Random-Walk Chemotaxis: Trial and Error as a Control Process

Richard S. Marken The Aerospace Corporation El Segundo, California William T. Powers Northbrook, Illinois

The biased random-walk chemotaxis of the bacterium *Escherichia coli* is a remarkably effective method of navigation based on random trial-and-error responding rather than steering. Humans restricted to the same mode of responding are able to navigate to target locations, just like the bacterium. This mode of navigation can be modeled as an input control process that selectively retains favorable and rejects unfavorable consequences of the random responses. The selection process is determined by the internal organization of the system rather than the external influence of the environment (as in natural selection or reinforcement).

Control theory is commonly applied to human goal-seeking behavior in situations where behavior has moderately predictable influences on environmental processes, and those processes simultaneously have regular influences on behavior (Powers, 1973). A few years ago, however, we became interested in an apparent goal-seeking phenomenon that takes place through a highly irregular, in fact random, link (Marken, 1985). This phenomenon was described by Koshland (1980); it is the method that the bacterium *Escherichia coli* uses to make its way up concentration gradients of attractants and down gradients of repellents. We have simulated this behavior of *E. coli* using a control-system model and have extended the principle to experiments with human beings.

The control-theoretic analysis may improve understanding of the trial-and-error phase of learning, the phenomenon that Campbell (1960) called "blind variation and selective survival" (p. 205). Others have explored similar phenomena under the name of "hill climbing" (Krisky & Shik, 1964) but have not demonstrated them as examples of control-system operation (in at least one case, the claim has been made that control theory cannot handle this sort of behavior [Fowler & Turvey, 1978]). We present a model and a sequence of simple experiments that show how several variants of *E. coli*'s basic method of navigation can be reproduced in human trial-and-error behavior.

The Phenomenon

E. coli has only two modes of locomotion. It can spin all its flagella in the same direction and swim at constant speed in a straight line, or it can briefly reverse some of the organic motors to spin some flagella the other way, creating a disorganized tangle of filaments and causing the bacterium's body to tumble in space. Koshland (1980) determined that the direction of swimming after a tumble (when all the motors are again in forward gear) bore no discernible relationship to the direction before the tumble. The tumble simply resulted in a new direction of swimming at random. E. coli has no other method of steering.

Nevertheless, E. coli can make its way up a gradient of attractants at least half as fast as if it could simply turn in one direction and swim that way. This extraordinary efficiency, given the method, is what caught our interest. How could a random process have such a markedly systematic effect?

Koshland (1980) found the basic phenomenological explanation. Because swimming and tumbling occur in a stereotyped way, there is only one degree of freedom in which the total behavior can vary: the interval between tumbles. When E. coli detects a positive time rate of change of an attractant, the interval between tumbles increases in proportion to the rate of change; when the rate of change falls off or becomes negative, the interval between tumbles shortens in proportion. The bacterium's constant speed of swimming converts the angle of swimming through a concentration gradient into a time rate of change of concentration, as detected by the bacterium. Koshland (1980) determined these relationships quantitatively, through perfusion experiments with individual bacteria. As a result of the arrangement that he found, the bacterium spends much less of its time tumbling and swimming the wrong way than the right way.

A Control-System Model of E. Coli

As a first step in understanding this phenomenon, we attempted to make a computer model that would behave as $E. \ coli$ does. The purpose was to see what might be required to make such a model behave systematically and, later, to try to apply it to human behavior. We began with a standard diagram of a behavioral control system that many control theorists now use (Figure 1).

The input quantity q_i is taken to be the concentration of attractant present at the cell wall of the bacterium. Koshland (1980) showed how chemical receptors in the cell wall respond to the presence of the attractant by liberating "messenger" substances that diffuse inside the cell, constituting perceptual signals that represent the external concentration. He also offered several possible chemical mechanisms through

Correspondence concerning this article should be addressed to Richard S. Marken, The Aerospace Corporation, 2350 East El Segundo Boulevard, El Segundo, California 90245-4691.



Figure 1. Canonical control-system model of behavioral organization.

which these signals could be converted to effects that are measures of the time rate of change of concentration, that is, second-order chemical signals.

In the model, we represent these sensory processes and subsequent signal processing as a single "input function" that takes the first time derivative of the concentration and represents it as an internal signal, the perceptual signal, p.

On the output side of the model, there is an "output function" that converts an internal signal, e, into variations of the observable behavior. In *E. coli*'s case, this output function is very strange, at least if one is used to working with continuous functions. The output quantity, q_o , produced by this function is a direction of swimming, an angle in the laboratory coordinate system. The angle, however, is changed at random. We confine the model to two dimensions for simplicity, although we have verified that it works just as well in three. Thus, this angle is measured in the plane of movement of the bacterium, as if the animal were swimming in a thin film of liquid (or across a computer screen).

The basic output relationship needed is one that converts a given magnitude of the internal signal into a delay between tumbles. Although many mechanisms could be proposed, we selected one that is easy to compute and has at least some resemblance to biological phenomena. We assume that the output signal causes the accumulation of some substance that gradually rises in concentration toward a fixed threshold: When the threshold is reached, the tumble occurs, and the substance is "used up," or reset to zero. In this way, the signal can influence the timing of tumbles continuously, although the influence is only manifested when the tumble actually occurs. This picture bears a generic relationship to phenomena in neurons and muscle cells. The tumbles that are triggered are implemented in the model by random selection of a new swimming angle.

Finally, to connect the perceptual signal to the signal that drives the output, we use the "comparator," which computes the difference between the perceptual signal and the reference signal, p^* , producing an "error signal," e, representing the difference. The error signal is the signal that activates the output function. The comparator allows us to use the reference signal to adjust the baseline rate of tumbling so it is not zero when the time rate of change is zero, as in the real

organism. Variations in the reference signal can also be used, incidentally, to model the behavior of "mutants" observed by Koshland (1980): always-tumbling and never-tumbling bacteria, which would nevertheless stop or start tumbling, respectively, in sharp enough gradients of the right sign. Aside from these properties, we ascribe no physical reality to the comparator.

The last part of the model in Figure 1 is the connection between the output quantity and the input quantity—the effect of action on the variable being sensed. The sensed concentration changes at a rate determined by the gradient times the swimming velocity times the cosine of the angle between the swimming direction and the direction of the gradient. Computing the concentration itself is not necessary because it is only an intermediate variable. The input quantity may also be affected by disturbances, d, unrelated to the organism's output. The effect of disturbances on the input quantity is nullified by the outputs of a properly designed control system.

The model is expressed as a series of program statements. In the following, the symbol :=, or colon-equals, is used as in the Pascal language to indicate replacement, not equality. Herein, the variables are defined, but the program steps that initialize them are not. An asterisk is used to indicate multiplication in program statements, because the variable names may consist of several letters. With these guidelines, no one should find the program no harder to grasp than a series of algebraic statements.

Definitions:

x, y = position of the bacterium

- xt, yt = position of the "target" (center of gradient)
 - v = velocity of swimming, fixed
 - a = swimming angle, radians
 - ϕ = angle from bacterium to target, radians
 - G = magnitude of radial gradient, fixed
 - p = perceptual signal
 - p^* = reference signal (adjustable parameter)
 - e = error signal, equal to $p^* p$
 - K = output scaling constant (adjustable parameter)
 - dq = size of "accumulation" increment for timing
 - q = current amount accumulated
 - Q = threshold of accumulation for tumble, fixed.

The program steps are as follows, with initialization omitted.

> 1. $x := x + v \cdot \cos(a);$ 2. $y := y + v \cdot \sin(a);$ 3. If abs [(x - xt)] > 0 then $\phi := \arctan[(y - yt)/(x - xt)]$ Else $\phi := \operatorname{arccot} [(x - xt)/(y - yt)]$ 4. $p := G \cdot v \cdot \cos(\phi - a);$ 5. $e := p^* - p;$ 6. $dq := K \cdot e;$ 7. q := q + dq8. If q > Q then



In Step 3, "abs" is the absolute value function. In Step 8, the "begin... end" pair contains statements that are executed only if q > Q; "twopi" is a constant set to $2 \times \pi$ (3.14159). The "random" function returns a number to between zero and 2π radians, which is substituted for the existing value of angle *a*. This program sketch should suffice to show anyone who can program in Pascal, C, or BASIC how to set up a working version, and to suggest how to set the model up algebraically to run it by manual calculations.¹

Model Behavior

The behavior of this model is shown in Figure 2. The model bacterium, starting on the left, finds its way up the gradient toward the source of attractant on the right in a series of random changes of direction. Because the speed of movement is always the same, the length of a segment indicates the delay between changes in direction. When the model bacterium arrives close to the source of attractant, it simply hovers in the vicinity, making a dense cloud of tracks. As can be seen, there is motion in all possible directions, but the tumbles occur the least frequently when the direction is toward the target.

By adjusting parameters, we have been able to make this process as much as 70% as efficient as a straight-line motion to the target, in terms of average velocity in the right direction. The model has also been tested in gradients in one and three dimensions and in planer and inverse-square concentration fields. Nothing seems to faze it.

The absence of systematic changes in direction gives this kind of arrangement the ability to work under a variety of external conditions. Even when there are local "pockets" of altered gradients, the model has a good chance of getting out of them simply by encountering a run of either good or bad luck in random direction changes, that is, up or down the gradient. This mode of action presumes little about the properties of the world surrounding it. Where a systematically behaving organism depends on the world's maintaining its properties reasonably constant, this randomly acting system can work even under radical changes of conditions. One's initial impression of this mode of behavior is that it is crude, wasteful, and stupid. After watching the model work for a while, however, one sees that it is a powerful way of reaching a goal, particularly for an organism that lacks higher computational abilities and spatial perception. The method is the only feasible way for an organism to maintain control over



Figure 2. Typical random-walk behavior of spot produced by control-system model.

important effects on itself, when its environment is totally beyond its comprehension.

Experiment 1

Having a basic model that works reasonably well, we were able to explore the same general principle with human subjects.

Method

Two male and 2 female adult humans between the ages of 14 and 60 years participated in the experiments. All had normal visual and manual ability. The subjects were restricted to the same means of locomotion as E. coli. For the swimming of the bacterium, we substituted a spot of light moving in a straight line at constant speed on a computer screen, and for the tumbles, random, computergenerated changes in direction were initiated by the subjects' tapping the space bar on the computer keyboard. Each tap caused a new direction of spot movement to be picked at random. We verified that the new directions were evenly distributed over 2π radians and were unrelated to old directions. The random change of direction constituted the subjects' only means of "steering" the spot toward the target, which was a small square at the right center of the screen. The moving spot was placed initially near the left edge of the screen. Subjects were instructed to tap the space bar to make the spot reach the target and remain near it.

Results

All 4 subjects, even without practice, were able to accomplish the task. The movements of the spot during one experimental run, which was typical, are shown in Figure 3. Because spot speed was constant, the length of a segment indicated the delay between changes in direction. Even while we watched the experiment progress, it was difficult to believe that the subjects' only influence on the spot's motion was to produce a random change in direction. Figure 3 does not show all changes in direction, for sometimes a previous

^{&#}x27;We will supply this program and other programs mentioned in this article ready to run on a small-screen Macintosh or IBM PC, XT, or AT with CGA or Hercules monochrome graphics on receipt of a formatted disk in reusable mailer with return postage. Turbo Pascal source code will be included.



Figure 3. Typical behavior of spot produced by human subject.

movement was partly retraced, and sometimes the movement in the wrong direction lasted so short a time that several changes in direction appear to be a single change.

The human subjects behaved in the same way as *E. coli*. The most important result is the finding that subjects were able to move the spot to the target *on every trial*. A model of this behavior must be able to move the spot to the target with the consistency displayed by the subjects (success on every trial), although it need not mimic precisely the route taken to the goal on each trial.

Discussion

Modeling the human behavior. As already illustrated, the approach taken by control theorists in proposing an explanation of behavior is, whenever possible, to construct a working model that will reproduce the same behavior. By requiring that the model actually run and that it be capable of generating predictions of behavior through time, we can be sure that the offered explanation is complete in itself, not requiring added interpretation to make the model capable of predicting a specific behavior. The behavior generated by the model can then be compared with the behavior of the subject.

The model used to reproduce the human behavior is identical to the model used for *E. coli*. Because there was no chemical gradient, we assumed that subjects simply observed the component of spot velocity toward the target (or some aspect of the situation giving equivalent information). Under that assumption, the computation of the perceptual signal pis exactly the same as shown in the program steps, with *G* set to 1. Then, because we could obtain the experimental data directly, we could estimate the two adjustable parameters, the reference signal p^* and the scaling constant *K* in the output function.

Evaluating K and p^* . The value of the Output Function Constant K determines how rapidly the time delay will change for a given change in perceptual signal. An increase in perceptual signal decreases the timer's accumulation rate and thus affects time delay inversely. The value of K is thus measured by the slope of the best line matching a plot of perceptual signal versus 1/delay. The perceptual signal is numerically the same as the velocity component toward the target. On the same plot, the timer's accumulation rate will just fall to zero when $p = p^*$. Thus, p^* is evaluated at the point where the best fitting line crosses zero.

Figure 4 is a plot of 1/delay versus perceptual signal (velocity toward the target), with the best fitting line of the data that are presented in Figure 3. Averaged over all 4 subjects, the value of K was -7.2, and p^* was 1.17. The maximum possible value of p is actually only 1.0 (when the spot moves directly toward the target); p^* may be greater than 1.0, so that e (and hence dq) never becomes 0, which would result in an infinite delay between responses. The value of K used in the model was 7.2 rather than -7.2 so that dq would increase and the delay between responses would decrease with an increasing error signal.

To save space, we plotted the results of a model run with these parameters in Figure 2. For that run, the constants were set to Q = 100, G = 1, and v = 1. The model behaved like the subjects, moving irregularly but steadily to the target on every trial.

Experiment 2

Test for the Controlled Variable

The controlled variable in this experiment is proposed to be the component of velocity toward the target. It is this component that should be stabilized against disturbances by the behavior if there is control (Marken, 1988). To measure this stability, we first measured the effects of undisturbed behavior on the controlled variable, then we calculated the effects on the controlled variable that would be observed with



Figure 4. Plot of 1/delay versus p (velocity toward target) with best fitting regression line.

a disturbance present if there were no compensating change in the system's action, and finally, we measured changes in the variable when disturbances were present during an experimental run. If control is present, the controlled variable should change in the last step much less than predicted by the middle step.

Method

The subjects were the same as in the previous experiment. We first asked them to keep the spot near the target, once it arrived there, for l min. Then a second run was done with the addition of a disturbance, d. The disturbance consisted of a 0.005-cm addition to the y position of the spot every 0.05 s. The x and y positions of the spot were determined by the following program statements:

$$x := x + v \cdot \cos(a)$$

$$y := y + v \cdot \sin(a) + d.$$

Results

The y component of radial velocity (scaled to screen units) was averaged over the entire minute after the spot arrived at the target. The same treatment was given to the behavior of the model. The average y velocity of the spot, averaged over all subjects, was 0.016 cm/s for the experimental run, with no disturbance added; for the model run, the average y velocity was the same, 0.016 cm/s. A typical subject run is shown in Figure 5a.

Next, we calculated the predicted effect of a steady 0.1cm/s disturbance added to the y velocity of the spot using the data from the preceding subject run. If the action of the subject remained the same, we would expect an average yvelocity (averaged over subjects) of 0.116 cm/s (the average from the no-disturbance run, 0.016 cm/s, plus 0.1 cm/s because of the effect of the disturbance). This calculated result is shown in Figure 5b. Note that the result of the disturbance, when unopposed, was to move the spot vertically below the target (in screen coordinates, increases in y velocity led to downward movements of the spot).

Finally, we found that the average y velocity (averaged over subjects) for experimental runs with the disturbance added was 0.018 cm/s. Thus, the control behavior cancelled some 90% (or more) of the effect that an unopposed disturbance of this kind would have had. The model showed the same effect. The result of a typical subject run with the disturbance added is shown in Figure 5c.

Discussion

We conclude that velocity toward the target is the controlled variable, or depends directly on the true controlled variable. The behavior fits the definition of control behavior in the respect that a disturbance is counteracted. The same result is observed for disturbances in any direction.

Experiment 3

Further Tests for the Controlled Variable

The model is ensured to respond only to the component of velocity toward the target, but there are many more pos-



Figure 5. (a) Subject run with no disturbance. (b) The computed effect of 0.1-cm/s disturbance to y velocity component of above run. (c) Subject run with 0.1-cm/s disturbance to y velocity added.

sibilities for the subject, who can see the full two-dimensional situation. To narrow the possibilities, we changed the display so that the radial distance of the spot from the target was shown as the vertical position of a short horizontal bar of light on the center meridian of the screen.

Method

The subjects were the same as in the previous experiment. The subjects saw a horizontal bar that moved vertically on the computer screen. The distance of the bar below the top of the screen was proportional to the radial distance of the spot from the target in two-dimensional space: $[(x - xt)^2 + (y - yt)^2]^n$. The two-dimensional target location corresponded to the top of the screen. The subject was instructed to move the bar as close to the top of the screen as possible and then keep it there for 1 min.

As in the previous experiment, each subject did one experimental run with no disturbance to the y velocity and a second run with a



Figure 6. (a) Plot of radial distance from target over time for subject run with no disturbance. (b) Twodimensional movements underlying radial distance plot. (c) Radial distance plot with computed effect of 0.1-cm/s disturbance to y velocity of above run (panel a). (d) Two-dimensional movements underlying radial distance plot. (e) Radial distance plot with 0.1-cm/s disturbance added to y velocity. (f) Twodimensional movements underlying radial distance plot.

0.1-cm/s disturbance added. The subject no longer saw the twodimensional display as before. All the computations were done in two dimensions, but the program extracted the radial distance to the target and presented only that to the subject's view. Thus, the subject could no longer see the spot's relationship to the target or its direction of movement.

In carrying out the instruction, the subject might attend primarily to the vertical velocity of the bar, the choice embedded in the model, or primarily to the momentary distance from the bar to the top of the screen. By altering the input function of the model, we can make it sense the distance from spot to target instead of the velocity. Thus, we have two versions of the model and can use the test for the controlled variable to choose between them.

Results

Panel a of Figure 6 shows the behavior of the bar plotted against time, and b shows the two-dimensional relationships that underlay what the subject saw. The two-dimensional view shows the same sort of biased random-walk approach to the target as before. Panels c and d show the effect that adding a 0.1-cm/s disturbance would have on the y position of the invisible spot, with the data from the previous run. This is the calculated effect of an unopposed disturbance (assuming the subject used the same response pattern as in the previous run). Panels e and f show the results when the disturbance was added during a run. The average y velocity (averaging over all subjects) when no disturbance was added was 0.013 cm/s. The expected y velocity with the disturbance added was 0.113 cm/s. In fact, the average y velocity (averaging over all subjects) was 0.016 cm/s. The test for the controlled variable was passed as before.

Discussion

Using the existing velocity-control model, we found that p^* (averaging over all subjects) was 1.14 and K was -10.0.

Resistance to the disturbance occurred as for the subjects' runs, and the plots were essentially like those of Figure 6.

To convert the model to a position-control system, we altered the input function so that $p = [(x - xt)^2 + (y - yt)^2]^{\frac{1}{2}}$. Then we generated a plot of perceptual signal versus 1/delay. From this plot, we obtained K = -0.8 and $p^* = 88.2$. The value of p^* makes no sense-it implies that the subject's reference for radial position was well away from the target. The reason for this absurd result is easy to see. There is no consistent relationship between delay and radial distance from the target. We were unable to find any values for the parameters that would let the position-control model move the spot to the target consistently. We were surprised at how completely the position-control model failed in this situation. It seemed plausible that a system could reach the goal by accepting or rejecting results that moved it closer to or farther from the target, respectively. The model convinced us that this was not the case. We conclude that radial distance from the target cannot possibly be the controlled variable in this experiment: subjects must be controlling the component of spot velocity relative to the target.

General Discussion

In these studies, we were exploring a kind of behavior that has been called trial-and-error behavior, that is, achievement of a result by essentially random movements. When the behavioral variable is seen to be the timing of the random changes, however, and when the behavior is put into the context of control theory, we find that highly systematic results are to be expected. The systematicity is not simply the kind expected from "natural selection," but it comes close to what we would expect from a well-organized and coordinated behaving system. This behavior can be seen as intermediate between the blind operation of chance and the systematic execution of well-learned performance. Recent discoveries in evolutionary biology (Cairns, Overbaugh, & Miller, 1988) suggest that this sort of "directed evolution" might be found in organizations even simpler than E. colieven in the genetic material itself. Coincidentally, Cairns et al. (1988) found the capacity in a bacterium to alter its rate of mutation to have systematic inheritable effects on its ability to metabolize lactose: The bacterium was E. coli.

The fact that human subjects fall quite naturally into using this mode of control suggests that is not unfamiliar to them. One of the puzzles of behavior is how organisms make the transition from fumbling through a new task to mastery of it. Perhaps the random-control mode is always the beginning mode of control, whatever is being learned. An organism can actually acquire some degree of control without any ability to generate a specific appropriate behavior; one can almost imagine how control that begins in this way gradually firms up into systematic control. The biggest problem in modeling learning is that a "teacher" always seems necessary to point the learning system in the right direction. It may be that pseudorandom control would provide an internal teacher.

One basic principle brought out by control theory is that control systems vary their actions to control their inputs. We have shown here than human subjects restricted to pure trialand-error behavior manage to control the sensed variable, defending it against disturbance and maintaining it in a specified state. The fact that output is varied but not controlled is doubly underlined in these experiments, for the human subjects have no means of choosing the direction of movement that will result from a tap of the key. They could not control the direction of movement even if that were asked of them.

What makes this kind of organization work is not the generation of an action that will have a predictable effect but the use of a mode of sensing that picks out of the effects of action from those that the system specifies—in our model, by setting a reference signal to some specific value. What is sensed determines what is controlled.

There could be no better illustration of the meaning of Campbell's (1960) principle of "blind variation and selective survival." The "survival" part of the principle results not from some external influence that strengthens a given tendency to respond but from the organization inside the organism that postpones random changes when the presenttime state of the sensed variable is favorable, as judged by the organism against internal criteria. We have thus completed Campbell's thought by showing how blind variation can depend on selective criteria. The blind variations supply a whole range of real consequences to the organism. The organism, by advancing or postponing the next blind variation, can reject or accept any consequence.

The principle we see here is the inverse of the concept of reinforcement (Skinner, 1981). Where reinforcement theory assumes some external influence that increases the probability of repeating an action, control theory (in this application) assumes an internal selection process that decreases the probability of changing to a different action. Both principles explain what we see. But there should be a way of picking one view over the other. It would help if reinforcement theorists could devise a reinforcement-based model that would generate behavior through time.

References

- Cairns, J., Overbaugh, J., & Miller, S. (1988). The origin of mutants. Nature, 335, 142-145.
- Campbell, D. T. (1960). Blind variation and selective survival as a general strategy in knowledge-processes. In M. C. Yovits & S. Cameron (Eds.), Self-organizing systems (pp. 205-231). New York: Pergamon Press.
- Fowler, C., & Turvey, M. (1978). Skill acquisition: An event approach with special reference to searching for the optimum of a function of several variables. In G. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 1-39). New York: Academic Press.
- Koshland, D. E. (1980). Behavioral chemotaxis as a model behavioral system. New York: Raven Press.
- Kriskiy, V. I., & Shik, M. L. (1964). A simple motor task. *Biophysics*, 9, 661–666.
- Marken, R. (1985). Selection of consequences: Adaptive behavior from random reinforcement. Psychological Reports, 56, 379-383.
- Marken, R. (1988). The nature of behavior: Control as fact and theory. Behavioral Science, 35, 167-178.

Powers, W. T. (1973). Behavior: The control of perception. Chicago: Aldine. Skinner, B. F. (1981). Selection by consequences. Science, 213, 501-

504.

Revision received January 5, 1989

.

Accepted March 10, 1989

Received April 8, 1988