

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

Based on a large body of neurophysiological, neuroanatomical, and behavioral data, it has been suggested that the hippocampal formation serves as a spatial learning and localization system. This spatial representation is *metric* in nature and arises as a result of associations between sensory inputs and dead-reckoning information generated by the animal. However, despite the fact that these two information streams provide uncertain information (e.g., recognition errors, dead-reckoning drifts, etc.), the hippocampal computational models suggested to date have not explicitly addressed information fusion from erroneous sources. In this paper we develop a computational model of hippocampal spatial learning and relate its functioning to a probabilistic tool used for uncertain sensory fusion in robots: the Kalman filter. This parallel allows us to derive statistically optimal update expressions for the localization performed by our computational model.

1 Introduction

A large body of neurophysiological, neuroanatomical, and behavioral data has pointed to the involvement of the *hippocampal formation* (HF) in *relational* (or *episodic, declarative*) memory^{1,2,3}. On tasks of a spatial nature, the hippocampus (HC) then learns *spatial scenes*, allowing the animal to recognize particular *places* and *localize* accordingly⁴. Based on a corpus of data from neuroscience and cognitive science, O’Keefe and Nadel suggested that animals (primarily rats) learn a *cognitive map* of their spatial environments and that the hippocampal formation serves as a site for these manipulations⁴. Such cognitive maps represent space in a *metric form* and are believed to arise as a result of associations between the sensory inputs and the self-generated estimates of position derived from the animal’s path-integration system^{a 4}. However, it should be noted that both these information streams provide uncertain and possibly erroneous information. For instance, sensory systems might include errors in the recognition of objects, estimation of distances to objects, etc., while path-integration is known to suffer from estimation errors and drifts. Thus, the animal must possess mechanisms to derive robust metric space representations by appropriately handling the uncertainties in these input streams.

^aPath-integration (or dead-reckoning) is the process of updating an estimate of one’s own position based on self-knowledge of direction, speed, and time of motion.

Although a number of computational models of hippocampal spatial learning have been proposed (see ⁵ for a survey), very few of them are based on a metric representation of space (called the *locale system* by O'Keefe and Nadel ⁴). Further, the models that address metric spatial representations, (primarily ^{6,7}), do not consider the input streams as noisy and consequently do not incorporate mathematically sound mechanisms for dealing with this uncertain information.

In this paper we develop a computational characterization of the hippocampal formation that learns places in the environment and associates them with position estimates derived from dead-reckoning. Since both these information streams are erroneous, we need mechanisms for fusing uncertain information. Kalman filtering (KF) has been used in robotics for fusing uncertain information, world modeling, and robot localization, and it is worthwhile to consider the use of this tool to characterize hippocampal function. This parallel is rather easily drawn, and the use of the KF framework allows us to derive expressions for *optimally* fusing information from the two streams. Simulation results showing the usefulness of this approach are presented and the plausibility of KF taking place in the hippocampus is discussed.

2 Anatomy and physiology of the hippocampus

The HF is one of the highest levels of association in the brain, receiving highly processed sensory inputs from the higher-order associational areas of the cerebral cortex ^{2,1}, 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*. 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 and interneurons, however, unlike the CA3 region, CA1 cells do not have recurrent projections to other CA1 cells. The CA1 pyramidal cells 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^{2,3}.

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 environment⁴. It is now known that place fields are learned quickly when the animal is introduced into a new environment. The size and shape of these fields appear to be controlled by visual cues and change when the configuration of visual cues are manipulated in a familiar environment⁸. 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 animal⁸, implying the possible role of dead-reckoning in place cell firing. This hypothesis derives further support from the observation of place cell firing in darkness provided the animal is allowed to first initialize its bearings in familiar environments^{6,8}. Cells with such location-specific firing have been found in the EC, Dg, CA3, CA1, as well as the Sb⁹.

Head-direction cells have also been found in several regions of the rat's brain. 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^{8,9}. Head-direction cells have been found in the posterior parietal cortex, retrosplenial cortex, dorsal presubiculum, and the anterior thalamus^{10,11}.

3 A hippocampal model of spatial learning

Based on the data presented above and results from behavioral experiments with rodents, it has been suggested that the HC functions as a *spatial localization* system by learning an internal map of *places* and associating these places with *position estimates* from dead-reckoning, leading to a metric representation of the environment (cognitive map or the locale system)⁴. Our model is based on this thesis and is shown in Figure 1.

Since the EC layer has been found to contain cells that respond maximally to specific landmarks appearing at particular positions relative to the animal (irrespective of the environment that the animal is in)¹², our model contains EC cells with similar properties. These cells function as *spatial filters* by matching recognized landmarks (for instance from the neocortex) and their spatial positions (possibly from the parietal cortex), with stored information

in module 1. Module 2, which corresponds to the Dg, contains place cells that associate multiple spatial landmarks to form *spatial scenes*. It has been suggested that the Dg serves to provide the *context* or *reference frame*⁸, possibly through *non-redundant* and *orthogonal* coding of the EC input³. However, our current implementation does not consider the computational role of Dg. Module 3 corresponds to the CA3 layer, with *place codes* being formed by a *weighted conjunction* of landmarks sensed from that place and reflected in the activations of the EC layer. Further, the CA3 recurrent collaterals are assumed to associate place codes with motion information and thus *predict*^{13,14} places based on current animal motion.

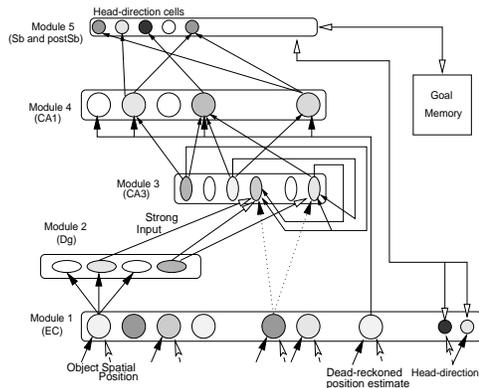


Figure 1: Computational model of hippocampus

centers of place fields in terms of *metric* coordinates from dead-reckoning. Since multiple places in the animal’s environment can produce similar sensory inputs (referred to as *perceptual aliasing* in robotics), we also suggest that the CA1 uses dead-reckoning information to *distinguish* between such perceptually similar places. Place cells have also been discovered in the Sb, however, we will not ascribe any roles to them in this paper.

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 EC⁹. Following this thesis, we assume that the head-direction cells in paraSb associate place codes in CA1 with the direction-reference derived from the head-direction dead-reckoning system, capturing a *directional map* of the environment. Finally, goals are assumed to be specified in terms of location information (in the dead-reckoning coordinate frame) and

Since the CA1 place cells fire in the dark and are silent when the motor system is restrained, it has been suggested that the motor system is involved in place cell firing⁸. Further, since animals (particularly rodents) commonly use dead-reckoning for goal-directed spatial navigation¹⁵, it is suspected that place cell firing is linked to the animal’s dead-reckoning mechanism. Following these ideas, the CA1 layer in our model (module 4) associates the place codes produced in the CA3 layer with position estimates derived from dead-reckoning, thereby learning the

are stored outside the HC.

Thus, EC cells are recruited if the exploring animal (or simulated animal) encounters a landmark not previously observed in that position (relative to the animal). Similarly, new place cells in CA3 are created when the animal visits a perceptually new place. CA1 cells are also recruited and associated with place codes of CA3 and position estimates from the dead-reckoning system^b. These constitute the *spatial learning* component of our model.

In addition, the animal recognizes places visited earlier and *localizes*, i.e., it updates the system parameters to better coincide with its observations. Since we assume that the estimate stored with a CA1 place code represents the position of the place field center, our model performs a *match* between the *predicted* place field center (from the dead-reckoning system) and the *observed* place field center (the estimate stored with the activated place code).

Based on this match the position estimate of the animal as well as the place field centers are updated as shown in Figure 2. If the animal is reintroduced into a familiar environment, it engages in *self-localization*⁷. Visual or sensory inputs activate place codes in CA3 and consequently in CA1, which are then used to initialize the position estimate of the animal (from the corresponding stored values in the place codes). In this system, place fields are formed quickly through exploration by the animal. Since these fields are driven by sensory as well as path-integration inputs, they can be manipulated by changes in the landmark configuration and yet continue to fire in darkness (CA1 cell firing is maintained by dead-reckoning while CA3 cells fire owing to the motion information in the recurrent collaterals).

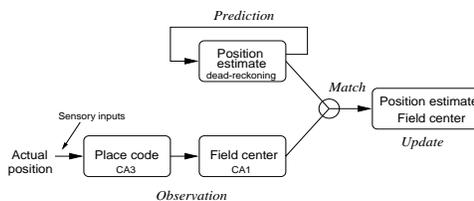


Figure 2: Hippocampal localization.

4 Probabilistic localization

As with animals, mobile robots also need mechanisms for spatial learning and localization, and much like animals, robot sensors and actuators are also noisy. This has led to many *probabilistic information fusion and localization* approaches for mobile robots. One such localization tool is the *Kalman filter*

^bWe have not implemented the Dg reference frames, paraSb head-direction cells, and topological learning in the CA3 recurrent collaterals

¹⁶, which maintains a *stochastic spatial map*, propagates and integrates sensory and motion uncertainties, and under certain conditions, provides *optimal* update rules for combining uncertain information¹⁷.

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:

$$\mathbf{x}_k = \Phi_{k-1}\mathbf{x}_{k-1} + \mathbf{u}_{k-1} + \mathbf{w}_{k-1} \quad (1)$$

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 :

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{v}_k \quad (2)$$

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

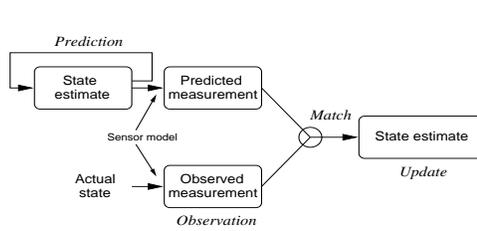


Figure 3: Kalman filtering.

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 Eq. 1, and the covariance matrix $\mathbf{P}_k^- = \Phi_{k-1}\mathbf{P}_{k-1}^+\Phi_{k-1}^T + \mathbf{Q}_{k-1}$. 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 Eq. 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 to¹⁶ 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. This process can be described by the schematic shown in Figure 3.

5 Hippocampal Kalman filtering

As can be observed from Figures 2 and 3, 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. Since the information streams leading to the formation of place codes in CA1 have been suggested to contain uncertainty, we are led to notice the close parallel between hippocampal function and KF. 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 characterized in terms of KF theory.

To do so, we first have to address the issue of *state vector representation*. As spatial localization in the HC appears to be based on *place recognition* and we have suggested that the *centers* of the places are learned in terms of dead-reckoning estimates, we define the state vector to be composed of the estimated centers of places encountered by the animal and represented in CA1. (To keep the discussion simple we will henceforth assume that place codes in CA3 and CA1 are represented by single units, although in reality, ensembles of units are known to code for place.) Thus, our state vector is given by:

$$\mathbf{x}_k = [x_{0,k}, x_1, \dots, x_n]^T \quad (5)$$

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.

5.1 The system and measurement models

We also need to specify the system model which captures the change of the state vector with the animal’s motion. In our model, we assume that the animal’s dead-reckoning system *predicts* the new position of the animal based on its previous position estimate and actual animal motion. Since dead-reckoning is error-prone, this estimate contains some error (in contrast to the system models used in KF where there is error in motion but the predictions are error-free).

Further, since the dead-reckoning system only predicts the changed position of the animal and not the place field centers, we can simplify the system models using $\Phi_k = \mathbf{I}$, $\mathbf{u}_k = [u_k, 0, \dots, 0]^T$, and $\mathbf{w}_k = [w_k, 0, \dots, 0]^T$, leading to:

$$\mathbf{x}_{0,k} = \mathbf{x}_{0,k-1} + \mathbf{u}_{k-1} \quad (6)$$

$$\hat{\mathbf{x}}_{0,k}^- = \hat{\mathbf{x}}_{0,k-1}^+ + \mathbf{u}_{k-1} + \mathbf{w}_{k-1} \quad (7)$$

where the w_{k-1} (with covariance matrix \mathbf{Q}_{k-1}) is the error on dead-reckoning rather than on the actual motion.

We also have 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: $\mathbf{z}_k = \mathbf{x}_{0,k} - \mathbf{x}_{i_k}$ 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. We circumvent this problem by stipulating that the measurement function always observes $\mathbf{z}_k = \mathbf{x}_{0,k} - \mathbf{x}_{i_k} + \mathbf{v}_k = 0$. This measurement function constrains the form of the random error to $\mathbf{v}_k = \mathbf{x}_{i_k} - \mathbf{x}_{0,k}$. Unfortunately, this error turns out to be *autocorrelated*, and hence is not a white sequence. However, if we assume that the animal either moves randomly or purposefully between place field centers, both the mean of the error and its autocorrelation can be shown to be almost zero (refer to¹⁸ for details of the proofs). Thus, using $\mathbf{H}_k \mathbf{x}_k = \mathbf{x}_{0,k} - \mathbf{x}_{i_k}$ and $\mathbf{v}_k = v_k$, we obtain the measurement model and predicted measurements as:

$$z_k = \mathbf{x}_{0,k} - \mathbf{x}_{i_k} + v_k = 0 \quad (8)$$

$$\hat{z}_k = \hat{\mathbf{x}}_{0,k}^- - \hat{\mathbf{x}}_{i_k}^- \quad (9)$$

Given these expressions for the system and measurement models, a minimum-variance derivation akin to KF leads to the following update rules:

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- - \mathbf{K}_k (\mathbf{H}_k \hat{\mathbf{x}}_k^-) \quad (10)$$

$$\mathbf{P}_k^+ = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k^- \quad (11)$$

with Kalman gain $\mathbf{K}_k = \mathbf{P}_k^- \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^T + \mathbf{R}_k)^{-1}$. For details regarding the exact form of these expressions, please refer to¹⁸.

5.2 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. In our

model, such perceptually similar places will activate the same place cell in the CA3 layer. We have suggested that the CA1 layer then distinguishes the places by using dead-reckoning information. Interestingly, we can extend our KF-based model to elegantly make such distinctions using the *Mahalanobis distance*¹⁷. Mahalanobis distance computes the difference between predicted and observed values and normalizes them by their covariance. Thus, given the current estimate of the animal’s position $\hat{x}_{0,k}$, if the sensory inputs activate a place cell in CA3 that is associated with a CA1 cell i_k with the estimated center \hat{x}_{i_k} , we perform the following test:

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

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

If the test in Eq. 12 is satisfied, then the place has indeed been visited before. However, if this test fails, we recruit a new place cell and include its parameters in the state vector. Thus, our system creates multiple units in the CA1 layer 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 some sensory input).

5.3 Implementation

The model described above was implemented 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 *allocentric* (with respect to the animal but independent of direction) spatial locations. The activations of CA3 units were computed through a linear combination of the EC activations. CA1 units were assumed to be connected to individual CA3 units and were associated with position estimates expressed in Cartesian coordinates. For more algorithmic and implementational details, the reader is referred to¹⁸.

Here we show the results on a simple task where the simulated animal follows a circular trajectory in a room with four identical landmarks. Figure 4 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 $3\text{-}\sigma$ boundary) keeps increasing.

Figure 5 shows how the hippocampal filtering process described earlier reduces the variances of the estimates when the animal revisits familiar places. Other results are presented in¹⁸.

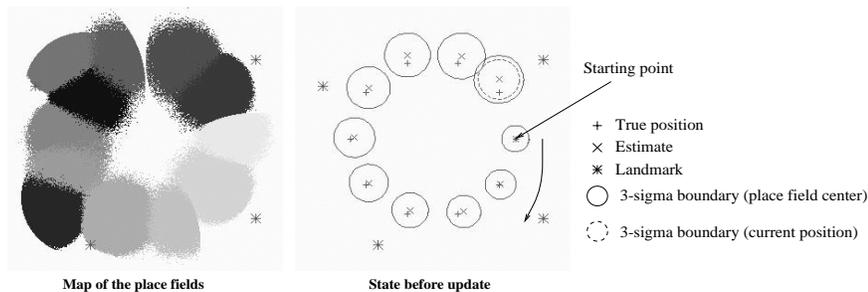


Figure 4: Place fields and the state of the system before updates.

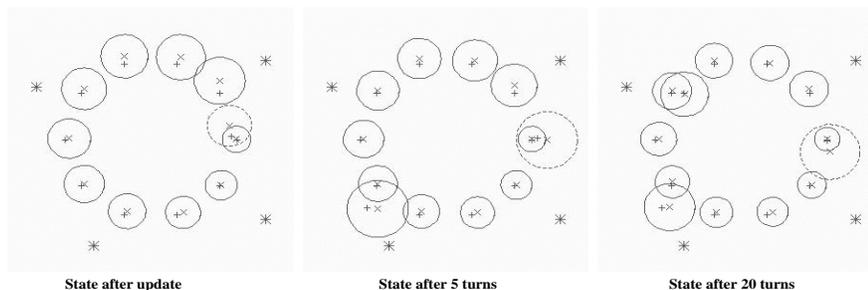


Figure 5: Decrease in the variance of state estimates with revisits.

6 Discussion

In this paper we have described a computational model of hippocampal spatial learning that acquires a metric map of the environment: places are represented in terms of their sensory features and their centers in terms of their metric positions in the dead-reckoning coordinate framework. Our model was inspired by the locale system of⁴ and is closely related to^{6,7}. However, it differs from the latter in a few significant ways. First, it attempts to map the computations to specific hippocampal regions. Second, it assumes that the two information streams carry uncertain information and provides a framework for the integration of uncertain information. Third, since the dead-reckoning information has uncertainty, it not only *learns* place field centers, but also *updates* them appropriately with subsequent visits by the animal. Finally, it uses dead-reckoning position estimates to distinguish between perceptually similar places.

Although Kalman filtering has been used in many robot applications, its use is restricted by the need for a *sensor model* in the measurement function.

Most current applications use the egocentric position of a landmark relative to the robot as the measurement, which leads to matching and localization problems if the sensor ranges are limited and the robot operates in an environment with many identical landmarks. Thus, the model developed in this paper can be seen as a *place-based* extension of KF for robots¹⁸.

Our model also leads to some testable predictions. For instance, we suggest that perceptually similar places in the *same* environment produce the same place code in the CA3 layer but different ones in the CA1 layer. Also, since the firing of CA1 and CA3 cells in darkness is maintained by different (although related) sources (dead-reckoning in CA1 and motion in CA3), they could drift in different ways and amounts. A behavioral prediction is that the animals will explore in slowly expanding trajectories from the starting point since this allows the animal to propagate reliable position estimates further and further away.

Although we do not know the exact neural substrates of the proposed KF computations, some plausible suggestions exist. O’Keefe has suggested mechanisms by which matrix inversions might possibly be performed by the CA1 layer (though an iterated update scheme)¹⁹. This reasoning applies to the matrix inversions required in KF. Also, Buzsaki has suggested that the 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 cells². It is possible that this consolidation phase provides a mechanism for the propagation (if not the computation) of the state updates and covariances. Much work has to be done to determine if HC indeed functions as a KF.

In conclusion, although we have developed a KF framework in the context of hippocampal spatial learning and localization, it is possible to use analogous mechanisms to describe the computations of other memory regions that behave in a similar predict-observe-match-update fashion and derive input from uncertain information streams.

Acknowledgements

This research was partially supported by the National Science Foundation through grant NSF IRI-9409580 to Vasant Honavar. Olivier Bousquet’s visit to Iowa State University during the summer of 1997 was funded by Ecole Polytechnique, France. Karthik Balakrishnan is supported by an IBM Cooperative Fellowship. The authors would like to thank the anonymous reviewers for their insightful comments and constructive suggestions.

References

1. N. Cohen and H. Eichenbaum. *Memory, Amnesia, and the Hippocampal System*. MIT Press, 1993.
2. P. Churchland and T. Sejnowski. *The Computational Brain*. MIT Press, 1992.
3. E. Rolls. In T. Ono, B. McNaughton, S. Molotchnikoff, E. Rolls, and H. Nishijo, editors, *Perception, Memory, and Emotion: Frontiers in Neuroscience*, pages 375–400. Pergamon, 1996.
4. J. O’Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, 1978.
5. O. Trullier, S. Wiener, A. Berthoz, and J-A. Meyer. *Progress in Neurobiology*, 1997. (to appear).
6. H. Wan, D. Touretzky, and A. Redish. In *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19. 1994.
7. D. Redish and D. Touretzky. In K. Ikeuchi and M. Veloso, editors, *Symbolic Visual Learning*. Oxford University Press, 1996.
8. B. McNaughton, C. Barnes, J. Gerrard, K. Gothard, M. Jung, J. Knierim, H. Kudrimoti, Y. Qin, W. Skaggs, M. Suster, and K. Weaver. *The Journal of Experimental Biology*, 199(1):173–185, 1996.
9. P. Sharp, H. Blair, and M. Brown. *Hippocampus*, 6:720–734, 1996.
10. L. Chen, L-H. Lin, C. Barnes, and B. McNaughton. *Experimental Brain Research*, 101:24–34, 1994.
11. J. Taube. *Journal of Neuroscience*, 15:70–86, 1995.
12. G. Quirk, R. Muller, J. Kubie, and J. Ranck. *Journal of Neuroscience*, 12(5):1945–1963, 1992.
13. B. McNaughton. In L. Nadel, A. Cooper, P. Culicover, and R. Harnish, editors, *Neural Connections, Mental Computation*, pages 285–350. MIT Press, 1989.
14. W. Levy. *Hippocampus*, 6:579–590, 1996.
15. C. Gallistel. *The Organization of Learning*. MIT Press, 1990.
16. A. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, 1970.
17. N. Ayache and O. Faugeras. *IEEE Transactions on Robotics and Automation*, 5(6):804–819, 1989.
18. K. Balakrishnan, O. Bousquet, and V. Honavar. Technical Report CS TR 97-20, Dept. of Computer Science, Iowa State University, 1997.
19. J. O’Keefe. In L. Nadel, L. Cooper, P. Culicover, and R. Harnish, editors, *Neural Connections, Mental Computation*, pages 225–284. MIT Press, 1989.