

Ensemble Dynamics of Hippocampal Regions CA3 and CA1

Minireview

John F. Guzowski,^{1,4,*} James J. Knierim,^{2,4,*} and Edvard I. Moser^{3,4,*}

¹Department of Neurosciences
University of New Mexico HSC
Albuquerque, New Mexico 87131

²Department of Neurobiology and Anatomy
University of Texas Medical School at Houston
Houston, Texas 77225

³Centre for the Biology of Memory
Norwegian University of Science and Technology
NO-7489 Trondheim
Norway

Computational models based on hippocampal connectivity have proposed that CA3 is uniquely positioned as an autoassociative memory network, capable of performing the competing functions of pattern completion and pattern separation. Recently, three independent studies, two using parallel neurophysiological recording methods and one using immediate-early gene imaging, have examined the responses of CA3 and CA1 ensembles to alterations of environmental context in rats. The results provide converging evidence that CA3 is capable of performing nonlinear transformations of sensory input patterns, whereas CA1 may represent changes in input in a more linear fashion.

The hippocampal formation has long been at the forefront of theory and research into the neurobiological mechanisms underlying learning and memory. This focus has its genesis in five lines of research (O'Keefe and Nadel, 1978; Eichenbaum, 2000). (1) Human patients with damage to the hippocampus display a profound anterograde amnesia in which they are unable to form new, long-term, declarative memories. (2) Animals with hippocampal damage also display distinct learning and memory deficits. (3) Hippocampal synapses exhibit long-term potentiation. (4) Anatomical connection patterns (in particular, the recurrent collateral circuitry of the CA3 subregion) have inspired numerous theoretical models of the hippocampus as an autoassociative memory network. (5) Hippocampal neurons display place-specific firing. Combined with the prevalence of spatial learning deficits in rats with hippocampal lesions, the spatial modulation of hippocampal neuronal firing provides one of the best opportunities to decipher how the brain creates high-order cognitive representations of the world from multimodal sensory input and to relate these representations to the behavior of the animal.

The hippocampal formation comprises multiple subregions, including the entorhinal cortex, the dentate gyrus (DG), the CA3 and CA1 fields of the hippocampus proper, and the subiculum. Differences in anatomical

connection patterns and synaptic physiology among these regions have inspired many models of the unique computational roles performed by each subregion (e.g., McNaughton and Morris, 1987; Rolls and Treves, 1998). Yet there has been limited experimental support for these proposed functions, despite numerous investigations of the DG and CA regions. Recently, a set of papers from three research groups has demonstrated striking differences in the ensemble activity of CA3 and CA1 neurons in behaving animals (Lee et al., 2004a; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004). These results offer tantalizing evidence in favor of a long-standing computational theory of the DG-CA3 network, namely, that it mediates a dynamic competition between two complementary processes of an associative memory system: pattern completion and pattern separation.

Pattern completion refers to the ability of a network to respond to a degraded input pattern with the entire previously stored output pattern. *Pattern separation* refers to the ability to make the stored representations of two similar input patterns more dissimilar, in order to decrease the probability of errors in recall. How might these processes be manifest in the ensemble activity of hippocampal place cells? One possible role of the place cell representation is to facilitate context-dependent learning (i.e., the ability to learn different adaptive responses to the same stimuli based on differing contextual information about the environment or about the organism's internal state). Place cells can "remap" when a familiar environment is sufficiently altered or when the animal's task is altered in a stable environment (for review, see Knierim, 2003). This remapping may allow the animal to create independent representations of the behavioral contingencies of stimuli in each context, thereby reducing the chances of producing an inappropriate response. Remapping may be an indication of pattern separation in the hippocampus, as changes to its sensory or cognitive inputs cause the hippocampus to create independent representations of the altered environment or context.

If every small change in an environment automatically caused the hippocampus to create new representations, such rampant pattern separation would be disruptive. Thus, pattern separation must be tempered with pattern completion, such that slight, behaviorally irrelevant alterations to the context (or small degradations in the quality of the input representations) can be ignored and the system can reconstruct the full, correct representation from the degraded input. This competition between pattern completion (or generalization) and pattern separation has been modeled as a sigmoidal function (McClelland and Goddard, 1996). Small changes to the input cause the hippocampus to perform a pattern completion operation, resulting in an output representation that is more similar than its inputs; as the inputs become increasingly dissimilar, the hippocampus switches to a pattern separation mode, and its outputs become even more dissimilar than its inputs (Figure 1).

Neural ensemble recordings by Lee et al. (2004a) and Leutgeb et al. (2004) have shown evidence for pattern

*Correspondence: jguzowski@salud.unm.edu (J.F.G.); james.j.knierim@uth.tmc.edu (J.J.K.); edvard.moser@cblm.ntnu.no (E.I.M.)

⁴These authors contributed equally to this work.

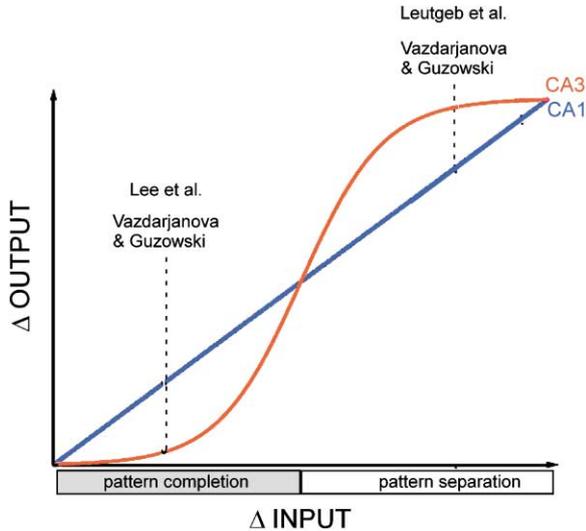


Figure 1. Hypothetical Nonlinear Transformation of Input Patterns in CA3 but Not CA1

The sigmoidal input-output function may reflect dynamic competition between pattern completion and pattern separation in the CA3 network (McClelland and Goddard, 1996; Rolls and Treves, 1998). Output signals are tolerant against small changes in input signals (pattern completion as in the Lee et al. [2004a] and Vazdarjanova and Guzowski [2004] studies) but shift radically in response to larger changes (pattern separation as in the Leutgeb et al. [2004] and Vazdarjanova and Guzowski [2004] studies). Although CA1 is shown here as responding linearly to changes in inputs, under other conditions, it may respond nonlinearly, perhaps reflecting a switch in its dominant input from the entorhinal cortex to CA3. Figure courtesy of J.K. Leutgeb.

completion and pattern separation in the CA3 subregion. Lee et al. (2004a) recorded simultaneously from CA3 and CA1 as an animal ran laps around a circular track in an environment with distinct, familiar cues on the walls and on the surface of the track. In probe tests, the cues on the track were rotated counterclockwise, while the cues on the walls were rotated by an equal amount clockwise. This double rotation of the proximal and distal cue sets caused a mismatch at each location on the track between the sensory input provided by the cues on the track and the input provided by the wall cues. When the magnitude of the mismatch was small (45°), both CA1 and CA3 ensembles output coherent representations that were similar to those of the original, familiar cue configuration. When the mismatch amounts were larger than 45° , however, the CA3 and CA1 ensembles reacted differently. The CA1 representation lost its coherence, as some place fields rotated with the proximal cues, others rotated with the distal cues, and a majority of the place fields remapped or displayed ambiguous behavior. In contrast, the CA3 representation was more coherent between the familiar environment and mismatch environments, as approximately half of the place fields rotated with the proximal cues and only a minority remapped or displayed ambiguous behavior. Thus, whereas CA1 reacted to the altered input patterns by creating similarly altered output representations, CA3 reacted in a way consistent with computational models of pattern completion or generalization (see also Kesner et al., 2000; Nakazawa et al., 2004).

In contrast, Leutgeb et al. (2004) provide evidence for stronger pattern separation in CA3 than in CA1. Rats were tested in enclosures with varying geometric similarity in three similar but distinguishable rooms. Enclosures used in different rooms were either identical (large box), moderately different (small box versus large box), or clearly different (large square versus small circle). In CA1, there was significant overlap between the populations of active neurons in the three rooms. With identical enclosures, the overlap was almost as large as during repeated testing in the same room, and place fields were correlated. With increasing differences between the enclosures, the amount of overlap decreased toward chance levels. In CA3, however, distinct subsets of pyramidal cells were activated in each room, even when enclosures were identical. The overlap of the subsets was no larger than the overlap expected by random shuffling of active neurons, indicating that activated neurons were as distinct as possible and suggesting that input patterns from the entorhinal cortex had been actively orthogonalized. These results show that overlapping input patterns are orthogonalized in ensembles of CA3 neurons whereas CA1 neurons remain more faithful to their inputs.

On first glance, the Lee findings (pattern completion in CA3) and the Leutgeb findings (pattern separation in CA3) appear to be in conflict. However, when viewed in terms of the differences in the testing environments (different rooms in the Leutgeb study versus altered cues in the same room in the Lee study), the results can be interpreted neatly in the framework of the competition between pattern completion and pattern separation processes (Figure 1). This interpretation is supported strongly by an experiment by Vazdarjanova and Guzowski (2004), using imaging of immediate-early gene (IEG) activity, rather than neural recordings, as a measure of ensemble activity patterns. This study examined the responses of CA3 and CA1 ensembles in rats exposed sequentially to (1) the same environment twice, (2) two similar environments in which either the proximal or distal cues were altered, or (3) two completely different environments. The activity history of neurons for each of the two context exposures was determined using the IEGs *Arc* and *Homer1a* as “genomic timers” of neural activation (Vazdarjanova et al., 2002). When presented with changes to either proximal cues or distal cues, ensembles in CA3 demonstrated greater overlap between contexts than in CA1, as in the Lee et al. (2004a) paper. The overlap in CA3 was only marginally less than during repeated testing in the same environment, suggesting that pattern completion occurred in CA3. When the rats were exposed to two dissimilar environments in which both proximal and distal cues were varied, the overlap of CA3 ensembles active in each exposure was at chance levels, whereas that of CA1 ensembles was higher, suggesting stronger pattern separation in CA3, as in the Leutgeb et al. (2004) paper. Like the two neurophysiological studies, the gene imaging study thus points to a discontinuous distribution of ensemble overlap in CA3, in contrast to a more continuous distribution in CA1.

The three studies suggest that cell assemblies in CA3 perform pattern completion under some circumstances and pattern separation under others. In agreement with theoretical predictions, pattern completion occurred

when the sensory inputs were changed to a small extent, whereas pattern separation occurred when the sensory inputs were changed more drastically, as predicted under the assumption that attractor networks respond nonlinearly to input patterns (Rolls and Treves, 1998). However, the combination of parameters determining whether CA3 responds with pattern completion or pattern separation are not known. It remains to be determined, for example, why a change in distal cues was sufficient to induce pattern separation in the Leutgeb study but not in the Vazdarjanova study. Whether the CA3 network is governed by nonlinear dynamics as predicted by the model needs to be tested more directly by changing the overlap between input patterns in a gradual and quantifiable manner within the same experiment. Lee et al. (2004a) attempted this test by varying the rotation mismatch in 45° increments from 45° to 180°, but it appears that even a 180° mismatch in that study was not a strong enough environmental change to trigger complete pattern separation in CA3. Preliminary data from other experiments suggest that CA3 ensembles do under some circumstances switch from pattern completion to pattern separation as an environment is transformed from one shape to another (J.K. Leutgeb et al., 2004, *Soc. Neurosci.*, abstract, volume 34).

One important caveat regarding all of these studies is that they rely on assumptions about the nature of the input representations to the hippocampus. Specifically, the interpretation that CA3 actively performs pattern separation or pattern completion relies on two assumptions: (1) when there are small changes to the environment, the CA3 representations of the standard and probe trials are more coherent than the input from the entorhinal cortex; and (2) when there are larger changes to the environment, the CA3 representations are more orthogonal than the entorhinal inputs. These assumptions must be tested directly in future experiments by measuring the ensemble properties of entorhinal cortex under these conditions, in order to determine whether CA3 is performing active pattern separation/completion on its inputs or is passively relaying the results of information processing that occurred in upstream structures. Similarly, the role of the DG in these processes (which is often modeled as performing a pattern separation function) remains to be determined.

By pointing to more heterogeneous response patterns in CA1 than CA3, the three studies also raise important issues about the functional relation between these hippocampal areas. The less coherent responses of the CA1 pyramidal cells may reflect the convergence of inputs from multiple cell assemblies in CA3, or the cells might be activated directly by layer III neurons in the entorhinal cortex. The differential time course of ensemble formation in CA3 and CA1 in the Leutgeb et al. (2004) paper speaks to this issue. When rats were tested in a novel room, the active subset of cells in CA3 changed over the first 20–30 min of recording. Some cells started to fire only after several minutes, while others turned gradually off. In contrast, the active subset in CA1 stabilized during the first 10 min, well ahead of the CA3 cells, suggesting that initial activity in CA1 depended more on direct entorhinal input. Another temporal dissociation between CA3 and CA1 was observed in a study by Lee et al. (2004b), in which the cue-mismatch conditions

first caused CA3 place fields to shift backward with experience, but not CA1 place fields; on subsequent exposures, CA3 place fields ceased shifting backward, whereas CA1 place fields began to demonstrate the phenomenon. These and other temporal dissociations between CA3 and CA1 (Nakazawa et al., 2004) are consistent with previous work showing that direct connections with the entorhinal cortex are sufficient for spatial modulation in CA1 pyramidal neurons after removal of afferent CA3 neurons (Brun et al., 2002).

The relative independence of areas CA3 and CA1 raises the intriguing question of how ensemble information from CA3 gains access to the neocortex, which is thought to be responsible for the long-term storage of hippocampal-dependent memories. Output from CA3 reaches the neocortex almost exclusively by way of CA1, so how can the codes of the CA3 network be transmitted to the neocortex if the firing pattern of CA1 pyramidal cells is determined primarily by the input from the entorhinal cortex? One possible answer is that the relative influence of CA3 and entorhinal input is dynamic (Hasselmo et al., 1996; Lisman and Otmakhova, 2001). While CA1 cells may respond strongly to entorhinal cortex under standard test conditions, tests that put larger demands on memory may reveal a stronger dependence on CA3, expressed perhaps as more coherent or divergent response profiles also in CA1. It is also possible that ensemble activity in CA1 varies with state. Sleep may be one condition during which CA1 neurons may transmit coherent or orthogonal signals from CA3 to target areas in the neocortex (Wilson and McNaughton, 1994). A third possibility is that ensemble organization in CA1 evolves slowly and that the CA1 response pattern becomes more CA3-driven with repeated testing. This would be consistent with the late differentiation of place cell maps in CA1 when rats are trained in geometrically distinct but otherwise identical enclosures, although the differentiation is frequently not complete after even several weeks of training (Lever et al., 2002; Leutgeb et al., 2004).

Finally, while the three studies provide evidence for attractor dynamics in CA3, the significance for behavior remains to be demonstrated. Pattern completion may allow the brain to recognize stored patterns when the retrieval cues overlap partially with the actual cues that were present as a memory was stored. Conversely, pattern separation may permit the brain to store events with common elements as separate traces or, in CA3, by the formation of nonoverlapping cell assemblies in order to minimize interference. While early studies suggest that the hippocampus is necessary for reducing memory interference (Jarrard, 1975) and later work provides evidence for pattern completion in CA3 (Kesner et al., 2000; Nakazawa et al., 2004), it remains a key challenge for future research to determine how pattern completion and pattern separation in CA3 contribute to specific memory processes.

Selected Reading

- Brun, V.H., Otnass, M.K., Molden, S., Steffenach, H.A., Witter, M.P., Moser, M.B., and Moser, E.I. (2002). *Science* 296, 2243–2246.
- Eichenbaum, H. (2000). *Nat. Rev. Neurosci.* 1, 41–50.

- Hasselmo, M.E., Wyble, B.P., and Wallenstein, G.V. (1996). *Hippocampus* 6, 693–708.
- Jarrard, L.E. (1975). *J. Comp. Physiol. Psychol.* 89, 400–408.
- Kesner, R.P., Gilbert, P.Y., and Wallenstein, G.V. (2000). *Curr. Opin. Neurobiol.* 10, 260–265.
- Knierim, J.J. (2003). In *The Neurobiology of Spatial Behaviour*, K.J. Jeffery, ed. (Oxford: Oxford University Press), pp. 226–239.
- Lee, I., Yoganarasimha, D., Rao, G., and Knierim, J.J. (2004a). *Nature* 430, 456–459.
- Lee, I., Rao, G., and Knierim, J.J. (2004b). *Neuron* 42, 803–815.
- Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.B., and Moser, E.I. (2004). *Science* 305, 1295–1298.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). *Nature* 416, 90–94.
- Lisman, J.E., and Otmakhova, N.A. (2001). *Hippocampus* 11, 551–568.
- McClelland, J.L., and Goddard, N.H. (1996). *Hippocampus* 6, 654–665.
- McNaughton, B.L., and Morris, R.G.M. (1987). *Trends Neurosci.* 10, 408–415.
- Nakazawa, K., McHugh, T.J., Wilson, M.A., and Tonegawa, S. (2004). *Nat. Rev. Neurosci.* 5, 361–372.
- O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Oxford: Clarendon Press).
- Rolls, E.T., and Treves, A. (1998). *Neural Networks and Brain Function* (Oxford: Oxford University Press).
- Vazdarjanova, A., and Guzowski, J.F. (2004). *J. Neurosci.* 24, 6489–6496.
- Vazdarjanova, A., McNaughton, B.L., Barnes, C.A., Worley, P.F., and Guzowski, J.F. (2002). *J. Neurosci.* 22, 10067–10071.
- Wilson, M.A., and McNaughton, B.L. (1994). *Science* 265, 676–679.