

Maze Navigation Using Optical Flow

Andrew P. Duchon

Department of Cognitive and Linguistic Sciences
Brown University
190 Thayer St.
Providence, RI 02912-1978
email: duchon@cog.brown.edu

Abstract

Some recent work with autonomous robots has focused on using optical flow for “direct” control of speed and rotation in obstacle avoidance and other simple behaviors. This work has been inspired by work with insects showing similar mechanisms. To extend these behaviors, three methods of maze navigation are investigated in a simulated robot modeled after a real one. A motor-based method places biases in the obstacle avoidance control law used previously. A perception-based method uses optical flow to detect possibilities for action (e.g., to turn left or right). Both of these require that the agent have a list of biases in order to navigate. The third method, called the Saliency Centroid Method, is based on a theory of the role of the hippocampus in rat navigation. This method trades off the memory of the first two for more advanced perceptual processing and allows the most flexible behavior.

1 Introduction

In the last few years, a number of researchers have explored the use of optical flow for directly controlling a robot’s behavior (Coombs & Roberts, 1993; Duchon & Warren, 1994; Sobey, 1994; Santos-Victor et al., 1995; Coombs et al., 1995). This research has concentrated on methods for obstacle avoidance without modeling the world (Aloimonos, 1992) and has even produced robots which can dock (Santos-Victor & Sandini, 1994) and play “tag” (Duchon et al., 1995). These implementations on a robot platform have been inspired by theories of how bees navigate through the world (e.g., Srinivasan, 1992). The crossover between these two fields indicates that it is possible to speak in general terms of an *agent*. Results in both fields apply to any

agent, be it biological or artificial, that can detect optical flow and act upon it.

This paper begins two extensions to the insect-based, on-line steering work cited above. The first is to develop methods using optical flow for goal-directed behaviors, still without using the information in optical flow for creating a 3-D model of the environment. The second purpose is to take inspiration from another animal literature, namely cognitive mapping in rats. As a first pass at these higher-level behaviors we use a simulated agent which is modeled after the real robot and “control laws” presented previously (Duchon & Warren, 1994; Duchon et al., 1995).

Three strategies for maze navigation are explored. The “motor-based strategy” is similar to the kinesthetic response chains of the early rat-in-a-maze literature whereby “the act of running a maze seemed to be more a series of concatenated reflex movements released by kinesthetic cues” (Munn, 1950, p.186). This conclusion was dismissed with further experimentation (e.g., Dashiell, 1930) and the simulation also quickly shows the limitations of this strategy. The “perception-based strategy” takes more advantage of the information in optical flow and allows the agent to make choices and learn a route-based representation of the environment. Although much more flexible, this strategy too has its faults. Finally, we explore the use of optical flow in a modification of O’Keefe’s (1989, 1991) model of hippocampal function. This method alone also has its problems which are addressed in the conclusion.

2 The Simulated Environment

2.1 The environment and the agent

The maze environment (Figure 1) has many possibilities for action including cul-de-sacs, and T-, L-, and U-junctions. The agent is an observation point with a

field of view (FOV) of $0^\circ - 330^\circ$, and although a point, a triangle is drawn around it for display purposes. For obstacle avoidance and for the first two strategies described below, the agent has 48 samples from a FOV of 120° . The agent has access only to the optical flow “seen” at the intersection of each sampled line of sight and the nearest maze wall.

2.2 Calculation of the optical flow

The instantaneous optical flow from translation only is used. No rotational components are added (a condition our robot approximates (Duchon & Warren, 1994) since it moves in a piecewise linear manner). As such, the optical flow can be calculated analytically. Because the simulations take place in a maze without other agents, the optical flow value (here, angular velocity in the horizontal dimension only) for a given line of sight is simply

$$\dot{\beta} = \frac{|\mathbf{h}| \sin \beta}{d} \quad (1)$$

where $\dot{\beta}$ is the angular velocity, \mathbf{h} is the velocity of the agent (here, the agent always moves along the optical axis), β is the angle of the sample with respect to \mathbf{h} , and d is the distance between the observation point and the sampled point on the wall. In this manner, the observer can obtain a speed-scaled distance between itself and the sampled points in the environment. Another piece of information provided by optical flow is the time-to-contact, known as τ (Lee, 1976). For small values of β , $\tau \approx \beta/\dot{\beta}$. In the maze, a low τ value signals that there is a wall in front of the agent. If τ is below a threshold, the agent will turn 180° (the “tau-reflex”).

2.3 Obstacle Avoidance

For obstacle avoidance, we used the Balance (or Centering) Strategy which has proven useful in robot implementations (Duchon & Warren, 1994; Coombs et al., 1995; Santos-Victor et al., 1995). With the Balance Strategy, the agent moves so as to equate the average magnitude of flow detected on each side of the optical axis (which is tied to the heading direction both in simulation and in our robot). At each time step in the simulation, the instantaneous optical flow values are obtained for each sample. The average for each side ($\overline{\beta_L}$, $\overline{\beta_R}$) is calculated and the agent rotates an angle r as given by

$$r_{bal} = k(\overline{\beta_L} - \overline{\beta_R}) \quad (2)$$

where k is a scaling constant. With $r_{bal} > 0$ indicating clockwise rotation, more flow on the left causes the agent to turn to the right.

The essential idea behind this strategy is that of motion parallax: when the agent is translating, closer ob-

jects will give rise to greater motion across the retina than farther objects. It also takes advantage of perspective in that closer objects will also take up more of the field of view, biasing the average towards their associated flow. In the end, the agent turns away from the side of greater flow, but only such that it does not turn into an obstacle on the other side. When the agent is in a hallway, it will tend to move down the center. Obstacle avoidance using the Balance Strategy has proven to be very robust with robots, but the question remains as to how to integrate this behavior with more goal-directed actions. Below, we explore a few possibilities.

3 A Motor-based Method

It is possible to induce other forms of behavior by placing biases in the Balance Strategy:

$$r_{bias}(\gamma) = k[(1 + \gamma)\overline{\beta_L} - (1 - \gamma)\overline{\beta_R}] \quad (3)$$

Since $r > 0$ indicates a right turn, a positive bias ($\gamma > 0$) makes the agent hug the right wall. With a constant bias, this control law will lead to wall-following, which can be used to map an environment (Mataric, 1992). With a variable bias however, say of a given sequence, the agent can navigate through the maze to a pre-determined location. We implemented this latter strategy.

Let’s posit a given sequence of biases (from some previous experience) which we know will lead the agent to the goal location, given that it always starts at the same start location (*home*). The agent will need to both make turns and be able to pass by openings that are undesired. To pass an opening on the left, it can use a small positive bias (to the right) that is enough to keep it from going into the left opening. To make a turn into an opening, a stronger bias is applied. For example, the sequence of biases used in Figure 1 is : $\Gamma = \{\text{PASSRIGHT}, \text{TURNLEFT}, \text{PASSRIGHT}, \text{TURNRIGHT}, \text{CUL-DE-SAC}, \text{STOP}\}$ where $\text{TURNLEFT} = -0.3$, $\text{TURNRIGHT} = 0.3$, $\text{PASSLEFT} = 0.15$, $\text{PASSRIGHT} = -0.15$, and $\text{CUL-DE-SAC} = 0.0$.

Since the agent only has access to the sequence of biases, it must have a way of determining whether or not it has actually made the turn or passed by the opening. This is done by keeping a moving average of the last few r values, denoted by σ , an “efference copy” of its movement. If it is not currently in a turn but the pointer in the sequence indicates the next action is a turn, then $|\sigma|$ moving above a threshold, θ_t , indicates it has entered the opening. When $|\sigma|$ again falls below θ_t , it has completed the turn. Similarly, passage by an opening can be detected as σ rises above and falls below another threshold, θ_p , characteristic of a slight turn into an opening and a straightening in the next

passage. A cul-de-sac can be detected by a low τ value which will also trigger the tau-reflex. Each time the agent completes a turn, pass or cul-de-sac, a pointer moves to the next bias in the list.

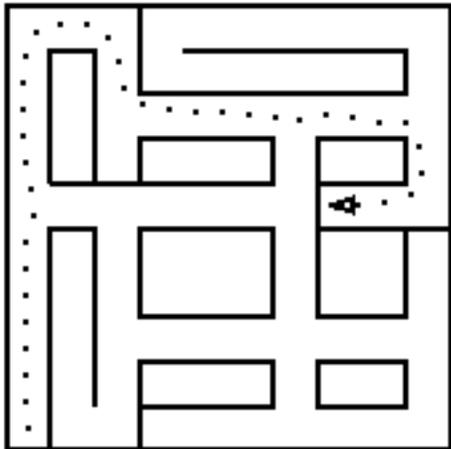


Figure 1: *Path of the agent using the motor-based method and the sequence in the text.* The agent starts in the lower left cul-de-sac. The dots indicate the agent’s location every 20 time steps. Notice that it does not travel down the middle of the passageways, but rather hugs the walls depending on the bias γ . For example, the first dot is in the middle of the passage, but the agent moves a bit to the left in order to pass by the opening on the right. It “dips” into the passage only slightly, the recovery from which signals that it has passed the opening so it moves to the next bias in the list.

There are a number of problems with this method of navigation. First and foremost, the agent must know the sequence beforehand and how this comes about is undetermined. Second, on the motor side there is a process that indicates it has completed a turn or passed by an opening, but there is nothing on the perceptual side by which the agent can “see” that it has a choice. To address these two issues we implemented a more perception-based method for navigating the maze.

4 A Perception-based Method

Given the rectilinear nature of the maze, in which all walls come together at right angles, the agent could “see” that there is an opening on the left or right by detecting a large drop in β from a given sample in the periphery to the next, more central, sample. An opening straight ahead is given by the central τ value. If τ is low enough, then a wall must be in front of the agent, if high enough, an opening. The agent must simply choose which among the possibilities (go straight at a cross-junction; turn or pass left or right at a T-junction) it will take and apply the appropriate bias

in Equation 3. In the motor-based method, the current bias in the list was applied constantly. With this method, the bias is applied only when at least two possibilities are available (a *choice point*) and only for a limited period (otherwise, the bias is zero). The latter condition is applicable because the agent will typically detect the opening at a fixed point in time before it must turn. Thus there is no need for an efference copy with this method.

The basic task for the agent is given that it starts in a cul-de-sac (*home*), it will move through the maze until it reaches another cul-de-sac, at which point it returns home (Figure 2). On the way out, at each choice point the agent makes a random choice among the possibilities and records the choice in a list. If, say, there is an opening to the right only, and a wall is in front of the agent, then there is only one way to go and no need to make or record a choice. Again, the list consists of a series of biases (Γ). For the return trip, the agent simply goes backwards through the list. At each choice point, the agent applies the negative of the current bias in the list and moves the pointer back one.

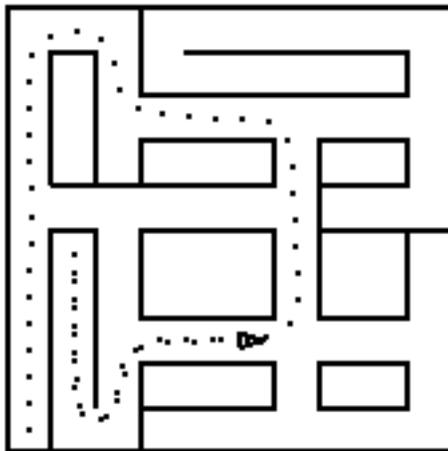


Figure 2: *Perception-based maze behavior.* The agent again starts in the lower left corner. It stays in the middle of the passage because there is no bias until it reaches a choice point halfway up the left side. There is a slight deflection at the first passage despite the PASSRIGHT bias. At the frame shown, it has made its way to the cul-de-sac at the middle left side and begun the trip home (doubled-up dots). It can consistently find its way back from any cul-de-sac.

There are a number of disadvantages to this method as well. First, as noted above, this method depends on the regularity of the maze. For instance, to determine if there is an opening, at least a 35% change in β is required in one of the peripheral samples given a 120° FOV. There is some flexibility in these values (in the percent change and the FOV), but not much. It also

depends on the size of the passageways and doorways and there is no contingency for entering a room.

Second, it is a route-based navigation system which means it is very “fragile” (O’Keefe & Nadel, 1978). The exact route taken on the way out must be used on the way back; no shortcuts can be taken. Moreover, if the agent is heading home and gets transported somewhere else in the maze, its list of biases is useless. A solution to this problem would be to make a connectivity graph of the possibilities at each junction and the choice made. The entire maze could then be represented in this way and standard AI techniques could be used to determine, say, the sequence of biases (a policy) to be applied at each choice point in order to get home from any location in the maze.

Third, a representation in terms of a list (or graph) is unsatisfying. It implies a symbolic memory mechanism, which behavior-based robotics is trying to avoid (however, see Mataric, 1992). Ideally, there would be a perception-based method, with minimum memory, which would use “higher-order invariants” in the environment for more flexible navigation.

5 The Salience Centroid Method

In many of the studies of maze-learning in rats, placed above the maze are a number of cues which can be seen at all times, termed “extra-maze cues” (Munn, 1950, p. 192). Most of these studies have shown that if the extra-maze cues are rotated a certain angle after the rat has learned where the goal is, the rat will now behave as if the goal location had also rotated by the same angle. That is, it is behaving with respect to a frame of reference defined by the extra-maze cues. This is apparent not only behaviorally, but also physiologically. It is now well known (O’Keefe & Nadel, 1978) that certain cells (“place cells”) in the hippocampus are active only when a rat is in a certain location in its environment (the “place field”). The place field of a cell also rotates with the cues (Muller & Kubie, 1987).

O’Keefe (1989) proposed that the hippocampal formation performs matrix transformations in a Cartesian mapping system. In his model, a current matrix \mathbf{C} contains the *egocentric* (x, y) coordinates of every object (as obtained by converting the distance and bearing of each object). Going to a goal requires multiplying the inverse of the current matrix by the goal matrix \mathbf{G} to get the transform, or movement, required: $\mathbf{C}^{-1} \cdot \mathbf{G} = \mathbf{T}$. O’Keefe proposed how various mechanisms in the hippocampal formation could carry out these operations, such that given \mathbf{C} (from vision) and \mathbf{G} (from previous experience), the hippocampus could generate the correct movement \mathbf{T} to move the rat to the goal.

In his second model (O’Keefe, 1991), coordinates are no longer egocentric, but make use of a “centroid-slope” construction. Once the bearing and distance of all the objects are obtained, the *centroid* is simply the average of the vectors from the rat to each object. The centroid is invariant to rotation, thus gives a radial distance for the rat’s location. The *slope* (for two objects) is simply the ratio of the difference of the y - to the x -coordinates, also in an egocentric frame. The slope is invariant with respect to translation, and specifies the rat’s orientation. The distance from the centroid and the angle from the slope are both invariant as long as the same objects can be seen. Movements are calculated with respect to these values, and not to egocentric coordinates as in the first model. However, one problem with this method is that the symmetry on the two sides of the “slope” means there is always an ambiguity in the rat’s location. Another issue is the computational expense of averaging the slope of *all* pairs of cues.

The method below is proposed for three reasons. First, it resolves the symmetry problem and scales linearly with the number of cues. Second, it exploits the information in optical flow (as O’Keefe (1991) actually suggests). And third, this method, like the previous two, integrates goal-directed movements with obstacle avoidance.

Three assumptions are required for the method to work. The first is that all the extra-maze cues must be within view simultaneously for the agent to “fix” its position. This “omnipresence” of the cues is assumed by most current computational models of the hippocampus since nearly all physiological studies have recorded place cell activity in an environment where all the cues could be seen¹. The second assumption is that at least one of the extra-maze cues must be perceptually distinguishable from the others. The third assumption is that the agent has some sense of its rate of motion, if only proportionally. Reasons for these assumptions will become apparent and methods for removing the first one will be discussed below.

To comply with these assumptions, the environment and the agent must be slightly modified. The environment is now equipped with four “extra-maze cues” which can be seen from anywhere in the maze. Each cue is assigned a random number as its “salience.” The agent now has a second set of optical flow samples, different from those described above which are still used

¹Once the rat is placed in an environment, place cell activity is maintained even as cues are taken away (O’Keefe & Speakman, 1987) or the lights are turned out (Quirk et al., 1990). This ability to use dead-reckoning to maintain place cell properties is mediated mostly by vestibular information modified by optical flow (Sharp et al., 1995). No dead-reckoning is used in any of the methods presented here.

for obstacle avoidance. This second set can be thought of as another set of horizontal motion detectors, but now in a plane above the rat. Thus, they only detect the extra-maze cues. Forty-eight samples are taken from a 330° FOV.

5.1 Saliency Centroid Algorithm

One flow value $\dot{\beta}$ at an edge of each extra-maze cue is detected (an edge is detected simply by a large difference between two adjacent samples). The (speed-scaled) distance to each cue is $d = \sin(|\beta|)/\beta$. For the purposes of computation only, we convert these (d, β) values to egocentric (x, y) coordinates (the y axis being the optical axis of the agent). The coordinates of the “centroid” (c) then, is simply the average of the (x, y) coordinates of the four cues: (\bar{x}, \bar{y}) . A directed vector (as opposed to just a slope) is required for an unambiguous coordinate system of the environment. Therefore, the original x and y values of each extra-maze cue are multiplied by the “saliency” of the cue and a new average obtained. At least one of the cues must have a different saliency than the others for the “saliency centroid” (s) to be different from c .

Given s and c , the saliency vector \mathbf{s} can be found (see Figure 3A). Through a couple of coordinate transforms, the location of the observer with respect to \mathbf{s} can be calculated, now back in polar coordinates: (r_O, θ_O) . We should note that crucial to this calculation is ψ , the bearing of c with respect to the agent’s heading (\mathbf{h}), as well as ω , the agent’s heading (\mathbf{h}) with respect to (\mathbf{s}). A little more vector algebra (Figure 3B) will determine the bearing β , with respect to the current heading, of a known location in the environment (r_G, θ_G) .

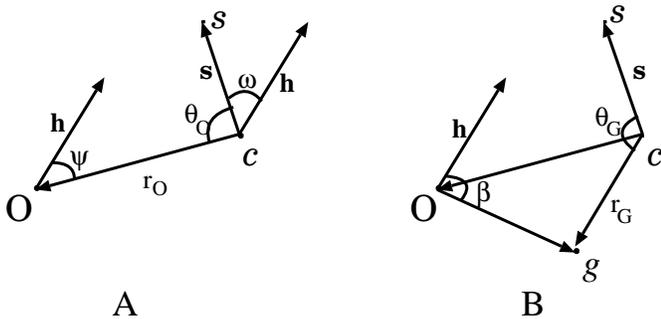


Figure 3: Vectors used in the calculations for the Saliency Centroid Method. See text.

5.2 Using the Saliency Centroid

At any given moment when all the extra-maze cues can be seen, the agent knows its location and the bearing

(β) of the goal point (if there is one). However, the agent need not turn instantly towards the goal. Instead, a certain percentage, $perc=0.2$, of β is used.

A few other factors need to be considered as well. The first is that if the goal is on the other side of a wall from the agent, the agent should not beat its head against the wall, as it were. This can be avoided by taking the total sum of optical flow that the agent sees (Σ), and reducing the bias accordingly:

$$\gamma = \beta \times perc \times \left(1 - \frac{\Sigma}{a}\right) \quad (4)$$

where a is an appropriately scaled constant. Other optic variables, like τ , could be used, but the sum takes into account the whole field of view in an appropriate and general way (Figure 4).

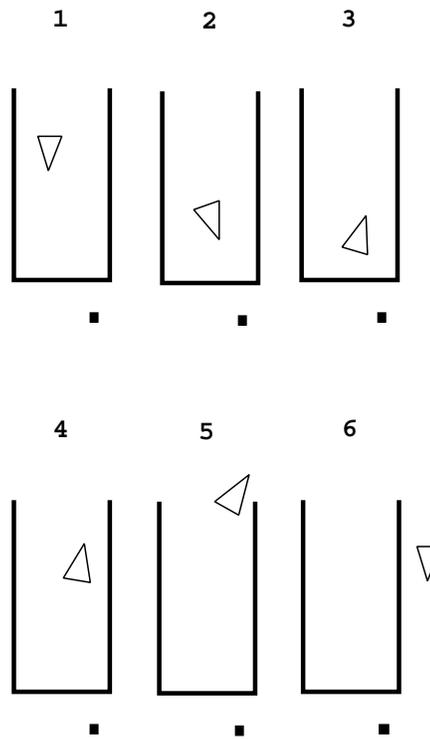


Figure 4: Escaping from a “goal-directed blind.” The sequence shows the agent (1) entering a “box” because the goal (square) is on the other side. As it nears the end (2), Σ increases, which reduces γ and so the Balance Strategy predominates. This allows the agent to turn itself around (3) and begin going up the wall (4) since being near the wall, Σ stays high. Without using Σ , the agent would stay in the cul-de-sac, continually using the tau-reflex to flip around and the bias to pull it back in. Once the agent gets to the corner (5), Σ becomes low, γ becomes high, and the agent begins heading towards the goal again (6).

The bias (γ) here is a function of the direction of the goal and is given in degrees. In the other two methods, it was related to an event (turn left or right) and did

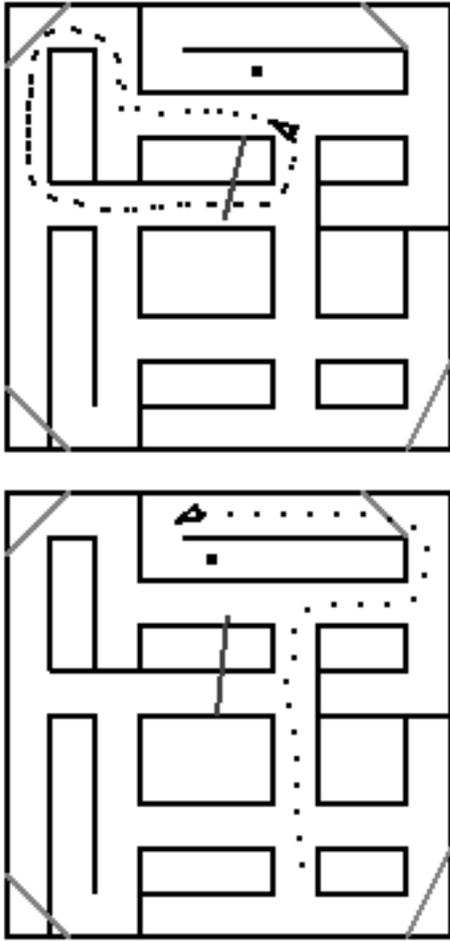


Figure 6: *Top*: The agent is caught in a loop (a local minimum) while trying to reach the goal. *Bottom*: With small Gaussian noise added to the observed flow values, the agent, by chance, makes a right turn at the intersection and can reach the goal.

5.4 Physiological Correlates

O’Keefe (1989, 1991) details much of the hippocampal physiology necessary to carry out the operations which we have proposed (which are in fact similar to his, but are based explicitly on optical flow). In brief, the theta rhythm in the hippocampus can be used to convert between polar and Cartesian coordinates (see O’Keefe & Recce (1993) for other use of the theta rhythm and a recent re-working of O’Keefe’s model). The place cells themselves would represent (r_O, θ_O) . Hippocampal place cells are sensitive to speed, either being modulated or tuned by speed (Wiener et al., 1989; Mizumori et al., 1990), and become nearly silent when the rat is not moving which may reflect the influence of ρ . Some cells in the postsubiculum are sensitive to head direction (Taube et al., 1990) and might represent ω , the an-

gle between the heading and the salience vector. Cells in the hippocampus proper are also sensitive to the direction of motion through the place field (McNaughton et al., 1983) and might be related to modulation by ψ , the angle of the heading with respect to the centroid. Or, they could be related to β , the direction of the goal. Recent evidence supports this idea. Markus et al. (1995) demonstrated that changes in task, i.e., changes in goal locations, resulted in changes in the directional selectivity of hippocampal place cells.

6 Conclusion

We have described a progression of methods an agent could use to navigate a maze. The most direct method was to place a bias in the control law for obstacle avoidance, but this required giving the agent a predetermined list of biases. The second method depended on the agent being able to detect openings, make a choice among them, record the choice (as a bias), and backtrack through the list of choices in order to return home. As mentioned above, this method could be used to create a connectivity graph of the environment. The final method we described required a nearly continuous perception of the agent’s place within the environment by means of which it could find the heading direction of another (goal) location. So, although it needed more on-line processing (a separate system to determine the agent’s relative location), memory was reduced to two numbers which could easily be considered as a kind of set-point for an error reduction system.

One might suspect that a memory system would also be required for the landmarks to be recognized, but this is not necessarily true. All that is needed is that the objects have the same “salience” (“quality” in O’Keefe, 1989) which is consistent across viewpoints. We have left “salience” essentially undefined for it could be any number of perceptual properties, or, more likely, a conglomerate of all of them. One intriguing idea is that a particular environment could be characterized by the length of the salience vector—a single value for which there is little chance of close matches among different environments.

One of the assumptions for the Salience Centroid Method was that all the cues be available at all times. This is true of most current models (e.g., O’Keefe (1991), Bachelder & Waxman (1994), Shapiro & Hetherington (1993); but see Touretzky et al. (1994), Worden (1992)). One of the properties of the hippocampus which is prominent in most models of its function is its autoassociativity. This is basic to many models of memory and allows for a “whole” memory to be recreated from a part. We have begun to look at using the autoassociative BSB-model (Brain-State-in-a-Box, see

e.g., Anderson, 1995, pp. 493-628) as a model of the hippocampus and hope to incorporate this work into our simulations. Such a model would allow the agent to determine the salience vector by only seeing a few of the objects. Once given this capability, it would not be difficult to “string along” a series of salience vectors, or create an extended one, which connected a number of distinct environments.

Moreover, the BSB network would be part of the perceptual process, which, if motor commands are also included, would create a single integrated system for both perception, action, and representations of the environment. That is, the agent’s actual perception of the environment would change as a function of learning in the BSB, which would change its behavior, which would change what it learned, etc. This would be an example of what Gibson called learning through differentiation (Gibson, 1966, p. 270). A number of criticisms of the ecological approach refer to the problems it has in explaining knowledge of things that are not currently perceivable. The autoassociative networks we have begun to work with indicate a way for that which can be perceived to *specify for the agent* that which cannot be seen. Once such a system is developed in simulation we will bring these methods back to the robots with which our work began.

7 Acknowledgments

This work was supported by a National Science Foundation Graduate Research Fellowship. Thanks to Leslie Pack Kaelbling for reading an earlier draft of this paper.

References

- Aloimonos, Y. (1992). Is visual reconstruction necessary? Obstacle avoidance without passive ranging. *Journal of Robotic Systems*, 9(6), 843–858.
- Anderson, J. A. (1995). *An Introduction to Neural Networks*. Cambridge, MA: MIT Press.
- Bachelder, I. A. & Waxman, A. M. (1994). Mobile robot visual mapping and localization: A view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, 7(6/7), 1083–1099.
- Buel, J. (1935). Differential errors in animal mazes. *Psychological Bulletin*, 32, 67–99.
- Coombs, D. & Roberts, K. (1993). Centering behavior using peripheral vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 440–445 New York.
- Coombs, D., Herman, M., Hong, T., & Nashman, M. (1995). Real-time obstacle avoidance using central flow divergence and peripheral flow. In *Proceedings of the 5th International Conference on Computer Vision*.
- Dashiell, J. F. (1930). Direction orientation in maze running by the white rat. *Comparative Psychology Monographs*, 7, 72 pp.
- Duchon, A. P. & Warren, W. H. (1994). Robot navigation from a Gibsonian viewpoint. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, pp. 2272–2277 Piscataway, NJ: IEEE.
- Duchon, A. P., Warren, W. H., & Kaelbling, L. P. (1995). Ecological robotics: Controlling behavior with optical flow. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, pp. 164–169 Mahwah, NJ: Lawrence Erlbaum Associates.
- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.
- Lee, D. N. (1976). A theory of visual control of braking based on information about time-to-collision. *Perception*, 5, 437–459.
- Markus, E. J., Qin, Y.-L., Leonard, B., Skaggs, W. E., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *The Journal of Neuroscience*, 15(11), 7079–7094.
- Mataric, M. (1992). Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, 8(3), 304–312.
- McNaughton, B. L., Barnes, C. A., & O’Keefe, J. (1983). The contributions of position, direction, and velocity to single-unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52, 41–49.
- Mizumori, S. J. Y., Barnes, C. A., & McNaughton, B. L. (1990). Behavioral correlates of theta-on and theta-off cells recorded from hippocampal formation of mature young and aged rats. *Experimental Brain Research*, 80, 365–373.
- Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12, 239–260.
- Muller, R. U. & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of

- hippocampal complex-spike cells. *The Journal of Neuroscience*, 7(7), 1951–1968.
- Munn, N. L. (1950). *Handbook of Psychological Research on the Rat*. Boston: Houghton Mifflin.
- O’Keefe, J. (1989). Computations the hippocampus might perform. In L. Nadel, L. A. Cooper, P. Culicover, & R. M. Harnish (Eds.), *Neural Connections, Mental Computations*, pp. 225–284. Cambridge, MA: MIT Press.
- O’Keefe, J. (1991). The hippocampal cognitive map and navigational strategies. In J. Paillard (Ed.), *Brain and Space*, pp. 273–295. Oxford: Oxford University Press.
- O’Keefe, J. & Recce, M. (1993). Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3), 317–330.
- O’Keefe, J. & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68, 1–27.
- O’Keefe, J. & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- Quirk, G. J., Muller, R. U., & Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat’s recent experience. *The Journal of Neuroscience*, 10(6), 2008–2017.
- Santos-Victor, J. & Sandini, F. (1994). Visual behaviors for docking. Tech. rep. TR 2/94, LIRA-Lab: DIST University of Genoa.
- Santos-Victor, J., Sandini, G., Curotto, F., & Garibaldi, S. (1995). Divergent stereo in autonomous navigation: From bees to robot. *International Journal of Computer Vision*, 14, 159–177.
- Shapiro, M. L. & Hetherington, P. A. (1993). A simple network model simulates hippocampal place fields: Parametric analyses and physiological predictions. *Behavioral Neuroscience*, 107(1), 34–50.
- Sharp, P. E., Blair, H. T., Etkin, D., & Tzanetos, D. B. (1995). Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *The Journal of Neuroscience*, 15(1), 173–189.
- Sobey, P. (1994). Active navigation with a monocular robot. *Biological Cybernetics*, 71, 433–440.
- Srinivasan, M. V. (1992). How bees exploit optic flow: Behavioural experiments and neural models. *Philosophical Transactions of the Royal Society of London B*, 337, 253–259.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, 10(2), 420–435.
- Touretzky, D. S., Wan, H. S., & Redish, A. D. (1994). Neural representation of space in rats and robots. In J. M. Zurada & R. J. Marks (Eds.), *Computational Intelligence: Imitating Life: Proceedings of the symposium held at the 1994 IEEE World Congress on Computational Intelligence*. IEEE Press.
- Wiener, S. I., Paul, C. A., & Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *The Journal of Neuroscience*, 9(8), 2737–2763.
- Worden, R. (1992). Navigation by fragment fitting: A theory of hippocampal function. *Hippocampus*, 2(2), 165–188.