



NORTH-HOLLAND

Intelligent Systems

A Reader for the Cognitive Map

ALLISTON K. REID

Department of Psychology, Wofford College, Spartanburg, South Carolina 29303, USA

and

J. E. R. STADDON

*Department of Psychology: Experimental, Duke University, Durham,
North Carolina 27706, USA*

ABSTRACT

A local diffusion model (Staddon and Reid, 1990) can reproduce exponential and Gaussian stimulus-generalization gradients. We show that a two-dimensional diffusion model, together with simple reinforcement assumptions, can reproduce many of the empirical properties of goal-directed spatial search, including area-restricted search, open-field foraging, barrier and detour problems, maze learning and spatial "insight." The model provides a simple, associationistic "reader" for Tolman's cognitive map. © Elsevier Science Inc. 1997

Edward Tolman spent much of his career devising clever experiments to show that stimulus-response accounts of rat behavior cannot be correct. Some of his most striking demonstrations involved spatial learning. One such example by Tolman and Honzik [11] utilized a maze which allowed three paths from the Start box to the box where food was available. The paths differed in length: path 1 shorter than path 2 shorter than path 3. In preliminary training, the rats were allowed to become familiar with all three paths to the food box. They also had experience with a block, which permitted access to the goal only via paths 2 and 3. In the test condition, the block was moved to another point—so that only path 3 was open. The question was: Would the rats choose path 3 as soon as they encountered the block in the new position, or would they choose path 2, which was normally preferred to path 3—indicating that they do not know that paths

INFORMATION SCIENCES 100, 217–228 (1997)
© Elsevier Science Inc. 1997
655 Avenue of the Americas, New York, NY 10010

0020-0255/97/\$17.00
PII S0020-0255(97)00042-X

1 and 2 share a common, blocked, segment? Tolman and Honzik's rats behaved intelligently, and usually went straight to path 3 after encountering the block. The experimenters took this as evidence that the rats knew something about the topology of the maze. The rats were not just operating on a fixed hierarchy of preferences ("path 1 better than path 2 better than path 3") nor were they responding reflexively to local cues. Tolman considered this behavior to be an example of "insight," although he did not specify exactly what that means. He did say that some kind of cognitive map must be involved; and he also argued that these results are incompatible with a blind S-R process. But he offered no satisfactory alternative, prompting Guthrie's [6] famous objection that his theory left the rat "buried in thought" (p. 143) at each choice point. Tolman provided a map, but no map reader.

We here propose a map reader that is based on the familiar behavioristic idea of stimulus generalization. We identify a class of spatial problems, including the Tolman and Honzik problem, that can be solved by a very simple local process, operating on an internal map. The process involves no overview of the problem as a whole, and hence nothing corresponding to "insight." We first describe our local-process model, and then show how it can simulate the behavior of animals in a variety of spatial situations. We conclude by pointing to some limitations on the simple model.

A LOCAL-PROCESS MODEL

MODEL STATE AND DIFFUSION RULE

Our map reader is based on a diffusion model for the dynamics of stimulus generalization (Staddon and Reid [9]). We assume that a stimulus dimension can be represented as a line of units, each with four connections: bidirectional connections to its immediate neighbors, an input from a perceptual mechanism that allocates a limited region (e.g., wavelength band, spatial location) of a single sensory dimension to each unit, and an output. When a given region is present in the stimulus, the corresponding unit is activated. Each unit has as its final output an expectation strength V_i that is its contribution to observed behavior. The key process is diffusion of expectation strength over the set of units.

ONE-DIMENSIONAL MODEL

The diffusion process works in the following way. We assume a set of connected units, each characterized by an expectation. Thus, $V_i(t)$ is the

expectation associated with unit i at time step t . At each time step, the expectation of each unit moves towards the average of its neighbors'. Formally, the change in the strength of the i th unit in a series at each time step is given by

$$\Delta V_i = \alpha [(V_{i+1} - V_i) + (V_{i-1} - V_i)], \quad 0 \leq \alpha \leq \frac{1}{2} \quad (1)$$

where ΔV_i is the change in V_i from one time step to the next (i.e., $V_i(t+1) - V_i(t)$) and $i+1$ and $i-1$ are the adjacent units. α is a diffusion rate parameter: larger values of α produce faster diffusion. Equation (1) is a discrete-time version of the standard diffusion equation that describes the flow of heat in a solid or molecules in a liquid (e.g., Berg [2]).

SPATIAL ORIENTATION: EXTENSION TO TWO DIMENSIONS

Spatial orientation has two logical parts, knowledge and action. Knowledge is the mapping between states of the organism and location in space—Tolman's cognitive map. The map may be more or less complete, more or less accurate and fine grained. An accurate map must at least have a unique state for each spatial position (within some spatial resolution). Action is the behavior taken when the organism is in a given state (i.e., at a given place on the map): stay, move forward or back, left or right, etc.

MAP STRUCTURE

Knowledge and action are logically independent, but may not be independent in practice. The program for action may expose the organism to new parts of the environment, and thus help it to acquire knowledge, for example. Nevertheless, we can focus on either component. Our model focuses on action. Our model organism is looking for food, and its behavior is guided at each instant by the expectations formed by its past experience in a given environment. We are *not* concerned with how the animal knows where it is. We assume sensitivity to landmarks, geometrical features of the environment, or whatever, *sufficient to locate the animal accurately*, i.e., to provide a unique mapping between the animal's real position and the active unit in the cognitive map. Our concern is not with learning about spatial layout, or even with how the organism uses cues, but with reward and what the animal does about it, given appropriate spatial knowledge.

We consider only the simplest class of map: we assume only that there is a unique element (termed a *unit*) corresponding to each spatial location, and that these units are connected in well-defined ways that accurately represent some of the features of the environment.

Each unit is characterized by three things: it is active or inactive, in the sense described earlier, it has a level of the expectation variable V_i , and it is connected to one or more adjacent units (termed *neighbors*). We assume that there are no isolated units (i.e., units that cannot be reached from some point in the map).

MODEL STATE AND DIFFUSION RULE

Equation (1) describes the diffusion process along a single dimension. At each step, the expectation of each unit moves toward the average of its neighbors—one on either side of the current unit. When the map has two dimensions, the form of the equation is the same, but the number of neighbors increases. Thus, the generalized diffusion equation is

$$\Delta V_x = \alpha \sum_{i=1}^N [V_i - V_x], \quad 0 < \alpha < 1/N \quad (2)$$

where ΔV_x is the change in V_x from one time step to the next. If the underlying map is a full lattice, each element has eight neighbors ($N=8$) (N, S, E, W, and the diagonals). The only difference between (1) and (2) is the number of neighbors N associated with each unit, which depends on the type of map and the location of the unit in the map. Note that diffusion rate into and away from any particular unit depends on the number of neighbors it has—which depends on the type of map, not on the diffusion process.

In this form, the model says nothing about the effects of reward and nonreward on the expectation surface. To include these effects, we make assumptions which describe the effects of stimulus presence or absence (unit activity), the effects of reward and nonreward, the effects of barriers, and an action rule:

1. *Activity.* A unit is either *active* or *inactive*, depending on whether or not its stimulus is present. If its stimulus is absent, the unit is inactive, and the expectation variable changes according to the diffusion equation. If its stimulus is present, the unit is active, and the expectation variable is

set to a fixed value, which is determined by whether reward or nonreward is also present.

2. *Reward.* Reward events set $V_i(t)$ for the active unit to a value directly related to the reward magnitude ($S(t) = 1$ in all our simulations) as long as reward is present (usually one time step). This assumption expresses the idea that as long as reward is available, reward expectation should equal reward value.

2. *Nonreward.* Nonreward (i.e., absence of reward in the presence of an active unit) sets $V_i(t)$ to 0. This assumption expresses the idea that any reward expectation that exists for a given stimulus should be zeroed as long as that unit is active and unrewarded. (Expectation will recover, owing to diffusion from neighboring units, when the unit is no longer active.)

4. *Barriers.* Barriers reduce the number of neighbors. This reduction affects both diffusion and the action rule: diffusion cannot occur between units separated by a barrier, and the model organism cannot move between them.

5. *Action Rule.* The action rule is very simple: at each time step, the system moves to the adjacent unit with the highest expectation (V value). If the current unit has a higher expectation than its neighbors, the model organism remains where it is; otherwise, it moves to the highest expectation neighbor. This rule is variously termed hill climbing or gradient descent, depending on the sign assigned to “good” events.

These assumptions allow the expectation surface to change in real time as a function of the temporal patterns of reward and nonreward in the presence of different stimuli. Notice that this model assumes only local processes: diffusion flow and the direction of movement from any unit are determined only with respect to immediate neighbors. The model involves no overview of the problem (or expectation landscape) as a whole.

We assume a cognitive map represented by a full lattice, so there are eight possible directions of movement. We assume that the organism's location at each time step is represented by a single “active” unit, which is accurately mapped on to the real environment in the sense that for every real spatial location, there is an appropriate active unit. Different locations need not activate different units (in a coarse-grained map, they may activate the same unit), but adjacent locations must activate adjacent units, in the same way that a fine-grained computer-graphic image can be mapped onto a coarse-grained screen image. Our “reader” specifies a pattern of movement within the map, driven by an expectation surface that is altered dynamically both by the passage of time and by the system's experience with reward and nonreward.

FORAGING PROBLEMS

We now briefly describe two simulations of simple foraging experiments. Reid and Staddon [8] provide more extensive demonstrations of how the diffusion model reproduces many additional movement patterns of animals foraging in the open field and within several types of mazes.

Area-Restricted Search

If rats are allowed to find buried food at a particular spot in a large enclosure, are removed, and then returned after a delay, they usually begin digging close to the rewarded spot (cf. Cheng [5]). If they fail to find food at once, the search is restricted to the immediately surrounding area before a more general, unlocalized search begins. Area-restricted search has been observed in a variety of species, including common shrews (Pierce [7]), pigeons (Cheng [4]), and even insect larvae (Banks [1], Bond [3]).

The two-dimensional diffusion model produces area-restricted movement patterns when the model organism is reintroduced to the enclosure after finding food at a particular spot. We begin our simulation with the assumption of a very low, uniform level of reward expectation at every unit in the map. This small expectation causes the organism to move around the enclosure in an irregular fashion, moving from spot to spot in a random-walk-like pattern as successive failures to find food “zero” expectation at each spot. Once the organism finds “food” at a particular spot, reward expectation for the active unit is set to unity and the organism is immediately removed from the enclosure (i.e., all units are made inactive). Reward expectation at that position immediately begins to diffuse according to (2) at a rate determined by α . Figure 1 shows the pattern of movement produced when the organism is subsequently reintroduced to the empty (no-food) enclosure.

The four panels of Figure 1 show the pattern of movement generated using four different values of α . The movement patterns were smoothed using splines to eliminate the abrupt changes in direction that result from a map which permits movement in only eight directions. In all panels, diffusion occurred for 30 time steps while the organism was out of the enclosure. This process creates a three-dimensional “landscape” of reward expectation (not depicted), centered on the rewarded location. For each panel, the organism was subsequently reintroduced at position 1,1 and allowed to move for 60 time steps. Equation (2) was iterated once for each step taken.

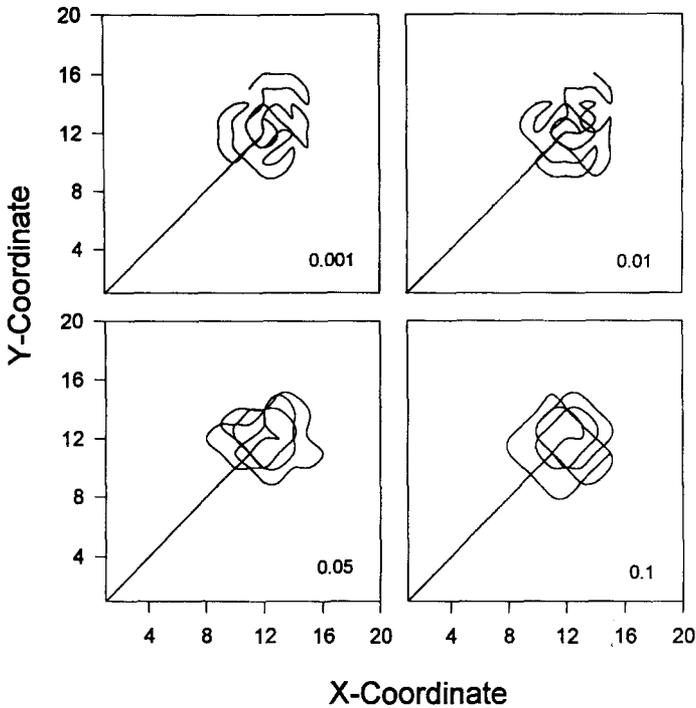


Fig. 1. Patterns of area-restricted search produced by the model organism with four different values of α , the diffusion rate parameter.

The hill-climbing action rule moved the organism directly to the peak of the expectation landscape. This movement has the appearance of goal direction (i.e., the organism moves straight to the rewarded spot), but, in fact, only local processes are involved: the model organism cannot sense the peak of the landscape from a distance. At each step, the value of the reward expectation for the active unit is set to zero since no reward is obtained. The diffusion process allows reward expectation to “flow” back gradually into the units the organism has already visited (much like the tracks made by walking through deep mud).

The area-restricted “search” pattern of movement is centered on the rewarded location. Movement in all directions from that spot is equally likely, as with actual animals (Cheng [5]). Even though the model has no stochastic component, the search path appears erratic and frequently intersects itself. With each unrewarded step, reward expectation is zeroed at that point. As more and more steps are made near the rewarded spot,

the overall level of reward expectation in the area declines, and the organism gradually moves farther away to areas in which “expectation” has not been depleted. This gradually expanding search pattern closely resembles descriptions of “focal” search followed by more “general” search in rats (Timberlake and Lucas [10]).

MAZES

Mazes have been widely used in the study of spatial learning. Adaptive behavior in mazes has often been used to justify large claims about cognitive processing. For example, Tolman *et al.* [12] argued that rats often take the shortest path to the goal, whenever such a path is available (i.e., not prevented by the design of the maze), and that this behavior provides evidence for a type of spatial insight. We have already shown that our model, which lacks any overview of the problem, is nevertheless capable of producing minimum distance paths under many conditions (e.g., Figure 1). We show in this section another example of adaptive, even insightful, behavior that can be simulated by the diffusion model.

Tolman's Starburst Maze

The shortcut problem of Tolman *et al.* [12] using the “starburst” maze is probably the best known example of spatial insight in animals. Rats were given extensive training on the elevated maze shown in Figure 2. The rats were first trained to run down the path labeled $F-G$, then down the longer path $C-D-E-F-G$. In the final phase of training, they began trials from position A . After five trials from starting point A , they were tested on the “starburst” configuration, shown in Figure 3. The subjects often chose path 6, which was the shortest distance from the tabletop to the goal. Tolman *et al.* concluded that training produces (in some rats, at least) a “disposition to take the shortest Euclidean Path to the goal” (1946, p. 20).

It is easy to see how our model produces the same choice of shortcut as in the Tolman *et al.* study. Training on the maze segments $F-G$ and $C-D-E-F-G$ allows the organism to identify the sides of the maze (the “cliffs”) as barriers to movement. Reward expectation does not diffuse through these learned barriers. Once reward is obtained, reward expectation diffuses outward along adjacent units representing open space and the maze. Diffusion is slightly faster along the runways of the maze than through open space because the number of adjacent units is reduced; so the value at any unit can “flow” only into units within the maze, rather

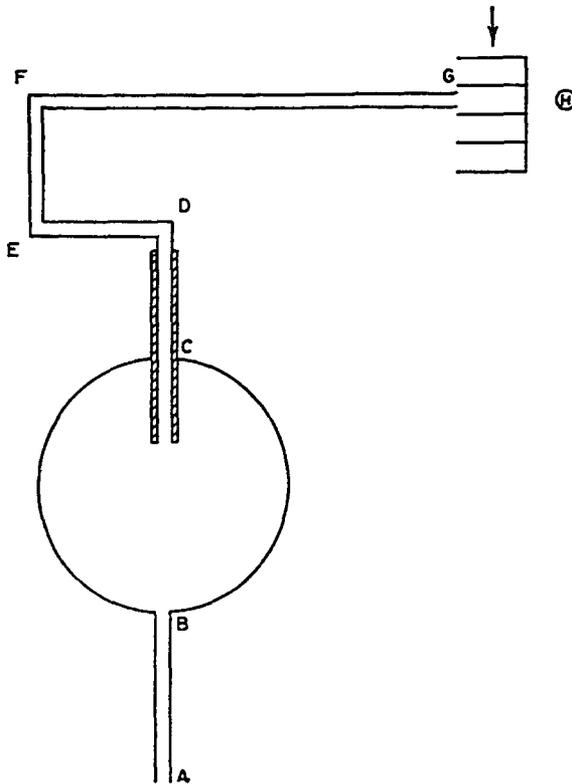


Fig. 2. Apparatus used during training in Tolman *et al.* [12].

than in all directions. Once the boundaries of the tabletop are discovered, the units representing the tabletop receive most of their value of reward expectation through the maze segment $C-D-E-F-G$, with reward expectation decreasing with distance from the goal box along the maze. Therefore, the organism traverses the maze from A to G .

The values of reward expectation outside the maze at point D and next to the upper right side of the tabletop are higher than the values of adjacent units within the maze, so the organism consistently "attempts" movement through the barriers at those points in consecutive trials. Since movement is prohibited in those directions, the unit with the highest value is attempted. When the maze topography is altered (cf. Figure 3), the barriers on the upper right of the tabletop are no longer present. Once their absence is discovered, expectation is free to diffuse directly from the

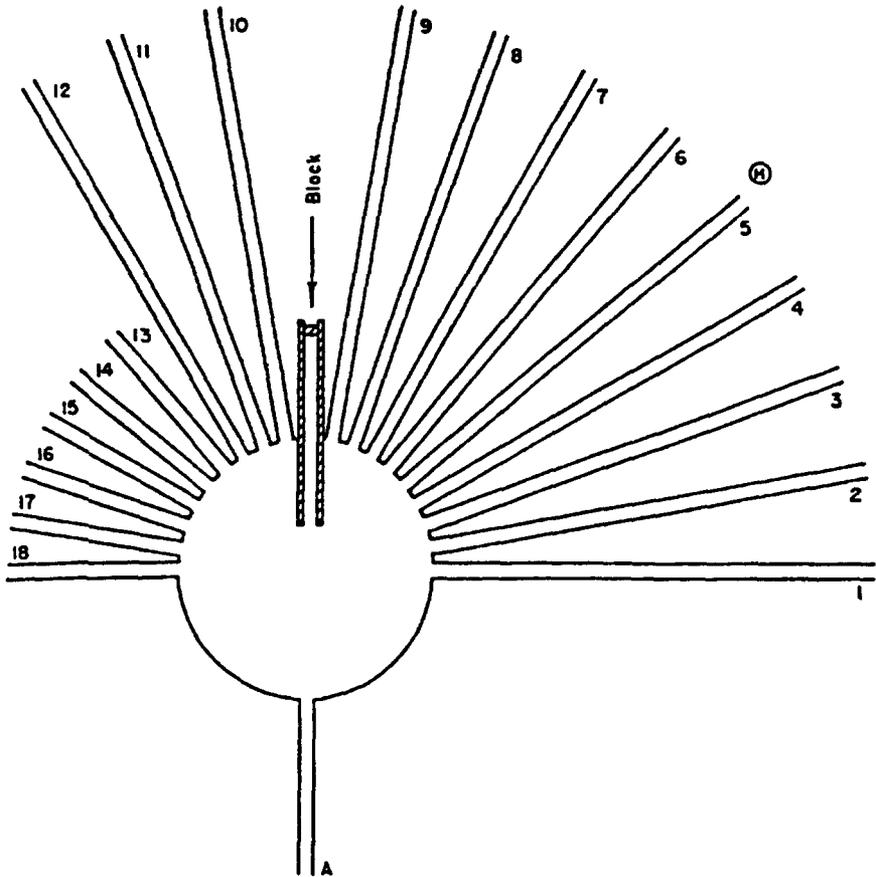


Fig. 3. Apparatus used during the test trials in Tolman *et al.* [12].

goal, which means that path 6, not the original (now blocked) path or those adjacent to it, will be the path with highest expectation. As in the detour problems above, prediction of this path does not depend upon the value of the diffusion rate parameter. It depends only on the accuracy of the cognitive map and the animal's ability to detect the absence of a barrier to the direct path.

The Tolman and Honzik "insight" problem with which we began can obviously be solved in the same way. Given an accurate map and knowledge of a barrier at point *B*, expectation can diffuse back to the start box only along path 3, which should therefore be the one chosen.

DISCUSSION

We have shown that a process in which expectation diffuses through a cognitive map can function as a map reader that guides action and adapts to reward contingencies. As Guthrie pointed out many years ago, a map by itself is not sufficient to account for spatial behavior. We show that a map plus a map reader that incorporates a response rule can account for the major features of a wide variety of spatial-learning experiments: in the open field, in enclosures, and in different types of mazes. Moreover, our map reader is nothing but a dynamic formulation of stimulus generalization, a basic behavioristic process. Thus, Tolman's contention that something more than *S-R* principles is required to account for maze behavior is only partly true. A map is required (in our scheme), but beyond that, no "insight," no internal cyclopean eye with an overview of the situation as a whole, is necessary.

There are several omissions in our account. Most glaring is the absence of a system for the formation of a map. Reid and Staddon [8] provide a fuller account of this model, as well as possible ways of overcoming these limitations.

Despite these limitations, the simplicity of our model and its ability to generate complex search paths whose general properties closely resemble those of real animals in a variety of environments recommend it as a useful property for any comprehensive model of spatial behavior.

Research support was provided by grants to Duke University from the National Science Foundation and the National Institute of Mental Health (JERS, Principal Investigator).

REFERENCES

1. C. L. Banks, The searching behaviour of coccinellid larvae, *Animal Behavior* 19:58-86 (1954).
2. H. C. Berg, *Random Walks in Biology*, Princeton University Press, Princeton, NJ, 1983.
3. A. B. Bond, Optimal foraging in a uniform habitat: The search mechanism of the green lacewing, *Animal Behavior* 28:10-19 (1980).
4. K. Cheng, The vector sum model of pigeon landmark use, *J. Experimental Psychology: Animal Behavior Processes* 15:366-375 (1989).
5. K. Cheng, Three psychophysical principles in the processing of spatial and temporal information, in: W. K. Honig and J. G. Fetterman (Eds.), *Cognitive Aspects of Stimulus Control*, Lawrence Erlbaum, Hillsdale, NJ, 1992, pp. 69-88.
6. E. R. Guthrie, *The Psychology of Learning*, Harper and Bros., New York, 1935.

7. G. J. Pierce, Search paths of foraging common shrews *Sorex araneus*, *Animal Behavior* 35:1215–1224 (1987).
8. A. K. Reid and J. E. R. Staddon, A blind reader for the cognitive map, unpublished manuscript (1995).
9. J. E. R. Staddon and A. K. Reid, On the dynamics of generalization, *Psychological Rev.* 97:576–578 (1990).
10. W. Timberlake and G. A. Lucas, Behavior systems and learning: From misbehavior to general principles, in: S. B. Klein and R. R. Mowrer (Eds.), *Contemporary Learning Theories: Instrumental Conditioning Theory and the Impact of Biological Constraints on Learning*, Lawrence Erlbaum, Hillsdale, NJ, 1989, pp. 237–275.
11. E. C. Tolman and C. H. Honzik, “Insight” in rats, *University of California Publications in Psychol.* 4(14):215–232 (1930).
12. E. C. Tolman, B. F. Ritchie, and D. Kalish, Studies in spatial learning. I. Orientation and the short-cut, *J. Experimental Psychol.* 36:13–24 (1946).

Received 3 November 1995