Is the Hippocampus a Kalman Filter?

Olivier Bousquet*, Karthik Balakrishnan† and Vasant Honavar‡
Artificial Intelligence Research Group, Dept. of Computer Science
Iowa State University, Ames, IA - 50011, USA
{bousquet,balakris,honavar}@cs.iastate.edu

Abstract

Growing evidence suggests that the hippocampal formation functions as a relational (or episodic) memory, and on tasks of a spatial nature, it serves as a place learning and recognition system. In this paper we develop a computational characterization of the hippocampal formation pertaining to spatial learning and localization, and show that the information processing posited for this structure closely matches an engineering tool for probabilistic spatial localization of robots: the Kalman filter. Based on this parallel, we derive expressions for information integration, propagation of covariances, and near-optimal update rules for the localization performed by our computational model.

1 Introduction

It has long been known that damage or lesions of the hippocampus (HC) and adjacent areas result in severe learning deficits, including the ability to learn new stimuli, recognition memory, object-place task memory, etc. [1, 2]. Hippocampal lesions have also been found to produce severe deficits in the ability of rodents to traverse complex mazes (cf. appendix of [3]). This lesion data, along with neuro-anatomical and cellular studies, has led to the hypothesis that the hippocampal formation (HF) functions as a *relational* or *episodic* memory by forming associations between spatio-temporally contiguous events or sensory inputs. Given this general understanding of hippocampal function, on tasks of a spatial nature, the HC learns *spatial scenes* thereby allowing the animal to recognize particular *places* and *localize* accordingly. In this paper, we will only consider the spatial learning aspect of hippocampal function, although the principles presented here may well apply to hippocampal function in general.

*Olivier Bousquet is a student at Ecole Polytechnique, France which supported his visit to Iowa State University.

†All correspondence concerning this paper should be directed to the second author.

‡This research was partially supported by the National Science Foundation through grant NSF IRI-9409580 to Vasant Honavar

A number of computational models of hippocampal spatial learning have been proposed in the literature (cf. [4, 5] for excellent surveys). Unfortunately, most of these models have been specified in terms of informal verbal descriptions. Thus, their implementation in neurobiological and cognitive simulations or mobile robots is far from straightforward. Some recent models have attempted to provide computational characterizations of the information processing and update mechanisms, which, though implementable, have been largely *ad-hoc*. This is probably due to the lack of precise information concerning actual information processing in the HC.

In this paper we propose a computational characterization based on the thesis that the hippocampal formation learns places in the environment and associates them with position estimates derived from dead-reckoning[1]. By analyzing the information transformation in this model, we relate it to a tool used for probabilistic localization of robots: the *Kalman filter*. Based on this parallel, we then derive expressions for *near-optimal* information fusion and updates in the model and present some experimental results.

2 Anatomy and physiology of the hippocampus

At the anatomical level the HF is one of the highest levels of association in the brain, receiving highly processed sensory inputs from higher-order associational areas of the cerebral cortex [2], which converge primarily in the *entorhinal cortex (EC)*. The HF itself consists of the *dentate gyrus (Dg)*, and regions *CA3* and *CA1* of *Ammon's horn*, as shown in Figure 1.

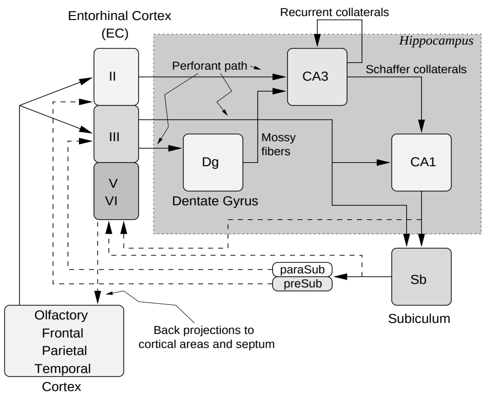

Figure 1: Anatomy of the hippocampal formation.

The Dg contains *granule cells* which respond to sensory information obtained from the EC through the *perforant path* and are known to be a site for *long term potentiation (LTP)*. The CA3 region consists of *pyramidal* (or complex-spike) cells, along with inhibitory interneurons like the *basket cells, chandelier cells, mossy cells*, etc. These cells receive input from the EC via the perforant path, the Dg via the *mossy fibers*, and from other CA3 cells via the *recurrent collaterals*. Both the mossy fiber synapses as well as the recurrent collateral synapses have shown evidence of Hebb-like modifiability. The CA1 region also contains pyramidal cells

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

and interneurons, however, unlike the CA3 region, CA1 cells do not have recurrent projections to other CA1 cells. The CA1 pyramidals receive input from the CA3 cells via the *Schaffer collaterals* and directly from the EC via the perforant path. Axons from the CA1 project via the *alveus* to the *subiculum (Sb)* and also back to the EC. Sb also receives input from the EC and projects to the *pre-* and *para-subiculum*, the deep layers of the EC, and to the hypothalamus, septum, anterior thalamus and cingulate cortex. Back-projections also exist from the EC to the cortical areas and the septum [2, 6].

At the cellular level, the HF of the rat contains *spatially responsive* cells, primarily *place cells* that fire maximally when the rat is in a particular portion of its environment. Since their discovery in 1971 [3], several key properties of place cells and their firing fields have been identified through systematic experiments. For instance, it is now known that place fields are learned quickly when the animal is introduced into a new environment. Further, the size and shape of place fields appear to be controlled by visual cues and change when the configuration of visual cues are manipulated in a familiar environment [7]. Researchers have also found that place fields resemble *radial-basis* functions. Another key property of place cells appears to be their *non-directionality* of firing in open environments, while in directionally-constricted environments like mazes their firing is direction-specific. Probably the most important property is the reliance of place fields on the motor system of the animal [7]. Finally, after visual initialization in a familiar environment, place cells continue to fire even in darkness, implying the possible role of dead-reckoning in the maintenance of place cell firing [8, 7].

Head-direction cells have also been found in several regions of the rat's brain, primarily the posterior parietal cortex, retrosplenial cortex, dorsal presubiculum, and the anterior thalamus [9]. These cells fire selectively based on the direction of the animal's head irrespective of its position in the environment. The firing of these cells appears to be dynamically alterable through complex interactions between the visual and angular motion signals, and it is strongly believed that rats update their head-direction by dead-reckoning [7, 10].

3 A hippocampal model of spatial learning

Based on the data presented above and results from behavioral experiments with rodents, it has been hypothesized that the HC functions as a *spatial learning* system [8, 7, 4]. Although the precise nature of hippocampal spatial learning remains unknown, it is our belief that the HC learns an internal map of *places* and associates these places with both motion information and dead-reckoning estimates, thereby maintaining *topological* as well as *metric* representations of the environment. In what follows, we present a computational characterization of hippocampal function from the point of view of spatial learning, and ascribe specific roles to the different hippocampal layers. Our model is based on the following:

1. The EC layer contains cells that respond maximally to specific landmarks appearing at particular positions relative to the animal irrespective of the environment that the animal is in [11]. Thus, in our model EC cells act as *spatial filters* by matching recognized landmarks (for instance from the neocortex) and their spatial positions (possibly from the parietal cortex), with stored information.

2. The granule cells of the Dg are primarily place cells and they appear to associate multiple spatial landmarks to form *spatial scenes*. Further, the mossy fiber inputs from Dg to CA3 cells is much *stronger* than the perforant path synapses between EC and CA3. Based on these facts, it has been suggested that the Dg serves to provide the *context* or *reference frame* [7], possibly through *non-redundant* and *orthogonal* coding of the EC input [6]. We use this interpretation of Dg in our model.

3. The CA3 cells produce a *place code* by forming a *weighted conjunction* of landmarks sensed from that place and reflected in the activations of the EC and Dg layers. The CA3 recurrent collaterals have led researchers to suspect an *autoassociation* (thus a pattern completion) circuit [6], however, recent results indicate that the CA3 system might be functioning as a *predictor* by associating places with *motion information* [12]. We believe that the CA3 forms associations of the form: $P_1 \times M_1 \rightarrow P_2$, where the P's represent place codes (hence places) and M_1 represents the animal motion that takes it from one place to the next, thereby learning a *topological map* of the animal's environment.

4. The CA1 layer too, contains place cells. However, since place cells continue to fire in the dark and are silent when the motor system is restrained, there appears to be motor system involvement in place cell firing. Further, since animals (particularly rodents) commonly use dead-reckoning for goal-directed spatial navigation [13], it appears that place cell firing is linked to the animal's dead-reckoning mechanism. Based on this intuition we suggest that the CA1 cells associate the place code produced in the CA3 layer with the current position estimate of the animal derived from its path-integration system. Thus, in our model, CA1 captures a *metric map* of the environment.

5. The Sb contains, among other cells, head-direction cells. It has been suggested that the thalamic nucleus integrates angular velocity signals, thereby providing a dead-reckoned estimate of head-direction to the Sb via the direct input from the EC [10]. We suggest that the head-direction cells in Sb associate place codes in CA1 with the direction-reference derived from the head-direction dead-reckoning system. Thus, Sb captures a *directional map* of the environment.

In our computational model shown in Figure 2, new sensory-input driven place cells are recruited in the CA3 layer as the animal explores its environment. These cells are then associated with the position estimate from the path-integration system through Hebbian-like learning in the CA1 layer. Further, motion information derived from a motor efference copy or ideothetic means is also used to associate the previous place code with the new place code in the CA3 recurrents. If necessary (for instance in constricted environments like maze-arms), the place code of CA1 is also associated with directional information from the head-direction path-integration system, making the corresponding place cells directional.

The hippocampal system performs *spatial localization* by *matching* the *predicted* position of the animal (from the dead-reckoning system) with the *observed* position of the place field center (dead-reckoning estimate stored with the activated place cell). Based on this match, the position estimate, head-direction estimate, as well as the place field centers are updated as shown in Figure 3.

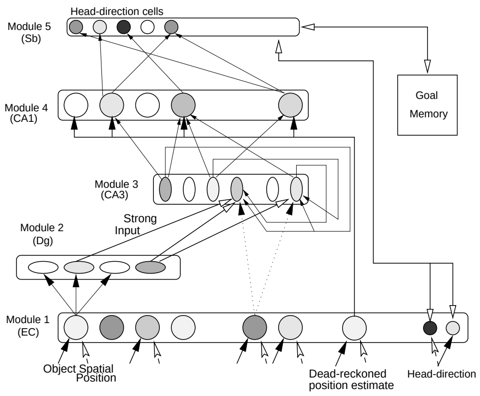

Figure 2: Computational model of the hippocampus

Upon reintroduction into a familiar environment, the animal distrusts its path-integration estimate and also possibly its head-direction estimate (if it has been intentionally disoriented). Visual (or sensory) inputs activate corresponding place codes which can then be used to initialize (or correct) the animal's position (from CA1) and head-direction (from Sb) estimates.

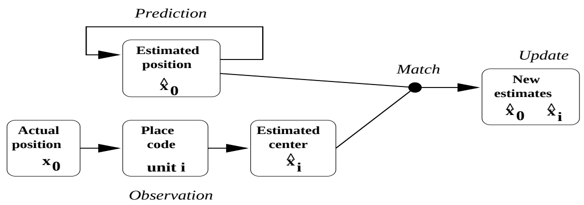

Figure 3: A schematic of hippocampal localization.

In a familiar environment, the animal places high confidence in its position estimate. Under such circumstances, even if some landmarks are removed, the high confidence in path-integration overrides the low sensory activation allowing the animal to still identify the correct place. We believe that synaptic changes at the EC-CA3 as well as EC-Dg synapses soon align the place code with the new (and deprived) sensory input. In our model, place fields are quickly formed through exploration by the animal. The place fields are driven both by sensory as well as path-integration inputs, thus they can be manipulated by changes in the landmark configuration and yet continue to fire in darkness. However, in darkness, without the appropriate updates of Figure 3, the position estimates accumulate drifts and the place cell firing deteriorates.

4 Probabilistic localization

In our model of the HC as a spatial localization system, we have suggested that the HC integrates information from two streams: the sensory inputs and the path-integration system. It should be noted that neither of these two information streams is perfect: sensory systems of animals accommodate considerable errors (for e.g., in the estimation of distance and direction to visible objects, recognition of objects, etc.) and path-integration is also prone to estimation errors and drifts. Thus, in order for the HC to perform robust spatial localization using these uncertain information sources, it must, in fact, be capable of *probabilistic localization.*

As with animals, mobile robots too need mechanisms for spatial learning and localization, and much like animals, robot sensors and actuators are also noisy. This has led to many probabilistic localization approaches for mobile robots. One such localization tool is the *Kalman filter* (KF) [14], which maintains a *stochastic spatial map*, propagates sensory and motion uncertainties, and under certain conditions, provides *optimal* update rules for combining uncertain information [15].

The KF technique for robot localization maintains a stochastic map of the robot's environment at each discrete time-step k (called the state vector $\mathbf{x}_k$), which includes an estimate of the robot's current position and possibly the estimated positions of other landmarks in the robot's environment. It is assumed that the *system model* denoting the evolution of the state based on robot motion, is specified:

$$\mathbf{x}_k = \mathbf{\Phi}_{k-1}\mathbf{x}_{k-1} + \mathbf{u}_{k-1} + \mathbf{w}_{k-1} \tag{1}$$

where $\mathbf{u}_{k-1}$ is the *movement command* and $\mathbf{w}_{k-1}$ is the motion error with variance matrix $\mathbf{Q}_{k-1}$. Also, a *measurement model* is assumed to be given, which denotes the measurements or observations the robot would make when in a given state $\mathbf{x}_k$:

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{v}_k \tag{2}$$

where $\mathbf{v}_k$ is the measurement noise with variance matrix $\mathbf{R}_{k-1}$.

Given these two models, the KF stores and updates an *estimate* of the current state $\hat{\mathbf{x}}_k$ and its associated covariance matrix $\mathbf{P}_k = E\{(\mathbf{x}_k - \hat{\mathbf{x}}_k)(\mathbf{x}_k - \hat{\mathbf{x}}_k)^T\}$, by making *predictions* and combining them with *observations*. Suppose the current state estimate is $\hat{\mathbf{x}}^+_{k-1}$ with the covariance matrix $\mathbf{P}^+_{k-1}$. Based on robot motion, the KF predicts the new state of the system $\hat{\mathbf{x}}^-_k$, using equation 1, and the covariance matrix $\mathbf{P}^-_k = \Phi_{k-1}\mathbf{P}^+_{k-1}\Phi^T_{k-1} + \mathbf{Q}_{k-1}$. Based on this state prediction and using the sensory model $\mathbf{H}$, the system predicts the measurement or observation $\mathbf{H}_k\hat{\mathbf{x}}^-_k$ using equation 2. This is the sensory input the robot is predicted to observe at its predicted position. Based on actual measurement $\mathbf{z}_k$, the KF then allows the state estimate and covariance matrix to be updated as follows (refer to [14] for details of the derivation):

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + \mathbf{K}_k(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_k^-) \quad (3)$$

$$\mathbf{P}_k^+ = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_k^- \quad (4)$$

where $\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}_k^T(\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \mathbf{R}_k)^{-1}$ is the *Kalman gain* and $\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_k^-$ is the *innovation*.

It can be shown that the KF updates are *optimal* (minimum variance, maximum likelihood, etc.) if the system and measurement models are linear, and $\mathbf{v}$ and $\mathbf{w}$ are assumed to be uncorrelated, zero-mean, white noise. (Note: uncorrelated means $E(\mathbf{w}_i\mathbf{v}_j) = 0$ for all i, j; zero-mean implies $E(\mathbf{v}_i) = 0$ and $E(\mathbf{w}_j) = 0$; and whiteness refers to $E(\mathbf{w}_k\mathbf{w}_l^T) = \delta_{kl}\mathbf{Q}_k$ and $E(\mathbf{v}_k\mathbf{v}_l^T) = \delta_{kl}\mathbf{R}_k$, where δ_{kl} is the *Kronecker delta function* ($\delta_{kl} = 1$ iff k=l, and 0 otherwise) and $\mathbf{Q}_k$ is the covariance matrix of the random noise $\mathbf{w}_k$ and $\mathbf{R}_k$ is the covariance matrix of the random noise $\mathbf{v}_k$.) This process can be described by the schematic shown in Figure 4.

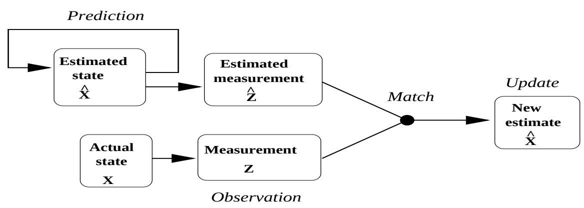

Figure 4: A schematic of Kalman filtering.

5 Hippocampal Kalman filtering

As can be observed from Figures 3 and 4, our computational model of hippocampal function and KF both share 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. Given this close parallel between our suggested hippocampal function and KFs, and considering the fact that we do not yet have a full understanding of the actual processes of hippocampal information fusion and update, it is interesting to explore whether hippocampal function could be described in terms of KF theory. If so, one can apply KF theory to characterize information processing in the HC and use it in practical realizations.

In considering hippocampal function as Kalman filtering, we first have to address the issue of *state vector representation*. In robot localization, the state vector usually contains the position of the robot and the positions of landmarks in the robot's environment. Since spatial localization in the HC appears to be based on *place recognition*, we have to use a state vector that contains the locations of these places. As the places are labeled with dead-reckoning estimates in our model, we consider the state vector to be composed of the estimated *centers* of places encountered by the animal and represented in the CA3. Remember that this notion of center is probabilistic, i.e., it is the expected position of the animal given that the corresponding place has been recognized. Mathematically, $x_i = E(x_0|i)$, where x_0 is the position of the animal and x_i is the center of place i. To keep the discussion simple we will henceforth assume that place codes in CA3 are represented by single units, although in reality ensembles of units are known to code for place.

5.1 The system and measurement models

For Kalman filtering, we also need to specify the system model which captures the change of the state vector with the animal's motion. In robots, the motion commands are known and hence the expected position of the robot can be easily predicted (it is simply the current position of the robot plus the motion command vector). In animal motion however, the exact motion command is rarely known (if at all), and the post-motion position of the animal has to be computed through the process of *active dead-reckoning*. Thus, we assume that the apparent motion of the animal was the intended one and the predicted motion is provided by dead-reckoning, leading to the following modified system models:

$$\begin{aligned} \mathbf{x}_k &= \mathbf{\Phi}_{k-1}\mathbf{x}_{k-1} + \mathbf{u}_{k-1} & \text{/* Animal motion */} \\ \hat{\mathbf{x}}_k^- &= \mathbf{\Phi}_{k-1}\hat{\mathbf{x}}_{k-1}^+ + \mathbf{u}_{k-1} + \mathbf{w}_{k-1} & \text{/* Dead-reckoning estimate */} \end{aligned}$$

where the $\mathbf{w}_{k-1}$ is a white noise (on dead-reckoning rather than on the actual motion) with covariance matrix $\mathbf{Q}_{k-1}$.

Since our hippocampal model recognizes places and associates places with position estimates, we would like to choose a measurement model that captures the spatial relationship between the estimated position of the animal and the center of the current place that it is in. For instance, we could choose the measurement model to be: $\mathbf{z}_k = x_{0,k} - x_{i_k}$ which is simply the vector from the position of the animal to the center of place i_k where the animal is at time step k. However, this quantity is not measurable (or observable) since we do not know the exact centers of the places or the exact position of the animal. This problem can be circumvented by specifying a measurement function that always observes $\mathbf{z}_k = x_{0,k} - x_{i_k} + \mathbf{v}_k = 0$. This measurement function constrains the form of the random error to $\mathbf{v}_k = x_{i_k} - x_{0,k}$. Although it can be shown that $\mathbf{v}_k$ still has zero-mean, it turns out that it is *autocorrelated* and hence is not a white sequence. Since this autocorrelation is difficult to measure, we choose to ignore it (in our simulations we found this autocorrelation to be negligible).

5.2 Update equations

Based on the above discussion, our state vector is given by:

$$\mathbf{x}_k = [x_{0,k}, x_1, \ldots, x_n]^T \tag{5}$$

where $x_{0,k}$ denotes the position of the animal, x_i denotes the center of place field i, and n is the number of places visited by the animal. 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.

It should be noted that the place field centers do not change with animal motion; only the position estimate of the animal changes. Thus, using $\mathbf{\Phi}_k = \mathbf{I}$, $\mathbf{u}_k = [u_k, 0, \ldots, 0]^T$, and $\mathbf{w}_k = [w_k, 0, \ldots, 0]^T$, the system model and the dead-reckoning estimates become:

$$\begin{aligned} x_{0,k} &= x_{0,k-1} + u_{k-1} & (6) \\ \hat{x}^-_{0,k} &= \hat{x}^+_{0,k-1} + u_{k-1} + w_{k-1} & (7) \end{aligned}$$

Similarly, using $\mathbf{H}_k\mathbf{x}_k = x_{0,k} - x_{i_k}$ and $\mathbf{v}_k = v_k$, we obtain the measurement model and predicted measurements as:

$$\begin{aligned} z_k &= x_{0,k} - x_{i_k} + v_k = 0 & (8) \\ \hat{z}_k &= \hat{x}^-_{0,k} - \hat{x}^-_{i_k} & (9) \end{aligned}$$

A minimum-variance derivation akin to KF leads to the following update rules:

$$\begin{aligned} \hat{\mathbf{x}}^+_k &= \hat{\mathbf{x}}^-_k + \mathbf{K}_k(\mathbf{H}_k\hat{\mathbf{x}}^-_k) & (10) \\ \mathbf{P}^+_k &= (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}^-_k & (11) \end{aligned}$$

with Kalman gain $\mathbf{K}_k = \mathbf{P}^-_k\mathbf{H}^T_k(\mathbf{H}_k\mathbf{P}^-_k\mathbf{H}^T_k + \mathbf{R}_k)^{-1}$.

5.3 Distinguishing perceptually similar places

Often, different locations in the environment produce the same sensory input (perceptual aliasing) and we need mechanisms to handle such cases. Fortunately, we can extend our model to elegantly make such distinctions based on the *Mahalanobis distance* [15]. Mahalanobis distance computes the difference between predicted and observed values and normalizes them by their covariance. This distance measure has a χ^2 distribution with q degrees of freedom where q is the rank of the covariance matrix. Since the errors in our system are Gaussian and the system equations are linear, we can use the Mahalanobis test to perform matches. Thus, given the current estimate of the animal's position $\hat{x}_{0,k}$, if the sensory inputs activate place code i_k with the estimated center $\hat{x}_{i_k}$ we perform the following test:

$$(\hat{x}_{0,k} - \hat{x}_{i_k} - 0)^T(\mathbf{H}_k\mathbf{P}_k\mathbf{H}^T_k)^{-1}(\hat{x}_{0,k} - \hat{x}_{i_k} - 0) < \epsilon \qquad (12)$$

where $(\hat{x}_{0,k} - \hat{x}_{i_k} - 0)$ is the difference between predicted and observed values and $\mathbf{H}_k\mathbf{P}_k\mathbf{H}^T_k$ is the covariance matrix (ϵ can be chosen such that, say, 95% of the correct matches pass the above test).

If equation 12 is true, then the place has indeed been visited before. However, if this test fails, it implies that the animal is now at a *new* place although it sensorily resembles some other place visited earlier. In the latter case, we recruit a new place cell and include its parameters in the state vector. Thus, our system creates multiple units that repond to the same sensory input but are tuned to different centers that correspond to the peaks of the multimodal distribution $P(x|s)$ (where s is a sensory input).

5.4 Implementation

The model described above was implemented in simulation with positions being represented in two-dimensional Cartesian coordinates. The EC cells were set up to compute a two-dimensional radial-basis function of the identity of the landmarks and their *ego-centric* (with respect to the animal) spatial locations. New EC units were inducted if a sensed landmark did not evoke an above threshold firing in any of the existing EC cells. Place cell (PC) activations of CA3 units was computed through a linear combination of the EC activations. If none of the PC units fired above a threshold, a new PC unit was inducted and connected to the ten most active units in the EC layer. However, if some PC units fired above the threshold *and* passed the Mahalanobis test, then the strongest activated unit was taken to be the winner and the filter updates presented above were applied to it. Other details can be found in [16].

We show results on a simple task where the simulated animal follows a circular trajectory in a room with four identical landmarks. Figure 5 shows the place fields generated by the animal for this task (left) and the state of the system at the end of its first turn (right). As can be observed, without any revisits or updates, the uncertainty in the location of places (denoted by the 3-σ boundary) keeps increasing.

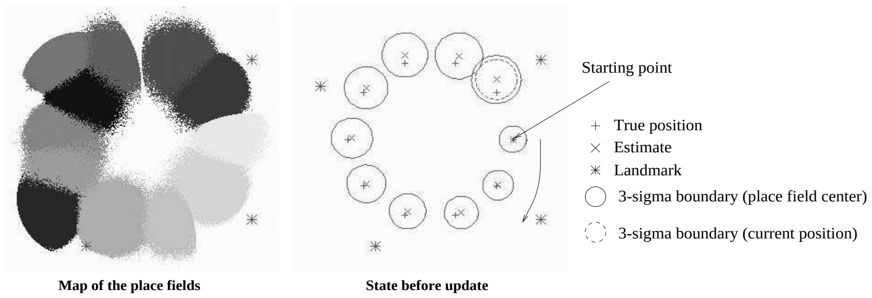

Figure 5: Place fields and the state of the system before updates.

Figure 6 shows how the hippocampal filtering process described earlier reduces the variances of the estimates when the animal revisits familiar places.

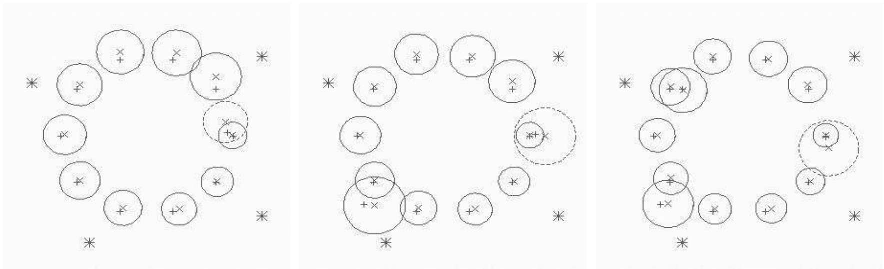

Figure 6: Decrease in the variance of state estimates with revisits.

6 Discussion

In this paper we have described a computational model of hippocampal function from the point of view of spatial learning. Although we have not rigorously characterized the computations of Dg and CA3 regions or their synapses, the nature of computations and associations in the CA1 region have been dealt with in more detail. The primary thesis is that the hippocampal system learns places and associates them with dead-reckoning estimates, thereby allowing the animal to localize itself in space, even in the presence of perceptual aliasing. Our model resembles the anatomy of the HC and is functionally consistent with results from behavioral experiments.

Although many models of place-based localization have been proposed [4, 5, 8, 7], they have failed to concretely specify the nature of information processing and the updates involved in the models. Some researchers have tried to adopt an information-theoretic approach for characterizing hippocampal function [17], but they have met with little success owing to the lack of information pertaining to actual information processing in the HC. In contrast, our approach draws the parallel between the overall nature of hippocampal function and probabilistic information fusion using a KF, thereby allowing us to use the KF framework for our computational model.

Finally, if the HC were indeed functioning as a KF, how are the computations (storage and update of place estimates and covariances) actually realized? Although we do not know the exact neural substrates of the proposed computations, some plausible suggestions exist. For instance, Buzsaki has suggested that CA3 cells in the rat hippocampus temporarily store information during exploratory behavior (by means of their recurrent collaterals), and during *consolidation*, i.e. when the rat rests, they discharge in *temporally-correlated* bursts thereby potentiating the CA1 cells [2]. It is interesting to note that the consolidation phase offers possibilities for the propagation (if not the computation) of state updates and covariances.

References

[1] N. Cohen and H. Eichenbaum. *Memory, Amnesia, and the Hippocampal System.* MIT Press/A Bradford Book, Cambridge, MA, 1993.

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

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

[4] N. Burgess, M. Recce, and J. O'Keefe. Hippocampus – spatial models. In M. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 468–472. MIT Press, Cambridge, MA, 1995.

[5] O. Trullier, S. Wiener, A. Berthoz, and J-A. Meyer. Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 1997. (to appear).

[6] E. T. Rolls. A theory of hippocampal function in memory. *Hippocampus*, 6:601–620, 1996.

[7] 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.

[8] H. Wan, D. Touretzky, and A. Redish. Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. Touretzky, J. Elman, and A. Weigend, editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19. Lawrence Earlbaum Associates, 1994.

[9] J. Taube. Head-direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15:70–86, 1995.

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

[11] G. Quirk, R. Muller, J. Kubie, and J. Ranck. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, 12(5):1945–1963, 1992.

[12] W. B. Levy. A sequence predicting ca3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, 6:579–590, 1996.

[13] C. Gallistel. *The Organization of Learning.* Bradford-MIT Press, Cambridge, MA, 1990.

[14] A. Gelb. *Applied Optimal Estimation.* MIT Press, 1974.

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

[16] K. Balakrishnan, O. Bousquet, and V. Honavar. Is the hippocampal formation a kalman filter?, 1997. (in preparation).

[17] W. Levy. A computational approach to hippocampal function. *The Psychology of Learning and Motivation*, 23:243–305, 1989.