A Computational Model of Rodent Spatial Learning and Some Behavioral Experiments

Karthik Balakrishnan (balakris@cs.iastate.edu)
Rushi Bhatt (rushi@cs.iastate.edu)
Vasant Honavar (honavar@cs.iastate.edu)
Artificial Intelligence Research Laboratory
Department of Computer Science
Iowa State University, Ames, IA 50011

Abstract

This paper describes a computational model of spatial learning and localization in rodents. The model is based on the suggestion (based on a large body of experimental data) that rodents learn metric spatial representations of their environments by associating sensory inputs with dead-reckoning based position estimates in the hippocampal place cells. Both these sources of information have some uncertainty associated with them because of errors in sensing, range estimation, and path integration. The proposed model incorporates explicit mechanisms for information fusion from uncertain sources. We demonstrate that the proposed model adequately reproduces several key results of behavioral experiments with animals.

Introduction

Animals demonstrate a wide range of complex spatial learning and navigation abilities (Gallistel, 1990). Considerable research effort has been expended in understanding different aspects of these spatial behaviors. These efforts have resulted in a large corpus of experimental data, a number of theories and models of animal spatial learning, and several implementations.

Among the known aspects of animal spatial learning is the suggestion that rodents learn *cognitive maps* of their spatial environments (Tolman, 1948). There is also a vast body of experimental data that directly implicate the hippocampal formation in rodent spatial learning (O'Keefe and Nadel, 1978). Based on this data, O'Keefe and Nadel proposed the *locale system hypothesis*, suggesting that the hippocampal place cells learn metric cognitive maps by integrating *sensory inputs* and *dead-reckoning* [1] position estimates generated by the animal.

In the two decades since the locale hypothesis was first proposed, a number of computational models of hippocampal spatial learning have been developed (Balakrishnan et al., 1997). Surprisingly, only a few of these models support metric spatial representations. Furthermore, the few models that are based on the locale hypothesis make an unrealistic assumption that the two information streams, namely, sensory inputs and dead-reckoning, are error-free. However, sensory and dead-reckoning systems of animals are prone to several sources of errors (e.g., errors in place recognition, distance estimation, dead-reckoning drifts, etc.), and therefore, any computational model of hippocampal spatial learning and localization must be capable of satisfactorily dealing with these errors.

In this paper we describe a computational model of hippocampal spatial learning that allows the animal to learn a metric place map and explicitly addresses information fusion from uncertain sources. Following a brief discussion of experimental data supporting the model, we present some key features of the model and some simulation results that demonstrate that the model satisfactorily reproduces certain behavioral experiments.

Hippocampal Spatial Learning

The *hippocampal formation* is one of the highest levels of association in the brain and receives highly processed sensory information from the major associational areas of the cerebral cortex. For anatomical and physiological details the reader is referred to Churchland and Sejnowski (1992).

Cellular recordings in the hippocampus have led to the discovery of place cells and head-direction cells which show highly correlated firings during the execution of spatial tasks. Pyramidal cells in regions CA3 and CA1 of the rat hippocampus have been found to fire selectively when the rat visits particular regions of its environment. These cells thus appear to code for specific places and have been labeled *place cells* (O'Keefe and Nadel, 1978). Cells with such location-specific firing have been found in almost every major region of the hippocampal system, including the EC, the Dg, regions CA3 and CA1, the Sb, and the post-subiculum.

In addition to place cells, *head-direction cells* in the hippocampal region respond to particular orientations of the animal's head irrespective of its location. These cells thus appear to function as some sort of an in-built compass (Taube et al., 1990).

A number properties of place cells and head-direction cells have been identified (see McNaughton et al. (1996) for details), primarily the fact that the firing of these cells is dependent on sensory as well as dead-reckoning inputs.

Hippocampal Cognitive Map

We have developed a computational implementation of the locale system hypothesis. Our model allows the animat (a robot simulating an animal) to learn its environment in terms of *distinct places*, with the *center* of each place also being labeled with a *metric position estimate* derived from dead-reckoning. A detailed treatment of this computational model

[1] Dead-reckoning or path-integration refers to the process of updating an estimate of one's current position based on self-knowledge of time, speed, and direction of self-motion.

can be found in Balakrishnan et al. (1997). Here we will only present a brief summary.

As the animat explores its environment, the model creates new EC units that respond to landmarks located at particular positions relative to the animat. Concurrent activity of EC units defines a place, and place cells in CA3 layer are created to represent them. These sensory input-driven CA3 place cells are then associated with position estimates derived from the dead-reckoning system to produce place firings in the CA1 layer. Thus, the firing of CA1 cells is dependent on sensory inputs from CA3 and the animat's dead-reckoning position estimates.

When the animat revisits familiar places, incoming sensory inputs activate a place code in the CA3 layer. Since multiple places in the environment can produce the same sensory input (called *perceptual aliasing* in robotics), the CA1 layer uses dead-reckoning estimates to disambiguate between such places and produces a *unique* place code that corresponds to the current place. The system then performs spatial localization by *matching* the *predicted* position of the animat (its current dead-reckoning estimate) with the *observed* position of the place field center (dead-reckoning estimate previously associated with the activated CA1 place code). Based on this match, the dead-reckoning estimate as well as the place field center are updated as shown in Figure 1.

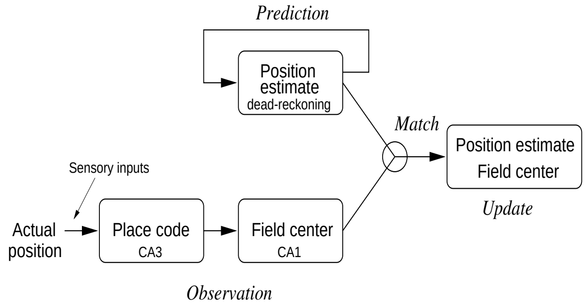

Figure 1: A schematic of hippocampal localization.

Thus, not only does our model learn a metric cognitive map of the environment, but it also permits the metric estimates to be *updated* when the animal revisits familiar places.

Hippocampal Kalman Filtering

It should be noted that the information provided by the sensory and dead-reckoning systems is uncertain because of possible errors in object recognition, distance estimation, and path-integration. Thus, if the hippocampus performs robust spatial localization, it must have adequate mechanisms for handling uncertainty in these information sources.

As with animals, mobile robots too have to deal with uncertainties in sensing and action. One of the probabilistic localization approaches for mobile robots is the Kalman filter (KF) (Gelb, 1974) (or some extension or generalization of it). KF allows the robot to build and maintain a *stochastic spatial map*, propagate sensory and motion uncertainties, and localize in *optimal* ways (Ayache and Faugeras, 1987). A schematic for a KF is shown in Figure 2.

As can be observed from Figures 1 and 2, the computational model of hippocampal function and KF both embody

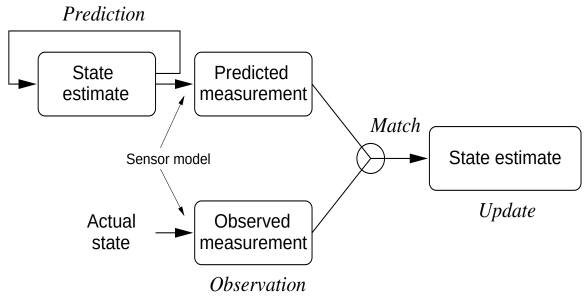

Figure 2: A schematic of Kalman filtering.

the same *predict-observe-match-update* principle. Further, KF provides a framework for performing *stochastically optimal* updates even in the presence of prediction and observation errors. Based on the similarities between the two, we have developed a KF framework for uncertain information fusion in the hippocampal localization model described above (Balakrishnan et al., 1997). In this framework, KF helps the animal in maintaining and updating an estimate of its own position as well as the estimates of the place field centers. These estimates are modeled by a state vector:

$$\mathbf{x}_k \quad = [x_{0,k}, x_1, \ldots, x_n]^T$$

where $x_{0,k}$ denotes the position of the animal at time instant k, x_i denotes the center of place field i, and n is the number of places that have been visited by the animal. These position estimates are assumed to be specified in 2D Cartesian coordinates, i.e., $x_i = (x_{i_x}, x_{i_y})$. The animal also computes and updates the covariance matrix associated with this state vector, denoted by $\mathbf{P}_k$, which is given by:

$$\mathbf{P}_k = \begin{pmatrix} \mathbf{C}_{00} & \mathbf{C}_{01} & .. & \mathbf{C}_{0n} \\ \mathbf{C}_{10} & \mathbf{C}_{11} & .. & \mathbf{C}_{1n} \\ . & . & . & . \\ . & . & . & . \\ \mathbf{C}_{n0} & \mathbf{C}_{n1} & .. & \mathbf{C}_{nn} \end{pmatrix}$$

where

$$\mathbf{C}_{ij} = \begin{pmatrix} C_{i_x j_x} & C_{i_y j_x} \\ C_{j_x i_x} & C_{j_y i_y} \end{pmatrix}$$

denotes the covariance between the 2D Cartesian representations of the state elements x_i and x_j.

When a new place is visited, the state vector is augmented by the center of this new place and the state estimate and its covariance matrix are modified accordingly. If the animal motions are assumed to be linear and the measurement function in Figure 2 is also a linear function of the state, the place field centers and the animal's position estimate can be updated in stochastically optimal ways. For details, refer to (Balakrishnan et al., 1997).

Frame Merging

We have also developed an extension of the computational model described above that permits the animat to learn separate place maps in different *frames* and to merge frames together in a well-defined manner.

Suppose the animat has learned a place map frame f_{old}. When the animat is reintroduced at another place, it stores away f_{old} in its memory, and begins a new frame f_{new} at the point of re-introduction. It also resets its dead-reckoning estimates to zero, thereby making the point of re-introduction the origin of its new dead-reckoning frame. Now it proceeds as before, learning places and creating EC, CA3, and CA1 cells using the algorithms described in Balakrishnan et al., (1997). At each step it also checks to see if sensory inputs excite CA1 cells residing in f_{old}. If this happens, the animat is at a place it has seen earlier in the older frame (f_{old}). It then *merges* the two frames, as follows.

Suppose CA1 unit c fires in f_{new} and m fires in f_{old}. Let $\hat{\mathbf{x}}_c^{f_{new}}$ and $\hat{\mathbf{x}}_m^{f_{old}}$ denote the estimated center of the animat's current place in the two frames f_{new} and f_{old}. We can update the place field centers of f_{old} to f_{new} via the transformation:

$$\hat{\mathbf{x}}_i^{f_{new}} = \hat{\mathbf{x}}_i^{f_{old}} - \Delta\mathbf{x} \qquad \forall i \in f_{old} \tag{1}$$

The covariances between units in f_{old} and f_{new} can be updated using the following expressions (details of the derivations can be found in (Balakrishnan et al., 1998b)):

Case I: If i and j were both units in f_{old}

$$C_{ij}^{f_{new}} = C_{ij}^{f_{old}} - C_{mj}^{f_{old}} - C_{im}^{f_{old}} + C_{mm}^{f_{old}} + C_{cc}^{f_{new}}$$

Case II: If i was a unit in f_{new} and j was in f_{old}

$$C_{ij}^{f_{new}} = C_{ic}^{f_{new}}$$

where C_{ij} refers to the covariance between units i and j in a particular frame.

The animat localizes to the first place that sensorily matches a place it has seen before. Therefore, if multiple places in the environment produce similar sensory inputs (*perceptual aliasing*), this procedure will lead to localization problems.

Goal Representation

Goals encountered by the animat can also be remembered in terms of their metric positions, which is derived from dead-reckoning. Goal position estimate is error prone because the dead reckoning system of the animat is corrupted by noise. We use equation 2 to update the goal location estimate whenever the goal is visited. These updates are stochastically optimal.

$$\begin{aligned} \hat{\mathbf{x}}_{\mathbf{G}} &= \frac{\sigma_0^2}{\sigma_0^2 + \sigma_G^2}\hat{\mathbf{x}}_{\mathbf{G}} + \frac{\sigma_G^2}{\sigma_0^2 + \sigma_G^2}\hat{\mathbf{x}}_{\mathbf{0}} \qquad (2) \\ \sigma_G^2 &= \frac{\sigma_G^2 \cdot \sigma_0^2}{\sigma_0^2 + \sigma_G^2} \end{aligned}$$

Here, $\hat{\mathbf{x}}_{\mathbf{G}}$ is the estimated goal position and σ_G^2 its variance, $\hat{\mathbf{x}}_{\mathbf{0}}$ is the current dead-reckoning estimate with associated variance σ_0^2.

We have also developed a mechanism to remember multiple goal locations in a single frame. When two frames are merged, a Mahalanobis distance test is performed between all pairs of goals in both frames. If two goals are found to be within a squared Mahalanobis distance of less than 4.61, the last visited goal is kept, while the other goal is discarded.

When the animat reaches a goal, a nearest remembered goal location is found from the locations the animat has already visited. If no goal location is nearby, a new goal position with it's goal variance $\sigma_G^2 = \infty$ is created. This ensures that the current dead-reckoning estimate is assigned to the new goal position when equation 2 is applied.

Since the animats can remember multiple goals, a goal selection mechanism is also required. We have implemented a heuristic strategy that involves choosing the most *recently* visited goal location first and navigating to it. If the goal is not found at this location, we use a mechanism in which the animat chooses to approach the nearest goal or a *confident/reliable* goal with equal probability.

The animats in our experiments navigate to goal locations in two ways. If the goal is visible, the animats directly move towards the goal (*goal approaching*). However, if the goal is not visible but the animat has a record of goal locations already seen, it selects a goal as discussed above and moves directly towards it (*goal seeking*). For the purposes of the experiments described in this paper the environments are assumed to be largely open and obstacle-free.

Simulation Details and Results

Experiments of Collett et al. (1986)

We simulated the behavioral experiments of Collett et al. (1986) using our computational model of hippocampal spatial learning described earlier. The experimental setup of Collett et al. consisted of a circular arena of diameter 3.5 meters placed inside a light-tight black painted room. Gerbils were trained to locate a sunflower seed placed in different geometric relationships to a set of visible landmarks. The floor of the arena was such that it prevented the gerbil from spotting the seed until it was very close to the gerbil (Collett et al., 1986).

In our simulations, we used a circular arena of radius 10 units. The walls of the arena were assumed to be devoid of any distinguishing sensory stimuli. The landmarks, on the other hand, were assumed to be visible to the animat from all points in the arena. Estimate of the relative position of the landmark was assumed to be corrupted by a zero-mean Gaussian sensing error with standard deviation σ_S = 0.01. Each landmark at a specific relative position caused an EC unit to fire. A simultaneous activation of EC units caused firing of CA3 and CA1 layers. The animat motion was corrupted by zero-mean Gaussians with σ_M = 0.5 units. The animats also possessed means for dead-reckoning with errors modeled as zero-mean Gaussians with σ_D = 0.05 units.

For each trial, the animat was introduced into the arena at a random position and was allowed to perform 500 steps of *sensing, processing* and *moving*. If the animat happened to see the goal, it was made to approach and consume it. Each animat was subjected to five such training trials. In each trial the animat learned places in a new frame by inducting EC, CA3 and CA1 units in appropriate ways, merged frames when required, and created goal representations.

Once training was complete, the animat was subjected to ten *testing trials,* in which the landmarks in the arena were manipulated in specific ways and the goal was absent. Here, the animat was released at random positions in the arena with its dead-reckoning variance set to ∞. Animats were only capable of localizing and did not induct any new units. A localized animat was allowed a maximum of 300 time steps to navigate to the estimated goal position. Since the goal was not found even after searching for 25 time steps at the goal location, animat chose another goal location.

For the training as well as testing trials, the trajectories followed by the animats were recorded. Also, the arena was decomposed into cells of size 0.33×0.33 and a count of the amount of time spent by the animats in each cell was kept. A normalized histogram for five animats was then plotted.

We simulated the one, two, and three landmark experiments of Collett et al. (1986), and the search distributions of our animats (Figures 3, 4, and 5) match rather closely with those of the gerbils. The large dark squares in the plots denote the landmarks.

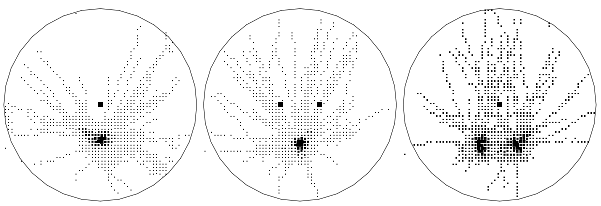

Figure 3: Left: One landmark experiment. Middle: Two landmarks experiment. Right: Two landmarks experiment with one landmark removed.

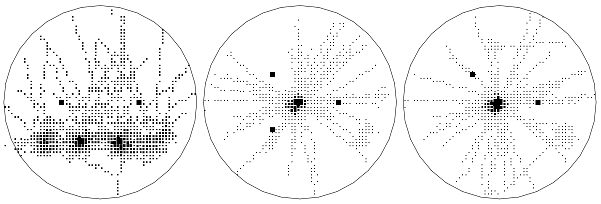

Figure 4: Left: Two landmarks experiment with landmark distance doubled. Middle: Three landmarks experiment. Right: Three landmarks with one removed.

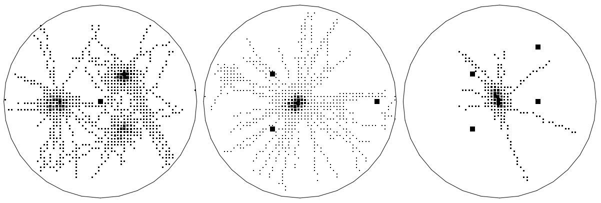

Figure 5: Left: Three landmarks with two removed. Middle: Three landmarks with one distance doubled. Right: Three landmarks with an extra landmark added.

Water-Maze Experiments of Morris (1981)

Morris (1981) experimented with male hooded rats of the Lister strain to demonstrate that rats are capable of rapidly learning to locate an object using distal cues.

A circular pool filled with opaque, milky water was used for the purpose. Objects present along the walls of the room served as distal cues. The pool was devoid of any objects except the escape platform. The escape platform was one of the following two kinds. First type was black colored, circular and protruding above the water, and therefore visible from a distance. Second type was white colored, circular and submerged in the water, thus virtually invisible.

The population of rats was divided into four groups of 8 individuals each. For *Cue + Place* group, the visible, black platform was used, which was always at the same location (NW, NE, SE, or SW) across all trials for a given rat. The second group was exactly like the first except that the white platform was used instead. This was designated the *Place* group. In *Cue-only* group, rats were trained using the black platform. However, in this case the platform was placed in one of the four positions, in an unpredictable sequence over trials. Finally, the *Place-Random* group was similar to the Cue-only experiment except in the use of the white platform instead of the black one.

For each trial, the rats were released in the pool, and their trajectories were recorded along with the time taken to find the platform. Following 20 such trials over 3 days, the groups were further divided into subgroups of 4 individuals each. Each of these subgroups was subjected to 4 testing trials, of type A or B.

In Test A the platform was removed and search behavior was observed for 60 seconds. For Test B, rats of groups Cue + Place and Place were tested with the platform now placed in the quadrant diagonally opposite the one used in training. Rats of groups Cue-only and Place-Random were tested with the platform position held fixed. The escape behavior of the animals was then observed.

In our simulations, we used a circular arena of radius 3.75 units inside a square room measuring 20 by 20 units. Consistent with the ratio of pool and platform sizes in Morris' experiments, we chose the radius of our simulated platform to be 0.65 units. It was assumed that the animat could see the platform from a distance of 0.325 units. Four indistinguishable landmarks were used, one along each wall of the simulated room.

The sensing, motion, and dead-reckoning errors were same as in the foregoing experiments. We also assumed that rats swam slower than their normal walking speeds, and hence the size of motion step was set to 0.4 units.

As in the case of the original experiment, we allowed our animats four pre-training trials in which they randomly explored the environment for 100 steps without the platform present in the pool. During this stage, our spatial learning system allowed the animat to acquire a spatial map corresponding to the environment. In the training trials, the animats engaged in the goal seeking behavior. If the platform was not found at the particular goal location, the animat searched for 15 time steps before selecting another goal location and navigated towards it.

Groups of eight animats each were used in experiments corresponding to groups Cue-only, Cue + Place, Place, and Place-Random as in Morris (1981). The escape latencies for the first 20 training trials and the last four trials of Test B are shown in figure 6.

As seen in Figure 6, the Place group quickly learned the goal position. Furthermore, the Cue group achieved very

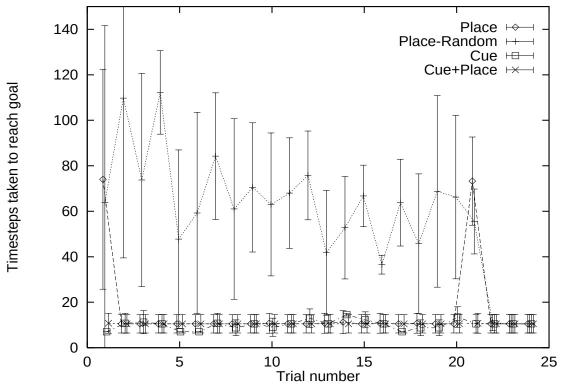

Figure 6: Escape Latencies while training and test B.

small escape latencies. One reason for this is the fact that our simulation had a built-in mechanism to directly approach visible platforms from the start. Actual animals may not have such direct approach behaviors preprogrammed but may learn them with experience. Further, as with rats, our animats too perform poorly in the Place-Random experiment.

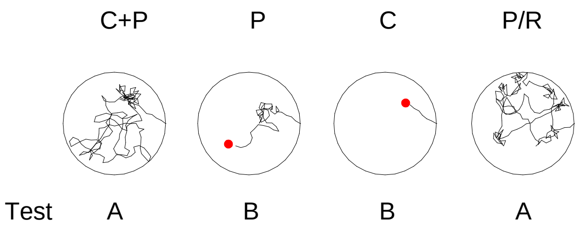

Figure 7: Trajectories followed by the animats (see text for detail).

Figure 7 shows the paths taken during the first test trial by representative animats in different groups. Labels C+P, P, C and P/R denote group Cue + Place, Place, Cue-only and Place-Random respectively.

Discussion

The primary goal of the simulations was to test whether our computational model of hippocampal spatial learning and localization was capable of reproducing the behavior of gerbils. We simulated a number of experiments conducted by Collett et al. (1986) and by Morris (1981) .

It should be pointed out that our animats did not remember goals in terms of *independent vectors* to individual landmarks, as suggested by (Collett et al., 1986). Instead, places were remembered as independent vectors to landmarks, while the goal was simply remembered as a place.

In the process of simulating behavior, we identified an important issue, namely, how do animals choose one goal to approach from multiple ones that they might remember? In order to simulate the Place-Random experiments of Morris, we had to incorporate a heuristic goal selection strategy. Our

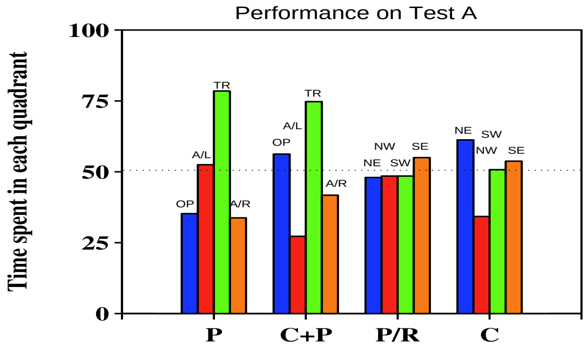

Figure 8: Performance on Test A. Histogram shows the duration of time spent by the animats in each quadrant. Here TR is the training quadrant, A/L and A/R are the adjacent quadrants to the left and right respectively, and OP is the opposite quadrant for groups Cue + Place and Place. The data from the other groups simply indicate the quadrant.

results using this mechanism closely parallel the behaviors observed by Morris. Indeed, our computational framework allows one to implement and test different hypotheses of goal selection. Such an approach can lead to a better understanding of goal selection processes in navigating rodents.

From Figure 8 it can be observed that Place and Cue + Place experiments indicate a strong spatial bias towards the training quadrant. While the former observation is consistent with the results of Morris, the latter is a surprise. However, this is a direct result of our spatial learning and navigation strategy, where we have assumed that the animat faithfully learns a place map. There is a possibility that in the presence of reliable visual cues (e.g., platform), place learning may not be as reliable, since it is not even necessary. This hypothesis regarding differences in place learning in the presence or absence of reliable cues, remains to be studied.

Related Work To the best of our knowledge the only other computational simulation of the experiments by Collett et al. (1986) is that of Redish and Touretzky (1996) (referred to hereafter as the RT model). Our computational model of hippocampal spatial learning is closely related to the RT model, since both models are based on the cognitive map concept of Tolman (1948) and the locale system hypothesis of O'Keefe and Nadel (1978). Finally, both simulations represent goals in terms of metric position estimates derived from dead-reckoning.

Despite these similarities, there are some significant differences between the two models and the behavioral results generated by them. Our computational framework explicitly addresses the issue of information fusion from erroneous (or uncertain) sources. Secondly, by using the framework of Kalman filtering, we have derived update expressions which are *stochastically optimal* for the given model. Thirdly, the RT model incorporates a mechanism for initializing the head direction. However, doing so makes the place cells *directional*, which appears to be at odds with experimental results that suggest the *non-directionality* of the CA3 and CA1 pyramidal cells. Our model assumes that the place cells are non-directional. This requires that the animats have reliable head-direction information. Fourthly, our model learns and remembers multiple goal locations. It is not clear how one could represent multiple goals in the RT model, considering that goals in their model correspond to the origin of the dead-reckoning system. Finally, the animats in our simulations were capable of actually moving in their environment, whereas the animats used in the RT simulations simply predict the goal location. A more detailed comparison between the two models can be found in Balakrishan et al. (1998a)

A number of researchers have developed computational models to simulate the behavioral experiments of Morris (1981). Blum and Abbott (1996) developed a model of hippocampal spatial learning, where the place cell activity encoded the place further ahead on the path. The animat navigated towards the position coded by the place cell activity at any given place. However, they only simulated the behavior of the Place group of rats.

Sharp et al. (1996) used a slightly different model to simulate the water-maze experiments of Morris. In their model the place and head direction cell firings allowed firing of unit in the nucleus accumbens which caused the animat to turn left or right by a certain angle. The experiments with this model matched closely with the Morris' experiments, but it was not clear how multiple goals can be represented in this framework.

Kalman filter-based localization approaches require a sensor model of the environment and often run into problems due to perceptual aliasing (Ayache and Faugeras, 1987). The hippocampal model, on the other hand, provides a place-based extension of KF and addresses these problems (Balakrishnan et al., 1997). A number of robot localization approaches based on *cognitive mapping* theories have also been developed (Kortenkamp, 1993). Although closely related to the hippocampal spatial learning model, they are not formulated to computationally characterize a specific brain region and differ in this regard. A number of neurobiological models of robot navigation have also been developed (Recce and Harris, 1996). A more detailed discussion about the different models is available in Balakrishnan et al. (1997).

Future Work As we mentioned earlier, our computational model assumes that the animat has an accurate head-direction estimate. We are currently exploring the possibility of such a head-direction reset mechanism being implemented by place cells in the subiculum with the correction being performed by the head-direction cells in the post-subicular region (Balakrishnan et al., 1998b).

Given the fact that Kalman filter based models of place learning and localization satisfactorily reproduce a non-trivial body of results from behavioral experiments in animals, it is natural to ask whether the hippocampus can perform KF computations. A discussion about this can be found in (Balakrishnan et al., 1997)

Acknowledgements

Olivier Bousquet's contributions to the theory of Hippocampal Kalman filtering is gratefully acknowledged. Karthik Balakrishnan and Rushi Bhatt are supported by an IBM Cooperative Fellowship and a Neuroscience Graduate fellowship from Iowa State University respectively. This research was partially supported by grants from the National Science Foundation (IRI-9409580) and the John Deere Foundation to Vasant Honavar.

References

Ayache, N. and Faugeras, O. (1987). Maintaining representation of the environment of a mobile robot. In *Proceedings of the International Symposium on Robotics Research*, Santa Cruz, California, USA.

Balakrishnan, K., Bhatt, R., and Honavar, V. (1998a). Spatial learning and localization in animals: A computational model and behavioral experiments. In *Proceedings of the Second European Conference on Cognitive Modelling*, pages 112–119, Nottingham, UK. Nottingham University Press.

Balakrishnan, K., Bhatt, R., and Honavar, V. (1998b). Spatial learning in the rodent hippocampus: A computational model and some behavioral results. (in preparation).

Balakrishnan, K., Bousquet, O., and Honavar, V. (1997). Spatial learning and localization in animals: A computational model and its implications for mobile robots. Technical Report CS TR 97-20, Department of Computer Science, Iowa State University, Ames, IA 50011. (To appear in Adaptive Behavior).

Blum, K. and Abbott, L. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93.

Collett, T., Cartwright, B., and Smith, B. (1986). Landmark learning and visuo-spatial memories in gerbils. *Journal of Neurophysiology A*, 158:835–851.

Gallistel, C. (1990). *The Organization of Learning*. Bradford-MIT Press, Cambridge, MA.

Gelb, A. (1974). *Applied Optimal Estimation*. MIT Press.

Kortenkamp, D. (1993). *Cognitive Maps for Mobile Robots: A Representation for Mapping and Navigation*. PhD thesis, University of Michigan, Electrical Engineering and Computer Science Department.

Morris, R. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–261.

O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford:Clarendon Press.

Recce, M. and Harris, K. (1996). Memory for places: A navigational model in support of marr's theory of hippocampal function. *Hippocampus*, 6:735–748.

Sharp, P., Blair, H., and Brown, M. (1996). Neural network modeling of the hippocampal formation spatial signals and their possible role in navigation: A modular approach. *Hippocampus*, 6:720–734.

Taube, J., Muller, R., and Ranck, J. (1990). Head direction cells recorded from the postsubiculum in freely moving rats: I. description and quantitative analysis. *Journal of Neuroscience*, 10:420–435.

Tolman, E. (1948). Cognitive maps in rats and men. *Psychological Review*, 55:189–208.