\theta = YX/OY = (RX + LR/2)/OY$; hence

$$\Delta I = -A_0 LR - 0Y2 \sin\theta/0L^2 0R^2, \quad \text{(N3.6)}$$

which is the stated relation. If the animal s far away from the source, so that $OY \gg LR$, then $OL \cong OR \cong OY$, and this relation reduces to

$$\Delta I = -A_0 LR \bullet 2\sin\theta / OY^3, \quad \text{(N3.7)}$$

so that receptor disparity is directly proportional to receptor separation, *LR,* and inversely proportional to the cube of the distance to the source.

4. Regulatory gadgets such as the furnace thermostat and the centrifugal governor used in early clocks have been around for some time. The modern, quantitative study of feedback devices dates from James Clerk Maxwell's paper (1868) on the centrifugal governor used in James Watt's steam engine. Feedback theory received its greatest impetus during the second World War when tracking devices of all sorts automatic gun sights, radar trackers, homing torpedoes, elicited a level of public interest lacking in more tranquil times. Since then the study of control systems has grown into a highly developed field of largely mathematical inquiry.

There are now so many texts on the theory of control systems that it is hard to give an objective survey. The classical source is MacColl's *Fundamental Theory of Servomechanisms* (1945), but this, although relatively easy to follow, is not explicitly directed at the novice, and gives no account of the now-standard Laplace transform method. Other books or papers aimed at biologists or psychologists are Toates (1975), a clear elementary account, McFarland (1971), Powers (1978), Grodins (1963), and Milsum (1966). Useful summaries of empirical and theoretical studies are to be found in McFarland (1974), the 1964 Symposium of the Society of Experimental Biology, Toates and Archer (1978), and Davis and Levine (1977). A comprehensive account of the mathematics involved is Kreider, Kuller, Ostberg, and Perkins (1966). Much recent material can now be found on the internet.

My account of control theory as applied to tropotaxis is, of course, greatly simplified. More comprehensive applications of the approach to orienting mechanisms such as the optomotor response (visual following of moving stripes, shown by insects and many other animals), are to be found in the German literature, e.g., Mittelstaedt (an account in English is 1964; see also 1972).

5. ***Linear systems analysis.*** The three fundamental assumptions of linear systems analysis have been given in the text. Here I describe the Laplace transform method that allows systems of linear differential equations to be solved in an almost automatic way.

Recall that the properties of any linear system are completely described by its Bode plot, which gives the gain and phase lag for all frequencies of sine-wave input. The *Laplace transform* is a function that incorporates all the information in a Bode plot. The transform method is similar to the use of logarithms, a technique familiar to most people with some mathematical experience. Both methods change quantities into a form where they are more convenient to work with. In the days before microchips, multiplication or division of large numbers was a tedious chore. The invention of logarithms in the seventeenth century by the Scottish mathematician John Napier allowed numbers to be converted to their logarithms, whereupon the sum of the logs yielded the log of the product, so that multiplication and division were replaced by the much simpler operations of addition and subtraction. The key to the simplicity of the method lay in the laboriously compiled tables of logarithms and antilogarithms that enabled the results of these additions and subtractions to be at once translated back into the raw number domain.

Most time functions that are of practical interest can be similarly converted into their Laplace transforms. So can the differential equations that describe the control function. These transformations allow the operations of differentiation and integration (necessary for the solution of differential equations) to be converted into arithmetic operations of multiplication and division, as I show in a moment. In this way any linear differential equation can be solved just using simple algebra. The solution takes the form of an expression in the complex variable s. By consulting a table of Laplace transforms (analogous to a table of antilogarithms), the time function corresponding to this expression can then be found.

The Laplace transform is denoted thus: $£F(t) = F(s)$ where s is a complex variable containing frequency, phase and amplitude information. Formally,

$$£F(t) = \int_0^\infty F(t)\exp(st)dt\,. \tag{N3.8}$$

For details of the derivation and the limitations that allow the integral to exist (i.e., be finite) see (for example) Toates (1975), Milsum (1966), or Kreider et al. (1966). The Laplace transform is closely related to the moment-generating function, familiar to students of probability theory.

The thing that makes the Laplace transform so useful in the analysis of linear systems is that the ratio of the output $F_0(s)$, and input, $F_i(s)$, is *constant* for all inputs; this ratio is the *transfer function.* Thus, the transfer function, $G(s)$, is defined by $F_0(s) = F_i(s) \cdot G(s)$. The transfer function behaves just like a multiplier, when input and output are expressed in terms of the complex variable s.

$G(s)$ can be derived in three separate ways. (a) Analytically, from the differential equations of the system, if these are known. This is very straightforward because the two operations of differentiation and integration correspond just to s and I/s, respectively, so that $£(d/dt)F(t) = sF(s)$, for example, where $F(t)$ is some function of time. (b) From the input-output relations derived from certain inputs, such as the "unit step" and "unit impulse," that contain a wide range of frequencies and produce responses with a "signature" characteristic of the system. This is known as *transient analysis.* (c) From the Bode plot obtained by empirically probing the system with a wide range of sinusoidal frequencies; this is termed *frequency analysis.*

As an example of the derivation of $G(s)$ from the differential equations of a system, consider again the simple model of tropotaxis described by Equations 3.5 and 3.6. First, since this is a dynamic analysis, we can replace a and b by $a(t)$ and $b(t)$, indicating that both these variables are functions of time. Substituting in Equation 3.7 and rearranging yields: $db(t)/dt = AB(a(t) - b(t))$. Taking Laplace transforms of both sides yields:

$$sb(s) = AB[a(s) - b(s)].$$

Rearranging so as to obtain the ratio output/input, $b(s)/a(s)$, yields:

$$b(s)/a(s) = G(s) = AB/(s + \mathrm{AB}). \qquad \text{(N3.9)}$$

The expression on the right-hand side is therefore the transfer function for this system. The output for any given input signal, $a(t)$, can now be obtained simply by substituting the Laplace transform for that signal for $a(s)$ in Equation 3.9 and rearranging to obtain an expression for $b(s)$ in a form suitable for consulting a standard table. The response explored earlier is equivalent to the "unit step," an instantaneous displacement of the target light from a value of 0 (the animal heading straight to the target) to 1. The Laplace transform of the unit step is just I/s, so that $b(s)$ is equal to $AB/s(s + \mathrm{AB})$. This expression does not appear in standard tables, but the expressions 1/s and l/(s - a) do appear there. AB/s(s + AB) can be written in this form using the method of partial fractions, which yields:

$$b(s) = 1/s - l/(s + AB).$$

From the tables this means that

$$b(t) = 1 - exp(-ABt), \qquad \text{(N3.l0)}$$

Which is just Equation 3.8 with $a = 1$, $b_0 = 0$, and $b = b(t)$. The response of this system to any other time-varying input can be obtained in the same way.

Equations N3.10 and N3.8 are termed *exponential lags,* and this example shows that an exponential lag in response to a step input is the characteristic signature of a transfer function of the form $a/(s + a)$, that is, of integral control. The same information can be derived from a Bode plot for this system, which has already been shown in Figure 3.11, using the vector properties of the complex variable s.

Exponential lags result when the rate of change of the dependent (controlled) variable is proportional to the level of the independent (controlling) variable — or, equivalently, when the controlled variable is proportional to the integral of the controlling variable. There is obviously a converse process, when the controlled variable is proportional to the rate of change of the controlling variable (rate control). Its properties can easily be deduced from what we already know about the Laplace transform method. Thus (using the same terms as before), rate control means that

$$b(t) = A \cdot d\,[a(t) - b(t)]/dt,$$

where A is a constant. Taking Laplace transforms and rearranging, as before, yields

$$b(s)/a(s) = s/(s + I/A).$$

Setting $a(s)$ equal to the Laplace transform for a step input (I/s) and rearranging yields

$$b(s) = l/(s + I/A),$$

so that (from tables)

$$b(t) = \exp(-t/A),$$

which is again an exponential with time constant A. This function, the exponential lead, should be especially familiar as the temporal change characteristic of effects variously labeled as adaptation, habituation, or fatigue; that is, rapid response to a sudden change in stimulation, such as the onset or offset of a light, followed by slow waning of response back to a low level.

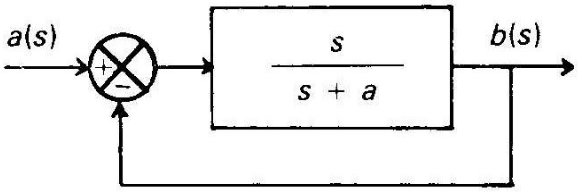

Figure 3.19. Effect of feedback on the exponential lead.

As an example of the use of the transfer function in learning about the temporal properties of a system, consider again the fundamental feedback equation (Equation 3.10). The transfer functions G and F are used just like simple multipliers in that equation, so we can use the equation to predict the effects of adding a feedback loop to one of the elements shown in Figure 3.13. Consider, for example, the exponential lead. Suppose we connect it in a feedback loop as shown in Figure 3.19; what will be the transfer function of the combined system? We can find out at once just by substituting the transfer function, $s/(s + a)$, for G in Equation 3.10, with a feedback gain (E) of unity. The result is the transfer function $s/(s + a/2)$; thus, the time constant, which was originally I/a, is now $2/a$ the system responds twice as slowly.

This is a general result: The flexibility conferred by feedback is bought at the cost of a slower response. Consequently, activities that must occur as rapidly as possible often do not involve feedback, as 1 describe in the text.

6. Sperry (1951) is a good review.

7. Lashley (1951). A now obsolete example: A rapid typist can verify immediately that his keystrokes are not guided by immediate feedback by noting how many keys he hits after the end of a line is reached and the carriage will not advance further (assuming that he has already ignored the bell). If he is at the same intermediate stage of incompetence as I, he will type fast enough to hit several keys after the carriage has reached the limit of its travel. This is a measure of the time it takes for feedback to affect the motor program. If you don't know what bells and carriages are…never mind.

8. There is a tendency to equate *unpredictability* with *randomness*. It may be worth pointing out that quite simple deterministic systems can lead to behavior that looks random to an observer who is ignorant of the generating process. Mathematicians have shown recently that a perfectly deterministic set of (nonlinear) differential equations can, nevertheless, lead to apparently chaotic, and certainly unpredictable, behavior (see, for example, May, 1976). Thus, unpredictability, from the point of view of an outside observer, does not imply absence of determinism: A system may be both deterministic and irreducibly unpredictable — unpredictable is not the same as random. (I return to this topic in Chapter 4 in connection with the concept of *observability.)*

9. Behavioral polymorphisms have been analyzed most thoroughly in connection with social behavior. Maynard Smith (e.g., 1976; see Krebs & Davies, 1981, for an elementary discussion) has used game theory to derive the idea of an *evolutionarily stable strategy* (ESS) a stable mixture of behaviors that coexist in a population because of frequency-dependent effects. ESS is more generally known as a *Nash equilibrium*, after the economist John Nash, who achieved notoriety unusual for a mathematician as the subject of the partly factual biopic *A Beautiful Mind* (2001). He also shared the 1994 Nobel Prize for Economics (with Reinhard Selten and John Harsanyi).

10. The argument here is closely related to W. R. Ashby's (1956) *law of requisite variety,* which states that if an organism is to maintain its internal environment constant in the face of environmental perturbations, the variety of its output must equal the variety of its input. This is easy to understand if you think of a simple tracking mechanism, such as the model of tropo-taxis discussed earlier. If the system is to continue to be aligned accurately with a moving target, then its correcting movements (responses) must exactly match those of the target. The ability of the tracker to do this is, of course, a function of its frequency response in relation to the frequencies in the pattern of movement of the target (which is directly related to the input *variety,* in the communications-theory sense).

Ashby's law (which derives from the theory of communication pioneered by Shannon & Weaver in 1949) refers both to the way an animal varies a single response and to its ability to come up with other responses if one fails. Linear systems theory is concerned more with the quantitative properties of a single response, such as shaft angle or voltage, to a single input, whereas real organisms cope with environments that require variation in the type, as well as the quantitative properties, of responding. The principle applies equally to both cases, however. The implication of this principle is that the reciprocity between systematic and unsystematic variation is likely to hold good only as between species that live in environments of comparable richness.

On a grander scale, the collapse of the Soviet Union at the end of the 1980s has been interpreted by some as confirmation of F. A. Hayek's prediction that the variety (in Ashby's sense) of the command structure in a Soviet-style command economy could never equal the variety of a free-market economy. Ashby and Hayek were probably not aware of each other's work (one was a physician, the other a political theorist), although they were active at about the same time.

11. The Yale psychologist Clark L. Hull (1884-1952) was one of the most influential figures in the history of behaviorism and stimulus-response psychology. His best-known books are the *Principles of Behavior* (1943), a theoretical and experimental account of how rats learn about mazes; the *Mathematico-Deductive Theory of Rote Learning* (1940, with several other authors), an account of how people learn nonsense syllables; and *A Behavior System (1952),* an updated version of the *Principles.* A good secondary account of his work appears in Osgood (1953). Hull, although not a mathematician either by nature or training, was nevertheless inspired by what he took to be Newton's mathematico-deductive method. Unfortunately, he failed to grasp the idea that the power of mathematical theory lies in its economy: Newton needed only three laws to explain the movement of all physical bodies, but Hull and his students, at various times, came up with dozens of axioms and corollaries to explain a set of observations hardly more varied than the theoretical constructs they applied to them. The corpus of arduously obtained experimental results that flowed from Hullian theory has yielded very little of permanent value, in large measure because the experiments were tightly bound to the testing of particular hypotheses, most of which are now of only historical interest. Nevertheless, Hull and his followers raised a number of issues that continue to be important, such as the relative contributions of reinforcement (reward and punishment) and practice to be "strength" of a habit, the distinction between learning and performance, and the similarities and differences between different kinds of motivational variables such as deprivation (e.g., time without food) and incentive (e.g., the attractiveness of food).