Francesco Savelli and James J. Knierim J Neurophysiol 103:3167-3183, 2010. First published Mar 31, 2010; doi:10.1152/jn.00932.2009

You might find this additional information useful...

- Supplemental material for this article can be found at: http://jn.physiology.org/cgi/content/full/jn.00932.2009/DC1
- This article cites 89 articles, 29 of which you can access free at: http://jn.physiology.org/cgi/content/full/103/6/3167#BIBL
- Updated information and services including high-resolution figures, can be found at: http://jn.physiology.org/cgi/content/full/103/6/3167
- Additional material and information about *Journal of Neurophysiology* can be found at: http://www.the-aps.org/publications/jn

This information is current as of September 17, 2010.

Journal of Neurophysiology publishes original articles on the function of the nervous system. It is published 12 times a year (monthly) by the American Physiological Society, 9650 Rockville Pike, Bethesda MD 20814-3991. Copyright © 2010 by the American Physiological Society. ISSN: 0022-3077, ESSN: 1522-1598. Visit our website at http://www.the-aps.org/.

Hebbian Analysis of the Transformation of Medial Entorhinal Grid-Cell Inputs to Hippocampal Place Fields

Francesco Savelli^{1,3} and James J. Knierim^{1,2,3}

¹Zanvyl Krieger Mind/Brain Institute, ²Solomon H. Snyder Department of Neuroscience, Johns Hopkins University, Baltimore, Maryland; and ³Department of Neurobiology and Anatomy, University of Texas Medical School, Houston, Texas

Submitted 16 October 2009; accepted in final form 25 March 2010

Savelli F, Knierim JJ. Hebbian analysis of the transformation of medial entorhinal grid-cell inputs to hippocampal place fields. J Neurophysiol 103: 3167-3183, 2010. First published March 31, 2010; doi:10.1152/jn.00932.2009. The discovery of grid cells in the medial entorhinal cortex (MEC) permits the characterization of hippocampal computation in much greater detail than previously possible. The present study addresses how an integrate-and-fire unit driven by grid-cell spike trains may transform the multipeaked, spatial firing pattern of grid cells into the single-peaked activity that is typical of hippocampal place cells. Previous studies have shown that in the absence of network interactions, this transformation can succeed only if the place cell receives inputs from grids with overlapping vertices at the location of the place cell's firing field. In our simulations, the selection of these inputs was accomplished by fast Hebbian plasticity alone. The resulting nonlinear process was acutely sensitive to small input variations. Simulations differing only in the exact spike timing of grid cells produced different field locations for the same place cells. Place fields became concentrated in areas that correlated with the initial trajectory of the animal; the introduction of feedback inhibitory cells reduced this bias. These results suggest distinct roles for plasticity of the perforant path synapses and for competition via feedback inhibition in the formation of place fields in a novel environment. Furthermore, they imply that variability in MEC spiking patterns or in the rat's trajectory is sufficient for generating a distinct population code in a novel environment and suggest that recalling this code in a familiar environment involves additional inputs and/or a different mode of operation of the network.

INTRODUCTION

An understanding of how information is processed by the circuits of the hippocampal formation can provide a mechanistic explanation for the role of this area in learning and memory (Eichenbaum et al. 2007; O'Keefe and Nadel 1978; Squire et al. 2004). Hippocampal principal neurons have spatially selective firing fields ("place fields") confined to a restricted region of a typical recording area (O'Keefe and Dostrovsky 1971) although they can display multiple fields in larger environments (Fenton et al. 2008). In contrast, one of the major inputs to the hippocampus, the medial entorhinal cortex (MEC), contains spatially specific firing patterns that are periodic (Hafting et al. 2005). Each MEC "grid cell" fires in multiple locations arranged as vertices of an equilateral triangular (or hexagonal) grid spanning the recording area. Hippocampal place cells are thought to build their place fields mainly by converting the many-location responses of grid cells into firing that is usually restricted to a single location (McNaughton et al. 2006; O'Keefe and Burgess 2005; but see Kropff and Treves 2008 for an alternative view). Understanding the rules that govern this transformation would be a major step forward in understanding hippocampal computation and function.

It has been proposed that place fields can be generated from grid cell inputs by a simple summation and threshold operation. That is, if a place cell receives input from an arbitrary set of grid cells, the activation of the place cell can be prevented everywhere except in the single region where the input to the cell is maximal (McNaughton et al. 2006; O'Keefe and Burgess 2005). This cutoff mechanism could be implemented by the postsynaptic spiking threshold. However, if the geometric alignment of the input grids is unconstrained, distinct subsets of coactive inputs will almost invariably generate similarly high levels of synaptic excitation in multiple regions covering the majority of the environment (Solstad et al. 2006). A simple threshold mechanism would not be able to single out one of these regions. The only circumstance in which this mechanism would work in a purely feedforward manner is if there was a single location where most of the afferent grid cells onto a place cell shared a common vertex, which would then summate to cause the place cell to fire at that particular location and fail to reach threshold at other locations (Solstad et al. 2006).

A subset of grids with overlapping vertices at a single location could be optimally selected from randomly aligned grid inputs by choosing a suitable synaptic weight vector, e.g., via Fourier analysis (Solstad et al. 2006), a fitting algorithm (Blair et al. 2007), or independent component analysis (Franzius et al. 2007). It is not known, however, how this task can be autonomously accomplished at a behaviorally relevant time scale with physiological mechanisms. In the present study, we investigated a Hebbian learning rule within a minimal, spiking network model of a layer of grid-cell inputs with modifiable connections onto an output layer. The model is not intended to simulate the output of a particular subregion of the hippocampus (e.g., dentate gyrus or CA1), any of which receive far more complex inputs than the reduced set of pure grid-cell inputs studied here. The model can be thought of as an explicit investigation of MEC grid inputs onto CA1 cells in the absence of the CA3 Schaffer collateral system (Brun et al. 2002) or perhaps onto individual dendrites of a DG granule cell (Leutgeb et al. 2007; Ujfalussy et al. 2009). However, we were primarily concerned with the problem of how a simple, feedforward network could accomplish the grid-to-place transformation problem in a general form, to generate insights that might be applicable at some level to all hippocampal subfields.

Address for reprint requests and other correspondence: F. Savelli, Zanvyl Krieger Mind/Brain Institute, Johns Hopkins University, 338 Krieger Hall, 3400 N. Charles St., Baltimore, MD 21218 (E-mail:fsavelli@jhu.edu).

METHODS

Computational model and simulations

Both place cells and their grid cell inputs were simulated as spiking units. Place cells of the hippocampus were modeled as integrate-andfire neurons. Their inputs were spike trains derived from a phenomenological characterization of the activity of MEC grid cells of varying phases, orientations, and scales, along the trajectory followed by a real rat recorded in our laboratory during a 15-min segment of a foraging session in a walled square box (60×60 cm). This design made the simulated course of activity of place cells and their inputs temporally and behaviorally relevant to typical experimental conditions.

A triangular grid template of each grid cell was defined by the position of one of its vertices (phase), the orientation of the vector



linking this vertex with one of its neighbors (orientation), and the intervertex distance (scale) (Hafting et al. 2005). The firing rate in the excitation map of this grid decreased exponentially with the square of the distance from the nearest vertex in proportion to the scale of the grid. In particular, the normalized firing rate (ranging from 0 to 1) at any position was computed as $\exp[-d^2/(k \times s^2)]$, where s is the scale of the grid, d is the distance of the position from the closest vertex of the underlying grid template, and k is a factor (held constant at 0.018 throughout the paper) chosen to make the simulated grids appear similar to published data (Hafting et al. 2005). The normalized firing rate was used to modulate the instantaneous rate of an inhomogeneous Poisson spike train according to the rat's trajectory. We implemented the generation of the spike train via dynamic thinning (Dayan and Abbott 2001). Briefly, at each iteration of the algorithm, the next spike interval was sampled from the exponential distribution of event intervals of a Poisson process of maximal rate (20 Hz, corresponding to the peaks of the grid); intervals shorter than 3 ms were set to 3 ms to create a refractory period. As the simulation proceeded, whenever the interval expired, the spike was accepted with a probability given by the normalized firing rate of the grid cell calculated for the current position. Positions were sampled at 30 Hz from the trajectory followed by the rat during an unconstrained foraging session lasting 37 min. Figure 1A illustrates the rat's trajectory (gray) and the spikes (red dots) for three simulated grid cells of different scales.

Two random number generators were used, both of the "Mersenne Twister" type (Gnu Scientific Library's implementation). The first assigned the geometric properties of phase, orientation, and scale to the grid inputs and set up the synaptic connectivity between these inputs and the postsynaptic cells (explained in the following text). The second was dedicated to the generation of the grid-cell spike trains. The separation of these two sources of randomness allowed us to run simulations that were identical except for the exact spike timing of the input trains by changing only the seed of the second random number generator. We used this type of manipulation when studying the remapping properties of the model.

Basic elements of the simulations and place-field generation by FIG. 1. synaptic plasticity. A: grid cells were simulated as inhomogeneous Poisson spike trains in which the momentary firing rate was modulated depending on the predetermined geometrical properties of the grid cell and on the position of the rat. A 15-min trajectory segment of a real rat foraging in a 60 cm² square box was used. Three examples of grid cells are shown. The gray lines show the rat's trajectories, and the red dots indicate the locations of the rat when the simulated grid cell fired spikes. Place cells were simulated as integrate-andfire units taking synaptic inputs from the grid cells. Each simulation contained 1,000 grid cells of varying scale, phase, and orientation and 500 place cells. Each place cell received inputs from 100 distinct grid cells that were randomly selected from the pool of 1,000 grid cells. B: uniform integration of randomly aligned grid inputs with constant and equal synaptic weights does not produce place fields. The gray line represents the rat's cumulative trajectory (15 min), and the red dots denote the location of spikes produced by the integrate-andfire output cell receiving 100 grid inputs (dot size is scaled to improve visualization). B1-B3: the outcome of the same simulation run with different weights (B1, $w = 7.0 \times 10^{-2} \mu$ S; B2, $w = 4.5 \times 10^{-2} \mu$ S; B3, $w = 3.7 \times 10^{ ^{2}$ μ S; weights have physical dimensions of a conductance; see METHODS). 10^{-2} Decreasing the weights causes a reduction of firing in the output cell but the firing does not become more spatially concentrated. C: postsynaptically gated synaptic plasticity supports the formation of a place field by clustering spatially coincident inputs. Trajectory and spikes during salient epochs of the field formation within the first 3 min of exploration are shown. The inputs, initial weights, and all parameters were the same as in B2 except for the inclusion of the synaptic plasticity rule of Eq. 2. D: temporal evolution (x axis) of the weights for the 100 grid inputs (y axis) during the 1st 3 min of foraging. Blue indicates a synaptic weight (w) of 0 and red indicates maximum synaptic weight. E: temporal evolution of weights as in D but for the whole duration of the foraging session. F: the resulting spatial response is consistent during the whole simulation.

The membrane dynamics of hippocampal place cells was modeled by

$$C_m \frac{dV}{dt} = -g_l (V - E_l) + \sum_{s} -g_s (V - E_s)$$
(1)

where V is the membrane voltage, $C_{\rm m}$ is the membrane capacitance (2 nF), g_1 is the leak conductance (0.2 μ S), and E_1 is the leak reversal potential (-65 mV). The synaptic contribution to membrane dynamics was modeled with dynamic conductances; for each synapse s, g_s is its conductance and E_s is its reversal potential. Whenever V crossed the firing threshold (-50 mV), a 1-ms spike event was superimposed. For the following refractory period (3 ms), V was kept at its resting potential (-70 mV), after which the integration of the voltage dynamics according to Eq. 1 was resumed. The membrane voltage V was never allowed to decrease below the lower bound -100 mV or increase above 100 mV by artificially clipping voltages outside these extremes (the dynamics of the conductances was not manipulated). Equation 1 was numerically integrated by the exponential Euler method with a 1-ms time step. When inhibitory interneurons were employed, they had the same physiological parameters as in the preceding text.

In this conductance-based model, the synaptic conductance g_s was 0 in the absence of presynaptic activity. To model synaptic activity, g_s was instantaneously incremented by a quantity w_s every time a presynaptic spike occurred and decayed exponentially to 0 according to the equation $dg_s/d_t = -g_s/\tau_s$. For excitatory synapses, the time constant τ_s was set to 2 ms, after the dynamics of AMPA receptors (Colquhoun et al. 1992). For inhibitory synapses, τ_s was set to 6 ms to follow the slower course of GABA receptor dynamics. The reversal potentials E_s of excitatory and inhibitory synapses were set to 0 and -70 mV, respectively. The value of w_s depended on previous potentiation or depression of synapse s by synaptic plasticity, as described below (Eq. 2). This variable, therefore played the role of synaptic weight (generally indicated with just the letter w in the rest of the paper) and was the subject of synaptic modification whenever synaptic plasticity was active in the simulation. The initial value of w_s at all synapses was $4.5 \times 10^{-2} \mu$ S, unless otherwise specified. The weight was artificially kept from becoming negative or growing above a saturation limit of 0.1 μ S.

Presynaptic and postsynaptic firing rates used in the synaptic rule (see Eq. 2 below) were computed by convolving the presynaptic and postsynaptic spike trains with an exponential kernel. The time constant of the kernel (τ_r) was generally set to 100 ms, but the effect of greater time constants was studied as well.

We used two types of networks in the simulations. The basic version contained 1,000 grid-cell spike-train generators as input to 500 integrate-and-fire units, representing potential place cells of the hippocampus (Fig. 1A). The phase, orientation, and scale of the grids were uniformly sampled as follows. The 1,000 units were first divided into 10 groups of 100 units that corresponded to 10 different scales of intervertex spacing, ranging from 30 to 53 cm by constant increments. This was the range of intervertex spacings observed in roughly the most dorsal 1 mm extension of the dorsocaudal MEC (Hafting et al. 2005). Each group was in turn split into 10 subgroups, each of 10 units, corresponding to 10 different orientations separated by 6° increments. The first orientation value was sampled randomly in the $0-6^{\circ}$ range independently for each different scale. Finally, the phases of the 10 units in each of these scale+orientation subgroups were uniformly sampled over the entire enclosure covered by the rat's trajectory. Each of the 500 integrate-and-fire hippocampal units received 100 excitatory synaptic inputs that were uniformly sampled without repetition from the available pool of 1,000 grid-cell spiketrain generators. Hence the same input unit was generally shared by different hippocampal cells.

The second network type included 50 additional integrate-and-fire units as feedback interneurons. These interneurons received excitatory inputs from the hippocampal cells and made inhibitory synapses (see preceding text) onto the same class of cells. The main results described in this paper could be qualitatively reproduced with many simulations varying in the number and strength of the synaptic connections. The quantitative results we report refer to the following parameters. Each excitatory unit projected to 40 randomly chosen but distinct interneurons (excitatory synapses); each interneuron projected to 300 randomly chosen but distinct excitatory units (inhibitory synapses). The strength of these excitatory and inhibitory synapses was not subject to plasticity: w_s for inhibitory synapses was set to 0.2 μ S and w_s for excitatory synapses was set to 0.8 μ S.

In each simulation, the computational model was run over a 15-min segment of the tracked rat's 37-min trajectory in accordance with the run-time of a typical physiological recording experiment of place cells. The output of the simulations mainly consisted of the times-tamps of every spike for each cell and the weight of every input to each cell sampled every 100 ms. These data were stored in files and analyzed off-line.

Simulations were implemented in C++ with use of the Gnu Scientific Library. Data analysis and plotting were implemented in Python using the NumPy and Matplotlib libraries as provided by the Enthought Python Distribution (Enthought, Austin, TX).

Data analysis

The average firing rate of a cell was calculated as the number of spikes it fired during the simulated session divided by the duration of the session. To compute firing rate maps, the recording enclosure was segmented into $\sim 3 \times 3$ cm bins. The firing rate in a bin was calculated as the number of spikes that occurred when the rat occupied the bin divided by the time spent by the rat in the bin. The rate maps were not smoothed. Bins that were visited for periods totaling <233ms were excluded from further analysis. The rate map was normalized by dividing every bin by the maximum value across all bins. For cells with an average firing rate ≥ 0.033 Hz, place fields were defined as sets of at least four contiguous bins (36 cm^2) if the bins all had value >0.15 (i.e., >15% of the maximal bin value) and at least one bin had a mean firing rate >1 Hz. When one or more place fields were detected on a rate map, their size and relative contribution to the total cell activity were computed. Size was simply computed by counting the number of bins composing the field. The proportion of "in-field" firing was computed as the ratio between the sum of rates in the bins that belonged to any detected field and the sum of the rates in all the bins of the map. High values of this ratio indicated that most of the cell's firing occurred within the boundaries of the detected place fields with sparse firing occurring outside the fields. The size and in-field firing proportion were conservative estimates in that the cutoff value (15% of the peak rate) for the boundaries of the place fields often tended to exclude the fringes of the field. Whenever a distribution or the statistics of peak firing rates is reported, it comprises only the cells that were included in place field analysis (i.e., with mean firing rate ≥0.033 Hz).

To determine whether place fields were distributed homogeneously throughout the environment, the cumulative rate map of a cell population was computed. For each bin of the cumulative rate map, the sum of the corresponding bins in the individual rate maps of all the cells in the population was first computed. Bins excluded from further analysis in individual maps were excluded from the cumulative map. Finally, the bins of the cumulative rate map were normalized with respect to the greatest bin value. To investigate whether the spatial biases represented by the cumulative rate map correlated with the initial segment of the rat's trajectory, an occupancy map for the first minute of trajectory was correlated with the cumulative rate map. The occupancy map had the same dimensions as the cumulative map and its bin values represented the time the rat spent in them within the segment of trajectory under consideration. This procedure was then applied to all the trajectory segments obtained by sliding a 1-min time window by 1-s steps along the entire available trajectory (37 min).

The resulting set of r values was used to assess the statistical significance of the r value corresponding to positioning the window at the starting point of the simulation.

RESULTS

A synaptic rule for generating place fields

The problem with integrating randomly aligned grid inputs (Solstad et al. 2006) is illustrated in Fig. 1B, which shows the trajectory of the rat (gray line) and the locations of the spikes (red dots) from a simulated integrate-and-fire unit receiving 100 different grid inputs. The three plots result from simulations that differed only in the initial weight assigned to all synapses. The synaptic weights were not subject to any modification during the session. Figure 1B1 shows the firing of the output cell when the synaptic weights were greatest. The output cell did not form a restricted place field. An increase in the firing threshold was simulated by decreasing the weights of all input cells, but the cell still fired in multiple locations (Fig. 1B2). Finally, the weights were decreased such that the cell only fired a few spikes (Fig. 1B3), but the cell still showed no spatial specificity. This example illustrates the general point that a simple summation-and-threshold model of random grid cell inputs will not generally cause a peak of firing of the output cell in a single location (Solstad et al. 2006). This was the common behavior of almost all 500 simulated cells that received different combinations of inputs from the same pool of 1,000 grid units. A qualitatively similar result was obtained by increasing an inhibitory current into the unit instead of decreasing the synaptic weights, as an alternative way of tuning the cutoff level of postsynaptic activation (data not shown).

As shown by Solstad et al. (2006), it is essential that a large proportion of input grid cells have a single location in which they all fire to create a single firing field of the output place cell at that location. Thus the inputs to a place cell must be either hard-wired in this fashion or a synaptic weight distribution must be learned to select such inputs for a given place cell. This selection can be accomplished by a postsynaptically gated Hebbian synaptic rule

$$\frac{dw}{dt} = k \left(\text{pre} - \theta_p \right) \text{post}$$
(2)

where w is the weight, pre is the presynaptic firing rate, post is the postsynaptic firing rate, k is the learning rate factor, and θ_p is a threshold on the presynaptic firing rate. Postsynaptic activity is required to trigger synaptic modification. The direction of change is determined by the presynaptic firing rate: if pre is greater than θ_p , the synapse is potentiated; otherwise, the synapse is depressed.

The rule in Eq. 2 permits place-field formation by detecting the spatial coincidence of vertices from multiple grid inputs onto a cell, potentiating such inputs, and depressing the rest of the inputs onto the cell. The process is illustrated in Fig. 1, *C*–*F*, for the same cell, grid inputs, and trajectory as in Fig. 1*B*. All the weights were initially set to the same value ($4.5 \times 10^{-2} \mu$ S, the same as in the constant-weight simulation in Fig. 1*B*2), θ_p was set to 5 Hz, and *k* was set to 4 nS/sHz² = 4 nS · s. This learning rate is high; for example, it allows an input firing at just 1 Hz above the threshold θ_p to increase its weight from the starting value of $4.5 \times 10^{-2} \mu$ S to the saturation value of 0.1 μ S in <14 s if the postsynaptic cell is firing at 1 Hz or in <4 s if the postsynaptic cell is firing at 4 Hz (or, alternatively, if the input is firing at 4 Hz above θ_p). The weights were updated every 4 ms; less frequent updates up to every 100 ms did not produce qualitative differences as far as the following results are concerned, provided that the learning rate was not altered in the process. In the remainder of the paper, the simulation parameters were the same unless otherwise specified.

Figure 1C shows the rat's trajectory and the cell's spikes during different salient intervals of the first 3 min of exploration. The concurrent evolution of the synaptic weights for the grid inputs to this cell is shown in Fig. 1D. The first postsynaptic spike was produced after the first minute of exploration (t1). As a result, the postsynaptic firing rate grew above zero and triggered the synaptic rule: the most active inputs (contributing the most to postsynaptic firing) were potentiated and the least active inputs were depressed. The change is represented by the transition between the uniformly green colored area and the region where horizontal stripes begin to appear after t1 in Fig. 1D. The discrimination between inputs was initially weak because the magnitude of synaptic change was proportional to the postsynaptic firing rate (Eq. 2), which at this point accounted only for a single spike. Successive spikes (t1-t4) reiterated this phenomenon, giving rise to a positive feedback loop between the potentiation of selected inputs and the postsynaptic firing rate in the area where the field was forming. After $\sim 2 \text{ min of exploration (t4), inputs}$ became maximally discriminated. A small subset of inputs with vertices in the place field (red/yellow stripes in Fig. 1D, denoting weight values approaching the saturation limit of 0.1 μ S) prevailed over the rest of the inputs, which did not have vertices in this location (blue background, denoting weight values approaching 0 µS). This discrimination remained relatively stable across the whole session-despite the synaptic rule staying active at all times—as reflected by the continuous horizontal pattern of yellow/red stripes over the blue background in the weight evolution plot spanning the whole 15-min simulation (Fig. 1E). The place field resulting from the activity of these inputs was also consistent and stable, as revealed by partitioning the position and spike data into three consecutive 5-min intervals (Fig. 1F). The final distributions of weights from each input grid onto this place cell and two other co-simulated place cells are provided in Supplementary Figs. S1-S3.¹

The same process produced place fields in most of the 500 co-simulated output cells (see Supplementary Fig. S4). Rate maps of five cells are illustrated in Fig. 2, *A* (without plasticity) and *B* (with plasticity). In each case, the addition of the plasticity rule caused the cells to form highly specific place fields. Population statistics for the entire set of 500 output cells are shown in Fig. 3. The mean firing rates ranged from 0.04 to 1.47 Hz (mean \pm SD = 0.39 \pm 0.2 Hz), so all the 500 cells qualified (0.033 Hz) for place field analysis (Fig. 3*A*). The peak firing rate was 14 \pm 5.8 (SD) Hz. Figure 3*B* shows the distribution of the number of fields per cell across the population. Most cells (403 cells) produced one field, 82 cells produced two fields, and 15 cells produced more than two fields. Hence cells with more than one field were less frequent than cells with single fields due to the lower probability that the

¹ The online version of this article contains supplemental data.

	cell 162	cell 163	cell 164	cell 165	cell 166
A Simulation 0	2	2	2	2	2
В	17	20	15	17	11
Simulation 1					• •
C Simulation 2	3	0.6	0.5	5	•
	30	20	 26	23	22
D Simulation 3	5	1		8	42
E Simulation 4	5		6	10	5
F Simulation 5		7	5	16	4
G Simulation 6	2	2		2	
			2.5.2	S	
H Simulation 7	22	35	14		21

same set of potentiated grid inputs would coincide in multiple spots of the recording enclosure. It is not known precisely how many place cells fire in multiple locations in a given environment, and the answer is almost certainly a complicated function of hippocampal subregion (CA1, CA3, DG) (Chawla et al. 2005; Vazdarjanova and Guzowski 2004), size of the environment (Fenton et al. 2008; Henriksen et al. 2009), behavioral

FIG. 2. Qualitative comparison of different simulations. Rate maps in the same column correspond to the same cell receiving the same inputs and were produced by simulations in which only the form or parameters of the synaptic rule changed (row H is an exception in that it receives different inputs; see text). The number above each rate map denotes the peak firing rate corresponding to the red color in the rate map. A: simulation 0: synaptic weights were fixed and no synaptic plasticity was present. Integration of grid inputs does not produce strong spatial specificity in the output cells. B: simulation 1: the postsynaptically gated rule (Eq. 2, $\theta_{p} = 5$ Hz) was active at all times. Place fields with strong spatial specificity are produced with the learning rule (see also Supplementary Fig. S4). C: simulation 2: same as B but with $\theta_{p} = 10$ Hz. Fewer inputs were potentiated because of the higher θ_{p} threshold, resulting in lower firing rates, but spatial selectivity was generally preserved. The 2 rate maps enclosed by the dashed rectangle did not meet the mean firing rate criterion (≥0.033 Hz) for inclusion in the data analysis. D: simulation 3: same as B but heterosynaptic depression was excluded from the rule ($\theta_p = 5$ Hz, $\theta_d = 0.05$ Hz). This exclusion caused the runaway potentiation of all or most inputs, leading to very high firing rates and a loss of spatial selectivity that was similar to A. E: simulation 4: same as C, but heterosynaptic depression was excluded from the rule ($\theta_p = 10$ Hz, $\theta_d = 0.05$ Hz). In this case, the higher θ_p threshold prevented the runaway potentiation observed in D but still caused a loss of spatial selectivity compared with C, mostly because of the emergence of multiple place fields per cell. F: simulation 5: the loss of spatial selectivity due to the lack of heterosynaptic depression ($\hat{\theta}_p = 5 \text{ Hz}$, $\hat{\theta}_d = 0.3 \text{ Hz}$) was recovered by the addition of intervertex background firing of the grid cell (0.5 Hz) and by a slower integration of the presynaptic spikes into the rate value used in the synaptic rule (see also Supplementary Fig. S6). This modification allowed the necessary function of heterosynaptic depression to be recovered in the form of homosynaptic depression. G: simulation 6: same as B but the rate of learning in the synaptic rule was 20 times slower. Place fields did not form within the foraging session duration (15 min). H: simulation 7: same as B but the grid inputs were sampled from a larger intervertex spacing range (70-93 cm) compared with all other simulations in the paper (30-53 cm). Using this larger range produces place fields closer in size to typical dorsal hippocampus recordings (see also Supplementary Fig. S5).

task/trajectories (e.g., open-field foraging versus linear track running), and other variables. We made no attempt to alter the proportion of cells with single place fields to match a predetermined proportion of single- versus multiple-field cells given our simulated environment and behavior, although subsequent simulations intended to analyze other aspects of the model were found to affect this proportion [see simulations with



FIG. 3. Quantitative verification of the quality of the place fields developed by 500 co-simulated place cells. The simulation was run with parameters as in Fig. 2B. A: average firing rates across the entire simulation. B: number of fields per cell. Most cells developed a single field. C: relative proportion of firing occurring inside the area of the place fields counted in B. Most of the cell activity took place inside these fields. D: size of the fields counted in B. All fields are 1 order of magnitude smaller than the recording box $(3,600 \text{ cm}^2)$.

inputs of larger spacing (Fig. 2*H* and Supplementary Fig. S5), a higher θ_p (simulation 2 of Fig. 2 and Table 1), or with added feedback inhibitory interneurons (see *Influence of early trajectory and feedback inhibition on spatial biases*)].

Figure 3*C* shows the distribution of the proportion of spiking inside the place field over the whole population (see METHODS); in the majority of the cells (454 cells) \geq 70% of the total firing occurred inside place fields. The typical size of the field was much smaller than the recording area (3,600 cm²); place field sizes ranged from 36 cm² (the minimum size requirement for field classification) to 288 cm² (mean ± SD = 101 ± 36 cm²; Fig. 3*D*). Although the place fields tended to be smaller than those typically found in CA3 and CA1 recording experiments, shifting upward the range of intervertex spacing of the inputs (from the 30–53 cm used in the preceding text to 70–93 cm) caused the place fields to resemble more closely those observed in typical dorsal hippocampus recordings in terms of size, hetereogeneity of shape, and number of fields per cell (compare Fig. 2, *H* with *B*, and Supplementary Fig. S5 with S4).

The process of field creation was qualitatively diverse across individual cells. Some cells built their place fields quickly, starting from the first traversal through the field (Fig. 4A). More often, due to a weaker presynaptic pattern or more rapid traversals through the location, multiple visits were needed before intense and consistent firing began (Fig. 4B; see also Fig. 1C). Some cells fired their first spikes in the eventual field (Fig. 4A), whereas other cells fired a few initial spikes elsewhere, thereby perturbing the synaptic weights in a way that later facilitated the initiation of a place field in a different location (Fig. 4C). In rare cases, remapping of a fully developed field emerged, reflecting interference between input patterns (Fig. 4D). In some cases, the cell did not fire for long periods during the initial stage of the exploration, spanning up to a few minutes, after which it started to fire consistently in a quickly formed field (Fig. 4, B and C). The initial lack of postsynaptic firing implies that visits to the future field were too brief for the temporal and spatial integration of input spikes to induce suprathreshold depolarization. Real place cells are known to display such a diverse range of individual behaviors in a new environment, as some place cells fire from the rat's first entry into the place field in a novel environment (Hill 1978), others begin firing abruptly after initial trajectories through the field produced no firing (Frank et al. 2004, 2006),

and there is a general sharpening of the place-field representation over time in a novel environment (Wilson and Mc-Naughton 1993). A number of cells also maintained a spatial pattern of firing that resembled their grid-cell inputs, such as Fig. 4*E* (see also Supplementary Fig. S4). Although such firing patterns have not been reported in CA1 or CA3 place cell recordings, we attribute their appearance in the simulations to the scaled-down, simple architecture of the model, and especially to the absence of spatially nonperiodic inputs; we did not make any attempt to artificially minimize the proportion of these cells.

The threshold $\theta_{\rm p}$ in Eq. 2 affects the number of synapses that are likely to undergo potentiation versus depression. A higher $\theta_{\rm p}$ favors depression, leading to greater discrimination of the inputs generating a field and yielding lower and more spatially restricted levels of activity in the output cells. In some cells, a potentiated input subset might not emerge at all because initial postsynaptic firing could induce depression on most of the synaptic inputs that were contributing to it, thus dramatically reducing or driving to extinction future postsynaptic activity. Figure 2C shows examples of rate maps from a simulation in which $\theta_{\rm p}$ was increased to 10 Hz from its previous value of 5 Hz (Fig. 2B). The firing rates of the cells and the field sizes were decreased, and two of the cells fell below the 0.033 Hz criterion for analysis (dashed outline). These two simulations are compared at the population level in Table 1. All 500 cells met the minimum firing-rate criterion for inclusion in the place-cell analysis (0.033 Hz) when $\theta_p = 5$ Hz, whereas slightly more than half of them did so when $\dot{\theta}_{p} = 10$ Hz (Table 1, simulations 1 vs. 2). On average, the peak firing rates and the field size in the population included in the place cell analysis were also lower when $\theta_{\rm p} = 10$ Hz (Table 1, simulations 1 vs. 2).

Constraints on the synaptic rule

FUNCTION OF HETEROSYNAPTIC DEPRESSION. The synaptic plasticity rule of Eq. 2 is input specific; that is, the change in weight is entirely specified by the current presynaptic and postsynaptic firing rates at each synapse. However, the rule produces a form of heterosynaptic depression in that when some strongly active synapses drive the postsynaptic cell to fire, any inactive synapses undergo depression. We asked whether heterosynaptic depression is necessary for place field formation in our model.

TABLE 1. Statistics of place field formation in different simulations

Simulation number	1	2	3	4	5	6
Simulation parameters*						
$\theta_{\rm p},{\rm Hz}$	5	10	5	10	5	5
$\theta_{\rm d}^{\rm P}$, Hz	0	0	0.05	0.05	0.3	0
$k, nS \cdot s$	4	4	4	4	4	0.2
Minimum intervertex firing of grid inputs, Hz	0	0	0	0	0.5	0
τ_r , ms	100	100	100	100	1000	100
Results						
Number of cells that met analysis criteria	500	255	500	500	394	499
Peak rate, Hz	14.0 ± 0.3	6.2 ± 0.2	26.5 ± 0.3	7.7 ± 0.2	9.2 ± 0.3	3.9 ± 0.2
Number of place fields/cell	1.22 ± 0.02	1.05 ± 0.02	3.87 ± 0.06	1.72 ± 0.04	1.19 ± 0.03	1.38 ± 0.03
Place field size, cm ²	102.0 ± 1.6	69.7 ± 1.2	297.6 ± 13.2	69.4 ± 1.1	89.3 ± 1.5	79.5 ± 1.6
In-field firing/total firing	0.79 ± 0.00	0.83 ± 0.01	0.79 ± 0.00	0.66 ± 0.01	0.75 ± 0.00	0.52 ± 0.01

Values are mean \pm SE. *Parameters that varied between simulations that are displayed in this table. Values that differed from those in simulation 1 are in bold. θ_p is the presynaptic threshold for switching between potentiation and depression; θ_d is the minimum presynaptic firing rate to induce plasticity; k is the learning rate; τ_r is the decay constant of the presynaptic firing rate for excitatory synapses.



FIG. 4. Examples of qualitatively diverse place field formation. A: a cell that quickly formed its place field during the 1st visit to the place-field location. B: a cell that remained mostly inactive for \sim 7 min before starting to fire consistently in the field. C: a cell that 1st fired in a location different from the future field, then was inactive for \sim 2.5 min, and finally established the place field. D: a cell that relocated its field after 10 min. E: a cell that developed 3 stable fields.

We ran simulations in which the plasticity at any synapse was disabled whenever the presynaptic firing rate was less than a threshold (θ_d) of 0.05 Hz. When the presynaptic firing rate was greater than θ_d , Eq. 2 was applied as before. The performance of the system-in terms of how many well-tuned place cells were produced-was severely undermined when heterosynaptic depression was disabled in this manner. The value of the threshold for inducing potentiation (θ_p) affected the outcome. With $\theta_p = 5$ Hz, the system lacking heterosynaptic depression $(\theta_d = 0.05 \text{ Hz})$ was drastically impaired in its ability to produce place fields (Fig. 2D, simulation 3). The reason for this impairment was that inactive synapses that were spared depression early in the simulation underwent potentiation at a later time. Eventually, runaway potentiation affected nearly all the weights, resulting in a loss of discrimination between input patterns. As a result, the spatial response of each output cell was as specific as the one obtained when all inputs had the same, fixed weights (Fig. 2*A*, simulation 0), albeit at a much higher rate. Across the 500 cells, the peak firing rate was 26.5 ± 0.3 (mean \pm SE; Table 1, simulation 3), the number of fields per cell was 3.87 ± 0.06 (Table 1, simulation 3), and the field size was 297.6 ± 13.2 (Table 1, simulation 3). Note that these values are respectively about 2, 2, and 3 times greater than the highest values from the other simulations shown in Table 1.

With a higher threshold for inducing potentiation ($\theta_p = 10$ Hz), the simulation lacking heterosynaptic depression did not lead to runaway potentiation, but it was nonetheless impaired in the ability to produce single-field responses (Fig. 2E). The average number of fields produced by each cell of the analyzed population was higher without heterosynaptic depression (1.72 \pm 0.04; Table 1, simulation 4) than with heterosynaptic depression (1.05 \pm 0.02; Table 1, simulation 2). In addition, more of the cells' firing occurred outside of the fields in the absence of heterosynaptic depression (Table 1, simulation 4 vs. simulation 2). These simulations showed that heterosynaptic long term depression is an essential aspect of Eq. 2 for the effective generation of place-cell-like behavior. Heterosynaptic long term depression (LTD) and depotentiation were the first experimentally observed examples of activity-dependent reduction of synaptic efficacy in the hippocampus (Bliss et al. 2007) both in vitro (Lynch et al. 1977) and in vivo (Levy and Steward 1979). However, recent experimental results at perforant path synapses suggest that low-rate, spontaneous activity that typically occurs in vivo might provide a background level of activity of presynaptic cells that allows depression of these low-rate synapses using homosynaptic mechanisms (Abraham et al. 2007). Alternatively, as the rat moves into the intervertex regions of a grid, the trace of presynaptic activity caused by the recent traversal of a more active area could allow this input to undergo enough homosynaptic depression, if the trace decay is sufficiently slow. To test whether there are conditions under which strict heterosynaptic depression was not essential for place-field formation, we ran simulations in which the rule lacking heterosynaptic depression was used, with different time scales for the presynaptic firing rate decay, with or without spontaneous activity of grid cells in the intervertex spaces. Figure 2F (simulation 5) shows that the combination of intervertex firing at ~ 0.5 Hz (which is similar to published figures of some grid cells in Hafting et al. 2005) and a longer time constant (τ_r 1,000 ms) can promote the formation of robust, unitary place fields, without the need for true heterosynaptic depression (see Table 1 for simulation results and parameters). To determine the relative influence of these two factors, their interaction was explored further in Supplementary Fig. S6. Thus if the argument that heterosynaptic depression in vivo may really be a disguised version of homosynaptic depression (Abraham et al. 2007) proves to be generally true, these results show physiologically plausible ways to preserve the essential *functional* role played by heterosynaptic depression in the abstract rule of Eq. 2—that is, the rapid depression of inputs that are relatively inactive when the place field is forming-without requiring true heterosynaptic depression of inactive synapses.

SENSITIVITY TO THE LEARNING RATE. The learning rate used in the synaptic rule permitted the formation of place fields within the first few minutes of exploration in most cells. Furthermore, it was not always necessary that the environment be thoroughly explored before a stable field was formed (e.g., Fig. 4A). Although these properties mimic the behavior of many real place cells (Frank et al. 2004, 2006; Hill 1978; Wilson and McNaughton 1993) and are consistent with the fast learning essential for the putative role of the hippocampus in episodic memory (Knierim et al. 2006; Manns and Eichenbaum 2006; O'Keefe and Nadel 1978; Squire et al. 2004; Vargha-Khadem et al. 1997), they are somewhat incompatible with the operating conditions of most analytical and computational studies of synaptic learning rules (see, for instance, Dayan and Abbott 2001; Gerstner and Kistler 2002). In these studies, the rate of synaptic change is assumed to be small compared with the rate of presentation of the input patterns. This assumption usually enables the mathematical derivation of the final steady state of the synaptic weights, which usually reflects some statistical properties of the stimulus space. For this purpose, the space of input patterns must be repeatedly experienced in its entirety for the learning system to reach the final state. In the case of place-field formation, this requirement would imply waiting until the animal has sampled the environment multiple times, in contrast with experimental observations.

To test the importance of a fast learning rate for place-field formation, we ran simulations with a range of learning rates that are an order of magnitude slower than that used so far. When learning was ten times slower ($k = 0.4 \text{ nS} \cdot \text{s in } Eq. 2$), visual inspection of firing rate maps of the output cells revealed fields that formed more slowly and were still weak (or absent) by the end of the 15-min session in about half the population. When learning was 20 times slower ($k = 0.2 \text{ nS} \cdot \text{s}$), the proportion of such cells was higher still (Fig. 2G). Inspection of the temporal evolution of the weights revealed that discrimination between inputs took much longer to build up under these conditions (not shown). Table 1 (simulation 6) shows the quantitative analysis of the fields of the whole population in this simulation. Note the lower average peak firing rates and in-field firing and the higher average number of fields per cell compared with simulation 1 of Fig. 2 and Table 1. It thus appears that fast learning is a critical aspect of the transformation of grid inputs to unitary place fields.

PRESYNAPTICALLY GATED RULE. A logical variation of the postsynaptically gated Hebbian rule (Eq. 2) is a presynaptically gated rule, in which the polarity of plasticity depends on the postsynaptic firing rate

$$\frac{dw}{dt} = k \left(\text{post} - \theta_p \right) \text{pre} \tag{3}$$

The postsynaptically gated and presynaptically gated rules are equivalent if the rate of learning is sufficiently slower than the presentation of the input patterns and if θ_p is set to the average presynaptic (for Eq. 2) or postsynaptic (for Eq. 3) firing rate across all the input patterns (Dayan and Abbott 2001). When these conditions are met, both rules amount to a covariance rule (Dayan and Sejnowski 1993; Sejnowski 1977). Because the postsynaptically gated rule (Eq. 2) with a slow learning rate does not generate place fields in a time frame consistent with experimental data, it is worth asking empirically if Eq. 3 could work as well as Eq. 2 in the latter's successful regime of fast learning ($k = 4 \text{ nS} \cdot \text{s}$). Simulations showed that the presynaptically gated rule with fast learning did not form strong, unitary place fields. For any cell, either positive or negative

runaway of all the weights occurred, leading to poor spatial specificity because of a lack of input selection or because of the loss of all excitatory drive onto the output cell. Figure 5 illustrates the resulting bimodal distribution of mean firing rates across the population ($\theta_p = 1$ Hz; higher values increase the proportion of cells that become inactive but do not qualitatively change the main result). Intuitively, this result can be understood as follows. The presynaptically gated rule either potentiates or depresses all active synapses at any given time (while leaving inactive synapses unchanged, since the rule is presynaptically gated). Potentiation occurs at all active synapses whenever postysynaptic firing is above θ_{p} , whereas depression occurs at all active synapses whenever postsynaptic firing is below θ_{p} . Because a grid input to a place cell with a single field will typically exhibit more firing fields outside of the place field than inside it, the mean balance of plasticity occurring at this input will be negative. In this case, runaway depression of the inputs will eventually occur. If instead the place field area covers the majority of the input's vertices, the mean balance of plasticity at that input will be positive. In this case, runaway potentiation of most or all inputs will eventually occur. Thus a synaptic weight distribution that generates single-peaked place fields from grid cell inputs is an unstable state when the presynaptically gated rule is active with fast plasticity.

The failure of simulations with the presynaptically gated rule and the poor performance of the simulations in which the postynaptically gated rule is applied with a slow learning rate suggest that successful place-field formation requires that the system work in the region of parameter space far away from the region where both rules converge appreciably to the covariance rule.

Effects on population coding

INFLUENCE OF EARLY TRAJECTORY AND FEEDBACK INHIBITION ON SPATIAL BIASES. As noted in the preceding text, cells can form stable fields before the exploration of the environment is



FIG. 5. Bimodal distribution of spatial response with the presynaptically gated synaptic rule. The application of this rule leads to either runaway potentiation or depression of all the weights of any cell; this in turn causes either elevated, indiscriminate firing or the loss of activity in the postsynaptic cell. This pattern is reflected in the bimodal distribution of mean firing rates in the histogram. *A*: rate maps of two sample cells that lost firing. *B*: rate maps of 2 sample cells that fired strongly over much of the environment. A simple threshold on these rate maps would produce activity reminiscent of the dentate gyrus, in which the majority of cells are silent and a minority fire in multiple locations (Chawla et al. 2005; Jung and McNaughton 1993; Leutgeb et al. 2007).

complete (e.g., Fig. 4A). The order in which different locations are initially visited might therefore be relevant for where cells will form their fields. This consideration raises the question of whether the initial stage of exploration exerts an influence over the final global distribution of place fields (O'Keefe and Conway 1978), which is not always uniform across the environment (Hetherington and Shapiro 1997; Wiener et al. 1989).

To address this question, we ran simulations with 15-min epochs of the rat's exploration starting from different time points of the entire 37-min foraging session. All other parameters and starting conditions of the simulations were identical to that of simulation 1. We calculated a cumulative rate map (CRM) from each simulation by summing the normalized individual rate maps of all the place cells. The CRMs showed that place fields did not spread across the recording enclosure uniformly; rather they tended to concentrate in certain regions and miss other regions (Fig. 6A). The regions that were over-represented were different in each simulation. An occupancy map was calculated for the first 60-s trajectory of each simulation, and these maps were positively correlated with the CRMs for their respective simulations (Fig. 6, A and C). It was important to determine whether the resulting r values were higher than those expected by correlating the CRMs with any arbitrary trajectory. A sample of r values was computed by sliding a 60-s window in 1-s increments along the entire 37-min trajectory and calculating the correlation of the CRM with the occupancy map for each increment. The r values of the initial trajectories shown in Fig. 6C were all in the top 6% of the simulated sample (P < 0.03, P < 0.04, P < 0.03, P <

Different simulations, trajectory epoch used (minutes):

0.002, P < 0.06), demonstrating that there is a significant relationship between the initial trajectory in an environment and the bias in the locations of place field formation.

The pronounced heterogeneity of the place-field distribution produced by the simple, feedforward model is greater than that typically observed in place-field studies. Many types of feedback interneurons are prevalent in the hippocampus (Freund and Buzsaki 1996). The spatial biases revealed by the CRMs were mitigated when a subpopulation of generalized, feedback inhibitory interneurons-receiving excitatory connections from the output place cells and making inhibitory connections back onto them-was added to the model (Fig. 6D). Intuitively, feedback inhibition reduced the gap between densely and scarcely populated areas by introducing mutual inhibition/ competition between place cells, which prevented too many cells from being active enough to start building a field at the same time. In this way, feedback inhibition made the global activity spread more uniformly in the environment. The interneurons fired in the range \sim 22–25 Hz, whereas the excitatory units projecting to these interneurons fired between 0 and ~ 0.5 Hz. An additional effect of feedback inhibition was a general reduction of average firing rates in the population of place cells. The mean (for all simulated cells) and peak (for the cells meeting the criterion of place cell analysis) firing rates without feedback inhibition were respectively 0.39 ± 0.2 (SD) Hz and 14 \pm 5.8 Hz compared with 0.07 \pm 0.07 and 5.6 \pm 2.7 Hz when inhibition was present (Fig. 6, column 1). Furthermore, all 500 excitatory cells in the simulation without inhibition fired at >0.033 Hz—the criterion for inclusion in the place-cell



FIG. 6. Effects of early trajectory on place field location and density. The 5 columns represent 5 different simulations that were identical except for the choice of the 15-min epoch of trajectory data. A: cumulative rate map of all cells showing inhomogeneous place field density across the simulated population. B: cumulative trajectory of each 15-min epoch. The spatial sampling was comparable across different simulations. C: trajectory data for the 1st 60 s of each epoch. Occupancy maps (not shown) were calculated from these early trajectories. The occupancy maps were correlated with the respective cumulative rate maps in A: place fields concentrated mostly in the locations first explored by the rat (r values: 0.34, 0.25, 0.33, 0.54, 0.21; occupancy maps were not smoothed prior to correlation). D: the biases in the distribution of fields shown in A are reduced when a subpopulation of feedback inhibitory interneurons is added to the place cells. E: different place fields developed by the same cell across the simulations in A. The change of spatial density in the cumulative rate maps is made possible by a similar individual "remapping" of the majority of the cells.

3175

analysis—but only 318 did so in the simulation with inhibition. Among these 318 cells, the percentage of cells with more than one field was drastically reduced (under 4%) compared with the simulation without inhibition (\sim 19%). Thus in addition to spreading out the distribution of place fields, the inhibition prevented a large number of cells from firing in more than one location in the environment.

REMAPPING. Across different environments (and sometimes within the same environment), place cells can change their



relative firing locations or change their average firing rates (Bostock et al. 1991; Knierim 2003; Leutgeb et al. 2004, 2005; Muller and Kubie 1987; Skaggs and McNaughton 1998). This phenomenon, called *remapping*, is thought to reflect the role of the hippocampus in contextual learning. We asked in what circumstances our simulations would produce different maps even if starting with the same grid cell inputs and the same uniform synaptic weights. As already shown, different place fields developed when the early trajectory was modified in the simulations described in the previous section. Indeed the spatial reorganization of place-field density (Fig. 6A) implied that the majority of the cells individually relocated their fields across simulations; an example of such a cell is shown in Fig. 6E. This result suggests that where a place field fires in a particular environment may depend in part on the rat's early exploration trajectories in that environment as speculated by O'Keefe and Conway in their pioneering studies of place-field formation and cue control (O'Keefe and Conway 1978).

An even simpler perturbation of the field-generating process that produced diverging maps consisted of small variations in the timing of spike input patterns. This variation was accomplished by repeating the simulation with a different realization of each grid cell's Poisson spike train. This was implemented by changing the seed of the random number generator (RNG) used in the grid-cell spike trains' generation process (Fig. 7A). All other stochastic aspects of the simulation—such as network connectivity and grid cells' geometric and firing propertieswere handled by another RNG with a seed that was not changed between simulations (see METHODS). Nonstochastic aspects, such as the initially uniform weights, were also identical. The spike-timing alterations had a minimal effect on the overall pattern of firing of the grid cells (Fig. 7B). These small alterations, which occurred independently for each grid input, caused similar small alterations in the exact spiking of postsynaptic place cells (Fig. 7C). The resulting differences in the momentary values of both pre- and postsynaptic firing rates in turn triggered plasticity differentially in terms of timing, location, magnitude, and direction of synaptic efficacy changes. Eventually, this caused the fast, nonlinear dynamics of the field formation to diverge, leading to the potentiation of different sets of synapses and thus to different field locations (Fig. 7D).

FIG. 7. Effects on place field location of exact timing and location of grid cell spikes. The 3 columns represent 3 simulations that were identical except for the initialization of the random number generator affecting the exact timing of the grid cell spikes. As a result, grid cells spiked at slightly different times and locations without altering their overall spatial firing pattern; by contrast, many place cells generated their place fields in distinct locations. A: an example of an input grid cell that spiked at slightly different locations and times during the 1st 60 s. The trajectory of the rat, as well as every other model parameter, did not change. B: the cumulative spatial firing patterns of this cell observed over the whole 15-min sessions were not affected appreciably. C: an example of an output place cell showing similar differences in early spiking. D: firing observed at the time scale of the whole session reveal that the same cell produced firing fields in 3 distinct locations. E: rate maps of a sample cell that changed location and number of place fields across the 3 simulations. F: rate maps of a cell that changed field location between the 1st and 2nd simulations and firing rate between the second and 3rd simulations. G: rate maps of a cell that did not appreciably change either field location or firing rate. *H*: cumulative rate maps computed from all cells of the simulated population that met analysis criteria. The spatial biases of place field density did not change across simulations in spite of the concurrent individual changes in place field locations (because the initial trajectory of the rat did not change across simulations—compare with Fig. 6A).

Not all cells produced different fields every time the RNG's seed was changed. Rather, different combinations of location and rate remapping (or lack thereof) were expressed throughout the population of place cells (Fig. 8, E–G), in a manner



FIG. 8. Varying place field locations and/or mean firing rate across simulations that differ only in the rat's early trajectory or in the input spike timing. A: scatter plot of mean firing rates of each simulated place cell in a reference simulation (abscissa, 0–15 min epoch of trajectory data) and in a simulation with a different start time (ordinate, 5-20 min epoch). Dashed lines indicate the firing rate threshold criterion for inclusion in data analysis (≥0.033 Hz). Cells that did not meet the criterion in one simulation could still meet the criterion in the other; the firing rates of these cells are denoted by the dots on the left of the vertical dashed line and above the horizontal dashed line, or by the dots below the horizontal dashed line and on the right of the vertical dashed line. Cells could have negligible activity in one simulation while being among the most active in the other. B: histogram of the distance between the place fields created in the reference vs. altered simulations as in A for each simulated cell that produced exactly 1 field in both simulations. The majority of cells produced different fields between the 2 simulations. C and D: these graphs are similar to A and B except that the altered simulation was obtained from the same reference simulation by changing the seed of the random number generator involved in the generation of the input spike trains (instead of the trajectory epoch starting point). Changes in place field locations and/or mean firing rates were similar in the 2 types of manipulations. E-H: same as A-D but in a simulation that does not contain feedback inhibition. The amount of remapping that occurs from both trajectory changes and spike-timing changes is similar in the simulations run with and without inhibition.

reminiscent of partial remapping of real place cells (Knierim 2003; Knierim and McNaughton 2001; Skaggs and McNaughton 1998; Tanila et al. 1997). (Note that we used the term remapping here to denote the phenomenology of the simulations producing different place fields for the same output neurons, analogous to the phenomenology of real place cell remapping. However, the simulations always started from the identical uniform synaptic weights pattern; there was no sequential history between simulations. See DISCUSSION for further elaboration.) The spatial biases of the population (as shown by CRMs) remained the same despite the underlying individual relocation of place fields (Fig. 7H) unlike the remapping obtained by changing the rat's trajectory (Fig. 6A). This result is consistent with the previous observation that the spatial biases reliably depended on the details of the initial trajectory, which was unaltered in the RNG seed simulations.

Figure 8 illustrates quantitative details of the remapping that occurred between simulations varying in trajectory (A and B, 0-15 and 5-20 min, respectively) or in spike timing (C and D); both sets of simulations included feedback interneurons (as in Fig. 6D). Figure 8, A and C, illustrates scatter plots of individual mean firing rates for all cells. Values for the same reference simulation are reported on the abscissa for both A and C, whereas values for the simulations with modified trajectory and grid spike timing are reported on the ordinates of A and C, respectively. In both cases, low-rate cells in one simulation could turn into high-rate cells in the other. Dashed lines indicate the 0.033 Hz threshold used for inclusion in the place-field analysis; many cells that do not reach threshold in one simulation are above threshold in the other. Figure 8, B and D, shows the distributions of distances between the field developed in one simulation and the field developed in the other. Only cells that produced exactly one field in both simulations were included: 196 such cells were produced in the remapping by change of trajectory (B) and 204 cells in the remapping by change of spike timing (D). (The distance between the 2 fields is the distance between the centers of mass of their areas in the rate maps in a common frame of reference.) In both cases, at least half of these cells experienced a relocation of their fields by a distance >15 cm. Conversely, a large fraction of cells changed their firing locations by only small amounts (<5 cm). Partial remapping was therefore widely present in the cell population to a similar extent in simulations varied by trajectory and input spike timing. Similar results were obtained from simulations that did not include feedback interneurons (Fig. 8, E–H). The only major difference between the simulations with and without feedback inhibition is the overall higher firing rate of the simulation without feedback, which eliminated the number of cells that fell below firing-rate threshold but had little effect on the proportion of cells that changed their place field locations.

DISCUSSION

Previous theoretical investigations of the grid-to-place field transformation hypothesized that some form of synaptic plasticity is involved in restricting the firing of a place cell to a single location (Blair et al. 2007; Franzius et al. 2007; Molter and Yamaguchi 2008; Rolls et al. 2006; Si and Treves 2009; Solstad et al. 2006; Ujfalussy et al. 2008). We have shown here that a physiologically plausible, Hebbian learning rule is capable of selecting inputs only from grid cells that share a common vertex location to accomplish this task (Solstad et al. 2006). Our investigations revealed that within the confines of our model fast learning dynamics and some form of heterosynaptic depression are required to transform the multipeaked input of grid cells into singlepeaked output of place cells. In addition, the model provides potential insights into the heterogeneous temporal dynamics of place-field formation, the prevalent remapping phenomena, and the potential separation of the roles of synaptic plasticity and network competition via feedback inhibition.

Plasticity

Our model assumes that plasticity plays a major role in place-field formation, an assumption that is made plausible by the ubiquity of synaptic plasticity in all subfields of the hippocampus. A challenge to this assumption comes from a study in which rats were injected systemically with CPP [a blocker of the *N*-methyl-D-aspartate (NMDA) receptor]. These animals formed apparently normal place fields in a novel environment, but they were unable to reactivate these same place fields when reintroduced into the same environment the next day (Kentros et al. 1998; Shapiro and Eichenbaum 1999). This result suggested that NMDA-dependent plasticity may be implicated in the long-term stability of place fields but not in their creation. In contrast, other studies using genetic knock-out techniques have shown decreased quality of place fields in novel environments when NMDA receptors were disrupted (McHugh et al. 1996; Nakazawa et al. 2004), a result that is more in agreement with our model. The reasons for the discrepancies among these studies are not clear, but they may be due to species differences and/or differences in recording or behavioral techniques. Alternatively, it is entirely possible that a form of non-NMDA receptor plasticity that operates in accordance with the present computational model could still support the grid to place transformation (Kentros et al. 1998). Further experiments and a more anatomically and biophysically detailed model of the hippocampus will be required to address the discrepancies between these studies and the respective roles on NMDAdependent and -independent forms of plasticity in place-field formation.

The properties of the learning rule are suggestive of the heterosynaptic depression/depotentiation of an inactive pathway when postsynaptic activation is driven by a second pathway (Abraham et al. 2007; Levy and Steward 1979; Lynch et al. 1977). We systematically tested our computational model to verify the extent to which the rule operates in agreement with this experimental phenomenon, especially in its more recent characterization (Abraham et al. 2007) as a form of homosynaptic depression "in disguise." In our simulations, successful generation of single place fields could only be accomplished with a faithful implementation of the heterosynaptic depression of silent inputs, unless the presynaptic activity trace was assumed to decay relatively slowly (500 ms to 1 s time constant; Fig. 2F and Supplementary Fig. S6). In the latter case, the resulting slower input integration provides a sufficient window of opportunity for depression to act homosynaptically as the rat moves from a peak to a trough of the grid. Determining the exact biophysical implementation of this function was beyond the scope of the present work, but a plausible candidate might be the binding of glutamate onto NMDA receptors; this permits Ca²⁺ influx into the spine at the occurrence of back-propagating action potentials. In this scenario, the dependence of the direction of plasticity on the level of the presynaptic rate in Eq. 2 would naturally reflect the long-held hypothesis that the level of $[Ca^{2+}]$ in the spine determines whether LTP or LTD is induced at the synapse (Artola and Singer 1993; Bear et al. 1987; Lisman 1989; Shouval et al. 2002). The time constant of the slow components of the decay of the glutamate binding to NMDA receptors (up to 600 ms) (Lester et al. 1990) appears compatible with the parameter range that succeeded in producing place fields in our simulations without proper heterosynaptic depression. Physiological realizations of the mechanism inherent in the postsynaptically gated synaptic rule could conceivably involve other forms of cellular plasticity as well. For instance, the global decrease in membrane excitability that accompanies input-specific LTP observed in CA1 Schaffer collaterals reduces the ability of nonpotentiated inputs to contribute to later postsynaptic firing (Fan et al. 2005; Narayanan and Johnston 2007). Thus this phenomenon-if present at perforant path synapses as wellcould have a functional net effect similar to that of heterosynaptic depression in Eq. 2.

It is possible that models of plasticity or parameter values that escaped our investigation could produce a grid to place transformation as effectively as accomplished by the rule in Eq. 2. Because this rule produces a crude form of synaptic competition, one might conclude that any synaptic rule with a similar property should be able to make place fields from grid inputs. For example, the BCM rule (Bienenstock et al. 1982) can refine the spatial response of a unit taking inputs from boundary-related cells into a place field (Barry and Burgess 2007; Barry et al. 2006; Lever et al. 2002). The BCM rule's rate of synaptic change must be considerably slower than the rate of presentation of input patterns to enable temporal competition between these patterns (Bienenstock et al. 1982; Dayan and Abbott 2001), thus requiring extensive spatial sampling of the recording enclosure. By contrast, our simulations suggest that the rate of synaptic change must be fast to generate fields at a realistic time scale, and some of these fields were formed before the environment was fully explored, as is often experimentally observed (Frank et al. 2004; Hill 1978). The present model does not exclude that an additional learning scheme like the BCM rule could affect the final place-cell response, possibly by eliminating additional fields and/or governing the integration of nongrid inputs into the hippocampal place representation (Barry and Burgess 2007; Barry et al. 2006), but this additional mechanism would need to operate far from the timescale of a foraging session that we were concerned with in this study.

Place field generation

The place fields generated by our simulations tended to be smaller than those typically found in CA3 and CA1 recording experiments. Incorporation of other inputs to the cells (e.g., boundary cells) (Lever et al. 2009; Savelli et al. 2008; Solstad et al. 2008), more complex network dynamics (Maurer et al. 2006), or temporal organization of the inputs (Molter and Yamaguchi 2008) might make the sizes of the fields more similar to those seen in experimental recordings. We also noted that the use of inputs with larger grids was an obvious solution to this problem (Fig. 2*H*, Supplementary Fig. S5). Whereas the dorsoventral topographical organization of grid spacing is now known to a good quantitative degree (Barry et al. 2007; Brun et al. 2008; Hafting et al. 2005), its translation to the dorsoventral axis of the hippocampus via the topographical organization of MEC-hippocampus projections is less quantitatively characterized especially with regard to the direct pathway that projects to CA1. Thus both the smaller and larger sets of grid inputs we tested are in principle compatible with available experimental data if the intent is to model place cells classically recorded from the dorsal hippocampus.

In single-unit experiments, although some place fields are present from the rat's very first visit to a novel location (Hill 1978), other place fields develop over time, and the ensemble code for location in CA1 improves its accuracy over the course of exploration (Frank et al. 2004; Wilson and McNaughton 1993). Our simulations reproduced both rapid and delayed formation of place fields, the latter occurring in a location that had been previously "ignored" by the cell. The diversity in our simulations was due to factors affecting each cell differently, such as the spatial distribution of suprathreshold synaptic activity, which was determined by the specific alignment of input grids to the cell, the relative degree of excitation at the peaks of this distribution; and the duration and order in which the peaks were first visited.

The simulations were initialized with equal weights on all inputs. Assignment of a heterogeneous distribution of initial weights is likely to be more physiologically realistic (de Almeida et al. 2009). This assignment would favor the influence of those inputs that start with higher weights, but because a reduced number of inputs still fails to produce single-peaked fields by nonplastic integration when grid phase is kept random (data not shown, and see Solstad et al. 2006), the nature of the problem would not be essentially affected. For similar reasons, we refrained from modeling the firing rate differences that are experimentally observed across different vertices of the grid (Hafting et al. 2005). In their model based on learning in a competitive network, Rolls and colleagues (2006) found that such differences helped reduce the number of output cells with multiple fields.

Remapping

Although an exploration of hippocampal place-field remapping was not an original goal of this work, the model offers potential mechanistic insights into this prevalent phenomenon. The model produced place fields that were stable within a single behavioral session (with rare cells changing their firing properties within the session, consistent with experimental data; Knierim 2002). However, in any pair of simulations that started from identical conditions-including the same uniform synaptic pattern-and differed only with respect to the initial exploration trajectory or the precise timing of input spikes, the model naturally produced different sets of place fields that mimicked the remapping seen experimentally under numerous conditions (Bostock et al. 1991; Jeffery and Hayman 2004; Knierim 2003; Kubie and Ranck 1983). The ease at which these independent sets of place fields develop from inputs with minimal or null spatial differences is an inevitable byproduct of the fast, nonlinear plasticity dynamics that were necessary for normal place field formation-a "butterfly" effect. The goal of these computational manipulations was to unveil the dramatic potential for the creation of orthogonal output representations inherent in the process of place field generation itself, even before taking into account the spatial changes that can occur in entorhinal inputs during remapping in real experiments (Fyhn et al. 2007). However, because these changes in the entorhinal inputs have not been extensively characterized, it is difficult at present to address how (or whether) the inherent pattern separation potential in the model might be utilized to reproduce remapping phenomena observed in response to environmental alterations in experimental data. Place cells typically fire in the same locations in repeated trips to an unaltered environment (Thompson and Best 1990). The learning rule we examined potentiates a set of synapses onto a place cell that produces a primarily single-peaked place field, while depressing all other synapses. However, when the animal leaves that environment, the same plasticity rule will tend to depress the previously potentiated synapses. This depression could occur, for example, during the exploration of a different environment. In this situation, it is not known whether the firing patterns of all grid cells always maintain the same spatial relationship to each other between environments (Fyhn et al. 2007). If clusters of grid cells can re-orient or shift independently of each other in a novel environment, the formation of a new set of place fields would erase the synaptic patterns previously learned in the original environment. Even if all grid cells remained coherent with each other across environments, the system could localize itself in a new environment to a different location on the two-dimensional sheet far away from the location active in the first environment (see supplementary discussion in Fyhn et al. 2007). In this case, the sets of grid cells coactive in the new environment would be different from those in the first environment, again promoting erasure of the previously learned synaptic patterns. When the animal is subsequently reintroduced into the first environment, even if the grid cells fire in the same locations as in the first exploration, the inherent remapping processes in the model would produce a different set of place fields.

This problem can be viewed as a case of the stability/ plasticity dilemma (Abraham and Robins 2005; Grossberg 2009) exacerbated by the very fast learning that we find necessary for building place fields. It is not limited to the specific plasticity mechanism studied here. Any feed-forward model that implicates plasticity in the formation of place fields from only grid-cell inputs (Molter and Yamaguchi 2008; Rolls et al. 2006; Si and Treves 2009; Ujfalussy et al. 2008) is unlikely to account spontaneously for the memory of many place field maps if investigated with spiking units and spatial sampling that is temporally and behaviorally realistic. One potential solution to this problem is that after a synapse from a particular grid cell to a place cell becomes potentiated, some biochemical marker or other neuromodulatory process might prevent plasticity at that synapse, thus protecting it from erasure by subsequent experience. For example, prion-like proteins have been proposed recently as candidates for longterm memory storage that would be resistant to subsequent reversal (Bailey et al. 2004). If the synapses are not protected, then a teaching mechanism would be necessary to retrain the synapses to fire the postsynaptic cell in the same location as previously learned. This could be done by a system that is sensitive to external landmarks or local cues-perhaps from

the lateral entorhinal cortex-providing input to DG or CA3 that drives the place cell to fire in the same location. Even a weak spatial bias might bootstrap the same set of grid cells and place cells to fire in synchrony as before (especially if coupled with hypothesized attractor dynamics provided by the recurrent collateral system in the CA3 region), driving the system to potentiate the same synapses as in the previous session. In the light of the experimental suggestion that long-term memory of place fields is NMDA dependent but their generation is not (Kentros et al. 1998), such a teaching process should be experimentally distinguishable from the field formation. Although this necessity for constant relearning may seem inefficient, there is experimental evidence that CA1 place fields undergo plasticity (revealed through backward expansion of place field size and location) every time an animal enters a familiar environment after a sufficient delay (<24 h) (Ekstrom et al. 2001; Mehta et al. 1997). Such plasticity in familiar environments is not evident in CA3 (Lee et al. 2004), suggesting that the mossy fiber and recurrent collateral inputs in CA3 may be part of the long-term memory system that allows CA3 cells to fire at the same location in repeated trips to an environment. The Schaffer collateral system may then teach the CA1 cells to fire in the same location each time relative to the incoming MEC input, thus providing the intersession placefield stability observed in most experimental studies.

Comparison with different mechanistic models

Rolls and colleagues (2006) proposed the first mechanistic model for the grid-to-place transformation. The model combined synaptic plasticity in place cells with a partial winnertake-all rule regulating their population activity to introduce competition between them. Plasticity was implemented as Hebbian potentiation of feed-forward projections from grid cells accompanied by fast weight renormalization and was shown necessary to generate place fields. This framework was largely shared by later studies that investigated how additional factors such as phase precession of grid cell firing (Molter and Yamaguchi 2008) or putative inputs from the lateral entorhinal cortex (Si and Treves 2009) could contribute to the properties of the transformation. Our study suggests decoupled roles for plasticity and competition: plasticity alone is sufficient for generating unitary place fields in single cells, whereas network competition may distribute fields more evenly across the environment and reduce the number of active cells. We also provided the first detailed analysis of the nature of the plasticity rule in relation to its success in producing place fields based on a spiking model.

A complementary hypothesis was investigated by de Almeida and colleagues (2009), who showed that the winnertake-all rule can produce multipeaked place fields similar to those recorded from dentate gyrus without requiring synaptic plasticity as long as the initial conditions encompassed a wide range of synaptic weight values. This hypothesis has the advantage of not requiring the inputs to be aligned at one vertex location because competition can in principle nonlinearly amplify even slight differences of postsynaptic activation across space. The stability of the place fields obtained with this mechanism, however, needs to be confirmed with a spiking model using nonphenomenological network competition, realistic spatial sampling, and noninstantaneous kinetics of synaptic excitation. For instance, the total momentary excitatory drive to a place cell firing in its field is a function of the recent trajectory, not only of the current position, due to the slow time course of the excitatory contribution of NMDA channels, which may be greater than the faster AMPA contribution (de Almeida et al. 2009). In our simulations, this effect proved beneficial for a less-than-ideal implementation of heterosynaptic depression (Supplementary Fig. S6), but in a model that resolves competition on the fast timescale of one gamma oscillation as proposed by de Almeida and colleagues (2009), it could cause different sets of winners to emerge across distinct traversals of the same location as a function of the approaching trajectory. Nonetheless, the existence of models based on different physiological mechanisms (plasticity vs. network competition) suggests that the problem of place field formation is still experimentally underconstrained.

Hayman and Jeffery (2008) have suggested that contextual inputs from LEC and grid inputs from MEC may be clustered onto individual branches of a DG granule cell dendritic tree such that the contextual input determines which branch is active and thereby produces context-specific firing of the cell. Ujfalussy and colleagues (2009) have investigated the same notion in a mechanistic model employing a plasticity rulewith similar heterosynaptic depression properties as in our study-applied separately to each postsynaptic dendritic branch rather than to the postsynaptic soma as in our model. Each branch received both visual and grid-cell inputs and could create a place field in one environment. The cell summated the different dendritic contributions thus giving rise to the independently modifiable, multiple firing fields typical of DG granule cells (Leutgeb et al. 2007). This may be a fruitful line of investigation into determining both context specificity of place fields as well as place-field stability in repeated trips to the same environment—a problem that we discussed in the previous section. A disadvantage of applying plasticity in a compartmentalized fashion, however, was that heterosynaptic depression could not work across branches. Manual intervention was hence necessary to silence those branches that did not reach sufficient activation for plasticity to depress their inputs (Ujfalussy et al. 2009).

Experimental tests

Some implications of the present model provide suggestions for future experimental research. The core prediction can be asserted as follows. While the field of a place cell is developing in a novel environment, every grid cell input that has an intervertex space at the place-field location must experience a fast reduction of its excitatory influence on the place cell. The reduction will need to be especially pronounced if the input at hand is not already depotentiated or depressed prior to the exposure to the new environment. Ideally, this prediction would be tested with intracellular recording of a single hippocampal place cell (Epsztein et al. 2010; Harvey et al. 2009; Lee et al. 2009), and (some of) its input grid cells during the animal's foraging session, by gauging the EPSPs provoked independently by each input via microstimulation before and after the session. This would be an extremely difficult experiment with current technology, however. Alternatively, the prediction could be tested in vitro by the quasi-simultaneous photostimulation of multiple dendritic sites (Gasparini and

GRID CELL TO PLACE CELL TRANSFORMATION

Magee 2006; Losonczy and Magee 2006) in the stratum lacunosum-moleculare (where the perforant path entorhinal inputs make their connections) of a proximal CA1 cell the soma potential of which is also recorded. The stimulation could be driven by the first stage of our simulations-which produces spike trains according to modeled grid cells and the tracked position of the rat. The spikes produced by the cell could be used with this positional data to build a virtual spatial rate map (based on the positions of the rat that correspond to the output spikes in the simulation). Assuming a virtual place field is thus produced, the EPSPs resulting from periodic independent stimulation of the different inputs could be used to monitor whether each input's excitatory influence on the postsynaptic cell evolves in accordance with the prediction. This experiment amounts to replacing the integrate-and-fire "place cell" unit in our simulations with a real cell in vitro. With the caveat that the spatial resolution of photostimulation does not currently attain control of single synapses (Gasparini and Magee 2006; Losonczy and Magee 2006), the results could provide preliminary evidence in favor or against the model. Note that strong evidence consistent with the model would require that the alterations in synaptic efficacy occur rapidly to reflect the requirement for fast plasticity shown by the model. Similarly, plasticity should occur when the postsynaptic cell fires, and the sign of plasticity should depend on the activity level of the presynaptic input. These predictions are specific to the current model rather than to a generic Hebbian plasticity model.

The emergent properties of the model have potential experimental implications as well. Some studies of place cells have demonstrated an over-representation of the boundaries of an environment compared with the center (Hetherington and Shapiro 1997; Wiener et al. 1989; but see Muller et al. 1987). We noted that minor changes in the initial exploration trajectory are sufficient to modify the location where place cells form their fields, and the locations visited the most in this early phase are more likely to eventually host place fields. This result suggests that one contributing factor to this boundary overrepresentation may be the tendency of rodents to first explore the edges of a novel environment before venturing out into the center (Drai and Golani 2001; Drai et al. 2001; Tchernichovski and Benjamini 1998; Tchernichovski et al. 1998), in addition to the possibility that the over-representation is partly inherited from boundary cells in MEC (Savelli et al. 2008; Solstad et al. 2008) or re-entrant subiculum outputs (Lever et al. 2009). Experimental manipulations could be performed to systematically constrain the early rat's trajectory (Hayman et al. 2008) and test the influence of the early trajectory on the inhomogeneities of field distributions. Moreover, in our simulations, we found that feedback inhibition considerably reduces the intensity of this influence. The prediction follows that if it were possible to inactivate feedback interneurons of the hippocampus (like OLM cells) (Freund and Buzsaki 1996) selectively without otherwise compromising normal place field formation, the global spatial distribution of place fields would become less uniform and possibly more concentrated in the areas first explored by the rat.

A C K N O W L E D G M E N T S

We thank M. Mauk and H. Shouval for helpful discussions, K. Zhang and S. Deshmukh for comments on the manuscript, J. Siegel for the rat trajectory data, and G. Rao for assistance with manuscript preparation.

G R A N T S

This work was supported by National Institute of Neurological Disorders and Stroke Grant P01 NS-038310 and by a postdoctoral fellowship to F. Savelli from the International Human Frontier Science Organization (LT00683/2006-C).

DISCLOSURES

No conflicts of interest are declared by the authors.

REFERENCES

- Abraham WC, Logan B, Wolff A, Benuskova L. "Heterosynaptic" LTD in the dentate gyrus of anesthetized rat requires homosynaptic activity. J Neurophysiol 98: 1048–1051, 2007.
- Abraham WC, Robins A. Memory retention-the synaptic stability versus plasticity dilemma. *Trends Neurosci* 28: 73–78, 2005.
- Artola A, Singer W. Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci* 16: 480–487, 1993.
- Bailey CH, Kandel ER, Si K. The persistence of long-term memory: a molecular approach to self-sustaining changes in learning-induced synaptic growth. *Neuron* 44: 49–57, 2004.
- Barry C, Burgess N. Learning in a geometric model of place cell firing. *Hippocampus* 17: 786–800, 2007.
- Barry C, Hayman R, Burgess N, Jeffery KJ. Experience-dependent rescaling of entorhinal grids. *Nat Neurosci* 10: 682–684, 2007.
- Barry C, Lever C, Hayman R, Hartley T, Burton S, O'Keefe J, Jeffery K, Burgess N. The boundary vector cell model of place cell firing and spatial memory. *Rev Neurosci* 17: 71–97, 2006.
- Bear MF, Cooper LN, Ebner FF. A physiological basis for a theory of synapse modification. *Science* 237: 42–48, 1987.
- **Bienenstock EL, Cooper LN, Munro PW.** Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2: 32–48, 1982.
- Blair HT, Welday AC, Zhang K. Scale-invariant memory representations emerge from moire interference between grid fields that produce theta oscillations: a computational model. *J Neurosci* 27: 3211–3229, 2007.
- Bliss T, Collingridge G, Morris R. Synaptic plasticity in the hippocampus. In: *The Hippocampus Book*, edited by Andersen P, Morris R, Amaral D, Bliss T, O'Keefe J. New York: Oxford, 2007, p. 343–474.
- Bostock E, Muller RU, Kubie JL. Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1: 193–205, 1991.
- Brun VH, Otnass MK, Molden S, Steffenach HA, Witter MP, Moser MB, Moser EI. Place cells and place recognition maintained by direct entorhinalhippocampal circuitry. *Science* 296: 2243–2246, 2002.
- Brun VH, Solstad T, Kjelstrup KB, Fyhn M, Witter MP, Moser EI, Moser MB. Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus* 18: 1200–1212, 2008.
- Chawla MK, Guzowski JF, Ramirez-Amaya V, Lipa P, Hoffman KL, Marriott LK, Worley PF, McNaughton BL, Barnes CA. Sparse, environmentally selective expression of Arc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus* 15: 579– 586, 2005.
- **Colquhoun D, Jonas P, Sakmann B.** Action of brief pulses of glutamate on AMPA/kainate receptors in patches from different neurons of rat hippocampal slices. *J Physiol* 458: 261–287, 1992.
- Dayan P, Abbott LF. Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems. Cambridge, MA: The MIT Press. 2001.
- **Dayan P, Sejnowski TJ.** The variance of covariance rules for associative matrix memories and reinforcement learning. *Neural Comput* 5: 205–209, 1993.
- de Almeida L, Idiart M, Lisman JE. The input-output transformation of the hippocampal granule cells: from grid cells to place fields. J Neurosci 29: 7504–7512, 2009.
- Drai D, Golani I. SEE: a tool for the visualization and analysis of rodent exploratory behavior. *Neurosci Biobehav Rev* 25: 409-426, 2001.
- Drai D, Kafkafi N, Benjamini Y, Elmer G, Golani I. Rats and mice share common ethologically relevant parameters of exploratory behavior. *Behav Brain Res* 125: 133–140, 2001.
- Eichenbaum H, Yonelinas AP, Ranganath C. The medial temporal lobe and recognition memory. *Annu Rev Neurosci* 30: 123–152, 2007.

- Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA. NMDA receptor antagonism blocks experience-dependent expansion of hippocampal "place fields." *Neuron* 31: 631–638, 2001.
- **Epsztein J, Lee AK, Chorev E, Brecht M.** Impact of spikelets on hippocampal CA1 pyramidal cell activity during spatial exploration. *Science* 327: 474–477, 2010.
- Fan Y, Fricker D, Brager DH, Chen X, Lu HC, Chitwood RA, Johnston D. Activity-dependent decrease of excitability in rat hippocampal neurons through increases in I(h). *Nat Neurosci* 8: 1542–1551, 2005.
- Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, Lytton WW, Ludvig N. Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. J Neurosci 28: 11250–11262, 2008.
- Frank LM, Brown EN, Stanley GB. Hippocampal and cortical place cell plasticity: implications for episodic memory. *Hippocampus* 16: 775–784, 2006.
- Frank LM, Stanley GB, Brown EN. Hippocampal plasticity across multiple days of exposure to novel environments. J Neurosci 24: 7681–7689, 2004.
- Franzius M, Vollgraf R, Wiskott L. From grids to places. J Comput Neurosci 22: 297–299, 2007.
- Freund TF, Buzsaki G. Interneurons of the hippocampus. *Hippocampus* 6: 347–470, 1996.
- Fyhn M, Hafting T, Treves A, Moser MB, Moser EI. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446: 190–194, 2007.
- **Gasparini S, Magee JC.** State-dependent dendritic computation in hippocampal CA1 pyramidal neurons. *J Neurosci* 26: 2088–2100, 2006.
- Gerstner W, Kistler W. Spiking Neuron Models: Single Neurons, Populations, Plasticity. Cambridge, UK: Cambridge Univ. Press, 2002.
- **Grossberg S.** Beta oscillations and hippocampal place cell learning during exploration of novel environments. *Hippocampus* 19: 881–885, 2009.
- Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. Microstructure of a spatial map in the entorhinal cortex. *Nature* 436: 801–806, 2005.
- Harvey CD, Collman F, Dombeck DA, Tank DW. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* 461: 941–946, 2009.
- Hayman RM, Donnett JG, Jeffery KJ. The fuzzy-boundary arena–a method for constraining an animal's range in spatial experiments without using walls. J Neurosci Methods 167: 184–190, 2008.
- Hayman RM, Jeffery KJ. How heterogeneous place cell responding arises from homogeneous grids–a contextual gating hypothesis. *Hippocampus* 18: 1301–1313, 2008.
- Henriksen EJ, Moser MB, Moser EI. Megaspace recording from hippocampal place cells and entorhinal grid cells. Soc Neurosci Abstr 101.8, 2009.
- Hetherington PA, Shapiro ML. Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behav Neurosci* 111: 20–34, 1997.
- Hill AJ. First occurrence of hippocampal spatial firing in a new environment. *Exp Neurol* 62: 282–297, 1978.
- Jeffery KJ, Hayman R. Plasticity of the hippocampal place cell representation. *Rev Neurosci* 15: 309–331, 2004.
- Jung MW, McNaughton BL. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* 3: 165–182, 1993.
- Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, Muller RV. Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280: 2121–2126, 1998.
- Knierim JJ. Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. J Neurosci 22: 6254–6264, 2002.
- Knierim J. Hippocampal remapping: Implications for spatial learning and navigation. In: *The Neurobiology of Spatial Behaviour*, edited by Jeffery K. Oxford, UK: Oxford, Univ. Press, 2003, p. 226–239.
- Knierim JJ, Lee I, Hargreaves EL. Hippocampal place cells: parallel input streams, subregional processing, and implications for episodic memory. *Hippocampus* 16: 755–764, 2006.
- Knierim JJ, McNaughton BL. Hippocampal place-cell firing during movement in three-dimensional space. J Neurophysiol 85: 105–116, 2001.
- Kropff E, Treves A. The emergence of grid cells: intelligent design or just adaptation? *Hippocampus* 18: 1256–1269, 2008.
- Kubie J, Ranck J Jr. Sensory-behavioral correlates in individual hippocampus neurons in three situations: space and context. In: *Neurobiology of the Hippocampus*, edited by Seifert W. London: Academic, 1983, p. 433–447.
- Lee AK, Epsztein J, Brecht M. Head-anchored whole-cell recordings in freely moving rats. *Nat Protoc* 4: 385–392, 2009.

- Lee I, Rao G, Knierim JJ. A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. *Neuron* 42: 803–815, 2004.
- Lester RA, Clements JD, Westbrook GL, Jahr CE. Channel kinetics determine the time course of NMDA receptor-mediated synaptic currents. *Nature* 346: 565–567, 1990.
- Leutgeb JK, Leutgeb S, Moser MB, Moser EI. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315: 961–966, 2007.
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309: 619–623, 2005.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305: 1295–1298, 2004.
- Lever C, Burgess N, Cacucci F, Hartley T, O'Keefe J. What can the hippocampal representation of environmental geometry tell us about Hebbian learning? *Biol Cybern* 87: 356–372, 2002.
- Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N. Boundary vector cells in the subiculum of the hippocampal formation. J Neurosci 29: 9771–9777, 2009.
- Levy WB, Steward O. Synapses as associative memory elements in the hippocampal formation. *Brain Res* 175: 233–245, 1979.
- Lisman J. A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc Natl Acad Sci USA* 86: 9574–9578, 1989.
- Losonczy A, Magee JC. Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. *Neuron* 50: 291–307, 2006.
- Lynch GS, Dunwiddie T, Gribkoff V. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature* 266: 737–739, 1977.
- Manns JR, Eichenbaum H. Evolution of declarative memory. *Hippocampus* 16: 795–808, 2006.
- Maurer AP, Cowen SL, Burke SN, Barnes CA, McNaughton BL. Organization of hippocampal cell assemblies based on theta phase precession. *Hippocampus* 16: 785–794, 2006.
- McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA. Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell* 87: 1339–1349, 1996.
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB. Path integration and the neural basis of the "cognitive map." *Nat Rev Neurosci* 7: 663–678, 2006.
- Mehta MR, Barnes CA, McNaughton BL. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94: 8918–8921, 1997.
- Molter C, Yamaguchi Y. Entorhinal theta phase precession sculpts dentate gyrus place fields. *Hippocampus* 18: 919–930, 2008.
- Muller RU, Kubie JL. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7: 1951–1968, 1987.
- Muller RU, Kubie JL, Ranck JB, Jr. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J Neurosci 7: 1935–1950, 1987.

Nakazawa K, McHugh TJ, Wilson MA, Tonegawa S. NMDA receptors, place cells and hippocampal spatial memory. *Nat Rev Neurosci* 5: 361–372, 2004.

- Narayanan R, Johnston D. Long-term potentiation in rat hippocampal neurons is accompanied by spatially widespread changes in intrinsic oscillatory dynamics and excitability. *Neuron* 56: 1061–1075, 2007.
- **O'Keefe J, Burgess N.** Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15: 853–866, 2005.
- **O'Keefe J, Conway DH.** Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* 31: 573–590, 1978.
- **O'Keefe J, Dostrovsky J.** The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34: 171–175, 1971.
- **O'Keefe J, Nadel L.** *The Hippocampus as a Cognitive Map.* Oxford, UK: Clarendon, 1978.
- Rolls ET, Stringer SM, Elliot T. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Network* 17: 447–465, 2006.
- Savelli F, Yoganarasimha D, Knierim JJ. Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus* 18: 1270–1282, 2008.
- Sejnowski TJ. Storing covariance with nonlinearly interacting neurons. J Math Biol 4: 303–321, 1977.
- Shapiro ML, Eichenbaum H. Hippocampus as a memory map: synaptic plasticity and memory encoding by hippocampal neurons. *Hippocampus* 9: 365–384, 1999.

3182

3183

- Shouval HZ, Bear MF, Cooper LN. A unified model of NMDA receptordependent bidirectional synaptic plasticity. *Proc Natl Acad Sci USA* 99: 10831–10836, 2002.
- Si B, Treves A. The role of competitive learning in the generation of DG fields from EC inputs. *Cogn Neurodyn* 3: 177–187, 2009.
- Skaggs WE, McNaughton BL. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J Neurosci 18: 8455–8466, 1998.
- Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI. Representation of geometric borders in the entorhinal cortex. *Science* 322: 1865–1868, 2008.
- Solstad T, Moser EI, Einevoll GT. From grid cells to place cells: a mathematical model. *Hippocampus* 16: 1026–1031, 2006.
- Squire LR, Stark CE, Clark RE. The medial temporal lobe. Annu Rev Neurosci 27: 279–306, 2004.
- Tanila H, Shapiro ML, Eichenbaum H. Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus* 7: 613–623, 1997.
- Tchernichovski O, Benjamini Y. The dynamics of long-term exploration in the rat. II. An analytical model of the kinematic structure of rat exploratory behavior. *Biol Cybern* 78: 433–440, 1998.
- Tchernichovski O, Benjamini Y, Golani I. The dynamics of long-term exploration in the rat. I. A phase-plane analysis of the relationship between location and velocity. *Biol Cybern* 78: 423–432, 1998.

- **Thompson LT, Best PJ.** 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, 1990.
- Ujfalussy B, Erős P, Somogyvári Z, Kiss T. Episodes in space: a modelling study of hippocampal place representation. In: *From Animal to Animats 10*, edited by Asada M, Hallam JCT, Meyer J-A, Tani J. Berlin: Springer, 2008, p. 123–136.
- Ujfalussy B, Kiss T, Erdi P. Parallel computational subunits in dentate granule cells generate multiple place fields. *PLoS Comput Biol* 5: e1000500, 2009.
- Vargha-Khadem F, Gadian DG, Watkins KE, Connelly A, Van PW, Mishkin M. Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* 277: 376–380, 1997.
- Vazdarjanova A, Guzowski JF. Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. J Neurosci 24: 6489–6496, 2004.
- Wiener SI, Paul CA, Eichenbaum H. Spatial and behavioral correlates of hippocampal neuronal activity. *J Neurosci* 9: 2737–2763, 1989.
- Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058, 1993.