TRACKING AND TRAILING:
ADAPTATION IN MOVEMENT STRATEGIES

Oliver G. Selfridge
Bolt Beranek and Newman Inc
Cambridge Mass

1 August 1978

^L

BOLT BERANEK AND NEWMAN INC

TRACKING AND TRAILING:
ADAPTATION IN MOVEMENT STRATEGIES

Oliver G. Selfridge
Bolt Beranek and Newman Inc
Cambridge Mass

1 August 1978

^L

TRACKING AND TRAILING: ADAPTATION IN MOVEMENT STRATEGIES

TABLE OF CONTENTS

TRACKING AND TRAILING: ADAPTATION IN MOVEMENT STRATEGIES

Foreword

This monograph forms part of a larger work concerned with the various expressions of adaptation in life and technology; and with understanding what adaptation is and how to use it. It is addressed to a wide audience in the biological and technological sciences, in the hope that the phenomena and mechanisms it discusses can be further studied.

Acknowledgments

This monograph was prepared specifically for the 1978 Cold Spring Harbor Summer Workshop on Biological Information Processing. One of its original stimuli was a 1976 NATO Institute on Man-Machine Interaction under the direction of Dr. Brian Shackel. Much of the material was discussed with a group of researchers at Otis Minot's in the summer of 1978. In addition, the manuscript has been read and criticized by many friends; MS, PS, Wallace Feurzeig, Steve Epstein, Michael Arbib, Neville Moray, and Scott Reed should be mentioned specifically.

O.G.S

TRACKING AND TRAILING: ADAPTATION IN MOVEMENT STRATEGIES

Introduction

> "Nature to [all creatures], without profusion, kind,
> The proper organ, proper powers assigned;
> Each seeming want compensated of course,
> Here with degree of swiftness, there of force
> All in exact proportion to the state ..."(1)

Scientists and engineers can, I suppose, take heart from Pope's optimism; and mathematicians can revel in his promise of linearity, "force/All in exact proportion." To discover what are the proper organs and the proper powers, and what has been the nature of the compensation, we need to deal with the complexity of organization and feedback. This may seem to fly in the face of Occam's razor, but simple strategies can produce complex behavior, and some simple behavior may in fact be the not so simple product of interacting strategies. Where there are common processes at work we should find and describe them.

This monograph studies the adaptive nature of tracking -- following tracks and trails. It asks what an organism needs to know in order to track, conjectures how the strategy of tracking can develop and evolve, and tries to show that complex behavior can follow in comprehensible ways from simpler behavior.

It also asks how an organism tracks its parameters and strategies of behavior to follow the moving optima of a changing environment full of other developing organisms, even as they respond to adaptations to them. In some way, I argue, this tracking by interacting organisms on a cosmic scale is responsible for much of the creativity we see.

Overview

In the first section various forms of guided movement are examined with respect to their effectiveness in heading the moving organisms to some kind of goal: for example, bacteria to their food, moths to their mates, ants to their nests. The path to be followed is a track or trail, and the organism must first be able to find it, and then to follow it; if the trail is lost, it must be recovered.

---

(1) Pope, Alexander, "An Essay of Man," from "The Best of Pope," ed. Sherburn, G., Ronald Press, N.Y., 1929, p. 121, lines 179-184.

In the second section, I discuss the equivalent actions with respect to the control devices themselves; in which, for example, a sensory organ or a steering wheel tracks its optimum setting. Evolution is a special case of constant tuning of fitness.

In the third section, the various aspects of trails and following them are considered in more detail, especially as they interact with the other perceptions and goals of the organisms. Trails are usually followed not just by using the scent of one kind of molecule, for example, but by using other clues as well, like gravity, sky polarization, and sight.

The fourth section discusses a number of related topics and examples in more detail, showing the relationships between the simple primitive strategies of trail following and the complex ones adopted by higher animals and man himself. I claim that many of the actions of intelligence are essentially tracking -- the optimum setting of an auditory filter, the control that moves a limb with an unknown load, the changing interactions of goals and subgoals.

1.0 Movement for a Purpose

Adaptation is such a fundamental feature of life that it has been considered as one of its distinctive features. There is observable in nature an enormous variation in the kinds and expressions of adaptive behavior.

This monograph discusses the process of adaptive behavior as expressed in strategies of movement by animals. In most animals, movement is concerned with nutrition, anti-nutrition, or procreation (by anti-nutrition I mean avoiding being something else's nutrition, by means of flight or schooling, for example). These higher purposes have guided the development of strategies of all kinds of movement by adapting those strategies to the changing expression of those purposes in highly complex environments. The purposes are usually implicit, especially in lower animals, and must be inferred from their behavior or deduced from analyzing the mechanisms of behavior. They express themselves in the behavior of single animals and interacting large groups, like mating swarms or packs of hounds pursuing game.

Here I want to discuss these strategies and how they have been adapted -- how the information the organisms get from their experiences change their behavior in future experiences. Much is known, of course, but far more is not; a great many questions seem not even to have been asked. Some conjectures and analyses will be related to models simulated on a computer.

I shall be constantly drawing a parallel between physical movement, with its adaptive characteristics, and movement in control or strategy space, with its adaptive characteristics. It seems to me that a control mechanism does not inherently care whether what it is controlling are organs of motion or other controls.

There is one special kind of adaptation that ought to have special attention, and that is the process of biological evolution. Evolution is presumably the adaptive mechanism by which adaptations arose. Analyses that treat evolution as a form of adaptation are rare; notable among them is a recent book by Holland.(1) I shall treat evolution separately in section 2.2.

---

(1) Holland (1975).

Most of the movement strategies I am concerned which are from phyla far below man. Indeed, the initial example is not even a metazoan. The starting point, therefore, is the adaptive behavior of creatures that occur lowest on the phylogenetic scale.

1.1 Elementary Strategies: Run and Twiddle (RT)

Most animals enjoy movement of some kind. Often this movement may appear to be random, but usually there are adaptive or purposeful elements to it. Movement itself may be directly adaptive, that is, responsive to a particular goal structure, like a search for food, or it may be regarded as part of a larger adaptive control strategy, like evolution or schooling..

In a paper in Nature in 1975, later summarized and extended in the Scientific American (1975), Berg at the University of Colorado discussed the motive habits of *Escherichia Coli*, a very common bacterium found in the human gut, and of *Salmonella Typhimurium*, another unpleasant bacterium. Each of those has a flagellum that rotates like a free shaft -- wheels and shafts had been thought to be absent from living organisms -- and each uses its flagellum in one of two modes of action: it is either moving, roughly in a straight line, or else it is 'twiddling', by which Berg means spinning around in the same place. The effect of twiddling is to change the direction of motion. The combination of the two modes makes the tracks of E. Coli look like steps in roughly straight lines connected by nodes which are the sites of twiddling.

In fact, there had been many observations before those of the late 1960's and 1970's that might have led to similar interpretations. Stanier and Cohen-Bazire described how purple bacteria (e.g., Bacterium photometricum) exhibit phototaxis:

> Their phototactic response consists of a sudden complete reversal of the direction of movement ('Schreckbewegung') when the light intensity is abruptly diminished ... In consequence, a local area of high light intensity acts as a light trap ...(1)

This kind of movement is more than random walk behavior,(2) which is not by itself adaptive. But Berg (and others) found out that the probability of twiddling varies with the environmental concentration of attractants (like light, or certain amino acids that might

---

(1) Stanier and Cohen-Bazire (1957), p. 84. The original discovery was apparently made by Engelmann (1882).

(2) Random walk is the term used by mathematical analysts to describe a step-by-step movement in which the steps have random characteristics, in length or direction.

indicate food). The initial observations concluded that the sequences of running and twiddling seemed to lead the bacterium into regions of high attractance, as if the creature could detect the gradient of the attractant concentration. But Berg showed that that probability does not depend on spatial distributions, but only on changes in time. If the attractant has been recently increasing in concentration, then the probability of twiddling is markedly less. That is, the bacterium appears to be dealing only with the temporal distribution of attractant as it perceived it, and not with the spatial distribution at all.

Bremermann (1974) reports on some analysis and simulations with a rather simplified version of the phenomenon. In particular, the behavior in his model is not probability driven, as Berg reports it is, but almost algorithmic:

> If the direction is one of increasing concentration, the organism continues until the concentration starts to decline... After sensing the decline, the organism tumbles again and as a result dashes off in a new random direction...
>
> If the direction is one of declining concentration ... then the swimming stops, the organism tumbles for a while and dashes off in a new random direction.(1)

The general effect is obvious - a tendency to keep climbing up a nutritional gradient, if there is one. We can think of this as the most simple control strategy - keep going the same way if things are getting better, and otherwise move around. By 'move around' is meant to vary the values of the control parameters implicitly or explicitly. The control parameters may be directions of motion, as with the bacteria; but they may also be other kinds of control.

When a drop of some attractant is placed on a cover slide or petri dish(2) filled with E. Coli, they soon cluster around, presumably tasting its delights, and growing enough to fission.

Computer simulation verifies that this simple strategy works, though in the real case there may be more to it than we can guess. I do

---

(1) Bremermann (1974), pp. 398-399. It might be noted here that in nature the change in direction is far from completely random, and certainly not uniformly distributed. See Berg (1975A) and Koshland (1957).

(2) The nature of the attractant need not concern us here, but typical ones represent sources of nutrient, like galactose, L-aspartate, and L-serine. Attractants are recognized directly, and not merely through metabolic action, for D-fucose is an attractant for, but not metabolizable by, E. Coli. Further descriptions may be found in Adler (1975).

not know of published work describing such simulation, beyond the modeling by Bremermann mentioned above; my own rough simulations (discussed in section 1.4) have used only two probabilities of twiddling, P1 and P2, and a threshold concentration at which the organism changes from P1 to P2. It may well be the case that the probability of twiddling is a continuous function of some recent history of attractant concentration, but the data do not seem to be rich enough to allow us to make that discrimination.

This strategy may be regarded as directly equivalent to the second kind of elementary continuous adaptive mechanism to be discussed in section 2.5. Essentially, the various possibilities are tried, and the successful ones tend to be kept up until they fail. If the possibilities are directions of movement, then the movement that results from this kind of control strategy can be considered adaptive.

I shall go further, and assert that this primitive control strategy forms one of the two basic strategies of movement and control that are found in living organisms. I shall refer to it as the RT strategy (for Run and Twiddle). The other strategy is more complex, and usually involves vision or hearing -- the long-range senses -- that is, seeing where one wants to go; a reasonable name for that is an AB strategy (from point A to point B).

Parenthetically, it seems not to have been pointed out how prevalent RT movement strategies are in life. In Scientific American(1) it was described how mosquitoes fly towards sources of blood using RT on the attractants of water vapor and carbon dioxide. For another example, Wigglesworth described how the louse adopts an RT strategy in order to approach some attractive nutrient.(2) One of his figures, taken from life, has a startling resemblance to figure 1.4A here, produced by the simulations described in section 1.4.

1.2 Other Expressions of RT -- The Use of Directionality

I suggest that many strategies of movement in animals will be found to be RT strategies. The essence of RT is determining a local gradient of attractiveness of an environment by taking sequential samples in time. In some sense, that is of course equivalent to taking samples simultaneously in time, but dispersed in space. If the organism is small, like a bacterium about a micron across, then its spatial sensitivity will not be very great, and it will no doubt be more profitable to integrate over time, as it were.(3) Since most motile protozoans move many body widths in a second, this

---

(1) See Wright (1975); and also Daykin (1967).
(2) Wigglesworth (1941), p. 105.
(3) In fact it would be essential, given the sensitivities reported in MacNab and Koshland (1972) and Adler (1969).

exchange of time and space affords a much larger baseline for measurement than a mere simultaneous spatial comparison.

Several researchers have asserted that some single cells can spatially derive directionality, but there is apparently doubt about exactly how. In a review on slime molds, Konijn remarks that "Probably an amoeba moves towards an aggregate by measuring a spatial gradient of the attractant over its total length (Kato, Losada, Nanjundiah and Konijn, unpublished)."(1) Wilkinson echoes this almost exactly in talking about leucocytes: "[leucocytes] probably sense gradients by detecting concentration differences across their own length."(2) There seems to be no published evidence for either claim.

As creatures become larger, the simultaneous sensing of attractants at comparatively remote sites enables the organism to respond faster and more accurately. The sensitivity in the dawn of life was probably primarily chemical; soon after primitive life became sensitive to photons, light directionality started to play a role in that sensitivity. Detecting this directionality requires distinct responsive sites, in effect. But since light quanta do not diffuse in the same way that molecules do, directions can be established in but a few microns; subject of course to the diffraction effects that are the optical equivalent of aperture-beamwidth relationships.

Even without directionality, lower animals (such as the purple bacteria mentioned above) use sensitivity to light to control their relationships with their environments. Nultsch describes a very simple example whose mechanism has still not been adequately understood or explained:

> Under diffuse light, blue-green algae of the family Oscillatoriaceae (Drews, 1957, 1959; Nultsch, 1961) and diatoms (Nultsch, 1956) display an alternating backward and forward movement without preferring any direction. On the onset of unilateral illumination the organisms do not change their movement direction actively. However, in individuals which are in a more or less parallel position to the light beam the movement toward the light source is prolonged while movement away is shortened resulting in positive photo-tropotaxis ... In organisms which are oriented perpendicularly to the light beam no phototactic effect can be observed. However, since the movement of the organisms is never completely straight, all individuals of a population come, sooner or later, for a long or short time into a position more or less 'parallel' to the direction of

---

(1) Konijn (1975), p. 148.
(2) Wilkinson (1975), p. 238.

light.(1)

This behavior is a kind of one-dimensional RT, with the additional complication that the step size is modified instead of the direction.

A partly directional intermediate step can be found in various maggots, which as a rule don't like light. In 1911, Mast described how a maggot, fleeing light, sweeps its head, which has a mildly directional light spot on its forehead, left and right so as to equalize the illumination on each sweep.(2) The directionality of the light spot is low, but it is aided by the opacity of the bulk of the body posterior.

Similarly, Euglena progresses towards a light source in a spiral, which moves the shadow of an opaque stigma around a light sensitive spot.(3) Euglena is mostly transparent, except for the embedded stigma, and if the shadow covers the spot, a correction is kicked in. The procedure is almost identical to that of conical scan radars:

> The beam was conically scanned by rotating the reflector about an axis which made an angle of 2.5° with its own. Commutators were mounted on the shaft of the antenna rotation mechanism so that signals from the upper half of the scan could be integrated and compared with those from the lower half; similarly, signals from the left half of the scan were compared with those from the right. The resulting error voltages were used to deflect the spot of a CRT used as an error indicator.(4)

This procedure represents an important modification of RT. The basic RT strategy invokes an essentially random twiddle(5) when things are not going as well as wanted. But here the geometry of the sensors allows the organism to make corrections in the right direction; figure 1.2A shows the operation in a highly schematic way, with one of Mast's maggots. The strength of each swing is very roughly proportional to the light that has just impinged on the light sensitive spot. This keeps the maggot pointing away from the light.

---

(1) Nultsch (1975), pp. 47-48.
(2) Mast (1911).
(3) The original observations were apparently first reported by Mast (1911); general reviews of the subject may be found in Feinleib and Curry (1971), and in Nultsch (1975), pp. 45-46.
(4) (Ridenour (1947), p. 207).
(5) Even if, as was remarked on page 3, the directions are not <u>uniformly</u> random.

The same thing happens with Euglena, except that, instead of a binary reciprocating sweep, there is a circular sweep resulting from the spiral or helical forward progress.

If one takes the radar analogy seriously, one might look for more deliberate scanning mechanisms in living beings; the scanning mentioned above involved essentially the whole beast, and the intent was apparently to track towards or away from some integration of the field of light as a whole. The copepod Copilia is reported to do that, at least for the females, which seem to have two eyes that scan towards the center of the beast.(1) I make some guesses about the operation of these scanning eyes in Appendix XXX.

Let me discuss the transverse sweep of the maggot, crustacean, or protozoan in more detail, because it is such a fundamental part of most strategies of following trails. The point of the sweep is to keep the peak of the concentration <u>in the middle of the sweep</u>. I defer discussing the problems of the noisiness of the signal, which can easily produce multiple peaks during a sweep, and of the actual detection of the signal that is the concentration.

For illustration, consider a clean signal and a smooth sweep. For the moment ignoring the forward progress, I can evaluate the sweep as the degree to which the peak does lie in the middle. If the value of the sweep drops, then some kind of change of parameter should take place -- a control twiddle -- to see if the value can be made to rise again. In fact, of course, that change is usually directed, by which side of the center the peak occurred, but that only makes the twiddle rather more constrained than it is with E. Coli. This twiddle is not a spatial twiddle, it should be understood, but a twiddle in 'control space,' so to speak. The way that twiddle is established and the parameters of its action may have a lot to do with the efficiency of tracking. It can <u>also</u> be regarded as some kind of spatial twiddle, for certain kinds of motion, and I shall argue that trail following and trail recovery procedures in moths have much more of an RT nature than the sweep of Mast's maggots or the rotation of Euglena. As I mentioned before, the use of parallel sensors, especially in larger animals, is only a convenience; to some extent, space and time are interchangeable in position finding.

Carlile makes a point of referring to the spatial sensitivities as "taxes" (plural of "taxis", a singular noun) and he regards them as essentially being more accurate than RT or klinokinesis.(2) The information about direction has to be gathered by sampling in space, whether simultaneously or not, and that accuracy is determined by

---

(1) Exner (1891), Wolken and Florida (1969), and Gregory et al. (1964). I am indebted to Professor Neville Moray for these references.
(2) Carlile (1975).

the amount of information that can be gathered in the ambient noise.

One can see that for a while now there has been a recognition that basic RT strategies have many modifications involving directionality. The confusing terminology is indicative of the complexity and variety of RT behavior. Thus Carlile describes how Fraenkel and Gunn(1) distinguish between 'kineses' and 'taxes':

> They classified the oriental [sic] movements of animals into kineses, taxes and transverse orientations ... Kineses are displayed when organisms are incapable of detecting the direction of a gradient or of the source of a stimulus. Instead, they respond to a change in the intensity of the stimulus by a changed rate of locomotion (orthokinesis) or turning (klinokinesis) in such a way as to lead to net movement towards or away from the source of the stimulus, even though the organism is incapable of precisely oriented movements.(2)

I emphasize these aspects and the terminology, because they refer to the earliest expression in life of what in man reaches its highest practice as free will. Berg's germs, acting individually entirely automatically, that is, not goal directed, are nevertheless in concert acting in a goal-seeking way. It must seem clear that a germ in which the probabilities have somehow been reversed (that is, which twiddles when things are getting good, and runs when they are not) is at an evolutionary disadvantage (see section 2.2). It must be stressed, however, that that conclusion depends on a particular kind of assumption about the distribution of nutrition. For example, if the clues to food do not form a smooth gradient, it might pay to twiddle on any contact with some attractant, knowing that nutrition is close by; that is, relying on a tight random walk to hit the real eatable concentrations. There are some botflies that apparently follow this strategy.(3)

Thus besides the basic RT strategy, there are modifications that incorporate uses of directionality of movement and of spatial distribution of attractant.

1.3 Search and Climb - Integration of Strategies

In this section, I shall concentrate on following the trail, in contrast to the previous section, where I emphasized the search or recovery aspects; that is, here I shall deal with the control of running instead of twiddling. The alternation between two different

---

(1) Fraenkel and Gunn (1961).
(2) Carlile (1975), p. 2.
(3) See a full discussion in Dethier (1976).

forms of "search" strategies, the first when there are no signs of attractants, the second when there are signs, is prevalent in the animal kingdom. (My emphasis on low forms of life, like germs, should not lead the reader to believe that these strategies are not biologically ubiquitous, for I believe they are.)

This kind of modification of RT strategies is extraordinarily prevalent: it is exemplified by a bloodhound. We can consider the general problem of the tracking of trails as a kind of anisotropic RT procedure. Researchers have been more impressed with the selectivity of dogs' olfactory faculties that with the actual process of tracking. Even the astronomer Carl Sagan is amazed by bloodhounds, which track a particular smell among "a bewildering and noisy background of other molecules."(1) Why 'bewildering'? In any case, silk moths do several orders of magnitude better. But Sagan never questions what <u>tracking</u> is. If a dog loses the scent, he enters a 'scent-recovery' mode until he finds it again. As he follows the trail, he adds a lateral jitter motion to his nose, to make sure that he is following the peak of the scent concentration. The case of hounds following the trail of, say, a fox is so interesting and complex that I will discuss it separately later, in section 4.2.

One might reasonably infer that silk moths do the same thing, while they are exercising their unbelievable sensitivity to the female pheromone (as low as one or two molecules a second!).(2) The procedure followed by a male silk moth as he goes a-courting is approximately:

1. Wait for arousal, which is triggered by the female pheromone.
2. Fly upwind, so long as the pheromone can be detected. Typically the path upwind is not straight, but zigzags. It is not known whether the zigzags tend to guide the insect towards the higher concentration of attractant, although it seems likely.(3)
3. If the trail is lost, fly back and forth across the wind until it is found again.(4) If it is not, give up, and return to step 1.
4. When the moth gets close to its target, other forms of behavior take over; sometimes the switch to visual guidance requires a continuing high concentration of attractant.(5)

Questions immediately present themselves from this account:

---

(1) Sagan (1977), p. 155. Many kinds of dogs, including bloodhounds, can follow the scent of individual people; see Kalmus (1955).
(2) See Schneider (1975) and Jacobson (1972).
(3) Daterman (1972).
(4) Kennedy and Marsh (1974).
(5) Brown (1972).

1. How does the insect know which way is upwind (or crosswind)? Answer: although it has been claimed that moths can fly in the blackest night, that is probably not true. Instead the moth orients his flight until the motion detectors in his visual system detect no sideways component.(1) If the mechanism of correcting his orienting is PT, it is a biased one, since the moth flies upwind and not downwind.
2. Why doesn't the insect just fly up the gradient of the attractant concentration? Answer: the gradient is far too small; it is corrupted by turbulence in the atmosphere; and furthermore, the molecular concentrations cause the signal from its antennae to arrive at the moth's brain in bursts -- each molecule can trigger an impulse or a burst.(2)
3. How does the moth know when it is close enough to mate, or to start to switch its approach procedures? Answer: when the attractant concentration exceeds some threshold, the insect becomes ready to switch tactics (e.g., to use sight, as above). Sometimes the forward progress diminishes also as the concentration rises.(3)
4. How often does the moth fail in following the pheromone trail upwind? Answer: nobody seems to know.
5. What happens when a male moth finds another male already there? Answer: The female stops emitting pheromone very quickly on being serviced by a male. Furthermore, males themselves emit a pheromone *in flagrante delicto* that seems to say "no room."(4)

As another instance, the following of trails by ants is well-known, but as far as I have been able to determine, the particular strategies have not been studied in detail.(5) Hangartner (1967) describes something rather like a lateral sweep across a solid pheromone trail, aided in the daylight by orienting with the aid of polarized sky light.

Michener and Michener claim that scent trails left by ants are directional:

> Of course, a simple, uniformly scented trail would give no guide to direction, and the ants must [sic] have some means of telling whether they are going or coming ... If you allow them to travel over a road made of narrow paper strips laid end to end ... [and then] turn one ... through 180° ... This will sometimes cause the

---

(1) Kennedy and Marsh (1974).
(2) Kaissling and Priener (1970)
(3) Farkas et al. (1974).
(4) Sebeok (1975) and Shorey (1976).
(5) Wilson (1971) reviews the older literature about the *mechanisms* of the trail-following: it is known that the antennae are the sensors, and that if they are crossed, for example, the ant has a much harder time. But the parameters of the strategies are unknown.

> ants from both directions to stop and show bewilderment ...(1)

Wilson does not believe that happens, at least in the ants he is concerned with:

> In a set of ingenious experiments, ... Hangartner (1967) was able to demonstrate the basis of osmotropotaxis in another trail-following ant species, *Lasius fuliginosus*. The method of following odor trails disclosed by these experiments makes it very unlikely that directional signals can be built into the trails. In other words, the odor streaks may or may not be tapered [I will return to this in section 3.4] or shaped in some way so as to point the way home -- as discovered for example in *Myrmica ruginodis* trails by Macgregor (1948) -- but it would be difficult for the follower ant to "read" this information.(2)

Wilson claims that ant trails typically have a duration of 100 seconds, so that a gradient of concentration along the path is obviously possible, although the geometry makes it seem extremely unlikely. In certain cases, where the length of the trails, etc., may be more favorable, it may be easier to credit gradients. Talbot (1967) maintains that trails laid by some slave-making ants can last for over half an hour.

The literature rarely reports actual data on polarization of trails. Willows writes:

> Information can be gleaned by a snail from a mucous trail ... [as] has been discovered by Hall (personal communication). When *Littorina* encounters a trail made by another member of its species, it turns to follow the trail. Surprisingly, however, it does not turn in both directions along the trail with the same likelihood. Instead, it turns to follow the trail-making snail in a statistically significant percentage ... Although the sensory mechanism underlying this process is unknown, observation ... suggests that an odour gradient along the 1-cm. distance sampled by the two tentacles may be used.(3)

Wells and Buckley (1972) find the same thing, although the mechanism for the directionality of the trail seems not to be known. No data are given.

---

(1) Michener and Michener (1951), pp. 17-18.
(2) Wilson (1971), p. 250.
(3) Willows (1973), p. 216.

The question of indicating direction on a trail may be examined theoretically. Regarding it as a problem in communication theory, one notices that it requires three symbols (not just two) to indicate direction -- a mere alternation, either of chemical or its concentration, is symmetrical, and a gap or length variation is another way of expressing or coding a third symbol. M. Minsky is preparing a general theory on this question (see Appendix 2).

Other arthropods follow trails in the same way. Hamner and Hamner (1977) show that certain shrimp follow trails of chemicals left in the water by descending food; they assert that they follow it down rather than up by being sensitive to gravity, but they present no data to support that assertion. They add:

> Probably ... the ability to track scent trails is widespread throughout planktonic taxa. Copepods (Calanopia elliptica) and a reef lagoon mysid ... both follow scent trails in our aquaria.(1)

Notice that the tracking here involves two dimensions rather than just one. That is, the shrimp cannot merely go back and forth across the trail or spoor, but have to check in both horizontal directions. It would be interesting to see whether the shrimp use conical scan, like the Euglena mentioned above, or some other tactic. It may also be that the sharply defined trails of amino acids afford a gradient easily detectable across the spread of the animals' sensors, allowing them to integrate over space as well as time. But there is always some diffusion, and, especially in the ocean, the disturbances due to waves and surges must ensure that trails have a comparatively short life. Other animals are known to follow scent trails in the open ocean, certain sharks in particular.(2)

The general procedure is an alternation of strategies:

1. Move in some direction, maintaining a jitter more or less at right angles to the direction of progress, to be sure of continuing to follow the maximum (usually of some chemical concentration).(3)

2. If, nevertheless, the concentration falls below some threshold, then perform a search for the trail, until the concentration rises above the threshold.

The anisotropy may be established in any of several ways: for ants, it may be primarily inertial -- that is, when they lose the trail, they look at right angles to it, at least initially, using some kind of orientational memory; silk moths track up wind, which they can

---

(1) Hamner and Hamner (1977), p. 883.
(2) Shulenberger, E. (1977).
(3) See Kennedy and Marsh (1974).

determine by sight; and shrimps track downwards, and search horizontally, by sensing gravity.

In these arguments there is an underlying assumption of continuity; the trail should have no breaks in it. In practice, all trails have breaks and irregularities of various sizes and frequencies. The trail following organism has to adjust his sensing filters and integration times so as to follow the trail over the breaks and irregularities; the exact ways that is done will have much to do with its success.

The complexities of argument that follow these simple strategies are typical of situations involving real beasts. There is a distinction elaborated on in the previous paragraphs that is worth making specific. I have drawn the distinction between hill-climbing and peak-tracking, as it were. But there are many instances where the primary job is to find a hill, for climbing it may be easy. In radar engineering, that is termed target acquisition, as opposed to radar tracking. There is no reason to suppose that the strategies for climbing a hill and tracking its (maybe moving) peak are necessarily any good for finding the hill in the first place.

There are at least three parts to trail-following:

1. Find the hill of attractant, that is, the trail; this can be termed trail acquisition.

2. Follow the trail by tracking the peak of the concentration, scanning across the trail, in either time or space.

3. Recover the trail when it is lost, by embarking on a search that usuallyprobably differs in detail if not kind from 1.

Higher forms of life often use a much more sophisticated form of movement using vision directly, with implied or explicit models of the environment. I have termed the long-range vision-guided strategy a point-A-to-point-B, or AB, strategy. Dogs, for example, usually switch between an RT strategy, used with their sense of smell, and an AB strategy when the prey comes into view.(1)

1.4 Run and Twiddle - Some Simulations

In this subsection I will report on some simple simulation experiments about RT strategies and their performance in a variety of environments. Consistent with an experimental approach, I will avoid analysis and theorems; not that they are impractically

---

(1) Though not always: note that retrievers as a class do exactly the opposite, using a general visual directioon and pinpointing with smell.

difficult, for they are not, but because the data do not allow us to measure the parameters of behavior and the environment with anything like enough accuracy, at least for the case of the protozoa. Furthermore, in this monograph the primary concern is in beginning with simple problems so as to be able to handle tasks that are beyond mathematical analyses.

The general effect of RT strategies is shown in figures 1.4A and 1.4B. The latter shows the first and final positions after 50 steps for 50 organisms, the former the actual tracks for 6 of them. Little is known of the particular dependency on parameter settings: in the figures, the two probabilities of twiddling were set at 0 and 1, the threshold at zero change in attractiveness, and the step at unity in a uniformly random direction. The operation of the program is shown in figure 1.4C.

Given:

```
Time t
Organism at location (X(t),Y(t))
         with speed  (Vx(t),Vy(t)), such that Vx**2+Vy**2=V**2
Organism Parameters
         Step Size V
         Threshold THETA
         Probability 1  P1
         Probability 2  P2
Organism States
         Twiddle
         Run
Attractant Source at (X0,Y0)
```

Procedure at Time t

```
In Twiddle: 1. Measure C(t)= 1/SQRT((X-X0)**2+(Y-Y0)**2).
            2. Get random angle A; set Vx=VcosA, Vy=VsinA.
            3. Change to Run.

In Run:     1. Make step X(t)=X(t-1)+Vx,
                         Y(t)=Y(t-1)+Vy.
            2. Measure C(t)= 1/SQRT((X-X0)**2+(Y-Y0)**2).
            3. If C(t)-C(t-1) > THETA,
               change to Twiddle, with probability P1;
               If C(t)-C(t-1) < THETA,
               change to Twiddle, with probability P2.
```

Figure 314C RT Strategy, Simulation Procedure

The point of these simple simulations is partly to verify a verbal argument -- that RT strategies do climb hills. Many questions can be raised about the assumptions made by the programmer. Do E. Coli really work with just two different probabilities, or is the probability of twiddling more or less a continuous function of the recent history? The latter would seem more likely on physiological grounds.

The RT procedure, even in this simple form, possesses a remarkable resistance to interference, both random and organized. The results of another simulation are shown in figure 1.4D; here the tracks between twiddles were not straight lines, but arcs of circles; the concentration of the attractant as the organisms perceived it had added to it a uniform random variable whose mean was ten times the difference in concentration experienced after a step in the right direction. Even so, as the figure shows, the 40 such organisms progressed uphill slowly and surely. In other words, the hill climbing power of these strategies is exceedingly <u>robust</u>, to use control theory jargon. The generalization above of RT to include running in paths that are not straight lines is not entirely frivolous. Pommerville has described the mating of fungal gametes (of <u>Allomyces macrogynus</u>); the male gamete is the seeker climbing the hill of concentration of 'sirenin,' emitted by the female. Some assymetry in the male causes it to travel in arcs of circles -- or so they appear under the microscope -- and the twiddling is much reduced in the presence of increasing concentrations of attractant from the much slower moving female.(1)

Nevertheless, the behavior has some degrees of sensitivity to the various parameters; and in section 2.2 I describe the result of a simple evolutionary experiment in which those parameters are changed.

Section 1.2 discussed the first simple modifications of RT into directionality, by introducing an anisotropy into the behavior. Here are some simple computer simulations showing some experiments with some of the parameters of such behavior.

I compared some simplified tracking strategies: a target moved in a sinusoidal curve leaving an attractant that disperses laterally and uniformly. Two organisms follow the track with different transverse sweeping policies. The first governs its step size by the recent change in concentration, like Nultsch' blue-green algae; the second governs the probability of reversing direction on the same grounds, as an expression of a pure 1-dimensional RT. If the perceived

---

(1) Pommerville (1976).

```
|
|                   ***************************************
|                   *  6 Organisms, 50 time steps         *
+                   *  * Starting points                  *
|                   *  Each integer marks a separate track *
|                   *  + Final points                     *
|                5  *  O Center of Hill                   *
|                5  ***************************************
|                *
+                5
|                5
|                5
|                5
|                5
|                5
+             +  5
|         5666   5             +3
|        66      5           3333       *
|         6      5         + 13333333333111     3
|        66      5       551111        1111333*3333
|       66       5   5555 111     111111
+       6        555+    +  101111
|       6      555          1  +2
|       6   4444                2222222222
|       6  144                           2
|     4 644                             22
|     446                               2
+   44444                             22
|     4444                             2
|       4                             2
|       6*                            2
|       66                           2
|       666                         22
+       6*                          2
|                                  22
|                                  2
|                                 22
|                                 2
|                                2
+                              * 2
|                              22
|
|
|
|
0---------+---------+---------+---------+---------+---------+
```

Axes in arbitrary units

Figure 1.4A 6 Organisms, RT Strategy

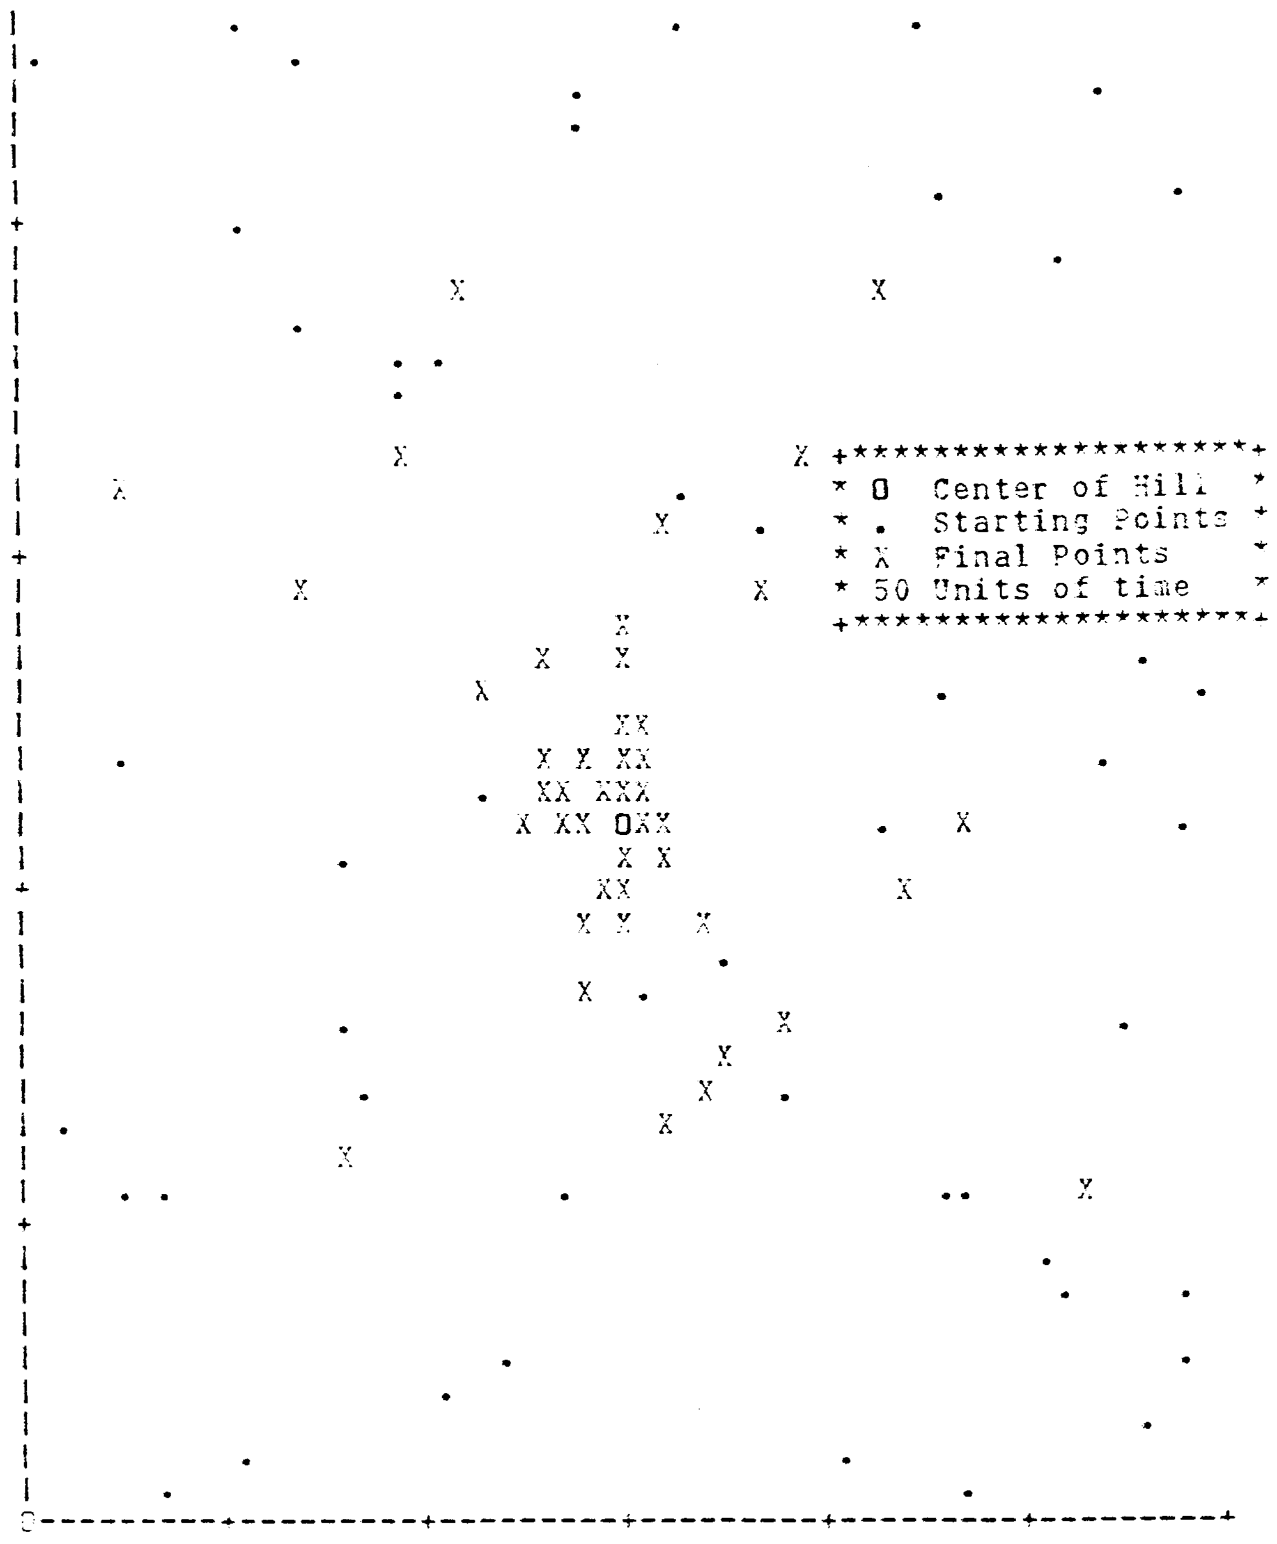

Figure 1.4B 50 Organisms, PT Strategy

```
+
|
|
|   Ordinate is Average of Y-values
|   Running Tracks are random circular arcs
|   Noise Mean is 10 times step size                    ***
|   40 Organisms, RT Strategy                        *** **
|   Uniform Slope up Y-Axis                      ****
|   Each X-step is 10 units of time           ****
|                                     *******  ***
+                                ******   * ****
|                              *****
|                          *****
|                   *********
|         *      ****   *
|******************
| *
|
|
|
0---------+---------+---------+---------+---------+---------+
   TIME --->
```

Figure 1.4D RT Strategy, 40 Organisms, Noisy System

---

concentration is not corrupted by interference or "noise," then the tracking seems to be fairly efficient, as is shown in figure 1.4E. Note that the target curves are the same, but displaced for the second target for clarity.

If the perceived concentration is corrupted, in our example by added Gaussian noise, the the tracking is still evident, as is shown in figure 1.4F. The mean of the added noise in that example is the same as the average step size.

One difference between the two strategies is that in the first, if the error is large, then the step size can be large too, while in the second the step size is always the same, and only the direction is changed. That is reflected in table 1.4A, which shows the average error for the two strategies as a function of the mean size of the added noise.

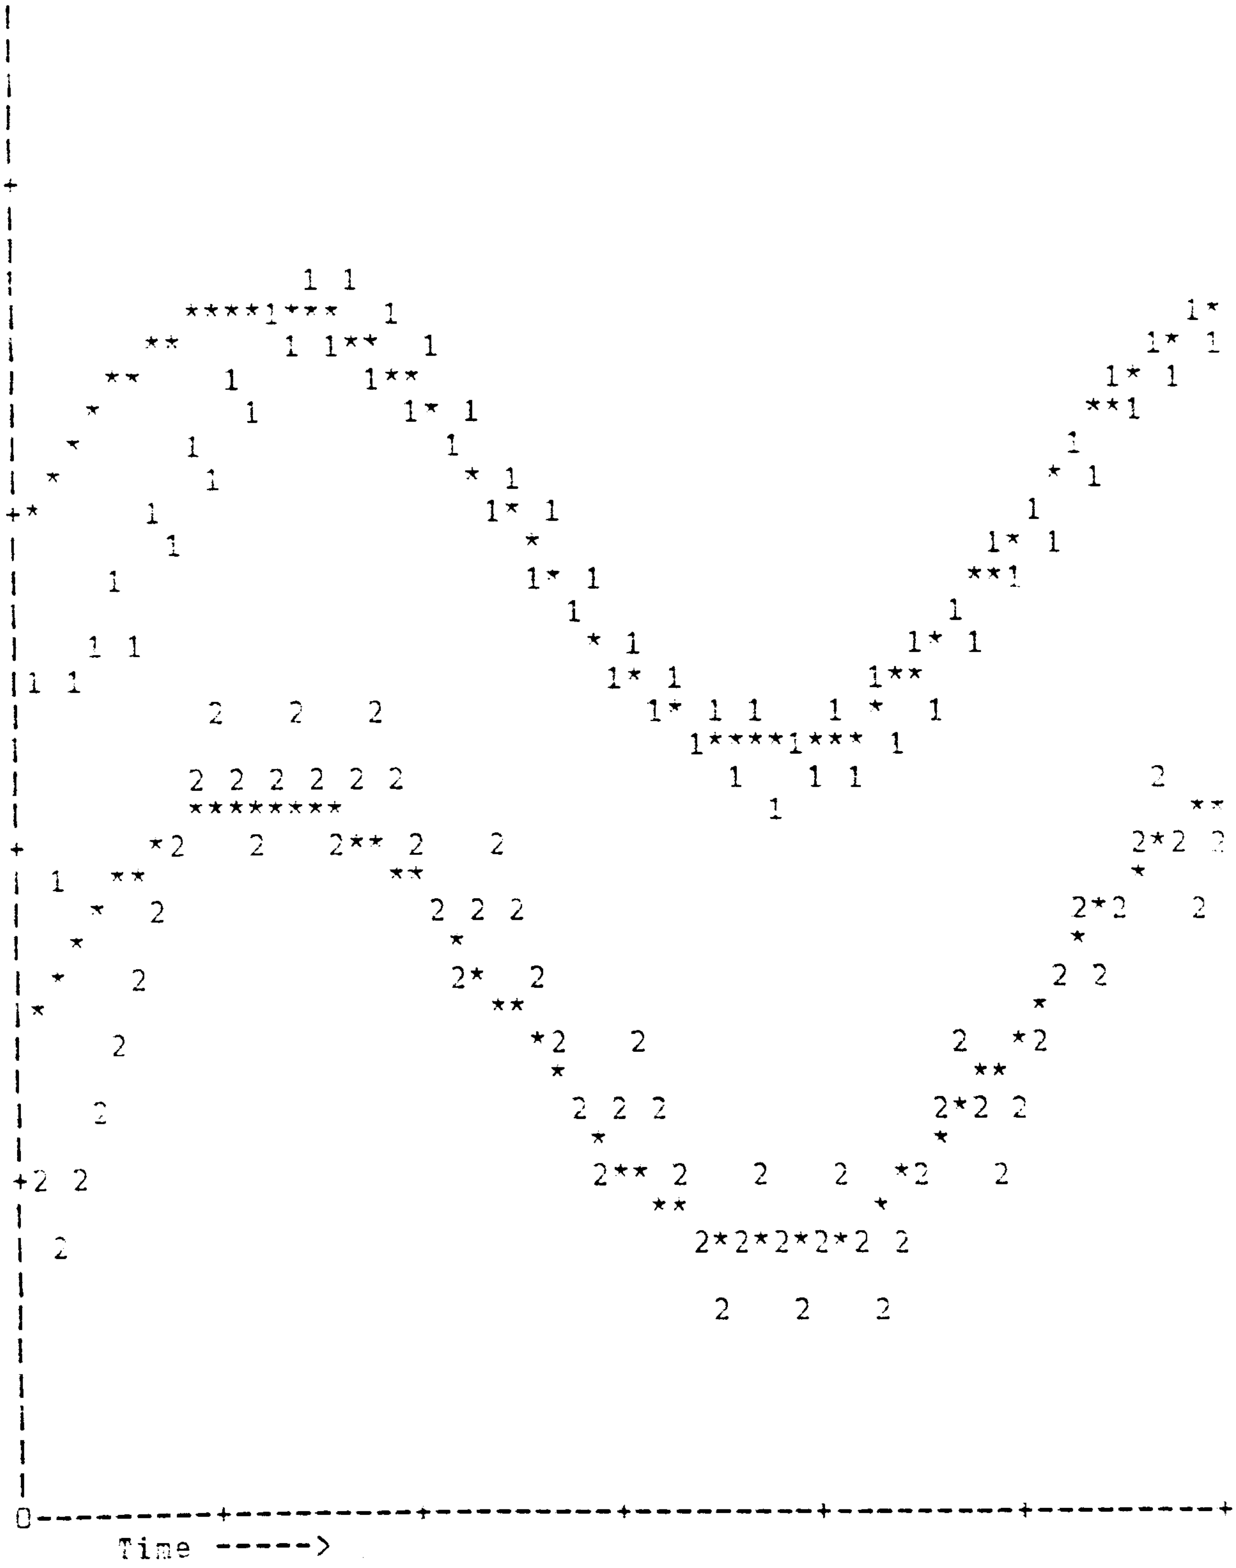

Figure 1.4E Noise-Free Tracking: Two Simulated Strategies

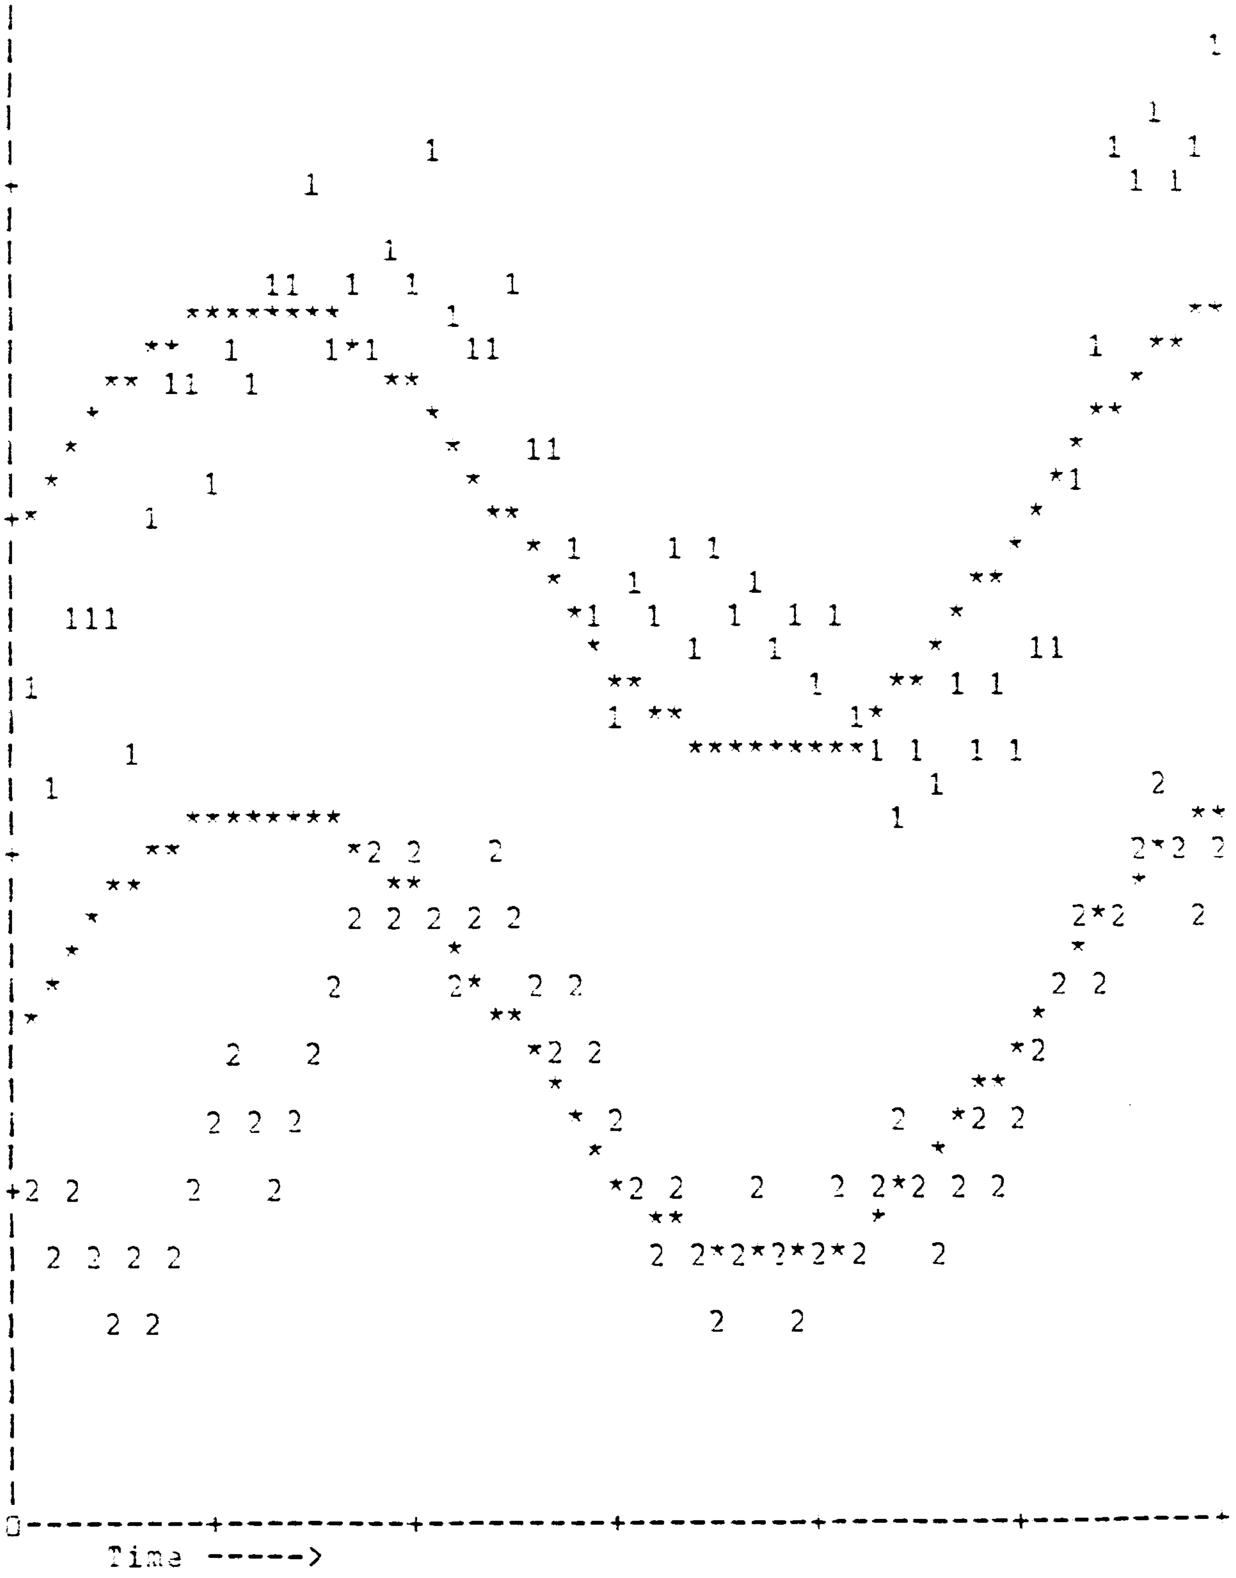

Figure 1.4F Noisy Tracking: Two Simulated Strategies

Mean Noise Amplitude	Strategy 1	Strategy 2
0.0	1.1	1.4
0.25	1.1	1.4
0.5	1.2	1.5
0.75	1.4	1.6
1.0	1.7	2.0
1.5	3.1	2.9
2.0	3.5	3.2
3.0	5.8	4.8
5.0	10.6	5.3
10.0	72.2	16.6

Table 1.4A Two Strategies: Average Error in Noise.

1.5 The Development of Strategies

It is very easy for people to inspect a creature's trail-following and hill climbing and attribute their management to intelligence in the creature. The true nature of its intelligence, and the diverse meanings we attach to the word, are beyond me here; but the intelligence surely depends on the interactions in wide domains of space and timei, and evolution is one of the prime mechanisms (see section 2.2). I believe anthropomorphic notions of intelligence are counter-productive. Lorenz writes:

> The process of 'learning by success' cannot evolve in unicellular or lower multicellular creatures which have no centralized nervous systems; for a system that is capable of exploiting the success or failure of a particular behavior pattern as a source of knowledge, and of using this knowledge as feedback to achieve an adaptive modification in the machinery of that pattern, obviously assumes the existence of various complex and highly organized subsystems.(1)

Apart from the gratuitous and suspect "obviously", it is not clear why inherently simple beasts cannot adapt their behavior, either by evolution or directly. Behavior may be simply directed by strategies, and one of my points here is that simple modifications of strategies can be easily responsive to need, and hence adaptive.

A simple adjustment of direction of movement can guide a bacterium to food. A simple and similar adjustment of strategies can guide any

---

(1) Lorenz (1973), p. 85.

organism to better behavior, no matter how complex the behavior. What does 'better' mean, and what are the ways of adjusting strategies? A strategy here is merely a formal set of rules for directing behavior.

The parameters of the organisms simulated in the previous section were four in number: probabilities P1 and P2 of twiddling when the change in perceived attractant concentration did or did not exceed the threshold, respectively; the threshold itself; and the distribution of step sizes, in the simulation uniform. Those were the explicit and obvious parameters. But there were a number of implicit parameters hidden in the program's structure. For example, in the simulation, twiddling produced a completely random new direction, contrary to what really happens in nature.(1) The change in behavior was governed by a sharp decision that some threshold had or had not been exceeded, but probably nature exhibits a more gradual change. Furthermore, the change in concentration was a straight subtraction with very little other time dependence; probably it would be more realistic to integrate the previous concentration with a kernel function of greater complexity.

---

(1) Berg (1975A), p. 391.

2.0 Movement and Tracking in Control Parameters

When we turn on a faucet or a car radio, we set the control to please us: if the water is too hot or cold, or the radio too loud or soft, we use corrections to set it where we want it. If we don't know which way to turn the control, we use what is obviously an RT technique to set it.

If other people are taking showers, or if the radio signal is fluctuating as we drive, then we track the optimum position of the control with small corrections every now and then. All living creatures use such techniques for control in one way or another.

In this section, I discuss the ways in which some simple controls can be classified and examined. A control is part of a feedback loop itself, in nature, and I shall try to show parallels between what is controlled in a movement, and what is controlled in a control.

2.1 Classical Adaptation

One of the primary controls of sensory systems sets the sensitivity in the range of input intensity being experienced at the current moment. In this section I try to relate the term 'adaptation' as I use it to the way it is used in the fields of psychophysics and neurophysiology. In those fields, it has a different meaning from that used here.

It is of course our common experience that we become accustomed to changing environments with grace and ease with all of our senses; a steady loud noise becomes less oppressive, a warm swimming pool becomes tolerably cool, offensive smells seem to vanish quickly, and so on. Adaptation in this sense thus serves to track the appropriate level of sensory input handling: we need to know the significance of our senses' messages, far more than, say, their absolute level. Since shapes are more significant to people than overall brightness, in dim light the human visual system, for example, becomes a great deal more sensitive to light, so as to be able to distinguish shapes at light levels which might otherwise be too low for the sensitivities that work in sunlight.

Geldard(1) attributes this first use of the word adaptation to

---

(1) Geldard (1953), p. 35.

Aubert (1865), who described adaptation to the dark. In this paper I use the term in the context of some overall evaluation of worth of behavior, so that I say that a system adapts when it alters its behavior (or position, or whatever) so as to improve its relation with its environment. In psychology and neurophysiology it is used with a much more limited meaning, generally without the imputation of any purpose or design; not only without purpose, but without any judgment of improvement. Indeed, Hilgard and Marquis define adaptation as 'The decrement in a response which is a consequence of its repeated elicitation.'(1) Thus, eyes 'adapt' to the dark by increasing their sensitivity many times. On the other hand, if the light intensity rises, then the output of the individual retinal ganglion cells rises, but then falls.

Such a phenomenon is well-nigh universal in nervous systems:

> If the stimulus is repeated in a regular, monotonous series, the evoked response diminishes to a low, stable level, often not even detectable. The response has habituated.(2)

How the mechanisms for this kind of adaptation have been half uncovered is a long, fascinating, and still developing story of research, full of excitement and heroes. One of them is clearly Adrian, who explored how nerves work:

> E.D. Adrian discovered this coding of stimulus intensity into impulse frequency in the 1920's at Cambridge. Working at first with apparatus that now seems rather primitive and cumbersome, he recorded from several kinds of sensory nerve fibre in different animals, while applying appropriate stimuli to the receptors. He stretched the receptors in the muscles of frogs, applied pressure to receptors in the paws of cats and displayed lights (and even himself) in front of the eyes of conger eels. What he recorded in each case was trains of impulses.
>
> He established that the codes in single axons from receptors and those sent to muscles were exactly the same. He also found two major varieties of receptor. One type responded to a stimulus by firing repetitively and steadily, reflecting the continued presence of the stimulus by a continuous sequence of impulses. In other types of receptor when the stimulus appeared (or disappeared) it was reported by a burst of impulses which died away. In these fibres the signal was said to undergo adaptation. In fact, even continuously

---

(1) Hilgard and Marquis (1961), p. 477.
(2) Hilgard and Bower, (1966), p. 443.

> responding receptors adapt partially in that the maintained presence of a stimulus is reported by a lower rate of firing than its arrival.(1)

Adrian himself describes it this way:

> All living cells tend to come into equilibrium with their surroundings ... and therefore a sudden change will cause far more disturbance than one which is established gradually. The cells of the nervous system show this property in a high degree. A slowly increasing current passed into a nerve-fibre will not excite it, whereas a sudden increase to a much smaller final value will do so; and at the other end of the scale we must all be aware of the way in which the whole organism can adapt its life to discomforts ... the adaptation of the sense organs is therefore an example of a general property of cells and organisms... in the eye particularly, and in all sense organs to some extent, the range of stimuli which can be distinguished is greatly increased by the power of adaptation, although the ability to signal absolute intensities is lost.(2)

Adaptation in this sense applies especially to perceptual processes; since there is no feedback loop, such a control is called an open loop process in control theory. It may appear, however, that the adaptation is for the purpose of restoring accuracy of perception of changes in the environment. Thus Wallach says:

> A modification of the perceptual processes ... , then, serves to compensate for the misinformation reaching the eye or ear ... There are two types of explanation ...:
>
> 1. Some sensory indication of the misinformation ... must be available to the subject, or
> 2. The misinformation must concern sensory data that are connected with a normative perceptual process, i.e., one that tends toward some distinct state ...(3)

It is widely recognized that this kind of adaptation is not unarguably 'adaptive,' in itself. Indeed, it is sometimes interchangeably termed 'habituation', 'equilibration', or 'perceptual compensation.' But of course it represents an adaptive evolutionary answer to a whole set of problems.

---

(1) Oatley (1972), pp. 58-59.
(2) Adrian (1946), pp. 83-84.
(3) Wallach (1968).

Rock considers this question quite explicitly:

> In the case of classical sensory adaptation, there is a gradual diminution of sensation based on a fatiguing of sensory cells stimulated. In the case of an optically altered image, what is presumed to occur is not so much a diminution of sensation as diminution of the perceived distortion. Classical adaptation is perhaps more nearly the same as the kind of normalization effects studied by Gibson ... The term "adaptation," therefore, might be appropriate in the sense of change in the direction of effectively coping with the environment.(1)

There is some overlap, therefore, in the usage of the term 'adaptation' in psychophysics and in this monograph; but here the roles of sensing the favorable directions of change and exercising the control are stressed over the actual sensory distortions. Of course one can see habituation as adaptation where the valued relation to the environment is maintaining some sort of status quo.

There is no doubt that in the larger evolutionary sense habituation is adaptive behavior (see section 3.5). Dyal (1973) writes:

> A worm which failed to habituate its reflexive withdrawal response would run the risk of starvation and [perhaps] suffocation. Habituation of the withdrawal response is thus a biologically adaptive technique to reconcile the need to eat against the threat of being eaten.(2)

In the same way,

> If the [mosquito] pupa is not disturbed, it remains quietly floating just below the surface [of the water], but if a shadow falls on it, it will quickly dive towards the bottom. If shadows repeatedly fall on it, however, it quickly adapts and stops diving.(3)

Presumably, a single shadow might be some predator, while repeated shadows probably represent a changing environment that does not justify the metabolic resources needed to dive continually to the bottom. That is, the adaptation might be considered as representing some sort of judgment of cost-effectiveness by the beast.

In fact, there are many researchers(4) who regard habituation as a

---

(1) Rock (1966), pp. 13-14.
(2) Dyal (1973), p. 248.
(3) Jones (1978), p. 138.
(4) See XXX and XXX.

kind of first step towards learning in its full glory. Our emphasis here is on habituation (and the other steps) considered as reactions of control loops, that is, of adaptation. Razran (1971) regards it as almost fundamental to life:

> Habituation is ... the lowest ... level of learning, functioning and prevailing in its functioning at the very dawn of life and under the most drastic experimental restriction of its neural substratum in higher animals."(1)

Further discussion and application of habituation as adaptation, and examples of its simulation in simple models will be found in section 2.4.

2.2 Adaptive Control of Behavior - Evolution

Regardless of whether individual organisms act adaptively or not in a behavioral sense, it is clear that they do in a genetic sense. The nature of this control is peculiarly hard to study, because, however much we are finding out about the genetic code, we know almost nothing of how that code is translated into behavior. It must be understood that genetic adaptation is mediated by biochemical mechanisms that we know almost nothing about in this context. Optimistic claims have been made: for example, Alloway says:

> It [appears] likely that the molecular processes underlying this form of learning may soon be discovered ...(2)

"A chicken is one egg's way of making another egg"(3) is a facetious but deep truth. Especially in lower animals, the adaptive control loops at the species level are exercised genetically. It is only in animals that take care of their young that there is much continuity of learned experience apart from the genetic code.

Evolutionary adaptation is a very attractive way of achieving progress, if only because it has already worked, albeit slowly. Fogel et al. (1966) had an intriguing title -- "Artificial Intelligence Through Simulated Evolution" -- but the results were unimpressive. Evolution is much more complicated a process than is commonly understood; the mathematics is difficult, and the applications are hard. Early promises(4) seem not to have been kept.

---

(1) Razran (1971), p. 45.
(2) Alloway (1973), p. 157. The form of learning referred to is conditioning of leg flexion in the cockroach Periplaneta Americana.
(3) Samuel Butler XXX
(4) See Box (1957).

The basic concepts and credibility of evolution are well established. Young (1971) comments:

> 4. The efficiency of selective death for adaptation.
>
> [There needs to be] an elaborate method for monitoring prospective genetic changes in order to determine their viability. There is indeed such a monitoring device, but it is of supreme simplicity. Any detailed plan for testing the probable effects of change in instruction would surely run into immense complications. But a perfect and simple test is always available -- viability ... all that are unsuitable will die.(1)

It is possible with some fairly simple genetic simulations to see some of the phenomena described in evolutionary studies. The bacterial behavior discussed in section 1.1 depends a great deal on the parameters, even if their exact nature and values are not known. One can easily enough imagine that they were set by evolution in random walk processes, rather like the PT strategies of the motion that they govern. The parameters need not be thought of as modifiable only by mutation or genetic variation, although that will be the basis for our abstract model. Spudich and Koshland (1976) even claim that some differences among bacteria are reflections of "non-genetic individuality," caused, say, by a Poisson distribution in the number of molecules of key behavioral enzymes. They argue that this may help preserve the species through short-lived violent changes in the environment better than genetic variation:

> Genetic variation would not accomplish the same purpose, since selection in a rare toxic situation would produce a mutant poorly adapted to the more common conditions of the environment. The population, however, was selected over evolutionary time for survival in all of the widely varying conditions of the environment. Thus non-genetic variability would be a preferred mechanism for accommodation to random fluctuations in the environment, and genetic variability the preferred mechanism for accommodation to long-lasting environmental changes.(2)

Presumably implicit in the argument is the notion that bacteria are too primitive to have been able to develop the hierarchical adaptive controls that might enable the single organism to alter its behavior in the required way.

But presumably, at least some aspects of the mechanism are genetically controllable, and those can be modeled very simply.

---

(1) Young (1971), p.333.
(2) Spudich and Koshland (1976), p. 470.

Suppose that P1 is the probability per unit time of twiddling when the local changes in attractiveness exceeds some threshold; and that P2 is the corresponding probability when it doesn't.

Figure 2.2A shows the results of a simple experiment along those lines. A population of fourteen organisms was modeled, consuming nutrition according to their distance from a source, dying when starving, and fissioning when replete. The two axes show P1 and P2, the probabilities of twiddling when the environment has been, and has not been, favorable, respectively. The offspring of the organisms have parameters, that is, P1 and P2, perturbed from their parents'. It will be noted that the drift is clearly and unmistakably towards the upper left, where a verbal analysis would suggest. The organisms in the lower right had less success in finding food, and hence fewer offspring. The figure shows ten generations.

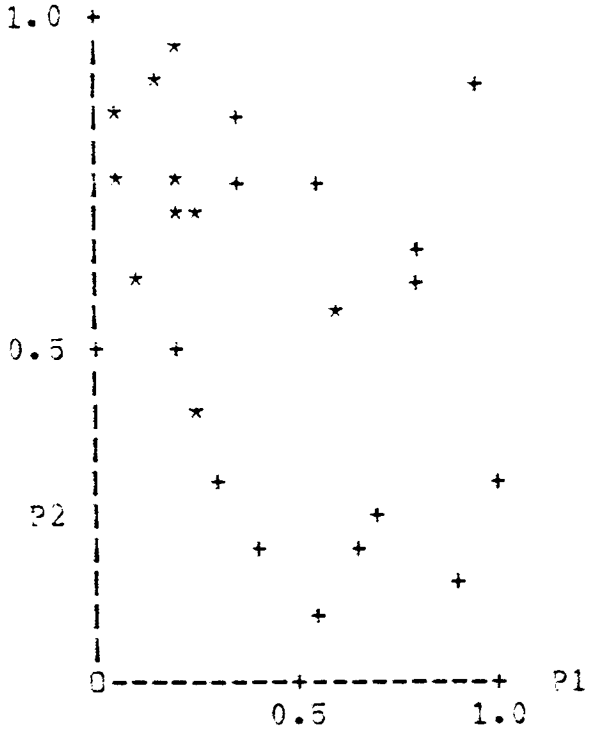

Figure 2.2A (see text)
\+ Ancestors
* Descendants

That example must be considered far more as an example than as an experiment. The perturbations in the probabilities that governed the behavior were small; the threshold of changing from P2 to P1 was set at zero, and there were other assumptions. Thus, in real life, it is not at all clear that 'small' mutations lead to small changes in the probabilities. But it is probably important to simulate enough examples to check the consistency and applicability of one's reasoning. If one can check against real data and real organisms, so much the better.

These points have been taken seriously among researchers in industrial control. In Box and Draper (1969), for example, the principles of evolutionary development are taken seriously enough, but in some sense the emphasis is on local planning and the design of efficient experiments, rather than also on the overall goals and directions. Actual industrial applications still seem to be infrequent. After a brave beginning, which stresses natural variation and the selection of favorable variations, the actual technique recommended to industrial engineers amounts to a local and few-dimensional gradient measurement; which is coupled with stress on statistics and the human relations difficulty of persuading real factory management to take it all seriously in practice.

The geneticists and the mathematical geneticists are more interested in modeling existing evolutionary processes than in using their mechanisms in the way we are here. Holland (1975) proposes an analysis of an evolutionary model, in which some fairly deep examination of the interaction of mechanisms can be carried out. The work is handicapped by a concentration on theorems and mathematical arguments, and by a complete lack of real models or simulations. Nevertheless, he discusses many of what now seem to be the crucial issues, and proposes important conceptual steps beyond those others researchers have taken. In a sense he is taking evolution seriously and constructively. Stebbins puts it:

> "Natural selection directs evolution ... by sorting new adaptive combinations out of a gene pool of variability ... built up ... over many generations. For the most part Darwin's concept ... fits in with our modern concept of interaction between evolutionary processes, because each new adaptive combination is a modification of an adaptation to a previous environment."(1)

Holland also complains appropriately about the difficulty of "apportionment of credit" to alleles or higher sub-structures.

> The greatest complexities come about because the effects of different enzymes are not additive - a phenomenon known as <u>epistasis</u> ... The main point is that the effect of each allele depends strongly upon what other alleles are present and small changes can often produce large effects ...
>
> Because of epistasis there is no simple way to apportion credit to individual alleles for the performance of the resulting phenotype. What may be a good allele when coordinated with an appropriate set of alleles for other genes, can be disastrous in a different genetic context. Thus adaptation cannot be accomplished by selecting among the alleles for one gene independently of what alleles appear for other genes.(2)

As Mayr puts it, "The fitness of a gene thus depends on and is controlled by the totality of its genetic background."(3) This is an absolutely key point in the practical analysis of complexity and the application of adaptive techniques. Indeed, it makes a mockery of the use of the word "complexity" in the field called "complexity theory," where the contribution of the simple components can be

---

(1) Stebbins (1966).
(2) Holland (1975), p. 10.
(3) Mayr (1963), p 295.

assigned with some uniformity and usefulness.(1)

It should be pointed out that 'survival,' or natural selection, is not a terribly directed principle; that is, the details of what has evolved are not pre-established or preordained. There is some feeling that evolution is an unreliable process, and that feeling can be traced back to the enormous public outcry at Darwin and his theory. Recently, Jacob (1977) has updated the argument:

> In contrast to ... evolution, the engineer works *according to a preconceived plan* ... [so that] objects produced by the engineer ... approach the ... perfection made possible by the technology. In contrast, evolution is far from perfection. Darwin emphasizes ... the structural or functional imperfections of the living world...
>
> Natural selection has no analogy with any aspect of human behavior ... natural selection does not work as an engineer works. It works like a tinkerer -- a tinkerer *who does not know exactly what he is going to do ... like a tinkerer who uses everything at his disposal*.[my stress](2)

To me, those words sound deprecatory in tone; but in content they carry encouragement and inspiration. I will argue that 'tinkering' is a needed and necessary source of richness and flexibility. Tinkering, one might say, is a form of conceptual twiddling.

In fact, of course, the real engineer, like the real programmer, rarely follows precisely his preconceived plan; whether or not current engineering and programming have too much tinkering, they certainly have their share of it. Moreover, the real engineer does not generate his plan in a preconceived 'correct' way either. Furthermore, the plans of any real engineer are not generated in a preconceived 'correct' way either.

One might describe the difference between debugging an information system and evolving one as the difference between tuning a system to meet a previously specified goal and adapting it to a new or newly realized goal or subgoal. I shall argue that "constructive serendipity" arises most readily when systems are built not with an over-tidy design, but with certain loose ends and discernible but ineffective interactions; so that there are always alternative modes of action and feedback available for adaptation as circumstances change.

---

(1) See Pippenger (1978).
(2) Jacob (1977), p. 1163.

In organic evolution, it is obviously hard to get a quantitative measure of progress, although rates of change can sometimes be usefully described by evaluating amino acid populations in the DNA.

* * * * * *

> Is there evil but on earth? or pain in every peopled sphere?
> Well, be grateful for the sounding watchword 'Evolution' here,
> Evolution ever climbing after some ideal good ...(1)

2.3 Mechanisms of Implementation

In comparing adaptation in life and in computer systems with their simulations of life, we are constantly reminded of the difference in how they do everything. Computer programs have an algorithmic simplicity and purity that make design comparatively direct; analysis often feasible; and modularity desirable and usually possible. Living systems are mostly modular only in a conceptual sense, since resources are always shared and metabolically dear; analysis almost never can be accurate, to say nothing of both accurate and precise; and design is so hard and uncertain as to require hundreds of millions of years without even a set of requirements.

The living system is thus not a tidy and self-contained example. But the untidiness, with its loose ends and interactions, gives the system a richness and flexibility of possibilities that can always surprise us. The mechanisms of bacterial hill climbing -- the RT strategy -- depend on molecular concentrations that affect the cell membranes and that thus interact with every process of the organism. In real life, for example, such chemicals may be dyes whose breakdown or enzymatic action may be affected by light density; and gratis, as it were, the bacteria have a sensitivity to light that can also be manipulated profitably. Variations in such concentrations can provide a sensitivity with a directional component, and so on. Thus the actual mechanisms of implementation in living system lead into possibilities for extensions of control, both inside and outside the actual modality originally practiced.

Evidence is arising how widespread this is in nature:

> So we start with an organism that senses temperature
> and has some form of thermoregulatory behavior... later

---

(1) Alfred, Lord Tennyson, "Locksley Hall, Sixty Years After," lines 197-199. But note Carruth's response in 1908: "Some call it evolution, and others call it God." ("Each in His Own Tongue, and Other Poems")

> on it develops another form of thermoregulatory behavior ... in parallel with the first. In changing its posture it accidentally develops a system for producing heat. Eventually the temperature sensors gain control over this new form of heat production and that is yet another integrating system... The same principle of new controls over an already existing mechanism for a new function can be used to understand the nervous organization of many forms of motivated behavior.(1)

That is rare, apparently, with computer systems. Living systems exhibit what we might term constructive serendipity; but most computer programmers do not find that new possibilities arise accidentally. Partly this is because a program is designed and written for a given purpose, and the specifications are generally limited to that purpose. 'Clean design' and economy of coding are not conducive to the loose ends and interaction of mechanism that we have argued enrich the evolutionary possibilities of living systems. This is of course not good or bad in itself; but it helps to explain why computer programs are often very hard to modify so that they have new faculties, or satisfy new ends.

The realities of the physical world are confusing enough to its inhabitants; trying to analyze them to determine the actual control variables may be harder. There are traps everywhere, often anthropomorphic ones, as we mentioned. For example, Bretherton and Rothschild (1961) claim that a range of mammalian spermatozoa exhibit 'rheotaxis' by swimming upstream in a uniform current. In human terms, swimming upstream is natural enough. But for protozoa, there is really no conceivable way for the organism to tell which way a uniform current flows if it is embedded in it. If there is observable behavior, it is not in response to the uniform current.

Similarly, Nultsch describes research of the last century:

> A different type of reaction was described by Engelmann (1882). He found that in the so-called _Bacterium photometricum_, probably a species of the purple sulphur bacterium genus _Chromatium_, an abrupt backward movement was caused by a sudden decrease in light intensity. Since the organisms behaved as if they were frightened, he called this reaction a phobic response ('Schreckbewegung').(2)

Other examples abound. The naturalist Lorenz is notorious in imputing to animals human ways of dealing with intent. Thus,

---

(1) Satinoff (1978), p. 21.
(2) Nultsch (1975), p. 30.

> "If an animal crawls from a cold to a warm zone, we can never determine objectively whether it is _escaping from the cold_ or _looking for the warmth_ ... In both cases, it is clear that the organism has been _irritated_." [my stress](1)

The objection to 'irritated' is not that Lorenz means it in the human sense rather than the biological, but that it begs the question.

I regard these points as important, because they are the natural science equivalent of certain fallacious views widely held in computer science.(2) It is of course partly my hope here that the control model aimed at elucidating adaptive mechanisms may help to provide insights into the morphology of behavioral phenomenology, and vice versa.

2.4 Adaptation, Habituation, and Learning

The reader will no doubt have already noticed the interesting parallel between the habituation discussed in section 2.1 and some of the attributes about guided movement in section 1. In section 1, we remarked that spatial hill-climbing probably made use originally of temporal (sequential) gradient detection; switching to parallel sampling techniques when organisms and sensitivities became large enough. In the same way, the "center-surround" process is a later spatial development of the notion that the temporal change in general is more important than its level on an absolute scale. It is time to examine these notions and their applications with more thoroughness.

In section 1.1, we described and discussed the adaptive behavior of microbes, which adapt their movement to the distribution of nutrition. It seems most appropriate that one think of microbes as having but a single strategy of behavior. For higher animals it is increasingly hard to think of the beast as following but a single strategy. It makes more sense, and is more convenient, as in the later parts of section 1.1, to think also of the behavior itself as adapting and changing.

In the last few decades, it has become almost a matter of competition among biologists to recognize and describe adaptation of behavior in phylogenetically lower and lower animals. For example, in 1952, Gelber asserted that Paramecia could be conditioned. That paper, and the subsequent ones she published, sparked a vigorous controversy over the next decade and a half, which is still not adequately resolved. Corning and Von Burg (1973) conclude:

---

(1) Lorenz (1973), p. 91.
(2) For a critique, see McDermott (1976).

> Attempts at demonstrating simple forms of learning ... [in protozoa] have been only moderately convincing, and associative learning remains a highly questionable phenomenon.(1)

But for our purposes this is an idle debate. It does not matter here where lines can be drawn, or even if they can be drawn at all.

There is clearly some progression of adaptive capabilities, in which responses are altered by exposure to more subtle clues, over longer time periods, and with more tenuous and global relationships to the stimuli. But in fact the concept of 'response' is itself equivocal. Shylock asks:

> ... If you prick us, do we not bleed? if you tickle
> us, do we not laugh? if you poison us, do we not die?
> and if you wrong us, shall we not revenge?(2)

Those four instances -- are they _responses_ in the behavioral sense? Shakespeare covers a nice spectrum with those four, each different in level and meaning, though for a sequence of phylogenetic behavioral consistency the second and third should be exchanged.

There seems to be little agreement among researchers about the meaning of these terms and how they should be applied. An extreme example is Bush and Mosteller:

> Any systematic change in behavior [is] learning whether or not the change is adaptive, desirable for certain purposes, or in accordance with any other such criteria ...(3)

The Shakespeare instances above are not meant to be merely dramatic. They illustrate important differences among reactions, as opposed to responses. It seems clear that a stimulus and a response must be elements of control loops to be meaningful, at least in the sense of the discussion here. After all, to extend Shakespeare's first instance, if an abominable pickpocket is to be punished by having his digits excised one at a time, nobody would suppose that ringing a bell before each severing cut would come to cause a finger to drop off at the sound of a bell without the cut.(4)

---

(1) Corning and Von Burg (1973), p. 117.

(2) Merchant of Venice, III, i, 60-62.

(3) Bush and Mosteller (1955) p. XXX.

(4) There _are_ beasts that respond to suitable provocation by shedding appendages; e.g., crabs their claws and legs, some lizards their tails. I do not know whether these responses can be conditioned; the difficulty would be in conducting enough conditioning trials, I suppose.

2.5 Some Theoretical Aspects of Adaptive Control

```
                +----------------+        +----------------+
                |                |        |                |
                |    CONTROL     |<-------|    SUCCESS     |
                |                |        |   EVALUATION   |
                |                |        |                |
                +----------------+        +----------------+
                        |                          |
                        |                          |
        Other   +----------------+                 |
     Factors-->|                 |  PERFORMANCE    |
                |     SYSTEM     |---------------->+
   External-->|                  |
     Stimuli    +----------------+
```

Figure 2.5A Single Adapting Unit; Conceptual Structure

The essence of *control* is the feedback loop, whose simplicity may be seen in figure 2.5A. The three blocks of figure 2A each need some separate description. The control itself is just that; it is the executor of some control plan built into it, either explicitly or implicitly. It is also itself possibly subject to controls of various kinds. Its assigned task is to manipulate its output, the *control setting*, so as, for example, to maximize *success* or to minimize *error*. (None of those terms is capable of precise definition at this point in the discussion.)

The system itself represents the interaction of the control with the entire environment. Note that we do not in general require *any* particular characteristics of the system; certainly not linearity, nor any of the other restrictions that are to be seen in current analyses of control theory.

The nature of the third block, which evaluates performance, is often confusing. The evaluation of success is the mediator of feedback, " ... in order to return a report to the one that sent him."(1) In simple systems, its output may be an actual error signal, as in the governor, or a feeling of hunger, as in people. In engineering applications, it represents a crucial decision. As Kalman points out:

> The choice of an optimal controller is largely
> arbitrary, depending on what aspect of system response

---

(1) From 'The Wisdom of Amenemope,' ca. 1750 b.c.

is to be optimized.(1)

But in complex systems, especially those with human components, it may represent a very complex computation indeed, and its outputs may be complex also. It should be noted that one of the parameters of control is the identification of the particular success evaluation function. That is, the evaluation function itself should be subject to adaptation, according to some higher level evaluation; in living organisms, that higher level evaluation might be survival.

I might propose a simple taxonomy of adaptive controls. This is not a Linnaean one; rather it differentiates on the basis of the form of the adaptive control action taken, or of the evaluation function or the range of responses.

For this purpose the notion of adaptation must be defined a little more precisely than has been done so far. I mean here a control whose function (or the purpose of whose builder) is to optimize (maximize or minimize) some function of the system that is being controlled. In a strict sense, we should imagine that this is an 'as if' definition, so that we do not worry about whose purpose it is to which bacteria adapt.

Perhaps the simplest kind of elementary adaptive control is the discrete selection of a 'good' choice, the simplest example being a binary choice. In some sense, of course, there is nothing to do but try them both, and pick the better one. In practice this is a great deal more difficult than might be imagined, because the 'simple' problem is invariably embedded in a host of harder ones. The binary choice can be illustrated in RT behavior with whether to twiddle or not.

Then there are continuous controls, that is, controls which involve setting a continuous variable or parameter, and not making a decision at all. The distinction I draw between servomechanisms and other kinds of controls is that in servomechanisms there is an explicit error signal, and it has a known sign.

A human example of a continuous adaptation that is not a servomechanism is the use of the Radio Direction Finder (RDF) by people on small boats to find the bearing of some radio station. The radio signal is usually a pure tone, or else a tone in separate beeps, and the RDF must be swung in a horizontal plane until the signal is at a maximum or minimum (the latter is called null-seeking, and usually turns out to be more sensitive). The human user does not know which way to rotate the RDF when he turns it on; his procedure is just to move it left and right far enough to detect the gradient of the signal strength. That is, he adds to the control some extra signal (the moving left or right), which he

---

(1) Kalman (1958), p. 444.

correlates with the signal returning around the loop to detect its sense. This is but a detailed description of the processes mentioned in the first paragraphs of this section on page 23.(1)

One may summarize these simple distinctions in the following primitive taxonomy of the spaces in which adaptive mechanisms work:

- Discrete
   - Independent (includes binary choices, perforce)
   - Non-independent, or quasi-metric
- Continuous
   - Servomechanisms (that is, with a signed error)
   - Continuous Hill-Climbers in general(2)

It must be confessed that that is not an entirely satisfactory taxonomy. For one thing, the distinguishing attributes are different for the two major categories. Between discrete units the distinction is exclusive, while among the continuous ones, servomechanisms are merely a special case of the general hill climber.

The term 'hill climbing' has been used to describe the entire class of general optimization, and derives from the image of climbing contours in some descriptive phase space. The idea is to maximize some performance; if we represent performance as some elevation gained, then there is a parallel with topography, and the sport of mountain climbing, in which we are trying to climb to the highest point.

That analogy can be fairly exact if we are trying to maximize a system that depends on the values of two variables, for then we can represent performance with contours on a graph in which the values of the two variables are the coordinates. Thus, in figure 2.5B, the contours of performance happen to be more or less straight lines. The system is at the starting point, and the hill climbing procedure is to take a step arbitrarily -- that is, to try the system at a different set of values for the two variables. If that step produces a better performance, repeat the process from the new values; if not, then return to the starting point and try a new step. In the figure, the first step goes downhill and is thereby rejected. The second merely stays at the same elevation and is likewise rejected. The third is successful.

The process of PT can be represented by *not* returning to the starting point if the step is not uphill; but if it *is*, then keep going in the same direction.

---

(1) There has been far less literature on this general problem than on problems more amenable to strict analysis. A comparatively recent review may be found in Young (1969).
(2) For example, see Minsky and Selfridge (1961).

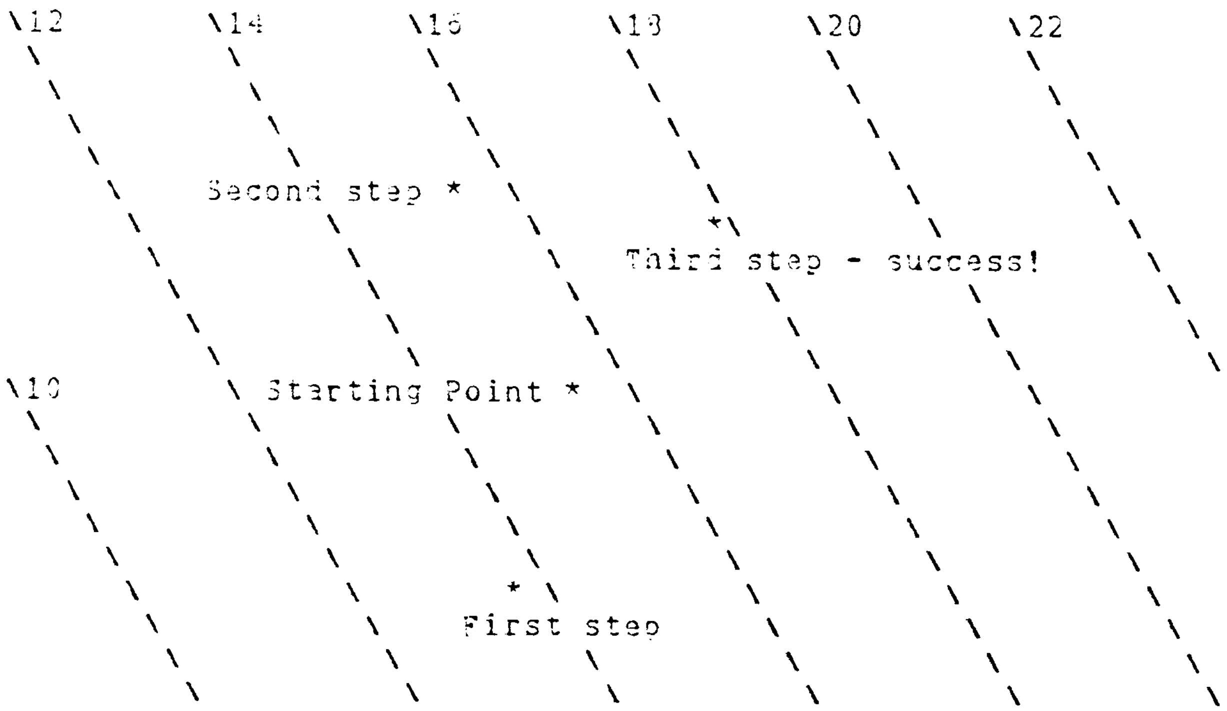

Contour lines of performance - arbitrary units

Figure 2.5B The Hill-Climbing Metaphor

It can be argued that all this is rather like a blind man trying to climb a hill or mountain. He cannot see which way to go, but he knows when a step he has taken is higher or lower.

In one dimension, that is, with just one variable, the process is easier, because steps can be taken in just two directions, instead of infinitely many. In three and more dimensions, measuring the effect of steps in different directions (that is, estimating the 'gradient') becomes increasingly difficult.

In examining simple cases of hill climbing, one can state some simple truths; for example, in one dimension, given certain properties like monotonicity away from a maximum, we can guarantee that a wide range of hill climbing units will find their way to that maximum.

The point can be simply illustrated. Suppose that we know, from other considerations, that the performance function of some system is monotonic decreasing away from the optimum position of its single controlling variable. See figure 2.5C. By what procedure should we test the system and find out the best position to set its control?

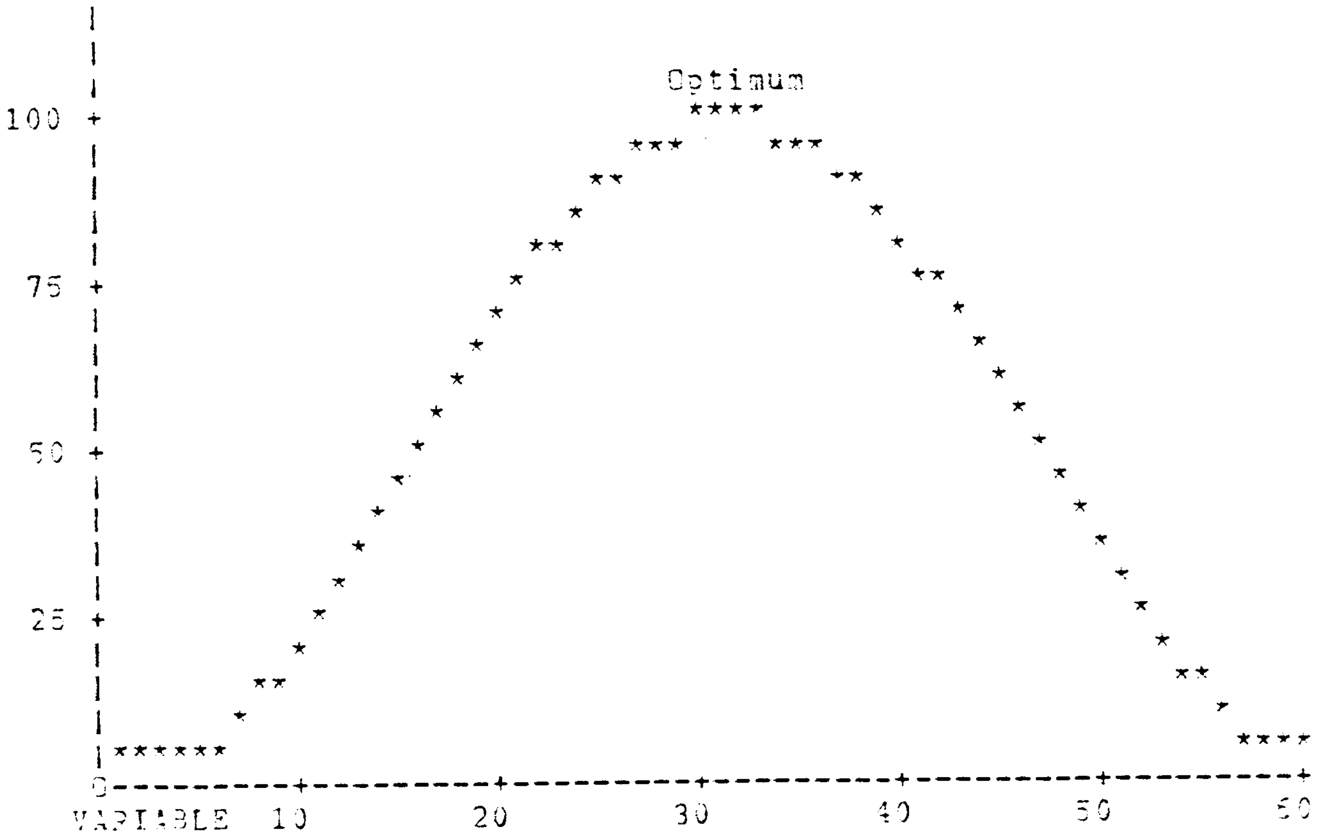

Figure 2.5C Performance of System as Function of Control
Performance Function is Monotonic away from Optimum

In the usual case, we have some notion, albeit vague, of the range of interest, and we start by selecting two points, trials 1 and 2, corresponding to the control values X1 and X2, in figure 2.5C. Then we evaluate the system for those settings, finding the values V(X1) and V(X2). If we suppose that we selected X1<X2, and we find that V(X1)<V(X2), then we know only that the optimum is not less than position X1 of the control. If we find that V(X1)>V(X2), then we know only that the optimum is not greater than position X2 of the control. Note that in the latter case, we do not know which side of X1 the optimum may lie. The two general possibilities are shown in the two curves in figure 2.5C. The next trial at X3 can resolve the ambiguity. We select X3<X1, say (the choice is arbitrary, except that it would be foolish to select X3>X2, for we know the optimum is not there). If V(X3)>V(X1), then we know that the optimum does not exceed X1. If V(X3)<V(X1), then we know that the optimum exceeds X3. And so on.

Even in one dimension, problems with false peaks arise easily and frequently, for we can rarely be assured of monotonicity away from

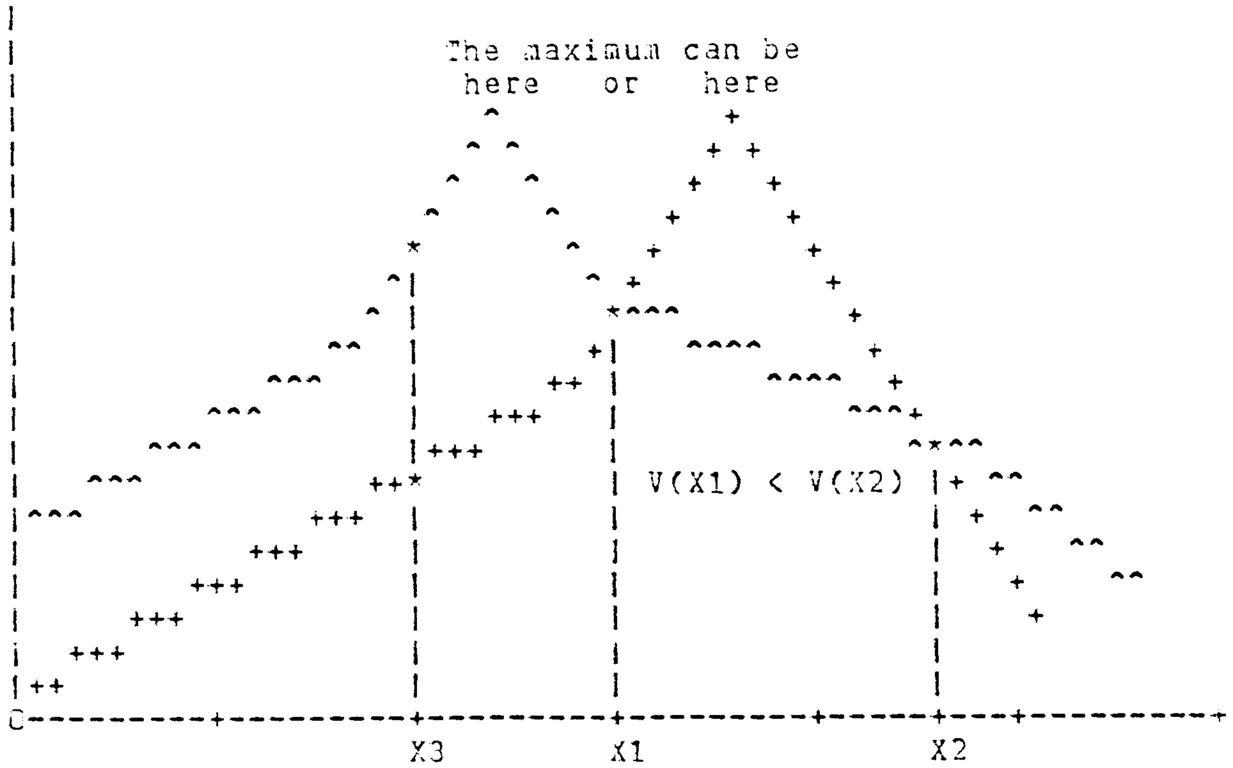

Figure 2.5D Converging on a 1-Dimensional Optimum

the optimum. Other pathologies may obtrude in the way of the optimum-seeker. If the slope is small, that might suggest that we take small steps, lest we overleap the summit; yet the summit might equally well be a long way off, so that the steps should be large. There is no general way to deal with such matters except by learning from experience what is the best thing to try.(1)

The crucial point about the hill climbing metaphor is that the surface of the performance function is reasonably smooth in the phase space; that is, that the mapping function from the variables to the performance is continuous or at least well-behaved.(2)

Also, the general problem is far harder. For example, in two dimensions, mere sampling cannot find a stationary maximum with any point-sampling techniques -- even if all the contours in the search

---

(1) Jarvis (1975) presents a review of many of the current optimization techniques in control theory.

(2) Discontinuities need not be bothersome by themselves, if they are few in number, and consistent in direction with the topography as a whole.

space are convex.(1)

2.6 The Controllable Parameters of Control

I regard all the trail following procedures, whether trail following spatially or in control space, as adaptive, in the sense that they are advancing the organism to some explicit or implicit goal. Again, some of the parameters that set the behavior are explicit, some implicit. Here will be discussed some of those parameters.

In most cases, it is possible to characterize adaptive units by their response time; that is, the delay around the control loop. Ascertaining this in practice may be hard. It may take ten or twenty seconds before the helmsman of a large oil tanker can observe any response to a change in the rudder position. The response of a racing car may be faster than most people can handle. Furthermore the response time depends itself very strongly on the speed. The racing car's control loop for steering reacts a great deal faster when the car is going faster, for example. For the animal examples that have been mentioned, the response times are far less; a bacterial twiddle takes a fraction of a second, ants make corrections in trail following in times on the order of a tenth of a second. Moths, on the other hand, flying upwind to find a mate using anemotaxis, make corrections whose effect must take tens of seconds to estimate, especially at low concentrations of the attractant.

Closely connected with the response time is the sensitivity or gain. That means, when an error occurs, what size of correction is to be imparted to the control? The easiest way to have a control system over-respond is to set the gain too high. That is true no matter whether we are dealing with a servo-mechanism, with the more general signless error-minimizer, or with any of the extraordinarily complex systems that proliferate throughout our societies.

In many of the examples the animals exhibited a constrained sweep, sometimes transverse, sometimes rotational, whose characteristics form another set of parameters of control. One can discuss what we might ask of them. First, the sweep should have enough amplitude to sense the real gradient through whatever noise might be present. Second, it should not, as it is in our example, be present more than is needed to sense; that is, if its environment is changing but slowly, it should be adding fluctuations to the control slowly also.

Most observations have some noise associated with them. That is, a single bit of observed information may contain a good deal less than one bit of information useful in correcting and guiding a system. If one needs a certain amount of information to make a decision -- like which way to move a control and by how much -- then one has to spend

---

(1) The simple proof is shown in appendix 1.

whatever time is needed to integrate what comes in. That integration time may have to be added to the loop delay mentioned above (that is, to the effective system response time), if one wishes to examine some equation representing the control mathematically.

A more subtle form of control consists of specifying or respecifying the evaluation function. In the long run it is probably the most powerfully able to direct what we call 'learning.' It is often not recognized as a form of control, and it has not been analyzed very deeply in the control literature.

Here is a list of these control parameters of a tracking or control system, with short remarks about them:

1. Gain or sensitivity; this is not a separable parameter, in the sense of its being measurable without considering other factors. In some way it measures the size of the correction, as I said above. With two systems, otherwise the same, if one applies larger corrections than the other, it must be said that the first has a larger gain. It should also be noted that the timing of corrections is crucially important.

If the gain is set too high, then instability will often result, causing divergence. This might be illustrated in tracking: an ant following a trail of pheromones accidentally deviates to the right, as in section 1.3. The left antenna is stimulated much more than the right, and so the ant corrects to the left. If the gain is too high, that means that the ant goes further to the left of the trail than he was originally to the right. Then a larger correction swings even more to the right, and so on.

This is not the place to go into control theory, which is concerned with the design and study of control systems so as to avoid divergence, inter alia.(1) It might be helpful, however, to reveal that the overall gain around the control loop should be less than one, to avoid divergence. What that means is this: consider a small deviation inserted at any point in the loop, in the control, say. That deviation will cause a deviation in the system, which will cause a resulting error signal, which will through the control mechanism cause another deviation in the control. The final deviation in the control must not be greater than the original one, or else, as the deviations go round and round the loop, they will constantly get larger.

2. Response time or gain; to some deviations from a track one must respond very quickly, and to some slowly. Notice that this is not the same as saying that the corrections must be large. For

---

(1) Two excellent books are Padulo and Arbib (1974) and Fortmann and Hitz (1977).

example, if one strays into the path of an oncoming car, one must act quickly, but the correction will probably be very small.

The response to a signal, an error, say, is always in practice some weighted average over a recent history. In real life, as opposed to computer simulations, events do not take place at points in time, whether evenly spaced or not, but over intervals, short or long. If the signal is corrupted by interference, or by some random process, then it may be necessary to listen to it for some time just to get a good idea of where it really is.

If the response time is long, it will mean that a perturbation takes a long time to go around the loop. If the gain is too high, then the oscillations will be slow, but growing. There are nice relationships between the setting of the gain and the setting of the response time. The response time is itself not a single delay, as I remarked above, but a weighted average. Mathematically, that is expressed by writing it as an integral of the signal (plus noise) multiplied by the weighting function, often called a 'kernel.' Properly, then, it is that function, the kernel, that is the parameter.

3. Evaluation function: this is conceptually a very complex parameter. Our analysis of its use emphasizes the fact that in a real system there is usually no single goal, but a hierarchy of them. In following a trail, for example, the purpose of the perpendicular swing is to keep the maximum of the scent concentration in the middle of the swing, so that the trail is not lost. But that purpose is followed only because it will lead to the end of the trail, and satisfy whatever higher level goal is to be found there.

In fact, the complex hierarchies of goals that manage human existence commonly exert their effects in this way. The principal control parameter of simple control loops is the 'set point,' that is, the particular value of the variable being controlled that we are trying to keep it on. In driving, it might be the lane of the turnpike; when we pass another car, the instruction, as it were, is to reset where we want the steering to head to the adjacent lane to the left.

4. Exploratory jitter: this includes the control motion added to explore the effects of the control. There may be several ways to specify such a motion, but the specifications must include the frequency and its amplitude if the error does not come with a sign; that excludes the classical case of servomechanisms.

3 Tracks and Trails

In this section I shall discuss physical trails in more depth, using the framework already set up, and relating the controls of the tracking procedures to those that can be inferrd from life. In section 1, I discussed various aspects of tracking a trail, by following a local (transverse) maximum of a chemical concentration, and I distinguished between 2-D cases (e.g., bloodhounds) and 3-D (e.g., shrimp).

In this section, I explore complex adaptive systems -- complex to different degrees -- whose structures incorporate what amount to rules of thumb, and hence are at least in part heuristic. Typically these rules have parameters that can be adapted. Such a rule might be the one that switches the mode of operation from acquisition/recovery to tracking. One can never be sure that the rules operate in a nice binary way. For example, the switch from acquisition to tracking often seems to be almost manifestly driven by pattern-matching of a sensor output; but the reverse switch is often a slow degeneration, a continuous modification of the parameters of behavior.

Some of the points will be illustrated by simple computer simulations.

3.1 Tracking a Trail - Elementary Heuristic Search

The essence of a trail itself is its one-dimensionality -- or at least an extreme anisotropy in one dimension; trails diffuse in the other two dimensions if they are chemical, although some kinds may not, like footprints. But all trails surely decay in one way or another in any case.

Trails are also often extremely uneven. When ants make trails, they do so by occasionally touching their abdomens to the ground, so as to deposit some chemical that can be scented.(1) An organism that is acting as a prey to a predator leaves its actual chemical spoor by accidental contacts, as well perhaps as direct chemical emanations from its body. All that means that the actual concentration of the chemical is going to be an uneven one, and the unevenness will be exacerbated by different exposures of different

---

(1) See, for example, Sebeok (1975).

parts of the trail to the diffusing medium. Furthermore, the chemical comprising the trail or spoor diffuses away; the rate at which that happens might be expected to affect the optimum tracking parameters.

Presumably, also, there is at least sometimes a background of identical or interfering chemicals that can provide some risk of losing the trail.

For organisms whose nutrition depends on tracking prey, much must depend on the efficiency with which tracks are discovered in the first place -- that is, on the search strategy. A key question is obviously how many dimensions are to be searched. If we consider a trail as a long thin hill with the peak at one end, in 2-space any random search path is going to cross the hill, by and large; that is not true in 3-space.(1) In any case, trails are to some extent ephemeral and greatly affected by unknown history and trends, and one cannot be certain _a priori_ what is the best search strategy for them.

The model for trail-following that was presented in section 1.3 on page 12 requires a transverse sweep to maintain the peak of the concentration in the middle of the sweep, coupled with a more or less steady forward movement. There are a number of parameters that control exactly how the model should be operated during the actual tracking, as opposed to the searching. First of all, how wide should the transverse sweep be? Wide enough, presumably, so that the declines in concentration at each end can provide the control signal to keep the trail centered on the sweep. Second, what direction should the tracker head in? Also, there must be a certain inertia in most tracking, so that the organism can continue over interruptions in the trail; yet that inertia must not be so powerful as to carry the tracker too far from the trail if it has been in fact lost. Hounds often overrun the scent of the fox, I am told. They must then find the trail again.

As I pointed out, instead of tracking by sweeping across it, it is equivalent to use spatial integration by taking advantage of simultaneous detection at sites remote from each other, that is, remote enough to be able to sense differences in concentration. Certain ants follow trails by balancing the sensory inputs to the two antennae. If the detecting capabilities of one antenna are reduced, say by amputating it, the ant will tend to one side of the trail.(2) There seem to be no data to estimate any of the parameters of trail-following in our sense.

---

(1) As a matter of fact, a random walk in 2-space will approach any target arbitrarily closely with probability 1; in 3-space that is not at all true.
(2) See Hangartner (1967) described in Wilson (1971), p. 252.

The trail recovery procedures need not look like acquisition procedures, at least to begin with. At some point, of course, if trail recovery proceeds unsuccessfully, it should degrade into normal search. Dethier (1976) describes how the blow fly recovery procedure is slowly modified into general search: the original tightly wound spirals become looser and wider. The actual performance is presumably adapted to the expected or experienced distributions of food or prey, the nature of the space, the physics of the trail and so on. Although in principle trail-following is trail-following, in practice the strategies tend to be different for prey catching and, say, mate finding. In the latter it can be reasonably supposed that the mate is not actively opposed to consummation; in the former it is well known that, for example, moths take evasive action on their hearing that bats are about to devour them.(1) One early review(2) suggests that the scale of turbulence found in moving air streams cannot provide useful directional guidance further away from the source than a few centimeters. In our terms, that is probably wrong, and would mean only that integration take place over distances long enough to smooth out the noise. The strategies discussed in section 1.3 -- involving optical cues to determine the direction of air flow -- have been observed in *Drosophila*,(3) mosquito,(4) and the moths *Plodia*(5) and *Spodoptera*.(6) There are few quantitative data on the hill-climbing tactics in most such instances. For mosquitoes, Wright (7) dismisses vision as operating at a range of more than a few inches, relying instead on the attractiveness of carbon dioxide and water vapor; but Jones claims that

> The adult mosquito has good vision and will quickly track a moving human hand. It is clear that the primary factor that attracts the *Aedes* mosquito to a human being from a distance, however, is odor ... Precisely what odors attract mosquitoes from a distance are unknown ..."(8)

Jones is, however, talking about the yellow fever bearing Egyptian mosquito, *aedes aegypti*, while Wright is talking about native North American mosquitoes, of several different species.

The use of optical cues to orient the animal into the wind has been termed 'anemotaxis' -- another taxis -- and it is clearly most appropriate in finding a stationary emitter, like a mate. Farkas

---

(1) See, for example, Treat (1957), and Roeder (1963), pp. 52-93.
(2) Dethier (1957).
(3) Kellogg et al. (1962).
(4) Kennedy (1939).
(5) Kennedy and Marsh (1974)
(6) Murlis and Bettany (1977).
(7) Wright (1975).
(8) Jones (1978), p. 140.

and Shorey point out that anemotaxis is often assumed and not proved:

> Contrary to previous beliefs, the mechanism by which the males steer toward the odor source does not require a sensing of wind direction.(1)

But they admit that "we still do not understand how the moth distinguishes the 'upwind' from the 'downwind' direction."(2) Of course, that is exactly where anemotaxis comes in. Grubb (1973) points out that their results do not mean that moths don't use anemotaxis, only that hill climbing is possible without it, a conclusion that Farkas and Shorey accept.(3) It will be clear from the simulations in the next few subsections that anemotaxis, when it is possible, is enormously advantageous, and hence would probably be adopted by evolutionary adaptation. The advantage comes from not having to integrate a non-uniform track along its path, so to speak, while the animal is simultaneously integrating at right angles to keep on the trail. After all, the essential directionality of a diffusion trail in air, whether or not it is turbulent, is very slight at any great distance downstream.

The question about the directionality of trails on the ground is another matter, which was mentioned in section 3.1. Given the noisiness of natural environments, it is hard to imagine that a smooth depositing and some kind of exponential decay would provide enough information for the directionality of direction discussed in Michener and Michener(4) and Willows.(5)

The next three sections will describe in somewhat more detail the three phases, search, tracking, and recovery, discussing how adaptive controls may have evolved, and how they might work. The fourth will describe some simulations and conjectures about the directionality of trails.

3.2 Searching for a Trail

It is probably obvious, and it can be shown rigorously, that in a completely uniform desert, the fastest way to find a trail, if there is a uniform distribution of them, is to travel indefinitely in a straight line. There are not very many assumptions needed for the proof, which involves showing that if you don't travel in a straight line a trail you might hit you would already have hit; that is, you are wasting some of your effort.

---

(1) Farkas and Shorey (1972), abstract, p. 67.
(2) Ibid, p. 68.
(3) Farkas and Shorey (1973).
(4) Michener and Michener (1951), pp. 17-18.
(5) Willows (1973), p. 216.

There is another and perhaps more profound point. The nature of a trail is that it divides space into two sets, success (being on the trail) and failure (not being so). If you are in the success set, you are near other points of success: that means success is correlated with success, and failure with failure. Hence, if you are not on a trail, and have not been recently, *you want to move as far away as possible*. This argument has been applied to adaptive controls and learning by Minsky and Selfridge (1961).

In practice, of course, regions where there are few clues about nutrition are seldom uniform. There is nearly always some anisotropy, even if it is only gravitation; and apart from that, then, one might expect that the mean free path of an organism will extend (through evolution) to be comparable to the implicit or explicit boundaries of uniformity in real life. I have found little enough analysis of this question in the literature. Dethier(1) discusses the problem in a general way for the blowfly. The general question of the adaptive possibilities for search strategies has indeed been hardly ever raised, let alone attacked. Dethier says:

> "Few quantitative studies bear on this topic. Casual observation suggests that flight [in insects searching for food] is random insofar as direction is concerned ... More precise studies ... support this conclusion in general ... Velocity and altitude are regulated by visual cues and are related to wind velocity. The crucial point is that the flight of food-deprived insects is basically nondirectional. This is a good strategy for searching because it is one way of ensuring that large areas will be explored."(2)

But, Dethier says, quantitative data are rare indeed, and it would be expected that search strategies would tend to avoid previously explored space. Certainly that is true in animals that graze, and that case will be discussed in section 4.3.

The reader will remember from section 1 that E. Coli has a positive probability of twiddling even if the environment is not changing in the slightest. One can imagine several reasons for this. First of all, there are variations in the perceived attractiveness of that environment due to variations as a consequence of the organism's size. Second, evolution may have allowed for a possible insensitivity to slow long-range changes in background attractiveness. And third, evolution has probably concluded that E. Coli is seldom indeed in an isotropic and uniform environment.

---

(1) Dethier (1976).
(2) Ibid, p. 17.

3.3 Following the Trail

The RT strategies are clearly the immediate precursors of trail searching and following. Indeed, if the trail is smooth enough, E. Coli can probably not only find it, but track it. At least, the program that generated the simple simulations in figures 1.4A and 1.4B is perfectly capable of finding and following a trail. Figure 3.3A shows the trajectories of three organisms following a simple trail, using the same program, modifying only some parameters and the attractiveness function.

Of course, the underlying assumptions about a trail are that either the source is moving through the medium, or the medium is moving around the source; and the average progress of the searching organism must exceed that motion. Probably organisms as lowly as bacteria do not find themselves customarily in such an environment, but, as I have already pointed out, beasts not much higher in the phylogenetic scale may.

We would expect the parameters of trail-following to have been evolutionarily tuned. Indeed, we can employ the biologically familiar conceptual reversal, and estimate the evolutionary forces from the current state of the adaptations. The difficulty is that the data on tracking abilities and strategies seem to be for the most part lacking in the literature.

The details of moths' strategies in following a trail upwind (like silk moths already mentioned) are not to be found in or deduced from the literature. The support for our description of the strategies is almost anecdotal. Kennedy and Marsh (1974) verified that some sort of anemotaxis was operating, when the moth flew upwind; when it lost the track its zigzags became perpendicular to the wind. It should be noted that of course the heading of the moth would not be perpendicular to the wind, if the description were accurate, because its upwind component would have to compensate for the downwind drift.

The problems of investigating the strategies used in trailing can be examined by first noticing whether the sensor is single or multiple. A single sensor must make some kind of a sweep across the track unless it is to perform RT tracking, like the bacteria or the moths. A multiple sensor (like an array) can direct the tracking so that the maximum response occurs in the middle of an array. Such a sensor needs no sweep at all, of course. Many beasts use two sets of sensors in parallel, especially eyes and ears.(1) I am not aware of any large class of animals that employ more than two parallel

---

(1) In some ways, of course, each ommatidium of a compound eye is a separate sensor; but so is each retinal rod or cone in vertebrates. The point is that such multiple sensors act as single integrated sensors.

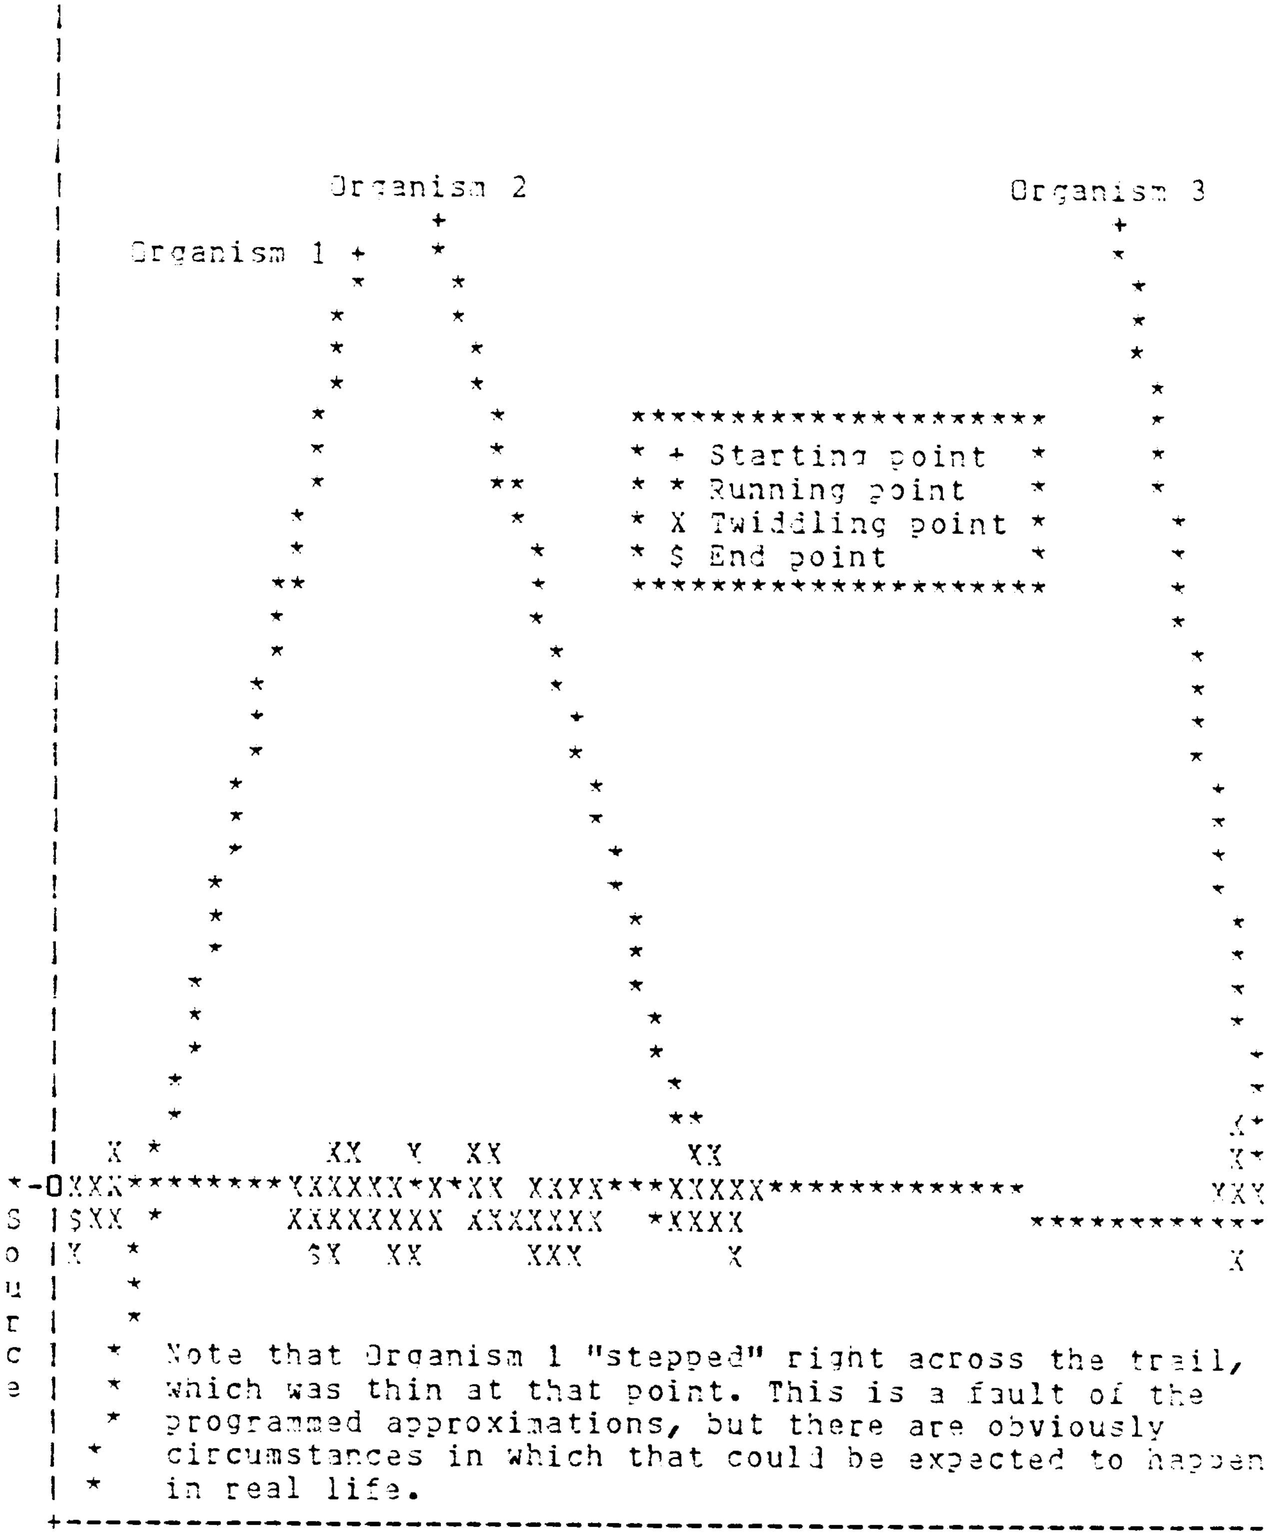

Note that Organism 1 "stepped" right across the trail, which was thin at that point. This is a fault of the programmed approximations, but there are obviously circumstances in which that could be expected to happen in real life.

Figure 3.3A Three Organisms and a Trail

sensors in that way. One example is of course animals that track along tunnels by means of whiskers that simultaneously brush against both walls and the ceiling and floor of the tunnel; but I doubt that a tunnel counts as a track or trail in the usual sense.

The two sensors may be antennae, as in moths or ants, whiskers, as in cats and catfishes, or ears, as in bats or seals. In the blowfly, Dethier points out that in tracking a thin trail of sugar water, consuming it as it goes, the fly uses the same technique, but the sensors are in the toes.(1)

The usual method of tracking, then, takes one of two forms. In the first, typified by hounds, a single sensor moves over the trail transversely as the hound progresses, maintaining the trail in the center of the sweep, with due regard for changes in intensity and direction; in the second, typified by ants, the sensory inputs from two sensors, the antennae, are balanced.

3.4 Recovering the Trail

Since the likely distribution of trails after one has just lost one is presumably different from when one is starting out, there is no reason to suppose that the best strategy to follow in order to recover it is the same as that to find it in the first place. Of course, after a certain duration of fruitless searching, the best thing to do no doubt degenerates to the initial case of search.

This can be illustrated by what a captain of a destroyer is supposed to do when the submarine he has been depth charging seems to have escaped his attack: he should spiral around the last known position of the submarine in such a way that his radial velocity is the same as that of the submarine, while of course his true speed and his tangential speed are far greater. In this way, if his estimate of the submarine's speed is correct, his course will intersect that of the submarine's, and he can resume the attack; this case is explored futrher in section 4.4.

There is in general no reason to suppose that such problems are susceptible to analysis, because in general the avoidance strategies are not even known in detail. For example, it is known that moths take fairly complicated avoidance on their hearing the beeps of a searching bat;(2) that avoidance takes at least two forms, whose efficacy we can hardly guess at.(3) The two cases, and their

---

(1) Dethier (1976), Fig. 14, p. 26.

(2) See Griffin et al. (1965) and Roeder (1963).

(3) Apparently, if both ears of the moth (there are two) are working well, the moth has some directional information about the bat, so that anisotropic avoidance strategies can be usefully selected. And if not, others turn out to be optimal. It is remarkable that the

relative frequencies, must surely influence the optimal strategy for the hungry bat, if he loses the sonar trail from the avoidance maneuvers. Strictly speaking, of course, the bat is attacking with the AB strategy, rather than the trail-following techniques I have been discussing here. There has been little discussion of trail recovery procedures in the literature, except for Dethier, and his work on the blow fly that I have referred to several times. It is simple enough to model such questions on computers; the programs and what they show are interesting but not startling. What is needed are data from life.

Anecdotal evidence can be found in abundance in cases interesting directly to people, like fox-hunting, considered in section 4.2. There are cases of economical importance that I can find no data in the literature about: for example, salmon in effect follow a chemical trail back up the stream at whose head they were hatched;(1) if they lose the trail, what recovery procedures do they follow, if any? Since salmon do find their way up salmon-less rivers, it must be far from perfect in any case.

3.5 On the Directionality of Trails

In considering the directionality of trails, we should distinguish between trails in a fluid like air or sea water, where diffusion and turbulence corrupt the trail, and trails deposited on the ground. The latter can be expected to maintain their shape of deposition for longer, and in addition are inherently motionless. In my initial simulations I had supposed that the trail was deposited uniformly, with variations that were smooth, and that the scent was diffused away at a more or less constant rate. Under those circumstances, it turned out to be hard to imagine how a gradient would exist along the trail that might indicate the direction in which it was laid. For a subsequent test, I supposed that the scent was deposited more realistically,(2) by the insect's, say, touching the ground with its abdomen, and laying a small blob of scent that then decayed exponentially as the abdomen dragged, making a tapered streak in Wilson's words. The trail thereupon slowly spreads out, in such a way that each unit of the chemical scent disperses into neighboring areas. Such a dispersal is illustrated in figure 3.5A. The scent blobs were given an unrealistic but computationally convenient

---

mites that render an ear of a moth inoperative never simultaneously occupy both ears of a moth (Frederick Webster, personal communication). I suppose that the new mite must follow some kind of pheronome trail to the correct ear.
(1) It has been shown that each stream has some kind of a flavor detectable by the salmon, and the flavor is apparently distinguishable all the way to the ocean (XXX); we know little about whether other senses aid this process, like geographical ones.
(2) XXX

Poisson distribution.(1) In the figure, the times of initiation of a blob of scent are indicated by "^".

We now suppose that an organism possesses a habituating sensor. In a crude way we write that the sensor output is the ratio of the local or temporary intensity to some average of recently received intensities, although it ought to maintain some nonzero level, to avoid dividing by zero. Such a sensor is clearly highly nonlinear. If it is used on the kind of trails generated above, and illustrated in the figure, the output is different depending on which way along the trail it is applied.

In such a case, we can be certain that if it is profitable for organisms to be able to detect the directionality of trails, they will have evolved their tracking mechanisms to be able to do so, since the basic detectable differences are there in the outputs. Notice that we do not have to worry even about which output is greater or less, because the adaptation will adapt to the evolutionary goal.

The directionality of the trail is detectable with such a filter(2) for some time, according to the very crude model in the program that generated the figure. It is true that the discrimination possible fades slowly, but even in the third curve of the figure the sensor makes effective discrimination for most values of its control parameter.

Table 3.5A shows the degradation in performance with time, expressed in arbitrary units. The score is the number of correct estimates made by the filter out of 50 trials, after the system has been tuned for that delay. That is, it shows the highest point on the success curve for the parameter that describes the strategy; in the case of this program, the parameter was the exponential weighting factor. In these figures, each decay step spreads just 1 percent of the amount of chemical at each location among the two neighboring locations.

The actual figures of success for the parameter of integration are shown in table 3.5B. That table shows the success of the filter with 20 different values of the gain, in arbitrary units. It suggests that a value of that gain around 0.1 is the best overall; it can determine the directionality of a trail like the third one in figure 3.5A in over 70% of the trials.

---

(1) It is true that I castigate work in which unrealistic assumptions are made for computational convenience, but here I am merely illustrating an argument.
(2) A habituating sensor of the kind already discussed can act as such a filter; it might work like an aircraft rate-of-climb indicator, which is usually a kind of leaky altimeter.

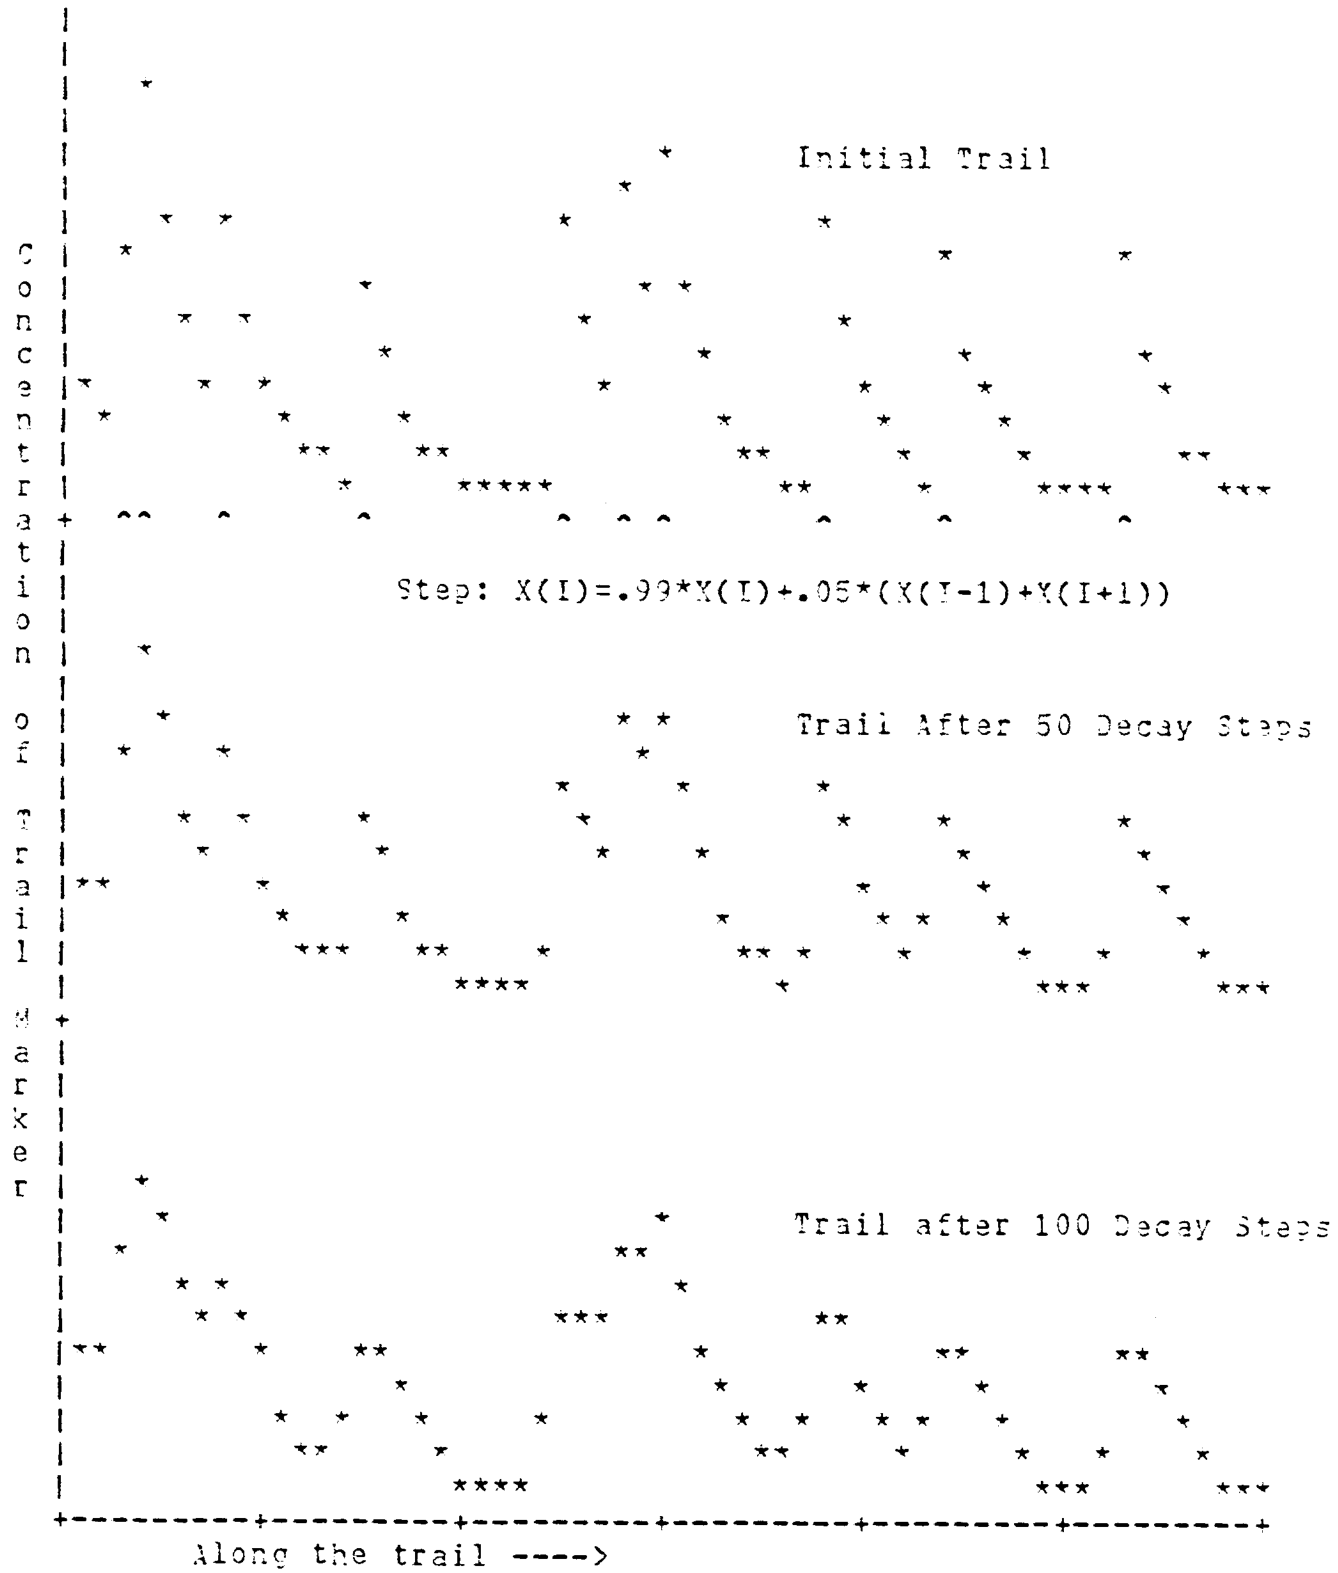

Figure 3.5A Scent Trail, Spreading Out.

After	0	decay steps,	best	filter scores	50	out of	50.
"	10	"	"	"	50	"	50.
"	20	"	"	"	48	"	50.
"	30	"	"	"	46	"	50.
"	50	"	"	"	42	"	50.
"	75	"	"	"	37	"	50.
"	100	"	"	"	36	"	50.

Table 3.5A Success of Best Filter for Decaying Trails

This illustration is of course over-simplified, and there is little to be concluded about how much further adaptation could go in discriminating the directionality of decaying trails. Other parameters and sensitivities are obvious, and they might provide far better performance. For example, the particular shape of adapted habituation might be optimally different from handling merely a ratio of sensation to current level (the logarithmic assumption). I should also point out that all this represents but a conjecture, for I have not found any data in the literature to support or controvert the mechanism proposed.

Successful Guesses out of 50 Trials,
For 0, 10, 20, 30, 50, 75, 100 Decay Steps

	0		.1		.2		.3		.4		.5		.6		.7		.8		.9		1.0
0	48	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50	48	45	39	29	29
10	31	50	50	50	50	50	50	50	50	50	50	49	49	49	48	45	45	41	36	29	30
20	26	46	48	46	45	45	44	44	43	43	42	40	39	39	38	37	33	30	23	23	16
30	29	43	46	43	42	41	40	40	40	38	38	37	33	30	28	27	29	28	26	25	28
50	29	41	42	40	38	36	36	36	37	36	34	32	30	30	30	31	31	29	26	29	32
75	27	37	34	33	31	23	23	22	21	20	20	20	20	23	25	27	25	24	26	24	25
100	29	33	36	34	31	31	28	28	25	26	27	27	27	27	28	28	29	30	29	29	29

Table 3.5B Polarity Estimation for Decaying Trails

Wilson's disbelief already referred to on page 12,(1) is not clearly directed at all trail-following by ants; and even if his argument is directed only at his favorite Solenopsis_saevissima, simulation

---

(1) Wilson (1971), p. 250.

with a modification of the program that produced the tables above will show that the polarization is easily detectable for appropriate choices of the parameters. This does not show, of course, that the trails actually are polarized, or that if they are the polarization is detected.

His argument is that the lateral sweep of the two antennae makes it "difficult for the follower ant to 'read' this information,"(1) meaning to ascertain the shape of the tapered drop. But the sharp rise of the attractant at the front of the drop, compared to the gentle slope at the back would be detected no matter which way the antennae swept over it, diagonally or not; at least, it is not clear that it could not be. One should be reminded that there are no data.

The arguments need a slowness of diffusion that seems not to apply to wind-borne trails. I have already in sections 1.3 and 3.1 alluded to anemotaxis, and the arguments about its mechanism.

The general conclusions about the directionality of trails can be summarized:

1. Deposited trails can be polarized by shaping the deposition droplets; this probably occurs in nature.(2) Theoretically, they can also be polarized by using different chemicals; I have not found any data that remotely suggest that such a thing happens.

2. Air-borne (or water-borne) trails are probably too noisy to allow for mere differential diffusion to serve as a reliable clue to directionality; some other directional guide is needed, like anemotaxis, or gravity.

3. It should be remembered that even if the directionality is not locally evident, there is always a 50% chance of making the right decision.

3.6 Trail Following - Integration of Strategies

There is a general problem that has already been mentioned in another context, that of determining the worth of a module in an assembly of modules. By and large, it seems reasonable to assume that better tracking will lead to better prey finding; but of course it may turn out that the effort expended in tracking better ought to be expended in better searching. That is, it is not enough just to optimize the parts, because there are nearly always some trade-offs among the parts.

---

(1) Ibid.

(2) See Macgregor (XXX) XXX

The assembly of a strategy that comprises several strategies of the kind discussed above -- search, tracking, track recovery, switching to other strategies -- is a complex process that has parameters of its own. Some of the obvious parameters are those that control the transitions from one strategy to another, as well as those that allocate resources to the various strategies that form something like modules in the whole. They are not entirely modular, because of the interdependence on common resources; furthermore, they may well make common use of certain other functions. Nevertheless, an initial launching point might be to treat them as independent.

The transition parameters can be examined if one considers the other elements as constant. There is obviously no sharp boundary of what is the best way to behave for, say, a moth bent on copulation. For one thing, if there are no other males around, then very slow progress upwind would perhaps be the best thing; making sure never to lose the trail. But if there are other males nearby, slow and steady may lose statistically against hasty and rash. As was the thrust of section 2, the parameters of the behavioral optima have to be tracked by evolution in the same way that the transverse position of the beast must track the attractant maximum.

More global features must also be treated. It is not enough for a bat to track any echo, no matter how clear, if it does not represent a target worth the effort. For example, it does not pay the little brown bat to catch drosophila, for the nourishment apparently does not provide the energy for the extra flight needed.(1)

A powerful consideration is always when to give up on being faced with failure, and what to do then. In human terms, there are two opposing temptations: to do what one has always done, and to try something new. The former is usually safer, while the latter may be the origin of new behavior and new capabilities.

---

(1) Griffin, et al. (1965), and Webster, F., personal communication.

4 Other Aspects of Tracks and Trails

In this section I discuss some more of the ways in which tracking can be extrapolated to exhibit more complex and integrative behavior. In one of these, concerned with foxes and fox hounds, the primary purpose of the tracking is still the chase; but in another, the tracking itself is an entirely subservient tool to be used by the intelligent organism or program.

In higher animals especially, the tracking of optimum parameter settings for control loops is all-pervasive; people are mostly not even conscious of it. An important parameter of an adaptive control loop is its evaluation function or goal. These days, the field of Artificial Intelligence is increasingly recognizing the need to improve and track optimum settings for goals. Schank and Abelson write:

> " ... From time to time, the definition of a goal changes... When the specified goal cannot be satisfied, he substitutes a different specification ... [and] finally substitutes a somewhat different goal...
>
> "After the professor succeeds with the substitute goal, ... he undertakes further actions [which] are not in the service of the original goal. Rather they help to improve the value of the higher-level goal ..."(1)

I shall discuss the role of such tracking in section 4.4; not just for goals, of course, but also for the many other duties it serves.

4.1 Aiming at a Goal - the AB Strategy

In section 1.1, I introduced the notion of RT strategies, which are adaptations that examine local gradients in some behavior space, changing behavior parameters (twiddling) if they seem to be leading downhill. I contrasted RT strategies with those that had a view, as it were, of distant goals, so that the organism could head directly from point A to point B; I termed the contrasting strategies AB strategies.

---

(1) Schank and Abelson (1977), p. 103.

In a practical sense, people usually can see what they are aiming at, either physically or metaphorically. Except at an unconscious level, that is, most people in their adult lives are not very aware of using RT strategies. There is a children's game in which the location of some hidden object is revealed to the child by telling him whether he is getting "hotter" or "colder" after some movement on the child's part. That is obviously equivalent to some local sensing of attractiveness gradient that is the essence of RT strategies.

But AB strategies are different in several ways. First and foremost, they can sense remote topography, both physically and in more abstract spaces. Physically, for example, once a hound has viewed the fox, he no longer relies on scent, but can cut the corners of the trail, heading directly for the fox (it might be remarked that foxhounds don't seem to have very good eyes). Previously, it paid the fox to complicate the trail with irregularities. But once the fox has been seen, the hound switches from RT trail-following to AB prey-following and the comparative values of the various avoidance strategies change drastically. In fact, in general, unless the fox can reach safety (like a burrow of sufficient complexity) or it can run faster, it is probably doomed. In that case, his best strategy, to delay the end as long as possible, is merely to flee in a straight line.

Second, apart from freeing the predator from the parallel or scanning process that maintaining the trail requires, the AB strategy enables higher speeds, because sudden changes in direction are called for only when the prey is very close.

I use the prey-predator terminology only because that is a common case, and the metaphor springs to mind. The reader should be reminded of other cases. If one is driving to some destination, one has probably learned that on a particular highway such and such a speed is the best, balancing speed limits, car's capabilities, cop densities, and so on. That is, in adapting one's driving to the destination goal, one need not experiment with speeds, but can accelerate directly up to the desired one.

The rapid movement of one part of the body to a particular location in space requires acceleration and deceleration of a rather careful kind. To catch a thrown ball is not an easy task for a child to learn. It is hard to see how that can be planned and executed with a pure RT strategy.

Probably this last point has the largest overall significance; by enabling planning in spite of local discouragement, the AB strategy in a very real way must have been the initiator of true intelligence. This possibility of global viewpoints, and then conceiving of the global viewpoint itself as only one of many, this aiming afar, indeed remembering the etymology of 'teleology', this modeling of and planning for the future that AB strategies require inherently -- all these make the intelligence and mind of mankind far more directive and purposeful than evolution, whatever certain poets, philosophers and theologians have to say.(1)

With our present cultural and conceptual heritages, it is hard to imagine how these faculties could have arisen without the newer, process-oriented sensory modalities of sight and hearing. The directional and action-at-a-distance properties are not only powerful in themselves in their most primitive forms, but extend easily in environments where propagation and discrimination are feasible. As I have already said; I do not believe in the slightest that evolution partakes of AS strategies at all, but eyes and ears grow better through evolution by working better, that is, by enabling profitable and protective adaptations.

It is of course well known that these adaptations reach their epitome in certain higher animals. Coupled with other social reinforcements, the planning encouraged and enabled by AS strategies are illustrated by cooperative predation by lions, cooperative protection by musk oxen against timber wolves, territoriality in birds, and the entire range of behavior in people.

But it must not be thought that AP and RT strategies are usually clearcut and separate in their application. Often they are jointly applied, and often sequentially, with complicated contingencies, depending on the larger contexts. Even insects make this very clear. (2) Indeed, Tinbergen reports on his own and others' work about the way the bee-wolf captures its prey the honeybee:

> " ... [Philanthus Triangulum] finds the bee with the aid of olfactory stimuli ... it can smell such a bee from up to 2 [meters] ... [Then] the wasp makes a quick dash or 'pounce' from appr. 10 cm; this pounce is guided by visual cues ..."(3)

Tinbergen reports that olfactory clues are used as a kind of trigger for the visually guided 'pounce'. Sometimes the pounce doesn't work, and the wasp alights predatorily and viciously on a twig; then the olfactory mismatch calls off the whole attack. It is this kind of integration of the inputs from the several sensory modalities that typifies what we term higher mental activity.

There is a substantial corpus of mathematical study about AS strategies under the title of Pursuit Theory. A splendid review of the current status is Hajek (1976).

---

(1) XXX

(2) The 'even' in that sentence is a little gratuitous, since insects are pretty advanced behaviorally: see Alloway (1973) for a review.

(3) Tinbergen (1972), p.144.

The complexities that arise when AB strategies are mixed with RT strategies can be startling and fascinating. Another example will be discussed in the next section.

4.2 Foxes and Hounds

> The oldest animal relationship is that of a man and a dog following a quarry. So, when you and your hound follow a fox, you go back to the beginning, the dawn age when man's own cunning was first supplemented by the nose of his new found friend, the dog. And the unsolved mystery of a trail's end is an old, old problem ..."(1)

Fox hunting to the hounds is a vestige of the medieval hunt whose purpose was to supply meat; it was only a secondary purpose to rid the countryside of predators. The hunt originally sought game like deer, but its efficiency was so high that the deer were nearly extinguished everywhere it was employed. Fox hunting survives as a sport in the U.S., England, and Ireland. "The unspeakable in pursuit of the inedible," said Oscar Wilde.(2)

Trail laying by foxes, and trail following by the hounds, are very complicated processes, about which there are almost no respectable data, and the generalities derive from tradition and anecdote. It is commonly imagined that foxhounds follow the trail of a fox primarily individually, the way a bloodhound notoriously used to follow escaped prisoners from the chain gang in the U.S. South a generation ago. Usually, however, the trail is followed collectively, by the hounds and the huntsmen; the latter play a vital role of higher level control and in some cases communication.

Even with these complications, this trail following can be fitted into the general structure I have been talking about, which I think will be apparent.

Fox hounds are large and swift dogs, carefully bred for the values that are thought to be needed for the hunt. They follow the fox primarily by scent, far more than wolves (which are closely related to hounds), and switch to the AB strategy only when necessary. The behavior of the scent depends on wind and weather, among other factors. The hounds smells the scent in the air, and the scent has a life varying up to about 30 minutes in ideal circumstances. The ideal circumstances are apparently a matter of profound debate and discussion: "Too much humidity deadened the scent, and made it lie

---

(1) Houghland (1933), pp. 1-2.

(2) Neville Moray tells me that fox-hunting has just been declared illegal in Great Britain.

heavily,"(1) I do not know what is the dynamic range of scent that a fox hound encounters, but it is probably not great. The trail left by a fox decays irregularly, and hounds frequently lose it, even if it is quite fresh, so that much of the hunt is taken up with trail recovery.

The fox lays a notoriously confusing trail, apparently deliberately. The evidence is all anecdotal, but it is very convincing:

> "The one thing to bear in mind about a fox is that he is actually smart. And he doesn't get drunk nor use tobacco, so his head is frequently clearer than is that of the Captains who pursue him...
>
> " ... The cry [barking of the hounds] was a mile away... The fox presently ran right by us into the pasture and then executed the most amazing maneuver. Unhurried, he doubled up in his line [ran back over his tracks], then made a circle of about a hundred yards in diameter which closed near the wall. Then he jumped on the wall, ran backwards a few yards, gave a great leap off on the other side, and ran off at right angles to his original course. It could only have been a premeditated move to confuse hounds."(2)

Similarly, the hounds are not stupid in following such a trail. Hougnland continues:

> " ... the pack was abreast of us with 'Big Stride' in a short lead. Without a check, he ran the circle just as the fox had traced it and reared upon the wall sniffing for the scent. At that instance, Ch. [Champion] "Cleo" leaped to the top of the wall, cried the line [barked at the scent track] along it, and leaped off where the fox had jumped. In a flash, that great pack ... were over and away in full cry ... they never caught that fox and never will."(3)

* * * * *

In field trials, the fox hounds are not provided with the human support of the huntsmen, but are on their own. Four phases of hunting can be distinguished:

---

(1) Mannix (1967), p. 130.
(2) Hougnland (1963), pp. 12-16.
(3) Ibid, p. 16.

1. CASTING: the hounds do not have the scent, and are wandering to try and find it. During casting, they may be led by the more adventurous hounds.

2. HUNTING: the hounds are not 'giving tongue', but are aware of the scent; they may be feathering their sterns, that is, wagging their tails. They do not have the line yet, that is, they have not found the actual trail, but only patches of the attractant.

3. TRAILING: the hounds have a positive scent, sufficient to speak to (bark, in a more or less special and recognizable way). If there are no clues about the direction of the line, they revert to 2. I have not found out what clues they might use to establish the direction of the line.

4. DRIVING: the line is strong enough to enable them to follow it while running; they tongue much faster, that is, bark more frequently.

Trail recovery can be complicated, because for one thing the fox lays a complex trail. Often hounds will execute a 'lollipop' maneuver, in which they stop dead and then circle as fast as they can with a radius of about a hundred yards. A coordinated pack will often regain the trail this way. This tactic clearly resembles that of the destroyer after the submarine that I mentioned on page 52.

* * * * *

"Yes, I ken John Peel, and Ruby too,
Ranter and Ringwood, Bellman and True,
From a find to a check, from a check to a view,
From a view to a death in the morning."

4.3 Inverse Tracking - Grazing

Grazing beasts should tend to avoid newly grazed territories, even if just because it is a needless waste of metabolic resources not to do so. By grazing I do not restrict myself to bovines, for examples, but include snails, whales, and termites; any creature, in fact, whose feeding depletes a slowly renewing resource. Even predators can be included under some assumptions. In higher animals other considerations often take priority over the use of metabolic resources, and grazing or browsing strategies may seem to be more complicated.

This section discusses only those aspects of grazing that seem to have similarities to the tracking and trailing that is our main topic. A good review of the state of the art of grazing models in general is Pyke et al. (1977).

Herbivores spend more of their time eating than carnivores, as most of us must have noticed, and efficiency will dictate that grazing tracks usually follow along the *edge* of old grazing grounds, for it would be inefficient to leave ungrazed spots to which access had to be gained by crossing grazed territories. I should repeat that this analysis, such as it is, cannot be expected to apply to what we ordinarily think of as herbivores, like horse or cattle; for these beasts have complex mammalian goal structures, and besides nutrition also seek shade, better tasting silage, and so on. Rather, it applies to lower animals like snails:

> " ... sediment feeding in benthonic invertebrates produces the most regular patterns ... best developed in deep-sea deposits (both Recent and ancient) because of the even distribution of food particles in most deep-sea sediments ... These requirements have been met by a multitude of two- and three-dimensional trail and burrow patterns. However, meander systems are used most commonly -- just as they are in human contour plowing and other agricultural activities... Richter ... pointed out [in 1928] that the trace fossil *Helminthoidea labyrinthica* had its movement controlled by a set of basic reaction: i) strophotaxis, that made the animal turn around 180° at intervals; ii) phobotaxis, that kept it from crossing other tracks, including its own; and iii) thigmotaxis, that made it keep close contact with former tracks."(1)

Such strategies make interesting patterns which can be seen in the literature as well as in nature -- spirals, zigzags, sweeps back and forth, and so on.

In some sense, one might expect that grazing strategies would be the exact reverse of trail-following ones, since the purpose of grazing is exactly the opposite: trail-avoiding. But in fact most of the considerations are the same -- follow the trail, but along the *edge* of the grazed trail. Sensing grazed ground may be done in ways not obvious to us, as can be seen by lifting up the bark on a fallen tree; larvae of various species will have eaten cambium in tunnels that make charming designs and hardly ever intersect.

The adaptive selection of non-crossing grazing strategies has been shown theoretically by Papentin (1973), who ran evolutionary experiments on a computer of a similar nature to the primitive ones shown above in figure 2.2A. These experiments derived from ethological analysis of Raup and Seilacher (1969), who showed that a wide variety of feeding and foraging trails could be simulated rather precisely with a handful of behavioral variables; like the turning radius, or the allowable lower limit of the distance between

---

(1) Raup and Seilacher (1969), p. 994.

the current and preexisting tracks. It was supposed that the simulated forager was really rather similar to the common "turtle" so frequently programmed with the LOGO programming language; the four commands are to move straight ahead, turn toward or away from a preexisting track, or make a full 180° turn.

Other forms of grazing are seen in the harvesting tactics of farmers and strip miners, and the appearances are similar, and for the same reasons; it is expensive to go twice over grazed ground.

4.4 Tracking a Submarine with Sonar

This section is not concerned with the deep technological details of tactical sonar operation, but rather with drawing parallels between that and other forms of tracking. The capabilities of sonar are in some way like those of the bat; the echo strength follows a fourth power law,(1) multiple targets need careful discrimination, and the predator is audible to the target at far greater ranges than vice versa. Furthermore, doppler effects -- changes in frequency of signals that correspond to relative velocities -- aid discrimination for both. The bat in general is subject to less interference than the destroyer with its sonar. For one thing, the propagation of sound in air is much lossier than it is in water. For another, the bat mostly modulates his squeak in frequency, so that targets at different distances will return echoes at different frequencies at a single instant in time. But most sonar signals are constant in frequency, so that signals reflected from a distant sea bottom will interfere with those from a closer but smaller submarine, in such a way that a gross frequency filter will not aid in discrimination.

The mathematical analysis of optimal methods of detecting targets is far advanced. The algorithms that decide where the target is have been proved mathematically to be the best, given certain reasonable assumptions about the environment, like the speed of sound and its variations, and the behavior of targets. Typically, a target is tracked, so that its position is known fairly accurately, and the current sonar signal is used only to update or improve that position. Since the update takes less information (in the sense of information theory), it can be accurately gauged when conditions are worse than would allow the same accuracy for the total position.

---

(1) Both in radar and sonar, the target reflected energy received by the sensors is generally sixteen times weaker if the range is doubled.

4.5 Tracking in Artificial Intelligence

> Most practitioners of the fledgling science of artificial intelligence would consider the simulation of guiding, balancing, and raw survival relatively simple compared to the problems about the structure of knowledge that they are currently tackling.(1)

And yet, I argue in this section, the functions of tracking optima are so pervasive in the problems that artificial intelligence is tackling that the procedures and tools for performing them deserve to be examined and analyzed. I shall not analyze them here; and it is not being argued that attempting to reach a profitable and manipulable representation language is not worth doing. Rather, the argument is that just as physical tracks and trails can examined, analyzed, and simulated profitably, so can and ought to be the function of tracking in concept and knowledge spaces. Winston has asserted that exploring this kind of adaptiveness can detract from the effort at representation:

> Infatuation with [questions of adaptation] often leads to *ad hoc* schemes for tinkering with awkward spaces rather than to finding spaces for which search is simple or not relevant.(2)

I confess that I do not see the strength of his argument; rather, the different attacks ought to work symbiotically. Indeed, there i a danger, manifest now, that looking solely at representation problems may tend to bias the kinds of representation that are to be studied, and the kinds of things that are to be subject to representations.

Some of the applications of tracking, both physically and in control space, will be immediately apparent. Many visual pattern problems, or conceptually similar, like the sonar problem above, require dynamic tracking, either of targets, or of optimum settings.

It seems likely to me that the applications can go far beyond such examples. I mentioned on page 59 the tracking and modifications of goals and goal structures that modern researchers in AI are exploring. In the same way, linguistic items have to be tracked, because in practice meanings drift:

> A language is
> ... like a biological organism, a bundle of interacting systems maintaining a dynamic equilibrium known as homeostasis (footnote: The term

---

(1) Mandler (1977), p. 1041.
(2) Winston (1977), p. 96.

> "homeostasis" seems to have been used first in linguistics by Martin Joos ...). There must be a network of interactions tying all parts together, and tying every pair of parts together in a number of ways. It is perhaps a reasonable speculation to insist that language, as we know it, with all its complexities, is very close to the simplest possible system which can maintain homeostasis in the face of the intricate set of external forces that impinge on it.
>
> But language must do more than this. It must continually adjust itself to new needs. No language can be a fixed system of words and patterns. It must be open to receive new words and new structures, and to change the old. Change is inevitable, since language functions in a society in ceaseless flux. The continual change in language produces inevitable maladjustments in the system. The language must repeatedly repair itself, restoring and maintaining equilibrium by additional changes to counterbalance those forced upon it by changed environment.(1)

I think it is not too fanciful to conceive of 'tracking' the meaning of a word, of acquiring its meaning, and narrowing down on it, as metaphorically equivalent to physical tracking. Obviously the process is embedded in a host of other ones, to which it relates in complicated ways, but often the meaning must be modified in only small ways. Indeed, some of the parallels in the processes are striking. The difficulty of a computer program's <u>unlearning</u> something that it has erroneously learned is very similar to the difficulty a computer has in correcting erroneous matching or targets and tracks.(2)

The aspects of tracking that arise in AI are ubiquitous but not prominent. In illustrative examples they tend to be omitted, thereby saving at least one extra level of control. In one way, it can be argued that their real power will be utilized in systems that are too complicated to understand as analyzable models. Consider the ARPA speech program, in which the purpose was to produce a computer based system that could understand continuous spoken speech. The basis for the interpretation was a

---

(1) Gleason (1965), pp. 105-106

(2) An example is a computer program designed to learn words and their meanings from using them and relating them to the behavior of the world. Given the strings GIVEMETHEBALL and PUTTHEBALLINTHEBOX, and some other similar ones, the program inferred that THEB was a proper substring or word. It proved extremely difficult to persuade the program eventually to reject that inference. (Selfridge, M., personal communication).

set of features that were assembled and analyzed at the semantic level, using a complex set of expectations and knowledge rules. The features themselves were never optimized for the resulting uses; in fact, one would suppose from the reports that very little experimentation was done with the features (which consisted of things like formant tracking, or the presence of a burst of high frequency noise). So there is no way of telling how the results are affected by the quality of the features, or whether better features would afford much better overall performance.

Appendices

Appendix 1: Maxima in 2-space

Appendix 2: Symbols for Trail Polarization

Appendix 3:

Appendix 1 Maxima in 2-space

On page 42, I asserted:

> " ... in two dimensions, mere sampling cannot find a stationary maximum with any point-sampling techniques -- even if all the contours in the search space are convex."

I shall justify that by presenting the following theorem, and then sketching its simple proof:

> Given any set of N points inside the unit circle, there exist two functions, each convex up in the unit circle, with identical values over the N points, and whose maxima are separated by a distance of 2-e, for any positive e, no matter how small.

To prove this, consider any set of N points in the unit circle; and notice that it is clearly feasible to place an ellipse so as to enclose none of the N points, and such that its foci are arbitrarily close to a diameter of the circle, and separated by a distance arbitrarily close to 2.

The class of ellipses with those two points as foci can be used as contours to define a suitable convex up function outside the ellipse. Inside the ellipse, -5c:kiconstruct conical surfaces, with vertices at each of the two foci, and extending to the ellipse; they can be made flat enough to satisfy the convexity requirements, so that the two functions generated are identical on the N points, and have maxima arbitrarily close to the edge of the circle, on opposite ends of a diameter.

Appendix 3 Scanning Eyes in Copilia

The copepod Copilia seems to have a specific visual scanning mechanism that does not involve the whole body. The observations reported in the literature(1) show that in a couple of species the eye of the female consists of

> "a highly refractile lens ... [and] a tenuous cone-shaped membrane extends backwards into the body ... about half the length of the thorax, [ending] in a small hgihly refractile element, which obviously corresponds to the crystalline cone of an insect eye ..."(2)

At the bottom of the cone is "a single narrow structure very like the rhabdomen of a single ommatidium or a compound eye"(3) with three photoreceptors.(4) In females of two of the species, C. quadrata and C. denticulata, the two sets of photoreceptors are mobile, and frequently "scan," that is, enjoy "a rapid movement towards the midline followed by a slower, probably passive recoil. The fastest observed sweeps have a sweep-time of about 50 ms., the slowest about 4 s."(5) Moray suggests, and Downing confirms, that this is a scanning system within the beast. Moray and Downing both point out the behavioral advantages such a system might have, and Downing claims that limb movements are specially likely to occur after visual stimulation. "The discovery that limb movements tend to be preceded by receptor movement corroborates the interpretation of the receptor movement in terms of scanning: it makes sense for overt behavior to be preceded by the acceptance of visual information."(6) But that observation serves little to elucidate what is really going on. It is from dynamic considerations, surely, that the eyes are moving in opposite directions; so that, whatever else obtains, the eyes are only looking in the same direction once per sweep.

---

(1) Exner (1891), Wolken and Florida (1969), and Gregory et al. (1964). I am indebted to Professor Neville Moray for these references.
(2) Moray (1972), p. 199.
(3) Downing (1972), p 243.
(4) Vaissiere (1961).
(5) Moray (1972), p. 201.
(6) Downing, op. cit., p 259.

Bibliography

Adler, J., "Chemotaxis in Bacteria," *Science*, Vol 166, 1969, pp. 1588-1597. "Adl1" [4]

Adler, J., "Chemotaxis in Bacteria," in Carlile, 1970 (see infra), pp. 91-100. "Adl" [3]

Adrian, E.D., *The Physical Background of Perception*, Clarendon Press, Oxford, U.K., 1946. "Adr" [25]

Alloway, T.M., "Learning in Insects Except Apoidea," in Corning et al., 1973 (see infra), Volume 2, pp. 131-171. "All" [27, 61]

Aubert, H., *Physiologie der Netzhaut*, Breslau, 1865. "Aub" [23]

Berg, H.C., "Bacterial Behavior," *Nature*, Vol 254, 3 April 1975A, pp. 389-392. "Ber1" [2, 3, 22]

Berg, H.C., "How Bacteria Swim," *Sci. Amer.*, Vol 233, No. 2, August 1975B, pp. 36-44. "Ber2" [2]

Box, G.E.P., "Evolutionary Operations, a Metnod for Increasing Industrial Productivity," *Appl. Stat.*, Vol 6, 1957, pp. 3-23. "Box1" [27]

Box, G.E.P., and Draper, N.R., *Evolutionary Operation*, John Wiley and Sons, New York, 1969. "Box3" [29]

Bremermann, H., "Chemotaxis and Optimization," *J. of the Franklin Institute*, Vol 297, No. 5, May 1974, pp. 397-404. "Bre" [3]

Bretherton, F.P., and Rothschild, Lord, "Rheotaxis in Human and Bull Spermatozoa," *Proc. Royal Soc.*, Vol 153, 1961, pp. 490-502. "Bre1" [33]

Brown, L.N., "Mating Behavior and Life Habits of the Sweet-bay Silk Moth," *Science*, Vol 176, 1972, pp. 73-75. "Bro2" [9]

Bush, R.R., and Mosteller, F., *Stochastic Models for Learning*, Wiley, New York, 1955. "Bus" [35]

Carlile, M.J., *Primitive Sensory and Communication Systems*, Academic Press, New York, 1975. "Car" [7, 8]

Corning, W.C., Dyal, J.A., and Willows, A.O.D., *Invertebrate Learning*, in three volumes, Plenum Press, New York and London, 1973. "Cor" [75]

Corning, W.C., and Von Burg, R., "[Learning in] Protozoa," in Corning et al., 1973 (see supra), Vol 1, pp. 49-122. "Cor1" [35]

Daterman, G.E., "Laboratory BioAssay for Sex Pheromone of the European Pine Shoot Moth *Rhyacionia_Buoliana*," *Ann. Entomological_Soc._of_Amer.*, Vol 65, 1972, pp. 119-123. "Dat" [9]

Daykin, P.N., "Orientation of *Aedes_Egypti* in Vertical Air Currents," *Canadian_Entomologist*, Vol 99, 1967, pp. 303-308. "Day" [4]

Dethier, V.G., XXX, in *Survey_of_Biological_Progress*, Glass, B., editor, Vol 3, Academic Press, N.Y., 1957, pp. XXX. "Det2" [47]

Dethier, V.G., *The_Hungry_Fly*, Harvard University Press, Cambridge, Mass, 1976. "Det" [8, 47, 49, 52, 53]

Downing, A.C., "Optical Scanning in the Lateral Eyes of the Copepod *Copilia*," *Perception*, 1972, Vol 1, pp. 247-261. "Dow" [73]

Dyal, J.A., "Behavior Modification in Annelids," in Corning et al., 1973 (see supra), Vol 1, pp. 225-290. "Dya" [26]

Engelmann, T.W., *Pflugers_Arch._ges._Physiologie*, Vol 29, 1882, pp. 387-400. "Eng" [33]

Farkas, S.R., and Shorey, H.H., "Chemical Trail-Following by Flying Insects: A Mechanism for Orientation to a Distant Odor Source," *Science*, Vol 178, October 1972, pp. 67-68. "Far" [48]

Farkas, S.R., and Shorey, H.H., response to Grubb, 1973 (vide infra), *Science*, Vol 180, May 1973, p. 1302. "Far1" [48]

Farkas, S.R., Shorey, H.H., and Gaston, L.K., "Sex Pheromones of Lepidoptera. Influence of Pheromone Concentration and Visual Cues on Aerial Odor-Trail Following by Males of *Pectinophora Gossypiella*," *Annals_of_Entomological_Society_of_America*, Vol 67, 1974, pp. 633-638. "Far2" [10]

Feinleib, M.E.H., and Curry, G.M., in Lowenstein, W.R., see infra, Vol I, pp. 366-395. "Fei2" [6]

Fogel, L.J., Owens, A.J., and Walsh, M.J., *Artificial_Intelligence Through_Simulated_Evolution*, Wiley, New York, 1966. "Fog" [27]

Fortmann, T.E., and Hitz, K.L., *Introduction to Linear Control Systems*, Marcel Dekker, New York, 1977. "For" [43]

Fraenkel, G.S., and Gunn, D.L., *The Orientation of Animals*, Dover Publications, New York, 1961. "Fra" [8]

Gelber, B., "Investigation of the Behavior of *Paramecium Amelia*," *J. Comp. Physiol. Psychol.*, Vol 45, pp. 58-65. "Gel1" [34]

Geldard, F.A., *The Human Senses*, John Wiley and Sons, New York, 1953. "Gel2" [23]

Gleason, H.A., Jr., *Linguistics and English Grammar*, Holt, Rinehart and Winston, New York, 1965. "Gle" [68]

Griffin, D.R., Friend, J., and Webster, F., "Target Discrimination by the Echolocation of Bats," *J. Exper. Zool.*, Vol 158, 1965, pp. 155-168. "Gri1" [52, 58]

Grubb, T.C., Jr. "Odor-Following and Anemotaxis," *Science*, Vol 180, May 1973, p. 1302. "Gru" [76]

Hajek, O., *Pursuit Games*, Academic Press, New York, 1975. "Haj" [51]

Hamner, P., and Hamner, W.M., "Chemosensory Tracking of Scent Trails by the Planktonic Shrimp *Acetes sibogae australis*," *Science*, Vol 195, No. 4281, 4 March 1977, pp. 886-888. "Ham" [12]

Hangartner, W., "Spezifitat und Inaktivierung des Spurpheromons von *Lasius fuliginosus* Latr. und Orientierung der Arbeiterinnen im Duftfeld," *Zeitschrift fur Vergleichende Physiologie*, Vol 57, No. 2, 1967, pp. 103-136. "Han1" [10, 11, 46]

Hilgard, E.R., and Bower, G.H., *Theories of Learning*, Appleton-Century-Crofts, New York, 1966. "Hil" [24]

Hilgard, E.R., and Marquis, D.G., *Conditioning and Learning*, revised by Kimble, G.A., Appleton-Century-Crofts, New York, 1961. "Hil1" [24]

Holland, J.H., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, Mich, 1975. "Hol1" [1, 30]

Houghland, M., *Gone Away*, Blue Ridge Press, Berryville, Va, 1963 "Hou" [62, 63]

Jacob, F., "Evolution and Tinkering," *Science*, Vol 196, No. 4295, June 1977, pp. 1161-1166. "Jac1" [31]

Jacobson, M., *Insect Sex Pheromones*, Academic Press, New York, 1972. "Jac" [9]

Jarvis, R.A., "Optimization Strategies in Adaptive Control: A Selective Survey," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol SMC-5, No. 1, Jan 1975. "Jar" [41]

Jones, J.C., "The Feeding Behavior of Mosquitoes," *Sci. Amer.*, Vol 238, No. 6, June 1978, pp. 138-148. "Jon" [26, 47]

Kaissling, K.E., and Priesner, E., "The Olfactory Threshold of the Silk Moth," *Naturwissenschaften*, Vol 57, 1970, pp. 23-28. "Kai" [10]

Kalman, R.E., "Design of a Self-Optimizing Control System," Trans. ASME, Vol 80, No. 2, Feb 1958, pp. 468-478, in Oldenburger (1966), pp. 440-449. "Kal" [36]

Kalmus, H., "The Discrimination by the Nose of the Dog of Individual Human Odours," *Brit. J. of Animal Behaviour*, Vol 3, 1955, pp. 25-31. "Kal1" [9]

Kellogg, F.E., Frizel, D.E., and Wright, R.H., XXX, *Can. Entomol.*, Vol 94, 1962, p. 884. "Kel1" [47]

Kennedy, J.S., XXX, *Proc. Zool. Soc. of London*, Ser. A, Vol 109, 1939, p. 221. "Ken1" [47]

Kennedy, J.S., Marsh, D., "Pheromone-Regulated Anemotaxis in Flying Moths," *Science*, Vol 184, 31 May 1974. "Ken" [9, 10, 12, 47, 50]

Konijn, T.M., "Chemotaxis in the Cellular Slime Molds," in Carlile (1975), see supra, pp. 101-153. "Kon" [5]

Koshland, D.E., Jr., "A Response Regulator in a Simple Sensory System," *Science*, Vol 196, 3 June 1977, pp. 1055-1063. "Kos" [3]

Lorenz, K., *Behind the Mirror*, Harcourt Brace Jovanovich, New York and San Francisco, 1973. "Lor" [21, 34]

Lowenstein, W.R., editor, *Handbook of Sensory Physiology*, Vol I, Springer, Berlin and New York, 1971. "Low" [75]

MacNab, R.M., and Koshland, D.E., "The Gradient-Sensing Mechanism in Bacterial Chemotaxis," Proc._Nat._Acad._Sci., Vol 69, 1972, pp. 2509-2512. "Mac" [4]

Mandler, G., review of Popper and Eccles, The_Self_and_the_Brain, vide infra, in Science, Vol 200, No. 4345, 2 June 1978, pp. 1040-1041. "Man1" [67]

Mannix, D.P., The_Fox_and_The_Hound, E.P. Dutton, New York, 1967. "Man" [63]

Mast, S.O., Light_and_the_Behaviour_of_Organisms, John Wiley and Sons, London, 1911. "Mas" [6]

Mayr, E., Animal_Species_and_Evolution, Harvard University Press, Cambridge, Mass., 1963. "May1" [30]

McDermott, D., "Artificial Intelligence Meets Natural Stupidity," SIGART Newsletter, ACM Special Interest Group on Art. Intell., April 1976, pp. 4-9. "McD" [34]

Michener, M.H., and Michener, C.D., Social_Insects, D. Van Nostrand, Toronto, New York, and London, 1951. "Mic" [10, 11, 48]

Minsky, M., and Selfridge, O.G., "Learning in Random Nets," Information Theory, Fourth London Symposium, Butterworth's, London, 1961. "Min2" [38, 49]

Moray, N., "Visual Mechanism in the Copepod Copilia," Perception, Vol 1, pp. 193-207. "Mor" [7, 73]

Murlis, J., and Bettany, B.W., "Night flight towards a sex pheromone source by male Spodoptera_littoralis (Boisd.) (Lepidoptera, Noctuidae)," Nature, Vol 268, 4 Aug 1977, pp. 433-435. "Mur" [47]

Nultsch, W., "Phototaxis and Photokinesis," in Carlile (1975), see supra, pp. 29-90. "Nul" [6, 33]

Oatley, K., Brain_Mechanisms_and_Mind, Thames and Hudson, London, 1972. "Oat" [25]

Oldenburger, R., editor, Optimal_and_Self-Optimizing_Control, M.I.T. Press, Cambridge and London, 1966. "Old" [78]

Padula, and Arbib, M., XXX "Pad" [43]

Papentin, F. "XXX," *Journ. Theor. Biol.*, Vol 39, 1973, p. 431. "Pap" [65]

Pippenger, N., "Complexity Theory," *Scientific American*, Vol 238, No. 6, June 1978, pp. 114-124. "Pip" [31]

Pommerville, J., "XXX," *Experimental Cell Research*, Vol 113, 1977, pp. 161-167. "Pom" [15]

Popper, K.R., and Eccles, J.C., *The Self and Its Brain*, Springer-Verlag, New York, 1977. "Pop1" [79]

Pyke, G.H., Pulliam, H.R., and Charnov, E.L., "Optimal Foraging: A Selective Review of Theory and Tests,", *The Quarterly Review of Biology*, Vol 52, No. 2, June 1977, pp. 137-154. "Pyk" [64]

Raup, D.M., and Seilacher, A., "Fossil Foraging Behavior: Computer Simulation," *Science*, Vol 166, 21 Nov 1969, pp. 994-995. "Rau" [65]

Razran, G., *Mind in Evolution: An East-West Synthesis of Learned Behavior and Cognition*, Houghton Mifflin, Boston, 1971. "Raz" [27]

Ridenour, L.N., *Radar System Engineering*, Vol 1 of the Rad. Lab. Series, McGraw-Hill, New York and London, 1947. "Rid" [6]

Rock, I., *The Nature of Perceptual Adaptation*, Basic Books, New York and London, 1966. "Roc" [26]

Roeder, K.D., *Nerve Cells and Insect Behavior*, Harvard University Press, Cambridge Mass, 1963. "Roe" [47, 52]

Sagan, C., *The Dragons of Eden*, Random House, New York, 1977. "Sag" [9]

Satinoff, E., "Neural Organization and Evolution of Thermal Regulation in Mammals," *Science*, Vol 201, 7 July 1978, pp. 16-22. "Sat" [33]

Schank, R.C., and Abelson, R.P., *Scripts, Plans, Goals, and Understanding*, Lawrence Erlbaum Associates, Hillsdale, N.J., 1977. "Sch2" [59]

Schneider, D., "The Sex-Attractant Receptor of Moths," *Sci. Amer.*, Vol 231, No. 1, July 1974, pp. 28-35. "Sch" [9, 59]

Sebeok, *How Animals Communicate*, Indiana University Press, Bloomington, Ind., 1975. "Seb" [10, 45]

Shorey, H.H., *Animal Communication by Pheromones*, Academic Press, New York, San Francisco, and London, 1976. "Sho1" [10]

Shulenberger, E., "Contaminants in the Navy Medium," *Naval Research Reviews*, Vol 30, No. 9, Sept. 1977. "Shu" [12]

Spudich, J.L., and Koshland, D.E., Jr., "Non-Genetic Individuality: Chance in the Single Cell," *Nature*, Vol 262, Aug 5 1976, pp. 467-471. "Spu" [28]

Stanier, R.Y., and Cohen-Bazire, G., "The Role of Light in the Microbial World: Some Factors and Speculations," in Williams and Spicer (1957), see infra, pp. 56-89. "Sta" [2]

Stebbins, G.L., *Processes of Organic Evolution*, Prentice-Hall, Englewood Cliffs, N.J., 1966. "Ste" [30]

Talbot, M., "Slave-raids of the ant *Polyergus lucidus* Mayr.," *Psyche*, Vol 74, No. 4, pp. 299-213. "Tal" [11]

Tinbergen, N., *The Animal in its World*, Vol 1, Harvard University Press, Cambridge, Mass, 1972. "Tin" [61]

Vaissiere, R., "Morphologie et Histologie Comparees des Yeux des Crustaceens Copepodes," *Arch. de Zool. Exper. et Gen.* Vol 100, No 1, 1961, pp. 1-125. "Vai" [73]

Wallach, Hans, "Informational Discrepancy as a Basis of Perceptual Adaptation," in *The Neuropsychology of Spatially Oriented Behavior*, ed. Freedman, S.J., Dorsey Press, Homewood, Ill, 1968. "Wal1" [25]

Wells, M.J., and Buckley, S.K.L., "Snails and Trails," *Animal Behavior*, Vol 20, 1972, pp. 345-355. "Wel" [11]

Wigglesworth, V.B., *Parasitology*, Vol 33, 1941. "Wig" [4]

Wilkinson, P.C., "Chemotaxis of Leucocytes," in Carlile (1975), see supra, pp. 205-243. "Wil5" [5]

Williams, R.E.O., and Spicer, C.C., eds., *Microbial Ecology*, 7th Symp. of Soc. for Gen. Microbiology, Cambridge Univ. Press, Cambridge, U.K., 1957. "Wil2" [81]

Willows, A.O.D., "Learning in Gastropod Mollusks," in Corning et al., 1973 (see supra), Vol 2, pp. 187-274. "Wil" [11, 43]

Wilson, E.O., *The Insect Societies*, Harvard University Press, Cambridge Mass and London, 1971. "Wil4" [10, 11, 46, 56]

Winston, P.H., *Artificial Intelligence*, Addison-Wesley, Reading, Mass., 1977. "Win3" [57]

Wright, R.H., "Why Mosquito Repellents Repel," *Sci. Amer.*, Vol 233, No. 1, July 1975, pp. 104-111. "Wri" [4, 47]

Young, J.Z., *An Introduction to the Study of Man*, The Clarendon Press, Oxford, England, 1971. "You1" [28]

Young, L.R., "On Adaptive Manual Control," *Ergonomics*, Vol 12, No. 4, July 1969, pp. 635-674. "You" [38]

BOLT BERANEK AND NEWMAN INC