



Does Complex Behavior Require Complex Representations?

James S. Magnuson (MAGNUSON@BCS.ROCHESTER.EDU)

Department of Brain and Cognitive Sciences, University of Rochester
Meliora Hall, Rochester, NY 14627 USA

Abstract

Models in cognitive science often postulate that individuals maintain complex representations of their environment when simpler explanations, based on simple behaviors interacting with each other and environmental constraints, would suffice. As an example, I consider representational approaches to animal behavior (e.g., Gallistel, 1990; Myerson and Miezin, 1980), which posit that complex group behavior results from complex representations of events within the central nervous systems of individual animals. For example, ducks feeding from two food sources distribute themselves proportionately to the density of food available at each source. This phenomenon, probability matching, is typically explained by attributing representations of the density of food available at each source within the central nervous system (CNS) of each duck. Are such complex representations required to explain this phenomenon? I will compare the results of two simulations of probability matching in groups. In one, individuals maintain and update representations of food available at each source. Although probability matching emerges, the organisms exhibit various unrealistic behaviors. In the second, each individual follows simple behavioral rules but has no representation of the food density at each source. Probability matching emerges and the behavior observed is more realistic than that in the first simulation. This adds to demonstrations in other domains that complexity at one level of analysis need not result from complexity at lower levels (e.g., Resnick, 1994; Sigmund, 1993).

Centralized versus decentralized analysis

It used to be said that our ability to build arches distinguished humans from lesser species -- arches require planning and engineering. It turns out that at least one lesser species -- termites -- builds arches. In fact, it appears that their arch building depends on highly organized cooperation. Scores of them deposit excrement, eventually forming several columns, which eventually join to form arches. The arches form the foundation for their nests. How is such an amazing joint activity organized? One possibility is that each insect is acting as part of a team. At first blush, it appears that each insect monitors and represents the activities of the others and their collective progress. As Braitenberg (1984) points out, we often prefer such explanations, even for simpler systems, and are quite willing to attribute complex representations and motivations when simpler explanations would suffice. Another possibility is

that the complex group behavior emerges from the interactions of large numbers of insects acting on the same small behavioral repertoire. Indeed, this is the case. Each termite develops a seasonal sensitivity to a particular pheromone secreted in their waste. Where termites deposit their waste is guided by gradients of pheromone concentration. Eventually, columns form where the concentrations are strongest. Competition between adjacent columns leads to their being joined into (see Kugler et al., 1989, for a detailed analysis). Resnick (1994) describes these two ways of analyzing problems as centralized and decentralized thinking (for many more examples of decentralized analysis, see Resnick, 1994; Braitenberg, 1984; Sigmund, 1993; and Langton, 1989).

Decentralized thinking is not new in cognitive science. Consider this list of demonstrations that decomposing a system into simpler components can reduce the complexity of representations required: Wiener's (1948) treatment of self-organizing systems, Gibson's analysis of optic flow (e.g., Gibson, Olum, and Rosenblatt, 1955), the "active" or "animate" approach to vision (e.g., Ballard, 1991), Reynolds' (1987) treatment of flocking behavior, and the literature cited above. Despite these compelling demonstrations, decentralized analysis has not received the attention it warrants. Many models in cognitive science start out "representation-heavy." I do not wish to imply that there is no place for representation in models of behavior and cognition. My thesis is that as a starting point, models of complex systems need not postulate complexity at the lowest level. To demonstrate this point, I will apply decentralized analysis to the phenomenon of probability matching in groups of foraging animals.

Probability matching

Various species adapt quickly to changes in the probability of important stimuli in their environments. For example, when presented with multiple sources of food in an environment, individuals quickly divide into groups with sizes proportional to the density of food available from each source (see Gallistel, 1990, chapter 11, for an overview). It is important to note that the overall density of food is what determines the size of the groups. Given two sources dispersing morsels of equivalent magnitudes at equivalent rates, the animals divide into two groups of approximately equal numbers clustered around each of the sources. If the amount of food available from one source increases, due to an increase in either its rate of dispersal or the magnitude of

the morsels of food, the sizes of the groups of animals change proportionately. Gallistel (1990) and others (e.g., Myerson and Miezin, 1980) have explained such "probability matching" behavior as the result of complex and specific representations of the food source characteristics in the central nervous systems of individual organisms:

[Results suggest] that birds accurately represent rates, that they accurately represent morsel magnitudes, and that they can multiply the representation of morsels per unit time by the representation of morsel magnitude to compute the internal variables that determine the relative likelihood of their choosing one foraging patch over the other. (p. 358)

Gallistel goes on to present this as a particular challenge to nonrepresentational models:

The challenge for nonrepresentational theories ... is to propound a nonrepresentational model of ... a system that is altered by its past experiences in such a way that it chooses patches in proportion to their relative [densities] ... from ... observations alone ... without the system's having any internal variable[s] ... or a fortiori any operations that appropriately combine these nonexistent representatives of number, time, and magnitude. Is it possible to propound a model of the internal causation of the duck's behavior that avoids postulating an isomorphism between a system of variables and operations inside the duck and the corresponding system of number, time and magnitude external to the duck? (p. 359)

Are the complex representations Gallistel and others describe truly crucial for explaining probability matching by, for example, groups of ducks? Could animals instead follow simple rules, with the result that the probability of food density would be represented in the distribution of the group, rather than the CNS of each individual? Indeed, the equations that describe a system do not necessarily reveal causation. Reaction-diffusion equations (see Winfree, 1980) predict patterns of diffusion in various systems (e.g., slime molds), but those systems can be simulated in cellular automata using simple, locally-defined rules (e.g., Resnick, 1994). In the next section, I consider a nonrepresentational approach to probability matching that does not require internal representations of environmental variables or past experiences.

A nonrepresentational approach

From the decentralized perspective, the problem is not a matter of constructing a model by which each organism can represent the complexity of the environment. Rather, it is a matter of determining the simplest possible representations and behaviors, and constraints of the environment which, when they interact, could result in the emergence of the larger group phenomenon. The approach taken here is influenced by the work cited in the first section, as well as Simon's (1957) notions of "satisficing" (rather than optimizing or maximizing) algorithms, as expressed by Gigerenzer and his colleagues (e.g., Gigerenzer et al., 1991).

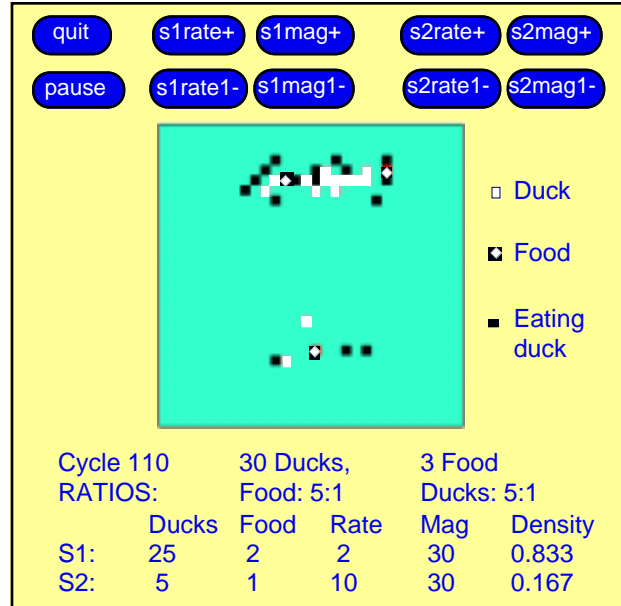


Figure 1: The simulator. In all simulations reported, 30 ducks were placed in a 30 x 30 "pond," with a food source at the top, and another at the bottom. Rate of dispersal and morsel magnitude could be specified for each source.

In the duck feeding example, rather than maintaining complex representations, the ducks may follow a rule as simple as "go to the closest morsel of food." Environmental constraints might also contribute to the emergence of the complex behavior. For example, morsel magnitude might be reflected in the size of the groups formed simply because it takes longer to eat large morsels, resulting in a larger number of visible morsels near a source of larger food. Suppose the morsels at one food source are larger than those at another source with an equal rate of dispersal, and, at a given moment, there are equal numbers of ducks at each food source. Soon, more food will be available at the source with the larger magnitude, because ducks at that source will require more time to consume each morsel.

Simulations

The general simulator

The environment was a 30 by 30 grid (the "pond") with two sources of food located at the top and bottom edges. Various characteristics of the sources could be specified (and changed interactively during a simulation -- although no interactive changes were made during the simulations reported here). For example, the rate of dispersal from the source (actually the interval between dispersals, so that a rate of "2" is slower than a rate of "1"), and the magnitude of the "morsels" of food could be specified.

In the simulations, when a duck made contact with a morsel of food, the duck would remain stationary until it

was finished "eating." Eating time, in update cycles, was defined as the magnitude of a morsel of food; given a morsel of magnitude 5, a duck eating that morsel would remain stationary for 5 update cycles. At the beginning of a simulation, a specified number of ducks was randomly distributed throughout the pond. At each time step, or cycle, food was distributed from a source (to a random, unoccupied location within a specified distance from the source) if the cycle number was divisible by its rate of dispersal. Subsequently, the locations of the ducks were updated in random order, in accordance with the constraints described in the next section. Simple obstacle avoidance procedures were implemented. If, for example, a duck en route to a morsel encountered a stationary duck, it could change direction to find an unimpeded path to the nearest morsel of food. An example simulator window is shown in Figure 1.

Nonrepresentational implementation

In the nonrepresentational implementation, duck behavior was determined by one rule and one explicit constraint of the environment. The rule was simply to approach the nearest morsel of food. The environmental constraint was that once a duck contacted a morsel of food, it had to remain stationary while it "ate." As mentioned in the previous section, a morsel of food of magnitude 5 would require the duck to remain stationary for 5 update cycles.

Certainly, other behaviors and constraints are possible. However, these two were chosen as the simplest possible. If others were used (e.g., ducks could maintain a representation of how much food they obtained at a source and probabilistically decide to change sources when they obtained too little), a behavior very much like the one used would still be needed to move the ducks to their desired source and to find food.

Representational approach

Gallistel (1990) presents a computational model (based on Myerson and Miezin's 1980 model) that accurately predicts the behavior of foraging animals. In the model, each duck represents the rate of food distribution at each food source and the magnitude of morsels being distributed. These variables are then combined neurally to represent food density at each source, yielding "relative patch affinities" (equivalent to the proportion of food at each source). Relative patch affinities multiplied by a "switchiness" constant (how often the animal is willing to change sources) are assumed to be equal to a Poisson process which predicts how much time a duck will spend at each food source.

In this implementation, each duck maintained an explicit representation of each source's density. Relative patch affinities were simply the proportion of food available at each source. However, it is not clear exactly how or when ducks should update their representations. This question becomes crucial if one attempts to implement the model as a simulation. If ducks constantly evaluate which source they prefer in a probabilistic fashion, how will they ever swim across the pond to the less dense source? If there is a probability of .70 of choosing one source (due to relative patch affinities), but the other source is 20 cells away

(requiring 20 update cycles' travel), a duck must choose the less likely source for several cycles running -- an unlikely event. Either complicated rules about when to update a representation or some sort of switchiness factor is also required, and one was implemented. For example, if switchiness was set to 80, this would mean that if the current preference (based on relative patch affinity) required switching sources, a second decision was made with an 80% chance of changing sources.

With switchiness set to 100 (so that it would have no effect), probability matching did not occur. Instead, all of the ducks would seek the densest source, since it was so unlikely that they would "choose" the less dense source enough times in one period of time to actually move there. In fact, simply setting switchiness substantially below 100 was not sufficient. For probability matching to emerge, I had to modify the switchiness factor to only come into play when a decision was being made to change preferences from least to most plentiful (i.e., it would never come into play if a duck decided to leave the more plentiful source). Trial and error determined that the optimal switchiness factor was 80.

Finally, although Gallistel's model is supposedly one of "internal causation," we must consider exactly what it is it causes ducks to do. "Choosing" a source is not enough. We must also consider how ducks implement their choice: how they travel between sources and arrive at a morsel of food. The model does not address this level. In order to give the ducks behavior in this implementation, they followed a slightly modified version of the behavior used in the nonrepresentational implementation: seek the nearest food *at the preferred source*.

General parameters

In every simulation, 30 ducks were used. As mentioned above, switchiness was set to 80 in all simulations run with the representational implementation. All simulations were run for 300 update cycles. The two implementations were compared by running two sets of simulations.

First, 100 simulations were run with each implementation with two food sources that differed in rate of dispersal. The rate of one source was set to 2 (i.e., a morsel of food was dispersed from that source every second cycle), and the rate of the other was set to 10 (i.e., a morsel of food was dispersed from that source every tenth cycle). Morsel magnitude was set to 30 at both sources (therefore, any morsel of food required 30 update cycles to consume). In this case, the expected distribution ratio of ducks would be 5:1, since the rate of dispersal at the first source was five times greater than that at the second source (i.e., the system should settle with 25 ducks at the source with the faster dispersal rate, and 5 at the other, slower source).

In the second set of simulations, rate was held constant, but magnitude differed between the two sources. 100 simulations were run with each implementation. Rate of dispersal was set to 5 at each source (i.e., a morsel of food was dispersed from each source on every fifth cycle). Morsel magnitude was set to 80 at one source, and to 40 at the other. Thus, the expected distribution ratio would be 2:1, since the density at the source with larger morsel magnitude was twice that of the other source (i.e., the system should

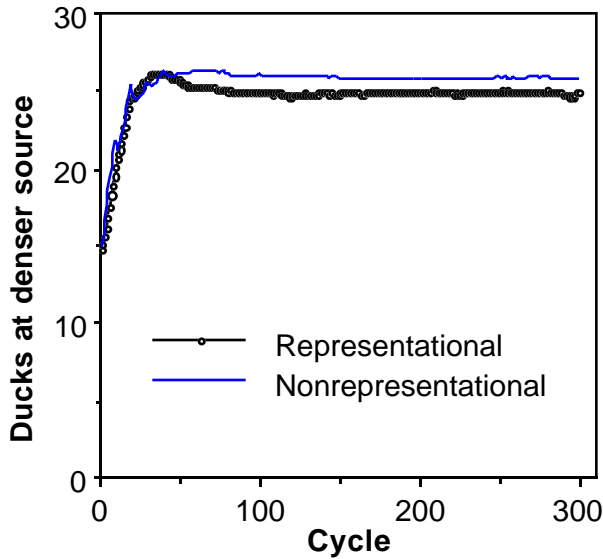


Figure 2: Probability matching to differences in rate in the two implementations. Morsel magnitude was set to 30 for both food sources. The dispersal rate of the denser source was set to 2. The dispersal rate of the other source was set to 10. With 30 ducks, the predicted number of ducks at the denser source is 25.

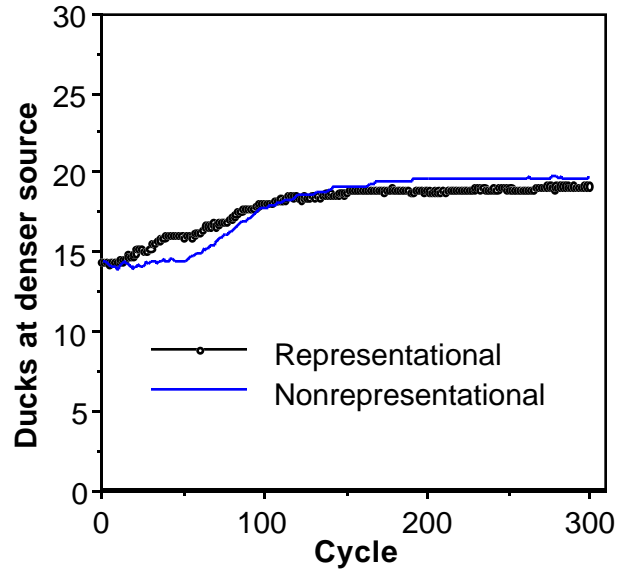


Figure 3: Probability matching to differences in morsel magnitude in the two implementations. Dispersal rate was set to 5 for both food sources. The magnitude of the denser source was set to 80. The dispersal rate of the other source was set to 40. With 30 ducks, the predicted number of ducks at the denser source is 20.

settle with 20 ducks at the denser source, and 10 at the other, less dense source).

Results

In general, the results of the both implementations fit the reported data -- that is, probability matching was observed. Ducks divided into two groups proportional to the food density at the two sources; they matched changes due to differences in rate and magnitude. However, not all combinations of number of ducks, rates and magnitudes resulted in probability matching. An examination of the environmental constraints -- the size of the pond, the distance between the sources, the size of the area a source disperses food into, the rate and magnitude values for the sources, and the number of ducks -- explains why.

For example, consider the relationship between magnitude and rate of dispersal in the simpler, nonrepresentational implementation. A change in magnitude will not have any effect unless the magnitude is approximately a multiple of the rate, since food dispersed at time t can be consumed by time $t + \text{magnitude} + 1$ (the minimum time for a duck to move to an adjacent morsel of food is one update cycle). For a duck to change sources, the only food available in the environment must be from a different source. This condition must remain true until the duck is closer to the new source's food than any new food from its original source, since the duck will seek the closest food. Therefore, for the system to converge on the distribution of ducks predicted by food density, food can be neither too plentiful -- in which case ducks have no reason to change sources -- nor too scarce -- in which case food will be consumed by a duck close to a source before other ducks can change sources

(Gallistel, 1990, notes that similar deviations from predicted probability matching have been reported in the literature; he explains them as a result of food being consumed before other ducks can update their representations). However, by varying only rate or magnitude, it is possible to observe behavior which closely resembles probability matching.

Differences in Rate Both implementations match the probability matching predictions for differences in rate. With the two sources' morsel magnitudes set to 30, one source with a rate of 2, and the other with a rate of 10 (5 times slower than the first), the group of 30 ducks split into two groups of approximately 25 (near the first source) and 5 (near the slower source) within 100 update cycles in both implementations (see Figure 2).

Differences in Magnitude Both implementations match the probability matching predictions for differences in magnitude, although effects of magnitude are more difficult to observe than those of rate. For the difference between sources' magnitudes to affect the distribution of ducks, it must be so large that food from the source with the larger magnitude remains available long enough that ducks at the opposite source can get closer to it before more food is distributed from the opposite source. So the "swimming speed" of the ducks and the distance between the sources also come into play. If magnitudes are set sufficiently high, however, the relative contributions of those secondary factors decrease, and probability matching emerges. With the two sources' dispersal rates set to 5, one source with a morsel magnitude of 80, and the other with a morsel magnitude of 40 (half the size of the first), the group of 30

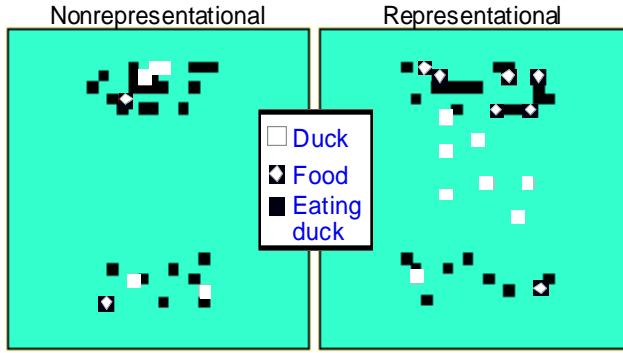


Figure 4: Examples of observed on-line differences between the implementations during simulations when magnitudes differed. The left panel is a snapshot of the nonrepresentational implementation, and the right panel is a snapshot of the representational implementation. In both cases, approximately 150 update cycles have taken place.

ducks split into two groups of approximately 20 (near the first source) and 10 (near the other source) within approximately 300 update cycles in both implementations (see Figure 3).

Discussion

Although it is not clear how to equate update cycles with real time -- with, e.g., the real swimming and eating rates of ducks -- the difference between convergence times for rate and magnitude differences are similar, in general, to those observed with real animals. Harper (1982) found that distributions of real ducks matched differences in rate within approximately 90 seconds. The same group of ducks matched differences in morsel magnitude within approximately 300 seconds. Compare these numbers with Figures 2 and 3: differences in rate were matched after




Figure 5: The number of ducks changing sources at each cycle in the two implementations when magnitude differed.

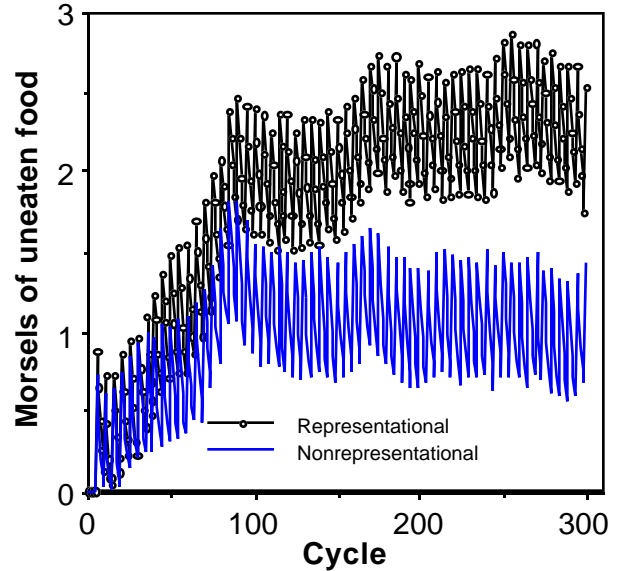


Figure 6: The amount of food left uneaten at each cycle in the two implementations when magnitude differed.

approximately 60 cycles, and the match to differences in magnitude did not asymptote until nearly 200 cycles had passed.

Both implementations provide similar simulations of the phenomenon. However, there were two major differences between them: the representational simulation was less realistic and less stable than the nonrepresentational simulation.

The difference in realism is illustrated in Figure 4. Snapshots of both implementations of the simulator are shown after approximately 150 update cycles when magnitudes differed. In the representational implementation (lower panel), there is much more uneaten food and there are several ducks that are not close to either source. During the representational simulations, ducks would often ignore the closest food -- even immediately adjacent food -- when their source preference had changed from their current location. Also, it was common to observe several ducks "stuck" in the region between sources for long periods, as their preferences oscillated between sources.

These differences between implementations are quantified in Figures 5 and 6. In Figure 5, the rate of source switching is plotted for simulations in which magnitude differed. On average, there were many more ducks changing sources at a given time in the representational simulations. In Figure 6, the amount of uneaten food available is plotted for simulations when magnitude differed. (The cyclic pattern in Figures 5 and 6 is due to new food being introduced every fifth update cycle, as dispersal rate was set to 5 for both sources in each implementation.) In the nonrepresentational implementation, a relatively steady state was reached after approximately 100 cycles, with approximately 1 uneaten morsel available. On average, in the representational implementation, there was more than twice as much uneaten food available after the nonrepresentational implementation reached its steady state.

While it would be simple to add rules designed to prevent the unrealistic behaviors in the representational implementation, the crucial point is that it becomes more difficult to simulate probability matching when calculations and representations are incorporated into the behavior of the individual ducks. Simulating a model like Gallistel's will require devising a set of behavioral rules in an ad hoc fashion to control behavior when, for example, the closest food is not at the denser source. In contrast, the simpler nonrepresentational approach succeeds in fitting the data with only one behavior and one environmental constraint.

Gallistel's (1990) model is an extension of a model of individual probability matching, and is meant to serve as model for both individual and group probability matching. However, the results of the simulations indicate that while it may account for individual data, it is not clear how to implement it in order to account for group data. It may well be that the two phenomena are not as directly related as is typically assumed. Another possibility is that the individual case could be modeled more simply. This is a question I plan to pursue in the near future.

Conclusions

By considering simple, local rules and environmental constraints, the nonrepresentational implementation successfully simulates probability matching in groups of foraging animals. This demonstrates that it is indeed "possible to propound a model of the internal causation of the duck's behavior that avoids postulating an isomorphism between a system of variables and operations inside the duck and the corresponding system of number, time and magnitude external to the duck." In addition, implementing the representational model reveals two critical weaknesses. First, the necessity of making animals move around requires that behavioral rules be added (in fact, the very rules that allow the nonrepresentational implementation to fit the data without complex representations). Second, relying on "relative patch affinities" leads to unrealistic behaviors which would require the ad-hoc addition of behavioral rules to prevent them.

The current simulations add to demonstrations in other domains that decomposing complex tasks into simple components can reduce the complexity of the representations required. Optimal individual representations are not necessary to explain complex group behavior. Quantitatively-equivalent, "satisficing" procedures, based on interactions between individuals, allow us to attribute less complexity at the individual level. While explanations based on such models may not turn out to be accurate, they should be considered along with more complex explanations. Finally, the simulator itself is a useful tool for preliminary tests of predictions about the behavior of real organisms, and studying the interaction of experimental and environmental variables.

Acknowledgments

Support was provided by an NSF Graduate Research Fellowship to JSM. I thank Gerd Gigerenzer and members of his lab, Inge-Marie Eigsti, Guy Brown, Stefan Schaal, Micheal-Spivey Knowlton, and Stephanie Forrest for stimulating discussions and suggestions, and two anonymous reviewers for their helpful comments, all of which substantially improved this paper.

References

- Ballard, D.H. (1991) *Animate Vision*. Artificial Intelligence Journal, 48, 57-86.
- Braitenberg, V. (1984). *Vehicles*. Cambridge, MA: MIT Press.
- Gallistel, C.R. (1990). *The Organization of Learning*. Cambridge, MA: MIT Press.
- Gibson, J.J., Olum, P., and Rosenblatt, F. (1955). Parallax and perspective during aircraft landings. *American Journal of Psychology*, 68, 372-385.
- Gigerenzer, G., Hoffrage, U., and Kleinbolting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98, 506-528.
- Harper, D.G.C. (1982). Competitive foraging in mallards: Ideal free ducks. *Animal Behaviour*, 30, 575-584.
- Kugler, P.N., Shaw, R.E., Vincente, K.J., Kinsella-Shaw, J. (1990). Inquiry into intentional systems I: Issues in ecological physics. *Psychological Research*, 51, 98-121.
- Langton, G., ed. (1989). *Artificial Life*. Santa Fe Institute Studies in the Sciences of Complexity, Proceedings VI. Reading, MA: Addison-Wesley.
- Myerson, J. and Miezin, F.M. (1980). The kinetics of choice: An operant systems analysis. *Psychological Review*, 87, 160-174.
- Resnick, M. (1994). *Turtles, Termites, and Traffic Jams*. Cambridge, MA: MIT Press.
- Reynolds, C.W. (1987). Flocks, herds, and schools. A distributed behavioral model. *Computer Graphics 21(4) (SIGGRAPH '87 Conference Proceedings)*, ACM, New York, 25-34.
- Sigmund, K. (1993). *Games of Life: Explorations in Ecology, Evolution, and Behaviour*. Oxford: Oxford University Press.
- Simon, H.A. (1957). *Models of Man*. New York: Wiley.
- Wiener, N. (1948). *Cybernetics, or Control and Communication in the Animal and the Machine*. New York: John Wiley.
- Winfree, A.T. (1980). *The Geometry of Biological Time*. New York: Springer-Verlag.