Simulation of place fields in the computational model of rodent spatial learning[*]

Rushi Bhatt
Dept. of Computer Science
Iowa State University
Ames, IA 50011
rushi@cs.iastate.edu

Karthik Balakrishnan
Allstate Research
& Planning Center
Menlo Park, CA 94025
kbala@allstate.com

Vasant Honavar
Dept. of Computer Science
Iowa State University
Ames, IA 50011
honavar@cs.iastate.edu

Abstract

Recent work (Balakrishnan, Bousquet, & Honavar 1997; Balakrishnan, Bhatt, & Honavar 1998) has explored a *Kalman filter* model of animal spatial learning the presence uncertainty in sensory as well as dead-reckoning estimates. This model was able to successfully account for several of the behavioral experiments reported in the animal navigation literature (Morris 1981; Collett, Cartwright, & Smith 1986). This paper extends this model in some important directions. It accounts for the observed firing patterns of hippocampal neurons (Sharp, Kubie, & Muller 1990) in visually symmetric environments that offer multiple sensory cues. It incorporates mechanisms that allow for differential contribution from proximal as opposed to distal landmarks during localization. It also supports learning of associations between rewards and places to guide goal-directed navigation.

Introduction

The computational strategies used by animals to acquire and use spatial knowledge (e.g., maps) for navigation have been the subject of study in Neuroscience, Cognitive Science, and related areas. A vast body of experimental data directly implicates the hippocampal formation in rodent spatial learning. The present model is based on the anatomy and physiology of the rodent hippocampus (Churchland & Sejnowski 1992) and on the locale hypothesis (which argues for the association of configurations of landmarks in the scene to the animal's own position estimates at different places in the environment as suggested by O'Keefe and Nadel(O'Keefe & Nadel 1978). A computational model of rodent spatial learning and localization was proposed in (Balakrishnan, Bousquet, & Honavar 1997). At the core of this model is a three-layer feed-forward constructive neural network. The first layer roughly corresponds to the Entorhinal Cortex (EC). Neurons in

this layer act as spatial filters, responding to specific types of landmarks at specific relative positions to the animat (simulated animal), producing a sparse code of sensory stimulus. The units have Gaussian activation functions, and the measurements of distances are corrupted by a zero mean, uncorrelated Gaussian noise. Neurons in the second layer (which corresponds to the CA3 layer) respond to a group of EC layer activations, activating an internal learned place-code. Neurons in the third layer, corresponding to the CA1 layer in the hippocampal system, compare this place-code with the dead-reckoning estimate (i.e., a position estimate that is generated by the animat by keeping track of its own movement) to establish a unique place code. A new set of neurons is recruited into the network if the animat finds itself at a previously unvisited place as it explores its environment. Such exploration can take place over multiple, separate episodes. Fragments of place-maps learned over several such episodes are integrated to form a coherent map whenever there is an overlap in incoming sensory information. The model is able to successfully deal with perceptual aliasing (i.e., when different places look alike) during a single episode (Balakrishnan, Bousquet, & Honavar 1997). This model of spatial learning uses a Kalman Filter like approach (Kalman 1960) to calculate and correct estimates of the animat's position in the environment in the presence of errors in sensing and dead reckoning (Balakrishnan, Bousquet, & Honavar 1997). This model is able to successfully account for a large body of behavioral results (Balakrishnan, Bhatt, & Honavar 1998), reported in the animal navigation literature (Morris 1981; Collett, Cartwright, & Smith 1986). In this paper, we report results of simulation based on extensions of the model in some key directions. The proposed extensions account for the observed firing patterns of hippocampal neurons in visually symmetric environments that include multiple sensory cues reported in (Sharp, Kubie, & Muller 1990). They allow for differential contributions of different landmark types during localization. They also support learning of associations between places in the environment and rewards through exploration thereby providing a basis for goal-directed navigation.

Variable Tuning Widths of EC Layer Spatial Filters

The Kalman Filter based model of spatial learning on which this work is based is described in detail in (Balakrishnan, Bousquet, & Honavar 1997). EC layer cells in the present model act as spatial filters, responding to individual landmarks at specific positions relative to the animat. We have extended these filters in light of work by O'Keefe and Burgess (1996), so that the tuning curves of such filters vary with the landmark distances, as discussed shortly.

As the animat explores its environment, it recruits a new EC cell if no existing cell fires above a threshold level to an observed landmark position (Balakrishnan, Bousquet, & Honavar 1997). A newly recruited EC cell has a Gaussian activation function:

$$EC_i = \frac{1}{2\pi \ \sigma_1 \ \sigma_2} \exp(-\frac{1}{2} \left(\frac{(x - \mu_x)^2}{\sigma_1^2} + \frac{(y - \mu_y)^2}{\sigma_1^2} \right))$$

$$\sigma_1 = \sigma_0 \ \left(1 + 4\mu_1^2/R^2 \right)$$

$$\sigma_2 =_0 \ \left(1 + 4\mu_2^2/R^2 \right)$$

Where μ_1 is the distance of the landmark in direction x_1 from the current position of animat, and similarly, μ_2 is the distance of the landmark in direction x_2 from the current position of animat. We have used a Cartesian cooridinate system where directions x_1 and x_2 are mutually orthogonal.

For the purpose of simulation, σ_0 was set to 1.0 and R was set to 20, equal to the diameter of the circular arena. As we shall see in what follows, this modified activation function has an interesting effect on the localization behavior exhibited by animats.

For the purpose of measuring the EC layer firing fields, animats were trained with one landmark present in the environment, and tested by adding another identical landmark in the environment. As can be seen in Figure , EC cell tuned to respond to a particular relative position of the landmark in training phase responded at two different locations, once for each of the landmarks during the testing phase.

All-Or-None Connections Between EC and CA3 Layers

While the animats are trained, if none of the CA3 layer cells fire above a threshold fraction of their peak firing level, a new CA3 layer cell is allocated. This newly created cell is then tied to the active EC layer cells. The connection weight assignment procedure was also modified to reflect a more all-or-none connection type. Rather than assigning weights proportionate to the activation of corresponding EC cells, a weight of $1/nconn$ was assigned to each of the links, where $nconn$ is the total number of EC layer cells firing above their threshold levels. Firing threshold of the newly allocated CA3 layer cell was arbitrarily set to 70% of its maximum possible activation level. During testing, however, the
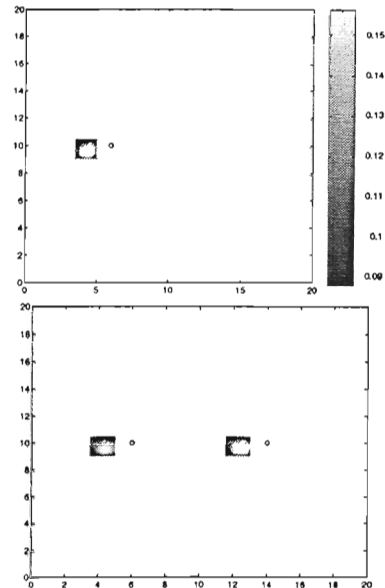

Figure 1: Activation field of an EC layer cell. Top: Trained with a single landmark. Bottom: Tested with an additional identical landmark.

threshold was reduced to 25% of the maximum possible activation level. This enabled activation of CA3 cells, which, in turn allowed animats to localize even in presence of partial sensory stimulus.

At this point it should be noted that CA3 cell activation is simply a threshold sum of incoming Gaussian activation functions from the EC layer cells. Such a method has been found to be successful in modeling place-cell firing characteristics in simplified environments (O'Keefe & Burgess 1996).

It has also been observed that rodents give more importance, while localizing, to landmarks physically closer to their actual positions. Sharp and colleagues (Sharp, Kubie, & Muller 1990) performed experiments on rodents in a cylindrical environment with a single polarizing cue present. After training, one more cue was added to the environment, producing a mirror symmetry in the environment. It was found that an overwhelming number of place-fields retained their shape and orientation with respect to only one of the two cues. In other words, instead of place-fields firing at multiple places, once for each of the cues, they maintained their firing characteristics with respect to either of the two cues. Also, in most cases, place-fields were fixed relative to the cue that was nearest to the animal when it was first introduced in the environment. As we shall see shortly, our model, with the aforementioned modifications implemented, also exhibits a similar behavior.

Association of rewards with places

We have also extended the model to incorporate mechanisms that result in an enhancement of response in
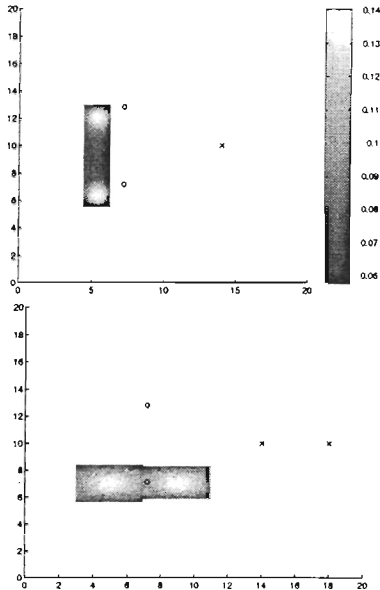
Figure 2: Activation fields of an EC layer cells in an environment with three landmarks. Top: a single EC cell responds to two landmarks. Bottom: EC cell firing before (distribution on left) and after (distribution on right) the landmark denoted by 'x' was moved

the EC layer to landmark types that are closer to the reward locations. Whenever the animat receives a reward upon visiting a pre-determined place in the environment, the maximum possible activations in EC layer cells are updated according to the following rule:

$$\delta w_j = \frac{1}{n-1} \ \Gamma \ \frac{\sum\limits_{i=1, i \neq j}^{n} d_i}{\sum\limits_{i=1}^{n} d_i}$$

where n is the number of types of landmarks present in the environment, $\Theta(i)$ is the total number of landmarks of type i present in the environment, and Γ is the amount subtracted from the landmark weights. Γ is computed as follows:

$\Gamma = 0$
for $i = 1$ to n
do
    if $w_i < \alpha \ \Theta(i)$
       $\Gamma = \Gamma + w_i$
       $w_i = 0$
    else
       $\Gamma = \Gamma + \alpha \ \Theta(i)$
       $w_i = w_i - \alpha \ \Theta(i)$
    endif
done

If multiple landmarks of same type are present, weights are altered by summing the distances of land-

marks of similar types to the estimated goal location. Here, the special case of $n = 1$ was handled separately. It is clear that the weights remain unaltered if all landmarks are of the same kind, or, if all landmarks are equidistant from the goal. For the purpose of our simulations, α was set to 0.05 and the weights were initialized to 1.0.

Briefly, the above rule gives more preference to landmark types that are near the goal location by removing a small amount from weights assigned to each of the landmark types, and redistributing it so that a landmark type gains weight if a landmark of the type in question is near the goal. On the other hand, if there are multiple landmarks of the same type, or, if landmarks are far from the goal, such landmark type loses weight. It is also clear from the above equations that the sum of weights assigned to all landmark types remains unaltered.

It should be noted that the activation level is modulated uniformly across all EC cells that respond to a particular landmark type, and not just for EC cells that are active at the time of reward presentation. Or, in other words, animat learns to give more importance to particular *types* of features present in the environment, rather than to particular features themselves.

Simulation Results

For the simulations presented here, all parameters and methods were identical to those in (Balakrishnan, Bhatt, & Honavar 1998). Briefly, the animats were introduced in an a-priori unknown environment that consisted of one or more landmarks. The landmarks could be identical or distinguishable from each other depending of the experiment being performed. Animats then explored their environments and allocated cells corresponding to different locations in the environment. They were also capable of updating their dead-reckoning position estimates and the place fields using a Kalman Filter like update rule. Animats could explore their environments over multiple trials and could merge these separate training frames to form a single, coherent place-map. The animats were also rewarded for visiting specific locations as they explored its environment. After a certain number of training trials, animats were removed from the environment, landmark positions were altered and the reward removed. When reintroduced in the environment, animats were able to re-localize, despite the change in configuration of the landmarks, using the available perceptual input, and moved toward the learned goal location.

Firing characteristics of units

In order to simplify analysis, for this part of experiments, animats were trained over a single training trial of 750 steps of random exploration.

As seen in Figure , animat consistently localized by giving more preference to the landmark physically closer to the point of entry into the environment. This

phenomenon was not guaranteed with the scheme used in (Balakrishnan, Bhatt, & Honavar 1998). It should be noted, however, that the overall behavior displayed by animat stays unaltered with these enhancements, and we get search histograms similar to those in (Balakrishnan, Bhatt, & Honavar 1998).
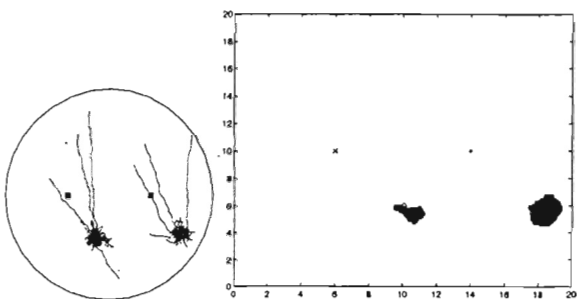

Figure 3: Left: Trajectories taken by an animat during test trials. Right: Superimposed place field firing regions during test trials.
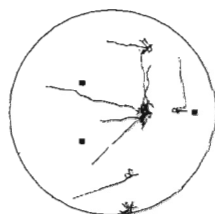

Figure 4: Left to right: Trajectories taken by animat when trained in an environment with three landmarks. The landmark on far right was distinguishable from rest. The landmark on far right was moved further while testing.

Figure also shows the activation of CA1 layer place-cells once animats localized. It is important to note that the place-cell in question fired only at one of the two places over a single test trial. The place-cell in question fired at the position cluster based around position $(12, 5)$ when the animat localized according to the landmark on the left, while the same place cell fired at places clustered around $(18, 6)$ when the animat localized using the landmark on the right. Interestingly enough, the right cluster is spread over a larger area, signifying a larger place-field when animat localized using the right landmark. We suspect an insufficient amount of training at those points of re-introduction. Figure shows the trajectories during test trials, when one of the landmarks was distinct from the rest during training. Also, in some of the trials animats were completely unable to localize. This was because of a lack of training at those points of introduction in the environment. It can

be seen in Figure that in one of the test trials, animat localized solely based on the position of the right most landmark. The reason for such a behavior is discussed in the next subsection.

Landmark prominence based on location and uniqueness

The extensions to the model that alter prominence assigned to landmark type is able to successfully replicate some of the behavioral results that were unaccounted for in (Balakrishnan, Bhatt, & Honavar 1998), namely, the experiments where an array of three landmarks with different types of landmarks was transformed, in figure 9 c of (Collett, Cartwright, & Smith 1986) as seen in Figure . The simulation parameters used here were identical to those in (Balakrishnan, Bhatt, & Honavar 1998).

In addition, simulations demonstrate that the proposed extensions enable the animat to acquire associations between rewards and places and use them for goal-directed navigation.
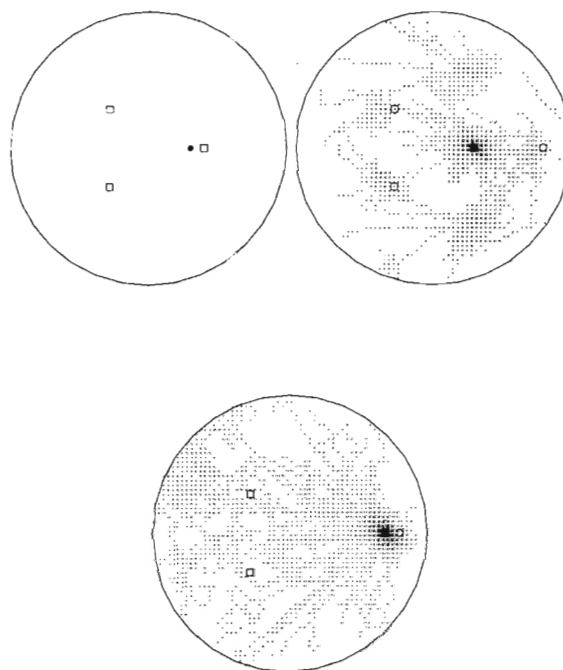

Figure 5: Top Left: Training Environment;Top Right: Normalized test histograms averaged over five animats with ten test trials each, when landmarks were indistinguishable from each other; Bottom: Right-most landmark distinguishable from the rest

As seen in Figure , during training the animats learned to give more weight to the type of landmark on the extreme right, due to its proximity to the goal as

well as the uniqueness of its type. During testing, the right most landmark, which was distinguishable from the rest, was moved further towards right. Animats localized based on this unique landmark. Hence, a simple rule to associate the landmark type to a goal location was learned. Obviously, if all landmarks are identical, no such rule was learned, and the animats localized using a majority vote, as seen in Figure , top right. In Figure , since only one visit to goal was allowed during training, the effect of weight update was not very pronounced.

Related Work

Burgess and colleagues have implemented a robotic simulation of rat navigation, which effectively reproduces place-cell firing characteristics (Burgess *et al.* 1997). However, it is unclear how a metric distance between any two places in the environment can be coded in their model. Our model, on the other hand, labels each place-cell with a metric position in the environment, thus, providing a basis for computing the distance between any two place fields. Also, as the animat navigates in its environment, the model uses a Kalman Filter like update procedure to reduce the effects of errors in the sensory and the dead-reckoning systems.

O'Keefe and Burgess have been successful in replicating place fields in an simple rectangular environment (O'Keefe & Burgess 1996). However, in their model, the EC layer cells parameters were preprogrammed. In the model discussed in this paper, on the other hand, EC layer cells are automatically assigned Gaussian activation functions of varying widths. Nevertheless, the characteristics of the EC layer cells in the present model are similar to those of O'Keefe and Burgess (1996) . Also, in the model discussed here, cells are allocated incrementally as new perceptual information arrives.

Conclusion and Future Work

We have presented several extensions of a *Kalman*-filter based model of animal spatial learning that was presented in (Balakrishnan, Bousquet, & Honavar 1997; Balakrishnan, Bhatt, & Honavar 1998). The model is capable of learning and representing metric places in an a priori unknown environment and localizing when reintroduced in an environment. The extensions presented here enable the model to learn to give preference to certain types of landmarks based on their uniqueness and proximity to the goals. This results in the usage of different strategies under different landmark configurations.

We have also demonstrated that the proposed model is able to successfully replicate the firing characteristics of cells in behaving animals in visually symmetric environments that offer multiple sensory cues as reported in (Sharp, Kubie, & Muller 1990).

The mechanisms that govern selection of landmarks in a dynamic environment are yet to be understood.

Also, if the hippocampus is capable of storing information about multiple environments, it would be interesting to find out how such representations can co-exist, and how they can be retrieved when the animal first enters a familiar environment.

Another question that can be asked is, do there exist multiple maps of the same environment, either at different granularities or using only subsets of the available sensory features? If such is the case, how is the path-integration system reset by the animal when it first enters a previously explored environment? Recent evidence on place-field firings in rodent hippocampus sched new light on the usage of sensory features. These findings suggest that neurons in the hippocampus encode places using distinct subsets of environmental cues. It is suggested that processing within the hippocampus tends to reduce the conflicting orientation information conveyed by these distinct subsets of cues under altered environments (Tanila, Shapiro, & Eichenbaum 1997). Authors are currently investigating the feasibility of such a mechanism for deciding on a correct orientation of the place-map in presence of conflicting environmental cues using the present model of the rodent hippocampal system.

Although several hypotheses about the hippocampal function have been put forth in the literature (McNaughton *et al.* 1996; Redish & Touretzky 1997), they remain to be verified through concrete realizations in terms of computational models that explain the relevant neurobiological as well as behavioral data. Against this background, precise computational characterization of the capabilities and limitations, as well as experimental comparison of different models of animal spatial learning is clearly of interest. It is also interesting to explore whether Kalman-filter based approaches offer a general framework for modeling the role of the hippocampus and related structures in settings other than spatial learning (e.g., episodic memories). Exploration of the relationship between such models and various hypotheses concerning the mechanisms underlying memory consolidation and the role of hippocampal system in learning (Nadel & Moscovitch 1997; Redish & Touretzky 1998), as well as further exploration of reward-based learning in spatial domains offer other promising directions for future research.

References

Balakrishnan, K.; Bhatt, R.; and Honavar, V. 1998. A computational model of rodent spatial learning and some behavioral experiments. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society*, 102–107. Mahwah, New Jersey: Lawrence Erlbaum Associates.

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 Com-

puter Science, Iowa State University, Ames, IA 50011. (To appear in Adaptive Behavior).

Burgess, N.; Donnett, J.; Jeffery, K.; and O'Keefe, J. 1997. Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society London B* 352.

Churchland, P., and Sejnowski, T. 1992. *The Computational Brain*. Cambridge, MA: MIT Press/A Bradford Book.

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

Kalman, R. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME* 60.

McNaughton, B.; Barnes, C.; Gerrard, J.; Gothard, K.; Jung, M.; Knierim, J.; Kudrimoti, H.; Qin, Y.; Skaggs, W.; Suster, M.; and Weaver, K. 1996. Deciphering the hippocampal polyglot: the hippocampus as a path-integration system. *The Journal of Experimental Biology* 199(1):173–185.

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

Nadel, L., and Moscovitch, M. 1997. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology* 7:217–227.

O'Keefe, J., and Burgess, N. 1996. Geometric determinants of the place fields of hippocampal neurons. *Nature* 381.

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

Redish, A., and Touretzky, D. 1997. Separating hippocampal maps. In Burgess, N.; Jeffery, K.; and O'Keefe, J., eds., *Spatial Functions of the Hippocampal Formation and the Parietal Cortex*. Oxford University Press.

Redish, A., and Touretzky, D. 1998. The role of the hippocampus in solving the morris water maze. *Neural Computation* 10:73–111.

Sharp, P.; Kubie, J.; and Muller, R. 1990. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic properties. *Journal of Neuroscience* 10:2339–2356.

Tanila, H.; Shapiro, M.; and Eichenbaum, H. 1997. Discordance of spatial representation in ensembles of hippcoampal place eells. *Hippocampus* 613–623.