

A Hybrid Model for Rodent Spatial Learning and Localization¹

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 [3, 2] has explored a Kalman filter model of animal spatial learning the presence uncertainty in sensory as well as path integration estimates. This model was able to successfully account for several of the behavioral experiments reported in the animal navigation literature [10, 6]. This paper extends this model in some important directions. It accounts for the observed firing patterns of hippocampal neurons [19] in visually symmetric environments that offer polarizing sensory cues. It incorporates mechanisms that allow for differential contribution from proximal and 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 data from lesion studies and cellular recordings directly implicates the hippocampal formation in rodent spatial learning[14]. The present model is based on the anatomy and physiology of the rodent hippocampus [5]. We draw inspiration 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. The system that generates the animal's own position estimate using the sensory as well as motor information is referred to as the path integration system[14].

The hippocampus (Cornu Ammonis, abbreviated as

CA), Dentate Gyrus (Dg) and the Subicular regions are collectively referred to as the hippocampal formation. The hippocampal formation receives majority of its sensory input from the Entorhinal Cortex (EC). The Entorhinal Cortex region itself is believed to be a major association area. Neuronal tract tracing experiments have discovered major projections originating from the inferior temporal gyri (higher processing area for visual sensory information), parietal and temporal lobes (higher processing area for auditory sensory information), olfactory bulb and pyriform cortex (primary processing areas for olfactory information) and from the frontal lobes of the cortex project into the superficial cortical layers of the EC. The EC can therefore be hypothesized to produce a sparse code of the scene based on the sensory information converging to it. Systematic and roughly topographic connections originating from amygdala have also been found to project into the entorhinal cortex as well as into the hippocampus [15]. Projections from medial septum, dorsal raphe nucleus, locus ceruleus and parts of the thalamus also arrive at hippocampus via the fornix. These subcortical projections are believed to carry modulatory signals that influence activity in the hippocampal formation[21]. The hippocampal formation also projects back to most of the areas that project to the hippocampal formation[21]

Units in the first layer of our model correspond to EC layer cells. In accordance with the experimental evidence [13], we model the EC layer units as spatial filters with Gaussian shaped receptor fields that fire in response to individual landmarks at specific relative positions from the simulated animal's (henceforth animat) current position. The sensory measurements are assumed to be noisy, corrupted by an uncorrelated Gaussian noise which varies in proportion to the distance of the observed landmark[3].

The EC layer cells project to the Dg via the perforant path fibers. The Dg in turn projects onto the CA3 layer via the mossy fibers. These mossy fiber synapses with CA3 are quite strong. It has been suggested that they provide the context for arriving information [12] or

¹This research was supported in part by grants from the National Science Foundation (NSF IRI-9409580) and the John Deere Foundation to Vasant Honavar, an IBM cooperative research fellowship to Karthik Balakrishnan during 1997-98 and a Neuroscience graduate research assistantship to Rushi Bhatt during 1997-98

reference frame in navigation [9] by transformation of the sensory input activity arriving at entorhinal cortex into a non-overlapping activity pattern of granule cells in Dg which is then conveyed to the pyramidal cells in the CA3 layer. It has been found that on an average a granule cell in Dg makes contact with around 14 pyramidal cells and each pyramidal cell is innervated by only about 46 granule cells[5] which suggests a sparsely connected network. Many pyramidal cells in the hippocampus as well as neurons in surrounding regions including EC and Subiculum have been found to display a place-dependent firing characteristic, that is, these cells fire in a complex spike burst only when the animal visits geometrically constrained regions of its environment. The cells that display such place dependent firing are commonly referred to as “place cells”[14].

Neurons in the second layer of our model (which corresponds to the CA3 layer) respond to a group of EC layer activations which activate an internal learned place-code. In the present model, a unit in the CA3 is connected to all EC layer units that were active at the time of incorporation of that unit. A unit that only responds to subsets of the active EC layer cells in accordance with observations by Tanila and colleagues [20] is also conceivable. We choose to connect the whole set of active EC layer cells to a CA3 cell under the assumption that due to processing in the Dg, contextual information, or in this case, the information about other landmarks might be available at CA3 layer. It is also possible that the recurrent connections amongst the pyramidal cells in CA3 as well as the widely spread lateral connections by the interneurons[1] provide information about activity in other pyramidal cells in the environment. It has also been suggested elsewhere that CA3 layer cells play a pattern-completion role using partially available sensory information[8, 18]. For the sake of simplicity, in our model CA3 layer cells respond to the complete set of EC layer cells that were active at the time of allocation of the CA3 cell.

Each unit in the third layer is connected to a CA3 layer unit which was active at the time of allocation of that unit. The newly allocated third layer unit, which corresponds to the CA1 layer in the hippocampus, is also labeled with the metric path integration position estimate generated by the animat. This newly incorporated unit is also labeled with a variance of the position estimate at the time of incorporation. A Mahalanobis test is performed between all CA1 units attached to the active CA3 units. The CA1 unit with the lowest Mahalanobis distance is taken as an active place cell. A new set of CA3 and CA1 units is recruited into the network if the animat finds itself at a previously unvisited place as it explores its environment. A previously

unvisited place is defined as a place where either no CA3 layer cell fires, or the Mahalanobis distances of a metric place-label of all active CA1 layer cells from the current path integration estimate are larger than a threshold distance. The model can learn such encodings of its environment over several episodes of random explorations which are then 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 [3]. The model uses a Kalman Filter like approach [7] to calculate and correct estimates of the animat’s position in the environment in the presence of errors in sensing and path integration by comparing and updating labels associated with CA1 layer units with the path integration estimates [3]. The model is able to successfully account for a large body of behavioral results [2], reported in the animal navigation literature [10, 6]. 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 [19]. 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

EC layer cells in the present model act as spatial filters, responding to individual landmarks at specific positions relative to the animat. O’Keefe and Burgess[13] showed that the place cells in hippocampus can be modeled as a sum of Gaussians of varying variances where each Gaussian function encodes the distance to an edge of the environment along one of the two orthogonal axes. We have extended the spatial filters in EC in the light of the aforementioned work so that the tuning curves of such filters vary with the landmark distances along two orthogonal axes.

As the animat explores its environment, it recruits a new EC cell if no existing cell responds to an observed landmark position [3]. A newly recruited EC cell has a Gaussian activation function:

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

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

$$\sigma_2 = \sigma_0 (1 + 4\mu_2^2/R^2)$$

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.

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

It should also be noted that EC cell firing is also based on the landmark type, so that an EC unit firing signifies a landmark of a particular type at a particular relative position from the animat.

All-Or-None Connections Between EC and CA3 Layers

In the training phase, if none of the CA3 layer cells fire above a predetermined fraction of their peak firing level, a new CA3 layer cell is allocated. This newly created cell is then connected to the active EC layer cells. We have modified the connection weight assignment procedure of the existing model to reflect an all-or-none connection type. Rather than assigning weights proportionate to the activation of corresponding EC cells, we assign a weight of $1/nconn$ to each links, where $nconn$ is the total number of EC layer cells firing above their threshold levels. We set firing threshold of the newly allocated CA3 layer cell to 70% of its maximum possible weighted sum of incoming activations. During testing this threshold was reduced to 25% of the maximum possible activation level in order to allow animats to localize even in presence of partial sensory stimulus, or in other words, partial activity in the EC layer. Such a method has been found to be successful in modeling place-cell firing characteristics in simplified environments [13].

It has been observed that rodents give more importance to landmarks physically closer to their actual positions while localizing. Sharp and colleagues [19] performed experiments on rodents in a cylindrical environment with a single cue card. 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. 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.

Association of rewards with places

We have also extended the model to incorporate mechanisms that result in enhanced response of the EC layer neurons to landmark types that are closer to the reward locations. Whenever the animat receives a reward upon visiting a location, the maximum possible activations in EC layer cells are updated according to the following rule:

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

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:

```

Γ = 0
for i = 1 to n
do
  if wi < α Θ(i)
    Γ = Γ + wi
    wi = 0
  else
    Γ = Γ + α Θ(i)
    wi = wi - α Θ(i)
  endif
done

```

If multiple landmarks of same type are present, weights are altered by summing the distances of landmarks of similar types to the estimated goal location. The degenerate case of $n = 1$ is 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.

The above rule gives more preference to landmark types that are near the goal location by removing a uniform amount α from weights assigned to each of the landmark types, and redistributing it so that a landmark type gains weight if such a landmark 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.

We hypothesize that such a computation, which gives more weight to a particular *type* of sensory stimulus, takes place in the EC-Dg layers, as these layers get sensory information from the cortical areas as well as feedback connections from the Subiculum. The Subiculum is strongly believed to be part of the path-integration system [17]. Assuming a population code in the subiculum, it is conceivable that units encode path integration information which is reset using subsets of landmarks. This information can be used to supply a modulatory feedback signal to the units in EC and Dg.

Simulation Results

All simulation parameters and methods were identical to those in [2]. 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 on the experiment being performed. Animats then explored their environments and allocated cells corresponding to different locations in the environment. 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 [2]. 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 [2]. 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 in one of the two clusters 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 lo-




Figure 1: Left: Trajectories taken by an animat during test trials. Right: Superimposed place field firing regions during test trials.




Figure 2: 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.

calized using the landmark on the right. Interestingly enough, the right cluster is spread over a larger area, signifying that the place-cell in question fired over a larger area of the environment when animat localized using the right landmark. This effect can be explained by conflicting CA3 layer cell firing pattern due to the incorrect binding of CA1 unit activity with the path integration system for some of the activated CA1 units, resulting in a greater path integration estimate variance which in turn causes the CA1 unit in question pass the Mahalanobis distance test over a larger area of the environment. Figure shows the trajectories during test trials, when one of the landmarks was distinct from the rest during training. In some of the trials animats were unable to localize because of lack of training in those regions of 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 extension 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 [2], namely, the experiments where an array of three landmarks with different types of landmarks was transformed, in figure 9 c of [6] as seen in Figure . The simulation parameters used here were identical to those in [2].

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.

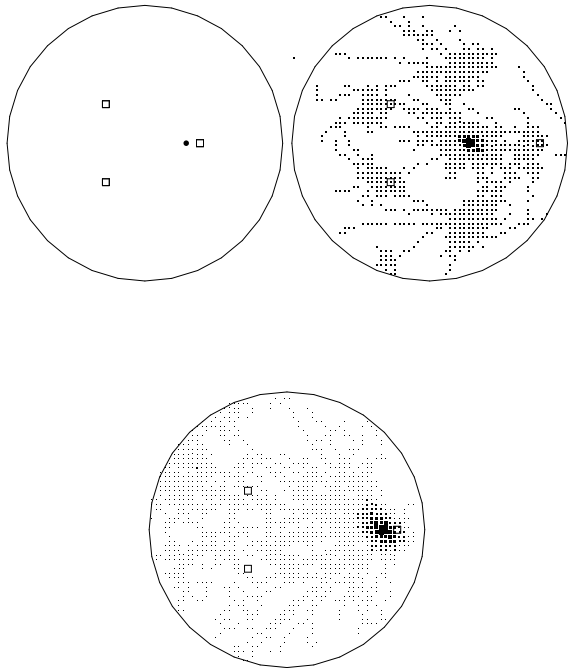


Figure 3: Top Left: Training Environment; Top Right: Normalized test histograms averaged over five animats with ten test trials each, when landmarks were indistinguishable from each other; Bottom: Right-most landmark distinguishable from the rest

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 local-

ized 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 prominence given by the animat to the unique landmark was not very pronounced.

Related Work

Burgess and colleagues have implemented a robotic simulation of rat navigation, which effectively reproduces place-cell firing characteristics [4]. 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 path integration systems.

O'Keefe and Burgess have been successful in replicating place fields in a simple rectangular environment [13]. However, in their model, the tuning widths as well as receptive fields of 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 [3, 2]. 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 [19].

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 one of the representations is retrieved when the animal first enters a familiar environment. Although several hypotheses have been put forth in the literature [9, 16], they remain to be verified through concrete realizations in terms of computational models that explain the relevant neurobiological as well as behavioral data.

A study of the EC-Dg layer neurons could shed more light on the merging of subsets of available multimodal sensory input streams which are then conveyed to hippocampus. Such a study could help understand how multiple traces of similar (and possibly related) events can be encoded into the hippocampal formation and then eventually consolidated and stored into other brain regions which function as long term memory storage devices. Recent theory suggested by Nadel and Moscovitch[11] justifies the need for such a research effort.

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 [11, 17] need to be further investigated. In order for hippocampus to perform as a temporary storage or an index of events, the encoding of information in at least the EC area must be similar to the ones used by other cortical regions. Further, there needs to be a switching mechanism that channels the memory traces to appropriate cortical regions based on modality of stimulus involved in the event replay during sleep. These suggest some interesting directions of research that are worth pursuing.

References

- [1]D. Amaral and M. Witter. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31(3):571–591, 1989.
- [2]K. Balakrishnan, R. Bhatt, and V. Honavar. A computational model of rodent spatial learning and some behavioral experiments. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society*, pages 102–107, Mahwah, New Jersey, 1998. Lawrence Erlbaum Associates.
- [3]K. Balakrishnan, O. Bousquet, and V. Honavar. Spatial learning and localization in animals: A computational model and its implications for mobile robots. *Adaptive Behavior Journal*, 1999. (to appear).
- [4]N. Burgess, J. Donnett, K. Jeffery, and J. O’Keefe. Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society London B*, 352, 1997.
- [5]P. Churchland and T. Sejnowski. *The Computational Brain*. MIT Press/A Bradford Book, Cambridge, MA, 1992.
- [6]T. Collett, B. Cartwright, and B. Smith. Landmark learning and visuo-spatial memories in gerbils. *Journal of Neurophysiology A*, 158:835–851, 1986.
- [7]R. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME*, 60, 1960.
- [8]D. Marr. Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London*, 176:23–81, 1971.
- [9]B. McNaughton, C. Barnes, J. Gerrard, K. Gothard, M. Jung, J. Knierim, H. Kudrimoti, Y. Qin, W. Skaggs, M. Suster, and K. Weaver. Deciphering the hippocampal polyglot: the hippocampus as a path-integration system. *The Journal of Experimental Biology*, 199(1):173–185, 1996.
- [10]R. Morris. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–261, 1981.
- [11]L. Nadel and M. Moscovitch. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, 7:217–227, 1997.
- [12]J. O’Keefe. Computations the hippocampus might perform. In L. Nadel, L. Cooper, P. Culicover, and R. Harnish, editors, *Neural Connections, Mental Computation*, pages 225–284. MIT Press/Bradford Book, Cambridge, MA, 1989.
- [13]J. O’Keefe and N. Burgess. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 1996.
- [14]J. O’Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Oxford:Clarendon Press, 1978.
- [15]M. Pikkarainen, S. Ronkko, V. Savander, R. Insausti, and A. Pitkanen. Projections from the lateral, basal and accessory basal nuclei of the amygdala to the hippocampal formation in rat. *The Journal of comparative neurology*, 403:229–260, 1999.
- [16]A. Redish and D. Touretzky. Separating hippocampal maps. In N. Burgess, K. Jeffery, and J. O’Keefe, editors, *Spatial Functions of the Hippocampal Formation and the Parietal Cortex*. Oxford University Press, 1997.
- [17]A. Redish and D. Touretzky. The role of the hippocampus in solving the morris water maze. *Neural Computation*, 10:73–111, 1998.
- [18]E. Rolls. Functions of the primate hippocampus in spatial processing and memory. In R. Kesner and D. Olton, editors, *Neurobiology of Comparative Cognition*, pages 339–362. Lawrence Erlbaum Associates, Hillsdale, NJ, 1990.
- [19]P. Sharp, J. Kubie, and R. Muller. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic properties. *Journal of Neuroscience*, 10:2339–2356, 1990.

- [20]H. Tanila, M. Shapiro, and H. Eichenbaum. Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus*, pages 613–623, 1997.
- [21]R. Traub and R. Miles. *Neuronal Networks of the Hippocampus*. Cambridge University Press, 1991.