Beta Oscillations and Hippocampal Place Cell Learning during Exploration of Novel Environments

Stephen Grossberg[1]
[1]Department of Cognitive and Neural Systems
Center for Adaptive Systems,
and
Center of Excellence for Learning in Education, Science and Technology
Boston University
677 Beacon St
Boston, Massachusetts 02215
steve@bu.edu
Telephone: 617-353-7858/7
Fax: 617-353-7755
http://www.cns.bu.edu/~steve

July 29, 2008
Revised: November 24, 2008

CAS/CNS-TR-08-003
Hippocampus, in press

Running head: Beta Oscillations during Hippocampal Place Cell Learning
Corresponding author: Stephen Grossberg, steve@bu.edu

Keywords: grid cells, category learning, spatial navigation, adaptive resonance theory, entorhinal cortex

Abstract

The functional role of synchronous oscillations in various brain processes has attracted a lot of experimental interest. Berke et al. (2008) reported beta oscillations during the learning of hippocampal place cell receptive fields in novel environments. Such place cell selectivity can develop within seconds to minutes, and can remain stable for months. Paradoxically, beta power was very low during the first lap of exploration, grew to full strength as a mouse traversed a lap for the second and third times, and became low again after the first two minutes of exploration. Beta oscillation power also correlated with the rate at which place cells became spatially selective, and not with theta oscillations. These beta oscillation properties are explained by a neural model of how place cell receptive fields may be learned and stably remembered as spatially selective categories due to feedback interactions between entorhinal cortex and hippocampus. This explanation allows the learning of place cell receptive fields to be understood as a variation of category learning processes that take place in many brain systems, and challenges hippocampal models in which beta oscillations and place cell stability cannot be explained.

The role of the hippocampal system in spatial navigation has been of great interest since O'Keefe and Dostrowsky (1971) showed the spatial correlates of pyramidal cell firing in the hippocampus. Many of these cells tend to fire in a specific portion of the environment independently of the head direction and movement speed, hence the term *place cells*. Such place cell selectivity can develop within seconds to minutes, and can remain stable for months (Thompson and Best, 1990; Wilson and McNaughton, 1993; Muller, 1996; Frank et al., 2004). Many models of hippocampal place cell formation have been proposed but, until recently, none has explained this combination of fast learning and stable memory, which is often called the *stability-plasticity dilemma* (Grossberg, 1980, 1999). How place cells are learned and remembered has attracted even more interest since the recent discovery of grid cells (Hafting et al., 2005) within entorhinal cortical circuits that project to the hippocampus.

Berke et al. (2008) have reported that beta oscillations occur during the learning of hippocampal place cell receptive fields in novel environments.. Paradoxically, beta power was very low during the first lap of exploration, grew to full strength as a mouse traversed a lap for the second and third times, became low again after the first two minutes of exploration, and remained low on subsequent days of exploration. Beta oscillation power also correlated with the rate at which place cells became spatially selective, and did not correlate with theta oscillations. Given the rapidity with which place cell learning occurred, and the sharp increase in beta activity during the second exposure to the environment, it would seem that a highly selective learning mechanism is at work. The present article explains these properties of beta oscillations as natural consequences of brain processes that enable place cell receptive fields to solve the stability-plasticity dilemma. This explanation unifies three parallel streams of modeling activity, and leads to testable predictions aimed at clarifying the underlying neural mechanisms.

Fast Learning and Stable Memory. The first stream of modeling activity concerns how the brain can quickly learn to categorize information in the world, and to remember it without experiencing catastrophic forgetting. As noted above, how the brain combines rapid plasticity with long-term memory stability is called the *stability-plasticity dilemma* (Grossberg, 1980, 1999). Adaptive Resonance Theory, or ART, predicts how the brain accomplishes this feat.

The development of ART included the discovery of how the brain learns recognition categories using *self-organizing maps* (Grossberg, 1976, 1978; Kohonen, 1984). In a self-organizing map, distributed patterns of input features are represented at a first processing level F_1. They activate adaptive connections to a second processing level F_2, whose cells represent recognition categories, or compressed representations, of the feature patterns (Figure 1a). These category cells compete with one another to choose one, or a small number, of winning cells that receive the largest total inputs. The winning cells then send teaching signals to abutting synaptic knobs, which learn the signals that reach them from the feature level via their adaptive connections. Grossberg (1976, 1978) mathematically proved that these associative and competitive mechanisms can together learn stable categories in response to a sparse series of input patterns whose statistics do not change through time. However, in response to a dense series of inputs whose statistics do change through time, catastrophic forgetting of previously learned categories can occur, during which the same input pattern can activate ever-changing categories when it is repeated in a series of intervening input patterns, because the intervening patterns can recode previously learned categories. Dense non-stationary inputs are, however, frequently encountered in the world in which we live. Thus, a self-organizing map does not solve the stability-plasticity dilemma.

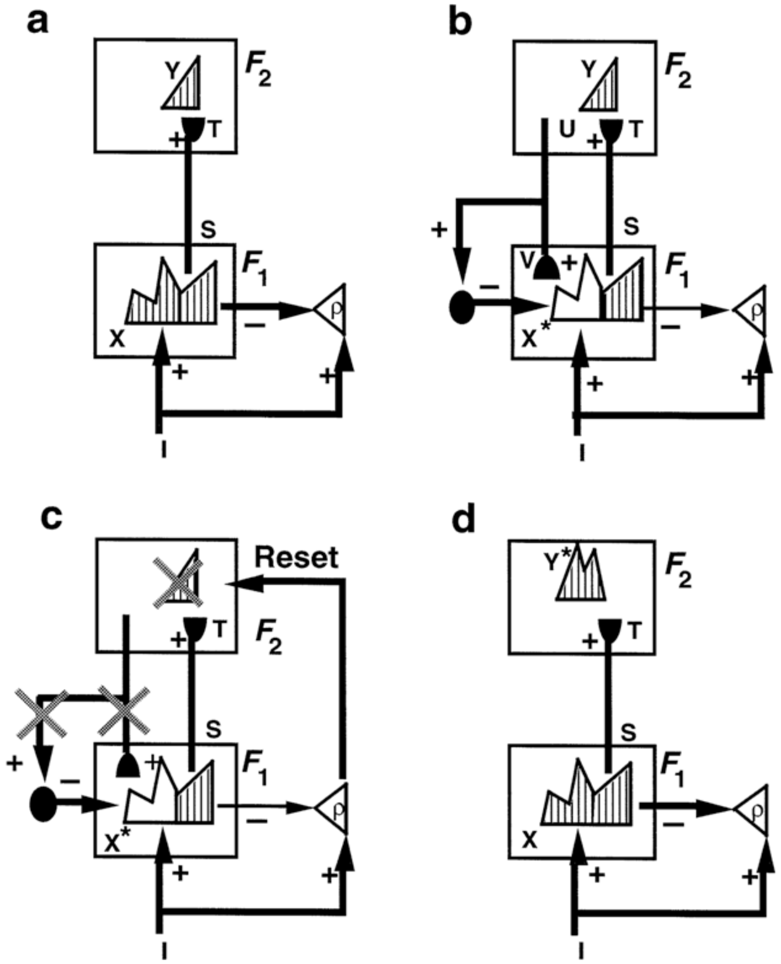

Figure 1. ART dynamics and learning: (a) Input pattern I is instated across feature detectors at level F_1 as activity pattern X (hatched pattern). I also nonspecifically excites the orienting system ρ, while X inhibits ρ. I excitation and X inhibition are balanced. I also generates output pattern S, which is multiplied, or gated, by learned bottom-up adaptive weights. These gated signals add to form input pattern T, which causes a recurrent on-center off-surround network to competitively select an activity pattern Y across winning recognition categories at level F_2. (b) Y generates top-down output pattern U. U signals are multiplied by adaptive weights and added at F_1 cells to form a *prototype* V that encodes the learned expectation of Y. Such a prototype represents shared features in all input patterns that activate Y. If V mismatches any I features at F_1, then a new activity pattern X* (hatched pattern) is selected at F_1, which is active at I features that match V. Matched features support gamma oscillations. Mismatched features (white area) are inhibited. Mismatched features support beta oscillations. When X changes to X^*, total inhibition decreases from F_1 to ρ. (c) If inhibition decreases sufficiently, ρ releases a nonspecific arousal burst to F_2; that is, “novel events are arousing”. Arousal resets F_2 by inhibiting Y. (d) After Y is inhibited, X is reinstated and Y stays inhibited as X activates a different category pattern Y* at F_2. Search continues until a better-matching or novel F_2 category is selected. When search ends, an attentive resonance triggers learning of the attended data. [Adapted with permission from Carpenter and Grossberg (1993).]

Attentive Matching, Resonant Learning, and Mismatch-Mediated Reset. ART showed how the stability-plasticity dilemma could be solved by using learned top-down expectations. These expectations are read out from the category level F_2 to the feature level F_1 via adaptive top-down connections (Figure 1b), and matched at F_1 against the bottom-up input feature patterns. A top-down, modulatory on-center, off-surround network realizes this matching process (Carpenter and Grossberg, 1987, 1991; Grossberg, 1995, 1999), and explains data about "biased competition" (Desimone, 1998; Kastner and Ungerleider, 2001). The modulatory on-center reads out learned weights corresponding to features in the prototype of the learned category. Top-down matching occurs with these primed features, while the off-surround suppresses mismatched features that are not in the on-center.

If the match is good enough, then selected features form an attentional focus. A feedback loop of mutual excitation then develops between the attended feature pattern and the active category, carried via their adaptive bottom-up and top-down connections. This mutual excitation causes a context-sensitive resonant state that synchronizes, amplifies, and prolongs cell responses, as it simultaneously inhibits mismatched features and categories.

Such a resonance drives fast learning within both the bottom-up adaptive weights that define recognition categories, and the top-down adaptive weights that define learned expectations; hence the name *adaptive* resonance. This is unlike learning in a self-organizing map, which is triggered by bottom-up events alone.

If the match is not good enough, then a mismatch state develops which causes *reset* of the current category and search for another category that can better learn to represent and predict the data (Figures 1c and 1d).

These predicted ART mechanisms have been supported by many behavioral and neurobiological experiments; see Grossberg (2003), Grossberg and Versace (2008), and Raizada and Grossberg (2003) for reviews. The main point for present purposes is that top-down attentive matching and mismatch-mediated reset can solve the stability-plasticity dilemma.

Every Input Pattern Can Initially be Matched by its Top-down Expectation. On the first learning trial, when a category is first selected by its bottom-up connections to represent a novel set of features, its top-down expectation must be able to match *any* input feature pattern. The category cannot know in advance what feature pattern it will end up representing. If it mismatched this input pattern, reset would be triggered, and learning could not get started. How a self-organizing neural system "gets started" is a critical issue for any type of learning.

In ART, all top-down expectations can initially match any input pattern because they have excitatory adaptive weights in their modulatory on-centers that start out large and are broadly distributed across the network (e.g., Carpenter and Grossberg, 1987, 1991). Subsequent learning trials *prune* these adaptive weights as they gradually select a more localized pattern of critical features in a category's modulatory on-center, while its other top-down weights approach zero. Partial mismatches occur as learning prunes a category's critical feature pattern. Larger mismatches reset active categories to search for better-matching ones (Figures 1c and 1d).

Mismatch and Beta Oscillations. Is there a relationship between mismatches and beta oscillations? Were this the case, then all the Berke et al. (2008) data could be explained: Beta oscillations are not seen during the first lap because, on the first learning trial, there are no top-down mismatches. Beta oscillations begin during the second lap and are correlated with the rate at which place cells became selective because mismatches occur when learning is refined. Indeed, Berke et al. (2008) reported that "the extent of beta-entrainment predicted the improvement in spatial specificity between the first 2 min.” Beta oscillations are attenuated after the first few trials because, when place cell learning is stabilized by top-down expectations, no more mismatches occur.

The Synchronous Matching ART (SMART) model predicts how beta oscillations occur during mismatch and reset events (Grossberg and Versace, 2008; Versace and Grossberg, 2006). More generally, SMART shows how ART mechanisms of attentive match-based learning and mismatch-mediated category pruning, reset, and search can be realized within the laminar circuits of multiple cortical areas as they interact with primary and higher-order specific thalamic nuclei and nonspecific thalamic nuclei. SMART hereby simulates multiple levels of brain organization, ranging from spiking dynamics to cognitive processing. In SMART, matches between bottom-up input patterns and learned top-down expectations can cause gamma oscillations that support attention, resonance, and spike-timing dependent plasticity (STDP) learning, whereas mismatches can cause slower beta oscillations that inhibit STDP at mismatched cells.

The SMART model simulates learning of perceptual and cognitive categories within thalamocortical and corticocortical circuits. However, the neural mechanisms that cause beta oscillations during mismatch states are not specific to the features that are categorized and matched. Thus, if place cells in the hippocampus are learned as spatial categories that obey ART laws, then an explanation of the Berke et al. (2008) data immediately follows.

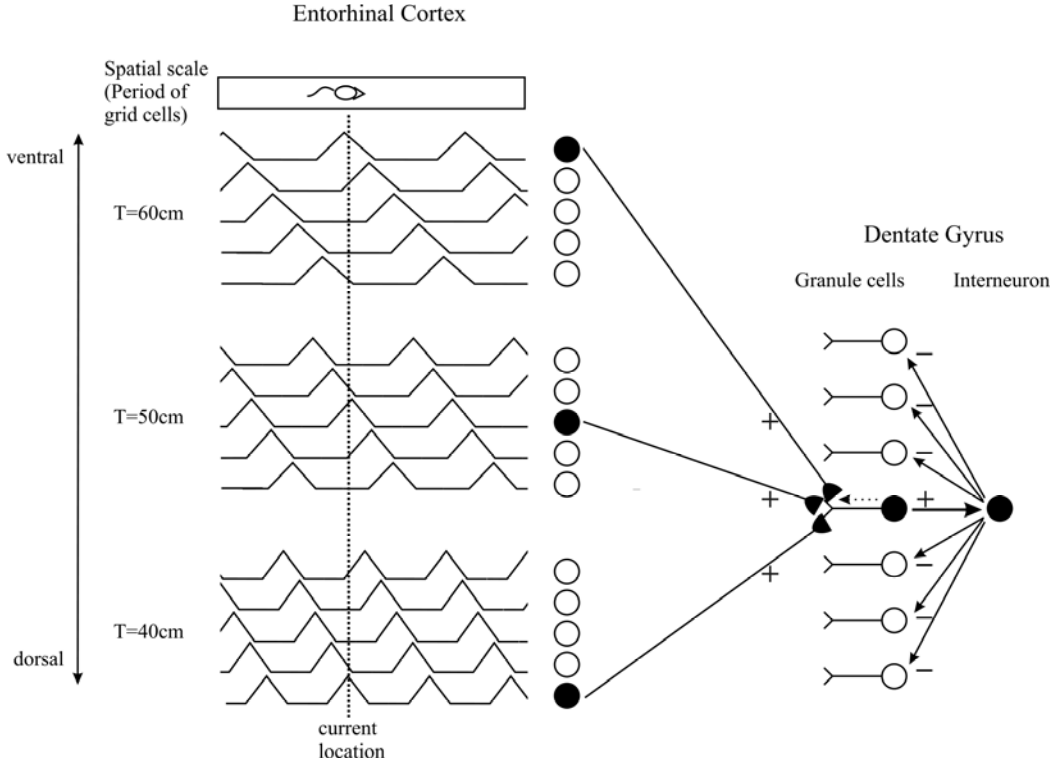

Figure 2. Place cell category learning: Grid cells in entorhinal cortex play the role of level F_1 in Figure 1, whereas place cells in dentate gyrus play the role of category cells in level F_2. Three populations of entorhinal grid cells of five cells each are aligned along the dorso-ventral gradient in entorhinal cortex and have different spatial scales. Their firing profiles are represented as peaks of activity that are aligned with the track along which an

animal moves. The animal's current location causes the corresponding grid cells to fire (filled circles). The dentate gyrus granule cell (population) that receives strong projections from all three active grid cells fires (filled circle) and activates a recurrent inhibitory interneuron (population) that suppresses other granule cells. A back-propagating action potential in the winning granule cell's dendrites (dotted arrow) carries a teaching signal that triggers associative learning at synapses receiving active grid cell inputs. Such back-propagating action potentials dissociate read-out of learned associations from their subsequent read-in after a competitive decision is made (Grossberg, 1975; Hasselmo et al., 2002). Only bottom-up projections from grid cells to hippocampal place cells are shown. Resonant learning dynamics also require top-down connections, as in Figure 1. [Reprinted with permission from Gorchetchnikov and Grossberg (2007).]

Place Cells as Spatial Categories of Multiple-scale Grid Cells. Gorchetchnikov and Grossberg (2007) have modeled how place cell receptive fields may be learned as spatial catgories in an ART system. In particular, Gorchetchnikov and Grossberg (2007) simulated how adaptive connections from entorhinal grid cells with multiple, but small, spatial scales (Hafting et al., 2005) may activate hippocampal place cells that can represent spaces of many meters during navigational behaviors (Figure 2). In this proposal, multiple-scale entorhinal grid cell activations form the level F_1 feature patterns that are categorized by hippocampal place cells acting like spatial categories in level F_2. The spatial scale of these place cells is the *least common multiple* of the incoming grid periods (e.g., see p.143 in Hartmann, 1997). By itself, however, these bottom-up grid-to-place cell learned connections cannot solve the stability-plasticity dilemma for the same reasons that no self-organizing map can do so.

Hippocampal-to-Entrorhinal Feedback, Novelty Sensitivity, and Visual Landmarks. The stability-plasticity dilemma may be solved, and catastrophic forgetting of place cell maps prevented, only if there is top-down attentive matching of grid cell activity patterns by top-down learned expectations that are read out by active place cells. Mismatches are predicted to cause beta oscillations. Large enough mismatches also activate novelty-sensitive cells that mediate reset of hippocampal place cells (Figures 1c and 1d). Many experiments have discussed novelty-sensitive cells in the hippocampus since Vinogradova (1975).

Recent neurobiological data from other experimental paradigms support the prediction that beta oscillations occur during mismatch states. For example, Buschman and Miller (submitted for publication) have reported beta oscillations in the frontal eye fields during spatial attention shifts. Buffalo et al (2004) have reported more beta oscillations in deep layers of visual cortex, and more gamma oscillations in superficial levels of visual cortex. These data are consistent with the SMART model prediction that reset is mediated through the deeper layers of visual cortex. This prediction can be tested by varying the number of reset-inducing events per unit time.

Given that visual landmarks can influence the firing of place cells, the place-to-grid feedback also clarifies how visual landmarks can modulate the activity of both grid and place cells, whose primary activation is derived from path integration signals (O'Keefe and Nadel, 1978; Hafting et al. 2005; Leutgeb et al., 2005; McNaughton et al., 2006). The model in this article may thus be tested by checking if beta oscillations are caused by inducing mismatch between path integration and visual estimates of an animal's position. In addition, the article summarizes many predictions about how grid cells give rise to place cells, and about the control circuitry that regulates refinement and stabilization of this learning process.

References

Berke JD, Hetrick V, Breck J, Green RW. (2008). Transient 23- to 30-Hz oscillations in mouse hippocampus during exploration of novel environments. Hippocampus 18: 519-529.

Buffalo EA, Fries P, Desimone R. 2004. Layer-specific attentional modulation in early visual areas. Soc. Neurosci. Abstr. 30: 717–716.

Buschman TJ, Miller EK. (2008). Covert shifts in attention by frontal eye fields are synchronized to population oscillations. Submitted for publication.

Carpenter GA, Grossberg S. 1987. A massively parallel architecture for a self-organizing neural pattern recognition machine. Comput Vis Graphics Image Proc 37:54-115.

Carpenter GA, Grossberg S. 1991. Pattern Recognition by Self-Organizing Neural Networks. Cambridge MA: MIT Press.

Frank LM, Stanley GB, Brown EN. 2004. Hippocampal plasticity across multiple days of exposure to novel environments. J Neurosci 24:7681-7689.

Gorchetchnikov A, Grossberg S. 2007. Space, time, and learning in the hippocampus: How fine spatial and temporal scales are expanded into population codes for behavioral control. Neural Networks 20:182-93 .

Grossberg S. 1975. A neural model of attention, reinforcement, and discrimination learning. Internatl. Rev. Neurobiol. 18: 263–327.

Grossberg S. 1976. Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors. Biol Cybern 23: 121–134.

Grossberg S. 1978. A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and plans. In Rosen R and Snell F, editors. Progr Theoret Biol. Vol 5. New York: Academic Press, p. 233–374. Reprinted in Grossberg S. 1982. Studies of Mind and Brain. Amsterdam: Kluwer/Reidel.

Grossberg S. 1980. How does a brain build a cognitive code. Psychol Rev 87:1-51.

Grossberg, S. 1995. The attentive brain. American Scientist, 83, 438-449.

Grossberg S. 1999. The link between brain learning, attention, and consciousness. Conscious Cogn 8:1-44.

Grossberg S. 2003. How does the cerebral cortex work? Development, learning, attention, and 3D vision by laminar circuits of visual cortex. Behav. Cogn. Neurosci. Revs. 2: 47-76.

Grossberg S, Versace M. 2008. Spikes, synchrony, and attentive learning by laminar thalamocorical circuits. Brain Res, in press.

Hafting T, Fyhn M, Molden S, Moser MB, Moser E. 2005. Microstructure of the spatial map in the entorhinal cortex. Nature 436(11): 801–806.

Hartmann WM. 1997. Signals, sound, and sensation. New York: American Institute of Physics.

Hasselmo, ME, Bodelon C, Wyble BP. 2002. A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. Neural Computation 14: 1-25.

Kastner S, Underledier LG. 2001. The neural basis of biased competition in human visual cortex. Neuropsychologia 12:1263-1276.

Kohonen T. 1984. Self-organization and associative memory. Springer-Verlag.

Leutgeb S, Leutgeb JK, Moser MB, and Moser EI. 2005. Place cells, spatial maps and the population code for memory. Current Opinion Neurobiol, 15:738-746.

McNaughton BL, Battaglia FP, Jensen O, Moser EI, and Moster MB. 2006. Path integration and the neural basis of the “cognitive map”. Nature Revs Neurosci 7: 663-678.

Muller RA. 1996. A quarter of a century of place cells. Neuron 17:813-822.

O’Keefe JM, Dostrovsky J. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Research 34(1): 171–175.

O’Keefe JM and Nadel L. 1978. The Hippocampus as a Cognitive Map. Oxford: Clarendon.

Raizada R, Grossberg S. 2003. Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. Cerebral Cortex 13: 100-113.

Thompson LT, Best PJ. 1990. Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. Brain Res 509:299-308.

Versace, M, Grossberg S. 2006. From spikes to inter-areal synchrony: How attentive matching and resonance control learning and information processing by laminar thalamocortical circuits. Society for Neuroscience annual meeting, 65.11/Z12.

Vinogradova OS. 1975. Functional organization of the limbic system in the process of registration of information: Facts and hypotheses. In: Isaacson R.L., Pribram, K.H., eds., The hippocampus. Vol. 2, 3-69. New York: Plenum.

Wilson MA, McNaughton BL. 1993. Dynamics of the hippocampal ensemble code for space. Science 261:1055-1058.

Acknowledgements

Stephen Grossberg was supported in part by the National Science Foundation (SBE-0354378)