A comparison of context-dependent hippocampal place codes in 1-layer and 2-layer recurrent networks

Simona Doboli[a], Ali A. Minai[a,*], Phillip J. Best[b]

[a]Complex Adaptive Systems Laboratory, ECECS Department, University of Cincinnati,
Cincinnati, OH 45221-0030, USA
[b]Laboratory of Cognitive Psychobiology, Department of Psychology, Miami University,
Oxford, OH 45056, USA

Abstract

Recently, it has been suggested that attractor networks may provide a mechanism for context-dependence in hippocampal place codes. We have proposed that context may be coded by "latent attractors" — mutually competitive and internally cooperative cell groups which channel the system's response to afferent stimuli. We have also argued that it is the disynaptically recurrent dentate gyrus–hilus (DGH) system which embodies these latent attractors in the hippocampus. Others have suggested the CA3 network — with its monosynaptic recurrence — is the site. While latent attractors (and, thus, context-dependent coding) can be implemented by a 1-layer or 2-layer recurrent network, we show that a 2-layer recurrent network can implement a sparse context-dependent place-code using latent attractors much more flexibly than a 1-layer network. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Hippocampus; Place cells; Spatial representation; Context; Attractor networks

1. Introduction and background

Recently, there has been considerable interest in the issue of hippocampal CA3 spatial representations in different contexts or frames of reference, and it has been suggested that attractor networks may provide a mechanism for context-dependence

---

[15,11,5]. The model we have proposed is based on the notion of "latent attractors" — groups of cells which tend to activate each other and inhibit cells in other groups, but which do not all fire at the same time. Thus, activity across the system is channelled into one of the groups in a metastable fashion. Cells in this group then become the only ones that respond to afferent input, creating a context-dependent representation with the selected cell group identifying the context [11,3–5]. We have shown through simulation [4,5] that latent attractors require only a very simple type of recurrent connectivity. Samsonovich and McNaughton [15] have proposed a qualitatively similar but more sophisticated scheme based on the "charts" hypothesis. However, this scheme requires a more complex pattern of recurrent connectivity (Recently, Samsonovich (unpublished) has proposed a simpler model closer to ours.)

One major issue in all this is the architecture of the underlying network. Samsonovich and McNaughton [15] have tentatively located the charts system in the CA3 — a 1-layer system with monosynaptic excitatory recurrence. In contrast, we have suggested [11,3–5] that the attractors are coded in the 2-layer dentate gyrus–hilus system, using the disynaptic recurrent connectivity between granule and mossy cells. In this paper, we address the issue of whether a 2-layer system has any advantage over a 1-layer system in supporting the appropriate type of place representations.

2. Theoretical issues

We have previously argued on the basis of known experimental results and anatomical considerations why the dentate gyrus–hilus (DGH) system is a likelier locus for latent attractors or charts than the CA3 [11,3–5]. In this paper, we focus on a more abstract question that goes to rather general issues in neural computation.

Most theoretical models of hippocampal function (Marr, 1971; [10,18,2,17, 15,13,7–9] hypothesize that the CA3 system serves as an associative memory for patterns or pattern sequences. In the rodent, such patterns can be seen as representations of place [19]. Two key requirements can be listed for these codes:

1. *Sparseness*: Successful storage of patterns in a recurrent network requires sparse coding [10,14]. Indeed, creating sparse codes for CA3 has been posited as the primary function of the dentate gyrus [10,14,18].
2. *Spatial Consistency*: Place codes must vary smoothly in space such that similarity in code reflects spatial proximity. At the same time, codes for distinguishable locations must be distinct. These requirements are met if the place code is correlated with afferent sensory input, which varies smoothly in space.

We argue that meeting these requirements while simultaneously maintaining a latent attractor is much easier for a 2-layer network than a 1-layer system.

We assume that, at time, t, the system is subject to two stimuli: The afferent sensory stimulus, $s(t)$, and the recurrent stimulus, $r(t)$, which reflects the system's own previous state, $x(t-1)$ (in a discrete-time formulation). The actual activity of the system's active primary cells must code the afferent information, $s(t)$. However, given a certain

magnitude of $s(t)$, the recurrent signal, $r(t)$ must have the minimum strength needed to confine activity within the cells of the appropriate latent attractor. Assuming that the recurrent connectivity of the system is fixed, there are two ways to achieve this strength: (1) Via a significant number of moderately strong active recurrent inputs, or (2) Via a small number of very strong active recurrent inputs. If the recurrent activity pattern is constrained to be sparse, only option (2) is available. However, this means that the net recurrent bias on cells within the attractor at time $t + 1$ due to activity at time t has a high variance. This, in turn, makes the activity of cells at time $t + 1$ much more dependent on the precise activity pattern at t, and, correspondingly, less informative about the afferent stimulus.

The key problem is that sparse patterns acting through sparse connections are not a good substrate for generating uniform biasing fields, which is precisely what the latent attractor scheme needs. However, good coding requires sparseness, thus creating a conflict. A 2-layer system can resolve this conflict because it has two extra degrees of freedom: Two instead of one set of connections in the recurrent path, and two instead of one level of activity (one in each layer). Thus, the layer whose output constitutes the system response (DG in our model), can have sparse activity and a sparse downstream projection, while the other layer, which acts only to bias the first, can have higher activity and a more diffuse projection to the first layer. This separation of the information carrying and biasing functions into two layers creates a much more flexible system than one which must operate within a narrow critical range of activity and connectivity needed to satisfy the requirements of both functions.

Interestingly, several features of the DGH system conform to our theoretical requirements: (1) The granule cells are very hard to fire and have low spontaneous activity [6]; (2) Granule cell axons (mossy fibers) project very sparsely [1]; (3) Mossy cells are very easy to fire [16]; (4) The projection from mossy cells to granule cells is highly divergent and relatively diffuse [16]. This is a system ideally configured to produce sparse place codes for storage in CA3.

3. Simulations and results

We simulated the environment as an $L \times L$ grid on which the simulated animal moves randomly. The 2-layer network has two functional layers — A and B, which can be seen as the DG and hilus, respectively — and an input layer, I, which can be seen as the entorhinal cortex. The I layer is modeled at a purely phenomenological level: Each cell is given a broad, noisy Gaussian place field with a randomly chosen center and random shape and orientation. As the animal moves on the grid, each I cell's activity just reflects its place field.

The A layer receives excitatory input from I and focused excitatory as well as global inhibitory input from B. A-layer cells, which are modeled as discrete-time threshold units, are fired by a K-of-N rule simulating competitive firing. Since the input from B represents a feedback, it has a delay of one time step. B layer neurons are also modeled as simple threshold units with competitive firing.
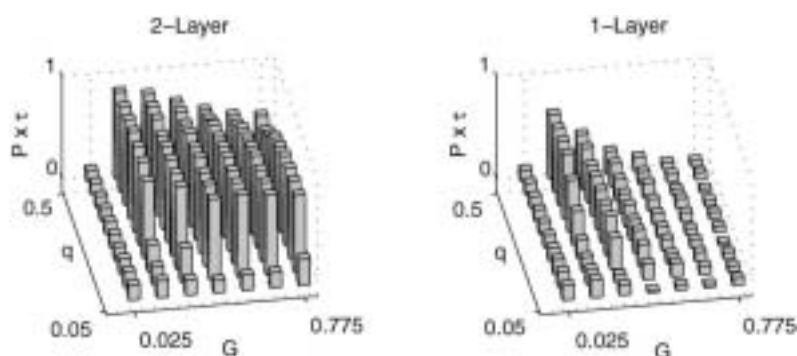
Fig. 1. Stability and quality of coding in 2-layer (left) and 1-layer (right) networks.

M cell groups are chosen from both the A and B excitatory cells, with no restrictions on overlaps. All groups in the same layer are of identical size. A-to-B and B-to-A cell connectivity is uniform random with a fixed probability of connection. Cell groups are implemented by setting all existing cross-group connections to low values and all within-group connections to high values.

The one-layer network is constructed as above, but with the A layer projecting directly to itself.

As the simulated animal traverses the environment, the activity of A cells represents a place code for the animal's current position, and active A cells exhibit place fields [5]. The first stimulus, $s(0)$, through a competitive process (due to the system connectivity), confines future activity within one A group (latent attractor), with cells in other groups firing only sporadically.

The architecture of the system as well as group size are fixed for all simulations. The strength of the input from I to A is also fixed, as is the activity level in layer B for the 2-layer system. However, the recurrent gain, G, is left as a potentially variable parameter (see below). The sparseness parameter, $q = K/N$ — the number of active A cells as a fraction of group size — is used as the independent variable. The system's performance is evaluated by two criteria:

1. Stability of the latent attractor, measured by the probability, P, of being in the correct group during steps 100 to 110 of a simulation.
2. Spatial consistency of the DG place codes, i.e., the coherence between pairwise similarity of place codes and spatial proximity of the corresponding sites. We measure it by using Kendall's tau (τ) [12].

Fig. 1 plots the composite $P \times \tau$ in $G - q$ space. It shows that it is possible to retain both high P and τ only in a narrow range within $q - G$ space for the 1-layer network. In the 2-layer network, however, a high value of $P \times \tau$ can be obtained over a wide area of $q - G$ space, indicating the robustness and flexibility of the 2-layer network.

4. Conclusion

Simulations done using reasonable connectivity parameters for 1- and 2-layer networks show that a DG-hilus-like 2-layer system can support sparse context-dependent place codes much more flexibly than a CA3-like 1-layer network. This does not, of course, preclude the existence of latent attractor dynamics in the CA3, but provides a compelling argument for the DG-hilus system as its primary locus.

Acknowledgements

References

[1] D.G. Amaral, Emerging principles of intrinsic hippocampal organization, Curr. Opinion Neurobiol. 3 (1993) 225–229.

[2] N. Burgess, M. Recce, J. O'Keefe, A model of hippocampal function, Neural Networks 7 (1994) 1065–1081.

[3] S. Doboli, A.A. Minai, P.J. Best, Generating smooth context-dependent neural representations, Proceedings of IJCNN'99 (1999).

[4] S. Doboli, A.A. Minai, P.J. Best, A latent attractors model of context-selection in the dentate gyrus-hilus system, Neurocomputing 26–27 (1999) 671–676.

[5] S. Doboli, A.A. Minai, P.J. Best, Latent attractors: A model for context-dependent place representations in the hippocampus, Neural Comput., in press.

[6] R.A. Fricke, D.A. Prince, Electrophysiology of dentate granule cells, J. Neurophysiol. 51 (1984) 195–209.

[7] M.E. Hasselmo, E. Schnell, E. Barkai, Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in hippocampal region CA3, J. Neurosci. 15 (1995) 5249–5262.

[8] W.B. Levy, A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks, Hippocampus 6 (1996) 579–591.

[9] J.E. Lisman, Relating hippocampal circuitry to function: The role of reciprocal dentate-CA3 interaction in the recall of sequences, Neuron 22 (1999) 233–242.

[10] D. Marr, Simple memory: a theory for archicortex, Phil. Trans. R. Soc. Lond. B 262 (1971) 23–81.

[11] A.A. Minai, P.J. Best, Encoding spatial context: A hypothesis on the function of the dentate gyrus-hilus system, Proceedings of IJCNN'98, Anchorage (1998) 587–592.

[12] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C: The art of scientific computing, 2nd Edition, Cambridge University Press, Cambridge, 1992.

[13] A.D. Redish, D.S. Touretzky, The role of the hippocampus in solving the morris water maze, Neural Comput. 10 (1998) 73–111.

[14] E.T. Rolls, The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus, in: R. Durbin, C. Miall, G. Mitchison (Eds.), The Computing Neuron, Addison-Wesley, Reading, MA, 1989, pp. 125–159.

[15] A. Samsonovich, B.L. McNaughton, Path intergration and cognitive mapping in a continuous attractor neural network model, J. Neurosci. 17 (1997) 5900–5920.

[16] H.E. Scharfman, Dentate hilar cells with dendrites in the molecular layer have lower thresholds for synaptic activation by perforant path than granule cells, J. Neurosci. 11 (1991) 1660–1673.

[17] D.S. Touretzky, A.D. Redish, Theory of rodent navigation based on interacting representations of space, Hippocampus 6 (1996) 247–270.

[18] A. Treves, E.T. Rolls, Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network, Hippocampus 2 (1992) 189–200.

[19] M.A. Wilson, B.L. McNaughton, Dynamics of the hippocampal ensemble code for space, Science 261 (1993) 1055–1058.

Simona Doboli received the Engineer Diploma (magna cum laudae) in 1995 and M.Sc. in 1996 from the "Politehnica" University of Timisoara, Romania, both in Systems Engineering. Since 1997, she has been enrolled in the Ph.D. program in Electrical Engineering at the University of Cincinnati. Her research focuses on the computational modeling of hippocampal function.

Ali A. Minai received his Ph.D. in Electrical Engineering and post-doctoral training in computational neuroscience, both at the University of Virginia. Since 1993, he has been an Assistant Professor in the Electrical & Computer Engineering and Computer Science Department at the University of Cincinnati, and directs the Complex Adaptive Systems Laboratory. His research interests include neural computation, nonlinear dynamics, and complex systems. Dr. Minai is a member of Sigma Xi, Eta Kappa Nu, Tau Beta Pi, IEEE, and INNS, and chairs the International Neural Network Society's Nonlinear Dynamics Special Interest Group.

Phillip J. Best is a Professor in the Psychology Department and the Center for Neuroscience at Miami University, Oxford, Ohio. He received his Ph.D. in Psychology at Princeton University, and has served on the faculties of the University of Michigan, the University of Virginia, the University of New Orleans, and Miami University. His research has focused on hippocampal cellular activity during spatial navigation, as well as the effects of lesions and ethanol on hippocampal activity and navigational behavior. Dr. Best has served on various review panels for NSF and NIH. He has been an associate editor of the Psychonomic Bulletin and Review since 1993, and was a consulting editor for Psychobiology from 1986 to 1997.