Modular Realignment of Entorhinal Grid Cell Activity as a Basis for Hippocampal Remapping

Joseph D. Monaco¹ and Larry F. Abbott²

¹Zanvyl Krieger Mind/Brain Institute, Department of Neuroscience, Johns Hopkins University, Baltimore, Maryland 21218, and ²Departments of Neuroscience, and Physiology and Cellular Biophysics, Columbia University College of Physicians and Surgeons, New York, New York 10032

Hippocampal place fields, the local regions of activity recorded from place cells in exploring rodents, can undergo large changes in relative location during remapping. This process would appear to require some form of modulated global input. Grid-cell responses recorded from layer II of medial entorhinal cortex in rats have been observed to realign concurrently with hippocampal remapping, making them a candidate input source. However, this realignment occurs coherently across colocalized ensembles of grid cells (Fyhn et al., 2007). The hypothesized entorhinal contribution to remapping depends on whether this coherence extends to all grid cells, which is currently unknown. We study whether dividing grid cells into small numbers of independently realigning modules can both account for this localized coherence and allow for hippocampal remapping. To do this, we construct a model in which place-cell responses arise from network competition mediated by global inhibition. We show that these simulated responses approximate the sparsity and spatial specificity of hippocampal activity while fully representing a virtual environment without learning. Place-field locations and the set of active place cells in one environment can be independently rearranged by changes to the underlying grid-cell inputs. We introduce new measures of remapping to assess the effectiveness of grid-cell modularity and to compare shift realignments with other geometric transformations of grid-cell responses. Complete hippocampal remapping is possible with a small number of shifting grid modules, indicating that entorhinal realignment may be able to generate place-field randomization despite substantial coherence.

Introduction

The locations of rodent hippocampal place fields (O'Keefe and Dostrovsky, 1971) can be randomly rearranged from one environment to the next during a process known as remapping (Bostock et al., 1991; Leutgeb et al., 2005; Wills et al., 2005). The freedom with which place fields remap suggests a link between the local spatial representations found in hippocampus and the global representation of grid cells (Hafting et al., 2005; Fyhn et al., 2008). Grid cells in medial entorhinal cortex (MEC) project to hippocampus (Witter, 2007b), and their periodic spatial responses realign during remapping (Fyhn et al., 2007). These shifts provide an attractive candidate mechanism for remapping in which grid-cell inputs cause large displacements in place-field locations. However, the realignment of colocalized grid cells during remapping is highly coherent (Fyhn et al., 2007). This apparent uniformity must be reconciled with the random reassignment of place-field locations during remapping.

In light of experimental evidence for modularity in MEC (Witter and Moser, 2006; Walling et al., 2006), including recent observations of modularity in grid-cell geometry (Stensland et al., 2010), we study whether grid-cell modules, within which grids realign coherently, can resolve this conundrum. Previous discussions (O'Keefe and Burgess, 2005; McNaughton et al., 2006) and models (Fuhs and Touretzky, 2006; Hayman and Jeffery, 2008) have considered place-cell remapping through independent realignment of grid-cell inputs, as well as partial remapping produced by less complete grid realignments (Fuhs and Touretzky, 2006). Our particular focus is on the following: (1) determining the number of independently realigning modules needed to produce statistically complete place-cell remapping; (2) studying the impact of assigning grid cells to modules either randomly or on the basis of their grid spacing (spatial-frequency-based modules); and (3) comparing the efficacy of different forms of grid-cell realignment, including shifts, rotations, enlargement of grid scale (Barry et al., 2009), and changes in grid ellipticity (Barry et al., 2007; Stensland et al., 2010). The second focus is inspired by the topographic organization of grid spacing along the dorsoventral axis of MEC (Hafting et al., 2005; Kjelstrup et al., 2008) and evidence for clustering of grid scales (Barry et al., 2007). In summary, our investigations provide a theoretical interpretation of clustering and modularity within MEC.

Our results are based on a model that transforms a periodic grid representation of space into one matching the sparse activity and high spatial specificity observed in hippocampus (O'Keefe and Dostrovsky, 1971; Wilson and McNaughton, 1993; Gu-

Received March 21, 2011; revised April 26, 2011; accepted May 4, 2011.

Author contributions: J.D.M. and L.F.A. designed research; J.D.M. performed research; J.D.M. analyzed data; J.D.M. and L.F.A. wrote the paper.

This work was supported by NIH Grant MH58754 and an NIH Director's Pioneer Award, part of the NIH Roadmap for Medical Research, through Grant 5-DP1-OD114-02 (L.F.A.) and NIH Grant P01 NS-038310 (to James J. Knierim, who supported J.D.M. during the writing of the manuscript). We thank Eric R. Kandel, Pablo Jercog, Francesco Savelli, Caswell Barry, Grace Hwang, and the anonymous reviewers for helpful comments and discussion.

The authors declare no competing financial interests.

Correspondence should be addressed to Joseph D. Monaco, Johns Hopkins School of Medicine, 720 Rutland Avenue, Traylor 407, Baltimore, MD 21205-2109. E-mail: jmonaco@jhu.edu.

DOI:10.1523/JNEUROSCI.1433-11.2011

Copyright © 2011 the authors 0270-6474/11/319414-12\$15.00/0

zowski et al., 1999). The model is meant to reproduce the first-pass activity of place cells in an unfamiliar environment (Hill, 1978; Frank et al., 2004; Karlsson and Frank, 2008) by combining fixed grid-to-place connectivity with global feedback inhibition among place cells (Buzsáki et al., 2007; Pelletier and Lacaille, 2008). This initial place-cell activity may determine the spatial representations that are ultimately learned with continued exploration (Savelli and Knierim, 2010). The simulated responses here based on randomly aligned grid inputs and uniformly distributed synaptic weights allow for flexible and independent remapping of place-field locations.

Materials and Methods

Place-network model. A simulated place network is defined by the grid-to-place weight matrix $\mathbf{W}$ that is created at the beginning of every simulation. For connectivity C , this matrix is constructed from a 1000-component reference vector that has $1000(1 - C)$ components set to zero, and the remaining $1000C$ components randomly sampled uniformly over the range $[0, 1]$. Each row of $\mathbf{W}$ is then set to a randomly shuffled permutation of this reference vector. Having place units with identical, but shuffled, afferent weights avoids the contribution of sampling effects to the heterogeneity of place-unit activity. The focus here is to allow the grid configuration and place-unit competition to drive the diversity of responses across the network. The model is integrated using the fourth-order Runge–Kutta algorithm over discrete 5 ms time steps. Place units are initialized to zero activity.

To determine the active place fields for each place unit, we find contiguous areas of activity $>20\%$ of the peak rate of the unit (Muller and Kubie, 1989). Then, putative fields with rates $>20\%$ of the population maximum rate and field size of at least 50 cm^2 are counted as active place fields.

Grid population model. Simulated grid-cell responses are constructed from interference patterns of three two-dimensional sinusoidal gratings oriented 60° apart. Similar formulations for simulating grid response fields have been used in previous theoretical and computational studies of grid cell function (Fuhs and Touretzky, 2006; Solstad et al., 2006; Blair et al., 2007; de Almeida et al., 2009). An exponential nonlinearity is used to shape the profile of the resulting subfields to be approximately Gaussian and to qualitatively match the shape of the firing fields of observed grid-cell responses.

The grid responses are normalized to the range $[0, 1]$. Each grid is characterized by three parameters, s , ψ , and ϕ , which determine the grid spacing, orientation, and spatial phase (relative offset of the peak nearest the midpoint of the environment), respectively. For a given position x in the environment, grid activity for these parameters is given by the following:

$$g(x) = \frac{1}{R[3]} \left[\sum_{i=0}^2 \cos \left(\frac{4\pi}{\sqrt{3}s} u(\theta_i - \psi) \cdot (x - \phi) \right) \right], \quad (1)$$

where $u(\theta) = \langle \cos(\theta), \sin(\theta) \rangle$ is a unit vector in the direction θ . The array of angles $\theta \equiv \{-\pi/3, 0, \pi/3\}$ determines the relative orientations of the component sinusoids, and $R[I] = [\exp(0.25I) - 0.75]_+$, where $[\cdot]_+$ specifies half-wave rectification that sets negative values to zero. The expression in Equation 1 defines an individual component of the vector $\mathbf{g}(x)$ in Equation 3.

Parameter search. We implemented a genetic algorithm to search the parameter space of fan-in connectivity C , inhibitory strength J , and threshold λ (Eq. 3). This was necessary because of nonlinear interactions among these variables (see Fig. 4). The fitness function was defined as the inverse variance from the target values of spatial map properties described in Results. Each generation consisted of 512 simulations of random grid/place-network pairs with parameters sampled from contracting hypercubes centered on the last-generation winners. A coarse-grained search (keeping the top 25% of winners) was performed starting from $C \in \{0.1, 0.9\}$, $J \in \{150, 2 \times 10^4\}$, and $\lambda \in \{0, 1.1 \times 10^3\}$ that converged in six generations. Based on those results, a fine-grained

search (keeping the top 10% of winners) was performed starting from $C \in \{0.2, 0.35\}$, $J \in \{1.8 \times 10^3, 4.0 \times 10^3\}$, and $\lambda \in \{0.0, 6.5\}$ that converged in five generations. Final parameter values (Eq. 3) are an average of neighboring winners from both searches.

Spatial map simulation. We simulate a 1 m^2 environment sampled over a 100×100 element array, so that each pixel represents 1 cm^2 . Spatial maps were constructed by setting x to the midpoint location of pixels in the environment and allowing the activity to converge. We found high correlations between this rasterization with fixed input and continuous-time simulations of the same environment using a naturalistic trajectory based on a random walk (see Notes). We use a checkerboard pattern consisting of every other pixel in the 1 m^2 area to improve the efficiency of our simulations; this does not significantly decrease accuracy because the scale of spatial activity is significantly larger than a single pixel. The first pixel is clamped for 10τ seconds, and all subsequent pixels are clamped for 5τ seconds, where τ is the rate time constant (Eq. 3). These dwell times are sufficient for place-unit activity to converge. Responses for unsampled pixels are interpolated as the average of their sampled neighbors. This interpolation can yield aliasing artifacts at the edges of place fields so, to mitigate this, we median filtered the rate maps with a 3×3 pixel kernel.

Autocorrelograms. Population autocorrelograms are computed by two-dimensional Fourier domain multiplication of a population rate matrix with its complex conjugate. The correlogram is obtained by taking the real part of the inverse Fourier transform of this product. We normalize the correlograms by dividing by the peak correlation.

Realignment. We tested three manipulations of grid responses, which we refer to as different types of realignment, to assess remapping. Unless otherwise specified, realignment parameters are randomly sampled as described here. Shift realignment is the spatial displacement observed by Fyhn et al. (2007) to be concurrent with complete remapping in hippocampus. For simulations of remapping with random modules, shift realignment is specified as a translation of the grid within the plane for a uniform random distance in the range 9–45 cm (or 10–50% of the maximum possible grid spacing) along a uniform random direction sampled from 0 – 360° . For simulations of remapping with frequency modules, distances are drawn from a range of 10–50% of the field spacing of the lowest spatial frequency grid in each module. Translations are applied equally to all grids within a module such that the location of any grid peak is shifted along the sampled direction by the sampled distance.

We also tested two other grid manipulations. First, changing grid ellipticity corresponds to a squeeze mapping in the plane, which is an equi-areal transform that preserves field size by magnifying the plane along one “longitudinal” axis and contracting proportionally along the transverse axis. This means that an ellipticity parameter l yields a primary elliptic flattening of $f = 2l/(1 + l)$. The ellipticity parameter is randomly drawn uniformly over the range 0.0 – 0.2 (or up to 20% magnification and contraction). The orientation of the longitudinal axis is drawn uniformly over the range $[-\pi/2, \pi/2]$. Second, grid rescaling is a uniform magnification of the plane. The scale is uniformly drawn over the range 1.0 – 1.2 (or up to 20% magnification).

Remapping measures. We present two measures of remapping between any two spatial maps. First, to quantify spatial differences, we determine the set of place units that are active in both maps. Then, for both maps, we compute all pairwise distances between peak firing locations. Spatial remapping is then calculated as 1 minus the Pearson’s product-moment correlation coefficient of the pairwise distances of one map with those of the other. Thus, spatial remapping is 0 for identical maps, ~ 1 for unrelated maps, and is sensitive to differences in the relative pairwise structure of the map. Second, to quantify the degree and randomness of turnover in the active subset of place units, we construct the three-element array α consisting of the proportions of place units active in neither, one, and both of the spatial maps. Then, activity turnover $T[\alpha]$ is computed by comparing the root-mean-square differences (RMSD) between α and two similar reference arrays:

$$T[\alpha] = 1 - \text{RMSD}(\alpha, \beta) / \text{RMSD}(\alpha_0, \beta), \quad (2)$$

where $\beta \equiv \{s^2, 2s(1 - s), (1 - s)^2\}$ is the expected activity array given random recruitment of active units, $\alpha_0 \equiv \{s, 0, 1 - s\}$ is expected given no

Table 1. Spatial map statistics

Type	Value	Mean $\pm$ 95% CI	SD
Map	Sparsity	$0.614 \pm 7.18 \times 10^{-3}$	0.0207
	Coverage	$0.988 \pm 1.82 \times 10^{-3}$	5.25×10^{-3}
	Representation	4.51 ± 0.0897	0.259
	Maximum rate	0.925 ± 0.0111	0.0319
Units	Number of fields	1.38 ± 0.0155	0.622
	Coverage	$0.0234 \pm 3.04 \times 10^{-4}$	0.0122
	Maximum rate	$0.418 \pm 4.81 \times 10^{-3}$	0.193
Fields	Area (cm ²)	169 ± 1.56	73.7
	Diameter (cm) ^a	14.3 ± 0.0665	3.13
	Peak rate	$0.399 \pm 3.95 \times 10^{-3}$	0.186
	Average rate	$0.208 \pm 2.35 \times 10^{-3}$	0.111

"Map" values are computed across sampled maps, "Units" across active place units in sampled maps, and "Fields" across active place fields in sampled maps. Mean $\pm$ 95% confidence interval (CI) and SD describe a sample set of 32 random grid/place-network pairs.

^aField diameter is computed from area assuming field circularity.

turnover, and $s = 0.614$ is the average network sparsity (Table 1). Thus, $T[\alpha] \sim 0$ indicates no change to the active subset and $T[\alpha] \sim 1$ indicates a fully randomized active subset. For comparison with these two measures, we computed population vector (PV) decorrelation as 1 minus the element-by-element Pearson's correlation of the original population rate matrix with the remapped population rate matrix, as has been used previously (Leutgeb et al., 2005, 2007; Wills et al., 2005).

Smooth densities. To visualize some statistical distributions (see Fig. 2), we created smoothed probability density functions (PDFs). This method was adapted from Karlsson and Frank (2008). For a given distribution, we computed a fine-grained cumulative histogram (1000 bins), extended its endpoints to reduce boundary effects, and convolved it with a Gaussian kernel for smoothing. To derive the PDF, we computed the differential of the smoothed cumulative data, cropped it to the original data range, and normalized the resulting densities to the trapezoidal integral. The SD and width of the smoothing kernel were 5 and 50%, respectively, of the data range.

Activity-dependent plasticity and recurrent excitation. To simulate activity-dependent plasticity of grid-to-place synaptic inputs, we implemented the weight changes $\Delta W_{ij} = \varepsilon g_j(x_i)$, where $\varepsilon = 0.01$ is the effective learning rate, g_j is the response of grid cell j , and x_i is the peak firing location of place unit i . We renormalized the weight matrix $W' = W + \Delta W$ so that the magnitude of synaptic gains remained the same for each place unit.

To simulate recurrent excitatory connections between place cells, we constructed a place-to-place weight matrix $P_{ij} = o_{ij}/\beta_{ij}$, where o_{ij} is the overlapping area between the primary place fields of units i and j , and β_{ij} is the smaller of the two place-field areas. Starting with the responses following the plasticity described in the previous paragraph, we applied the firing-rate change $\Delta r = rP$ twice (to include feedforward and feedback activity) to the responses at each location.

Software. We developed all modeling and analysis code as a package of custom Python libraries. These libraries extensively use NumPy for its ndarray implementation of numerical arrays and array operations. Plots and graphs were created from simulation and analysis output with Matplotlib and saved in the vector-based Portable Document Format. Two-dimensional arrays were first converted to red–green–blue image data using Matplotlib color maps and then saved in the lossless Portable Network Graphics format using the Python Imaging Library.

Results

The place-cell network model

We constructed a model of hippocampal spatial map formation in which grid inputs drive a recurrently inhibited network of nonlinear place units (Fig. 1A). The responses of 500 place units receiving input from $N = 1000$ grids are described by a vector r that obeys the time-evolution equation

$$\tau \dot{r} = -r + [\tanh(\alpha Wg(x) - J(r) - \lambda)]_+, \quad (3)$$

where $[\cdot]_+$ indicates rectification, $\tau = 50$ ms is the rate time constant, $\alpha = 100/(NC)$ normalizes the strength of the grid input, J sets the strength of inhibition, and λ is the threshold (values given below). Note that place-unit output is restricted to the interval $[0, 1]$, so we refer to the responses described here as a normalized firing rate. The first term in the argument of the tanh is the grid input: $g(x)$ is a 1000-component vector describing grid responses for the current location x in the environment, and W is a matrix describing the connections from the grid inputs to the place units. W is constructed randomly (see Materials and Methods) so that each place unit receives input from a fraction C of grid cells (value given below). The second term in the argument of the tanh, where $\langle r \rangle$ is the population average of place-unit activity, defines the global inhibition. Thus, we model this inhibition as a global feedback interaction for the place-unit population rather than by including inhibitory interneurons explicitly.

Grid responses, $g(x)$, are modeled phenomenologically as a regular hexagonal grid that tessellates the plane of the simulated environments (Fig. 1B), emulating MEC grid-cell activity (Hafting et al., 2005). Although we use an empirical method based on oscillatory interference to construct grid activity (see Materials and Methods), our results here do not depend on whether grid-cell activity derives from oscillation-based or attractor-based mechanisms (Fuhs and Touretzky, 2006; Giocomo and Hasselmo, 2008; Burak and Fiete, 2009). The spatial structure of each simulated grid depends on three parameters: field spacing, spatial phase, and orientation. For each grid, the spacing is drawn randomly, independently, and uniformly from 30 to 90 cm, which represents the severalfold range of spatial scales observed for grid cells found at the dorsal pole of MEC (Hafting et al., 2005). Similarly, the spatial phase for each grid is set independently to a random point within an origin-centered circle of a diameter equal to half of the grid spacing. This corresponds to the observation that even colocalized grid cells exhibit unrelated spatial phases (Hafting et al., 2005). Grid orientation is also drawn randomly and uniformly over the range $[0, \pi/3]$ (i.e., full angular sampling attributable to the sixfold rotational symmetry of the grids), but the same orientation is used for all of the grids (Hafting et al., 2005; Sargolini et al., 2006; Fyhn et al., 2008).

Each simulation starts with a random selection of the grid parameters (spacing, spatial phase, and orientation) and of the grid-to-place connection matrix W . We refer to this set of parameters, which fully define the grid inputs and place network for a given simulation, as a grid/place-network pair. The model is then run to determine place-unit activity across the simulated environment, from which we determine place fields and compute several spatial map characteristics (see Materials and Methods). Network sparsity is defined as the fraction of place units without active place fields. Coverage is the proportion of the environment covered by at least one place field. Representation is the average number of overlapping place fields at any point in the environment. For each active place unit, we compute the number of fields, peak firing rate, and peak firing location, and, for each place field, we compute the size, peak firing rate, average firing rate across the field, and center-of-mass location.

We used a genetic search algorithm (see Materials and Methods) to set the fixed parameters of the model: the connectivity C , inhibitory strength J , and threshold λ . We specified model performance targets on the basis of experimental data. First, several studies report that ~ 60 and 75% of place cells in hippocampal areas CA1 and CA3, respectively, are silent in any given environment (Guzowski et al., 1999; Lee et al., 2004; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004); accordingly, we specified a

Figure 1. The competitive network model for forming spatial maps with example grid inputs and place outputs. **A**, Model schematic (left) indicating the feedforward inputs from the 1000 simulated grids and a population of 500 recurrently inhibited place units. Competitive dynamics are facilitated by a saturating output nonlinearity (right). **B**, Simulated grid responses of various spacing are shown across the virtual environment. **C**, Example place-unit responses from the same simulation. Responses are place like and typically exhibit one or sometimes two distinct fields within the 1 m² environment. Numbers below the figures are normalized peak firing rates (see also Table 1).

target of 65% network sparsity. Second, to be able to test remapping, we required full representation of the environment by setting 100% coverage as a target. Third, we required a large dynamic range across place responses by specifying a target of 0.99 for the population peak rate. Fourth, to approach the spatial specificity of hippocampal activity, we specified that the model should minimize the number of place fields per active place unit. Finally, so that sparsity is not achieved at the cost of field size, we specified a target of 200 cm² for the average field size. Based on the search results, simulations reported here use inhibition strength of $J = 2250$ (the value is large because the sparseness of network activity makes $\langle r \rangle$ small), threshold $\lambda = 2$ (Fig. 1A, right), and connectivity $C = 0.33$. Thus, each place unit receives input from 330 grid cells, consistent with estimates on the range of MEC inputs to dentate granule and hippocampal pyramidal cells (100–1000) (Amaral et al., 1990). Example responses (Fig. 1C) demonstrate that this search-optimized model produces hippocampal-like place activity.

Spatial map statistics

To characterize the spatial activity generated by the model, we analyzed a sample of 32 spatial maps from different grid/place-network pairs with a total of 6178 active units and 8354 place fields. The results are summarized in Table 1. The model successfully produces spatially representative maps: on average, 98.8% of the environment is overlapped by at least one place field, and each point is represented on average by 4.51 fields. The average network sparsity of 61.4% approximates observations of CA1 activity levels in novel environments (Wilson and McNaughton,

1993; Guzowski et al., 1999; Lee et al., 2004; Karlsson and Frank, 2008). The peak firing rate of 0.925 indicates that the model uses most of its dynamic range. Active place units have an average of 1.38 place fields, intermediate between reported values for dentate gyrus and CA3 (Leutgeb et al., 2007). Most nonspecific activity is attributable to the 25% of active units with two fields; only 6.2% of active units have three or more fields (Fig. 2A).

The average peak rate of 0.418 indicates that responses are not saturating the output nonlinearity (Fig. 1A, right). Place-field peak and average firing rate distributions are positively skewed (Fig. 2B, left) and qualitatively match CA1 firing rate distributions observed in novel environments (Karlsson and Frank, 2008). The average field size of 169 cm² is smaller than typical hippocampal place fields (Muller and Kubie, 1989; Kjelstrup et al., 2008), and the size distribution is positively skewed such that only 5.2% of place fields are larger than 300 cm² (Fig. 2B, right). Network mechanisms beyond recurrent inhibition can be added to the model to produce larger place fields (see Fig. 8C and below, Alternative place-cell models).

To assess multifield responses, we classified the strongest place field of each active unit as primary and all others as secondary. The distribution of distances

between the centers of mass of every secondary field and its respective primary field (2356 secondary fields: 65.1 ± 15.8 cm, mean $\pm$ SD) may reflect the mean spacing of the grid-cell population (60 cm; one-sample $t = 15.8$, $p < 10^{-52}$). To examine grid-like periodicity in place fields, we computed two-dimensional spatial autocorrelograms of the grid and place population responses (see Materials and Methods) from an example grid/place-network pair (Fig. 3A). The grid correlation shows a strong central peak and weaker radial arms corresponding to the periodic hexagonal structure of the grids. The orientation of the radial correlations matches the grid orientation chosen at the start of the simulation. The radial arms are smeared out because the correlogram averages over grids of different spacing. The place-unit correlation has a narrower central peak and includes radial correlations similar to the grid autocorrelogram but much weaker. This is another reflection of the underlying grid input. The difference between the grid and place correlograms (Fig. 3B) reveals much stronger suppression of the center surround than of the radial grid correlations. This indicates that the model is better at local sparsification than at reducing the more remote redundancies of the periodic inputs.

Network competition

The inhibitory strength J and threshold λ of the nonlinearity are critical for the competition that generates place fields in this model (Eq. 3). To assess their impact, we simulated a 16×16 parameter grid using a single grid/place-network pair and holding all other parameters fixed at their usual values. To show the parameter dependence of key spatial coding properties, we com-

Figure 2. Distributions of place-field characteristics across a sample of 32 spatial map simulations. **A**, Distributions across 6178 active place units of the number of place fields (left) and proportional coverage of the environment (right). Most multifield responses consist of two fields, whereas 68.7% of place units have a single field. **B**, Distributions across 8534 place fields of peak and average firing rates (left) and area (right). Gray dotted line is field size minimum; black dotted line is parameter search target for field size.

Figure 3. Population autocorrelograms of grid and place responses from a spatial map simulation. **A**, Autocorrelograms of the grid (left) and place (right) population responses are visualized up to 25% of peak correlation to emphasize the structure of nonlocal correlations. The coherent orientation of both populations is evident in the radial arm-like correlations outside of the central peak, which are weaker but still present for the place responses. **B**, Difference map showing subtraction (grid – place) of the full correlograms from **A**.

puted bilinear interpolated maps of network sparsity, number of fields, field peak firing rate, and field area across this parameter grid (Fig. 4). As expected, increasing λ or J generally leads to increases in sparsity and reductions in field number, rate, and area. Holding λ fixed and increasing J yields monotonic changes to spatial coding. However, for fixed J , sparsity is U-shaped, and field number and rate are inverse-U-shaped with increasing λ . This is somewhat paradoxical because increasing λ decreases excitatory drive, but, up to a point, it yields more active place units with more and stronger place fields. This reflects a balance of the competitive dynamics in which the network is most efficiently

Figure 4. Dependence of spatial map characteristics on network competition parameters. Place-network competition was studied on a 16×16 grid of inhibitory strength J and threshold λ . Results are shown as a bilinear interpolation across the parameter grid. Increases in J lead to monotonic changes in spatial coding; increases in λ produce a non-monotonic dependence for some measures, suggesting a competitive balance between recurrent inhibition and afferent excitation.

driven with a certain amount of input but past which the global inhibition suppresses activity.

Remapping

To test the ability of changes in grid configuration to elicit place-field remapping, we compared two spatial maps generated by a single grid/place-network pair: the first as described above, and the second in which we have realigned the grid inputs. Grid inputs were divided into equal-sized subsets or modules, such that realignment was coherent within each module but different across modules (see Materials and Methods). Initially, modules were selected randomly with respect to the spatial metrics of the grids. In later simulations (see Fig. 7B,C) using “frequency modules,” the grids were sorted by field spacing (spatial frequency) before being partitioned into modules. We characterized the quality of remapping using two measures: the “remapping strength,” which characterizes the decorrelation of the distribution of pairwise inter-place-field distances before and after remapping, and the “activity turnover,” which is a measure of the randomness in the selection of active place units for the two maps (see Materials and Methods). Both measures range from ~ 0 to 1, in which values ~ 0 indicate little or no remapping. Remapping strength near 1 indicates randomization of the pairwise spatial relationships of place units active both before and after remapping. Activity turnover near 1 indicates that the subsets of place units active before and after remapping show no more overlap than expected by chance.

Remapping examples based on two modules and the three grid realignment types (shifts, changes of ellipticity, and uniform rescaling) are shown in Figure 5 for typical realignment parameters (cf. Barry et al., 2007, 2009; Fyhn et al., 2007). Remapping simulations here compute responses based on the initial environment A and then the realigned environment B. Restricting analysis to place units that were active in both environments, we

Figure 5. Modular remapping examples using two random modules for three types of grid realignment. Columns correspond to shift (left), ellipticity (middle), and rescaling (right) realignments. **A**, An example grid response of 0° orientation, 30 cm spacing, and no spatial phase offset (i.e., it has a peak at the origin) is transformed according to the realignments tested here: 15 cm shift, 20% ellipticity (along the 45° axis), and 20% rescaling. **B**, Schematic illustration of three modular remapping experiments. Arrows indicate relative direction (shift and ellipticity) and magnitude of the realignment between modules. Blue squares illustrate the changes in grid response for each module. **C**, Arrows indicate remapping of peak firing locations of active place units between environments. Circle markers indicate no change. **D**, Box-and-whisker plots showing distributions of remapped distances for active place units. Boxes show ranges of data from first to last quartiles, red lines show medians, whiskers show range of data up to 1.5 times the interquartile range (25–75%), and blue + markers show data points beyond the whiskers. **E**, 2D histograms of the distance between peak firing locations of coactive units in environment A against the same quantity in environment B. The maxima of these 64 × 64 histograms are 18, 77, and 115 for shift, ellipticity, and rescaling, respectively.

examined changes in the location defined by the peak firing rate (Fig. 5C). For shift realignments (Fig. 5C, left), quite a few long-range jumps in response maxima occur. These jumps typically occur near the edges of the environment, reflecting the appearance and disappearance of activity peaks across the borders. The other transformations (Fig. 5C, middle and right) show more place units that remap by small or even zero distances (median remapped distances are 3.2 and 4.5 cm for ellipticity and rescaling, respectively). The rescaling example (Fig. 5C, right) shows that units that remap short distances follow the expansion of the grid inputs about the origin. The distributions of remapped distances (Fig. 5D) confirm that shift realignment more effectively remaps place fields than ellipticity and rescaling in these two-module examples.

Remapping requires not only that place fields shift by variable distances but that they do so independently of each other. To test this, we computed the distances between the peak firing locations of pairs of place units before (distance A) and after (distance B) remapping and show two-dimensional histograms of these distances in Figure 5E. The remapping strength measure we use is based on the decorrelation of these pairwise data (see Materials and Methods). The off-diagonal activity evident for the shift example (Fig. 5E, left) reveals incoherence in the remapping, but the diagonal band indicates that the spatial map is not fully disrupted by this realignment. The high diagonal correlations evident in the ellipticity (Fig. 5E, middle) and rescaling (Fig. 5E, right) examples indicate that the spatial maps remained essentially intact across realignment.

Although most evidence argues against incoherent grid-cell orientations, preliminary findings by Stensland et al. (2010) indicate the possibility of small changes in orientation at discrete positions along the dorsoventral axis of MEC that may be associated with changes in grid ellipticity. To investigate the effect of grid rotations and compare them with shift effects, we simulated concurrent shifts (up to the maximum possible shift of 45 cm) and rotations (up to 30°) of one module against the other across a 16 × 16 grid of realignment parameters (Fig. 6A). Both modular shift and rotation can elicit strong remapping independently, with 15° of rotation approximately equivalent to 15 cm of shift. Whereas remapping saturates with ~20 cm of shift, rotations under 30° re-

quire shifts to fully remap (divide between top two remapping quintiles) (Fig. 6*A*, left). Thus, with just two grid-cell modules, rotations without concurrent shifts do not generate full remapping, but sufficiently large shifts do not require rotations to accomplish full remapping.

To observe continuous remapping from environment A to environment B, we performed two-module remapping simulations for an incremental series of realignments. We simulated realignments based on shift, ellipticity, and rescaling with 30 intermediate grid configurations between the two environments (Fig. 6*B*). Remapping strength and activity turnover, computed relative to environment A, are closely correlated, indicating that these realignments similarly affect the relative spatial map structure and the active subset of place units. These remapping curves show that shift realignment can be more effective at remapping than the other transformations for randomly chosen realignment parameters.

Previous experimental remapping studies have used PV correlations to characterize the similarity of spatial responses (Leutgeb et al., 2005, 2007; Wills et al., 2005). For the data in Figure 6*B*, we computed PV decorrelation relative to the spatial map of environment A (Fig. 6*C*). PV decorrelation shows qualitatively similar remapping trends as the other measures. Across the progressive realignment, PV decorrelation is generally smaller than remapping strength for all realignment types (Fig. 6*C*, inset). The exceptions are the strongly remapping shift realignments for which both measures have saturated at values ~ 1 . In general, PV decorrelation provides a smoother but less sensitive measure than pairwise remapping strength.

Modularity

Remapping may benefit from increasing the number of modules used for realignment. We tested this by running sweeps of progressive realignment simulations using from 1 to 16 modules. Realignment parameters were preserved as much as possible by drawing from a single set of 16 parameters for each sweep instead of resampling each module on every sweep. The progressive remapping curves (Fig. 6*D*) show that shift realignment benefits substantially from having two modules and only incrementally from additional modules. For ellipticity and rescaling (Fig. 6*D*, middle and right), remapping improves with two modules but shows no consistent effects with additional modules. The overall magnitudes are below the remapping achieved by shift realignment (Fig. 6*D*, left).

To assess the effects of modularity and realignment types on remapping, we simulated sets of 64 independent remapping experiments across a number of conditions. For each realignment type, we tested 1, 2, 4, 8, and 16 modules. These sets are called *sNN*, *eNN*, and *zNN* for the shift, ellipticity, and rescaling (zoom)

Figure 6. Progressive modular remapping as grids are realigned from environment A to B. Using two modules, we measure remapping strength and activity turnover to track the effects of grid realignment. *A*, Remapping (left) and turnover (right) are shown for a 16×16 grid of remapping simulations based on concurrent shift (up to 45 cm) and rotation (up to 30°) of grid cells. Data were linearly interpolated and colored by quintile (color bar). *B*, Progressive remapping (thick lines) and turnover (thin lines) across 30 intermediary configurations with randomly sampled realignment parameters: $[-24, -2.5]$ cm and $[40, 21]$ cm shifts; 16 and 5.8% ellipticity on 7.3° and 72° axes, respectively, and 1.4 and 6.7% rescaling. *C*, PV decorrelation are shown for the same data in *B* to compare the different remapping measures. PV decorrelation is plotted against remapping strength for every point of the realignment (*C*, inset) showing that PV decorrelation is less sensitive than remapping strength. *D*, Remapping strength for progressive realignments of each type with 1 (dark blue) to 16 (dark red) modules. Increasing modularity beyond two modules enhances remapping in simulations of modular shift realignment but does not have a consistent effect for ellipticity and rescaling.

realignment types, respectively, in which *NN* is the number of modules. In addition, we simulated conditions using 1000 modules (i.e., every grid-cell input is realigned independently) as an upper bound on realignment incoherence. These sets are called *srnd*, *ernd*, and *zrnd*, respectively. Finally, as an absolute upper bound on remapping, we simulated a condition called *rnd* in which an entirely new set of grid inputs was sampled to define the realigned representation. Means and SEM error bars are shown for all of these conditions in Figure 7*A*. The biggest gains in remapping occur between coherent (one-module) conditions and two-module conditions. For all realignment types, improvement in remapping levels off after the introduction of four to eight modules. We were interested in which conditions remapped sufficiently to not be significantly different from their respective incoherent sets or the *rnd* set. We computed Kolmogorov–Smirnov two-sample tests between all pairs of sample sets for remapping strength and activity turnover (Fig. 7*A*, horizontal brackets indicate nonsignificance at $p > 0.05$). Based on these tests, similarity here means we do not reject the null hypothesis (at 5% significance) that two samples result from the same distribution. For shift, both *s16* and *srnd* are similar to the *rnd* upper

Figure 7. Statistics for sets of 64 remapping simulations under a variety of realignment conditions. The first letter (s, e, or z) of the label indicates shift, ellipticity, or rescaling, whereas the number indicates modularity; srnd, ernd, and zrnd signify fully incoherent realignments; rnd signifies fully resampled grids. **A**, Mean and SEM error bars for activity turnover and remapping strength. Horizontal brackets indicate conditions for which Kolmogorov–Smirnov two-sample test is not significant ($p > 0.05$). **B**, The 95% confidence ellipses on remapping and turnover for modular shift conditions comparing random modules (no border) with spatial-frequency modules (dashed border). Numbers indicate modularity. Black ellipse is the rnd condition. **C**, Comparison of random and frequency modules. Remapping (circles) and turnover (triangles) are computed for matched sample sets of simulated remapping experiments for each realignment type and number of modules (2, 4, 8, and 16). Scatter plot (top) shows remapping and turnover averages for random versus frequency modules. Most points lie near the diagonal, indicating that the two module types produce similarly effective remapping. However, the stem plot (bottom) of the difference between module types ($\Delta = \text{random} - \text{frequency}$) shows a consistent remapping deficit for frequency modules. Of 12 modular conditions, 9 and 12 frequency-based realignments show lower remapping and turnover, respectively, than the equivalent random realignment. The largest differences occur for turnover in the case of shift realignment.

bound on remapping. Both s4 and s8 were similar to rnd for remapping but not for activity turnover. For ellipticity, both e8 and e16 are similar to ernd. For rescaling, only z16 was similar to zrnd. Thus, 16 modules is generally equivalently effective at remapping as total incoherence. Shift realignment with 16 modules is the only modular condition similar to our upper bound for both remapping strength and activity turnover.

In the Introduction, we suggested that grid cells could be assigned to modules on the basis of their grid spacing, that is, spatial frequency. To study the effects of such frequency-based modules, we ran all the modular sample sets again as above but using frequency modules (see Materials and Methods) instead of the random modules that we have used to this point. Frequency

modules are constructed by first sorting the grid-cell population by spatial frequency and then dividing it into equal-sized subsets. The spatial frequencies themselves remain uniformly random. Trends across the modular remapping conditions are similar to and highly correlated with those for random modules (Pearson's correlation, remapping, $r = 0.998$; turnover, $r = 0.992$). The shift condition shows that frequency modules are generally less effective at remapping. For example, shift realignment with 16 frequency modules yields less randomized activity turnover than eight random modules (Fig. 7B). Overall, random and frequency modules are similarly effective (Fig. 7C, top), but there is a consistent turnover deficit for frequency modules that is prominent under shift realignment (Fig. 7C, bottom).

Alternative place-cell models

The place-cell model used up to this point relies heavily on recurrent inhibition (Fig. 1). Feedforward inhibition imposed as a global activity threshold has been discussed and modeled previously as an alternative mechanism (O'Keefe and Burgess, 2005; Fuhs and Touretzky, 2006; McNaughton et al., 2006; Rolls et al., 2006; Solstad et al., 2006). In contrast to our model, spatial specificity in these sum-and-threshold models relies on place cells receiving input from grid cells with correlated spatial phases, not the random spatial phases we use. Nevertheless, it is interesting to see whether feedforward instead of recurrent inhibition can be used in our model. Using the same grid/place-network pair as in Figure 1, we computed linear responses with a threshold calculated to provide a sparsity of 61.6% (192 of 500 active place units). The place-field distributions for the recurrent (Fig. 8A, left) and feedforward (Fig. 8A, right) responses show that feedforward inhibition produces irregularly spaced clusters of place fields. Furthermore, threshold-based place fields are much smaller

(mean $\pm$ SEM, recurrent, $156 \pm 3.98 \text{ cm}^2$; feedforward, $116 \pm 3.20 \text{ cm}^2$). Correlated clusters of place fields do not provide a useful population code for position; they simply follow the peaks in the magnitude of the grid-cell inputs (Fig. 8B). The recurrent mechanism allows for a uniform representation of the environment despite variations in input strength.

The place fields generated by our model are atypically small and insufficiently sparse (Table 1; Fig. 8A, left). We explored two mechanisms for improving these features. First, noting that experience within an environment can produce larger firing fields (Mehta et al., 1997, 2000; Wallenstein and Hasselmo, 1997; Frank et al., 2004), we introduced activity-dependent plasticity on the connections between the grid and place cells (see Materials and

Figure 8. Comparison of place-field distributions generated by variants of the recurrent inhibition model. All place fields in a simulation are shown schematically (**A**, **C**) as a circle of equivalent diameter and plus sign centered at field center-of-mass. Simulations are derived from the same grid population and place-cell network shown in Figure 1. **A**, The recurrent inhibition model (left) is compared with a feedforward inhibition model (right). The feedforward activity threshold was chosen to match the network sparsity (192 of 500 units active) of the recurrent simulation. Unlike the recurrent model output, place fields in the feedforward model are distributed unevenly in several clusters across the environment. **B**, The population vector magnitude $|g(x)|$ of grid-cell activity shows relative peaks that match the locations of the place-field clusters in the feedforward output (**A**, right). **C**, The recurrent inhibition model is extended with associatively modified grid-place connections and recurrent excitation among the place units (for details, see Results, Alternative place-cell models). The spatial map is sparser (114 of 500 units active) and place fields larger (mean, 325 cm²) than the recurrent inhibition model (**A**, left) (Table 1; Fig. 2B). Place fields formed local clusters: 103 of 113 units active in **A** (left) and **C** shifted peak firing <10 cm.

Methods). After these activity-dependent weight changes, the place-unit responses demonstrated higher sparsity (71.4%, 143 of 500 active place units) and slightly larger place-field size (156 ± 3.98 cm² before and 166 ± 5.67 cm² after; Kolmogorov–Smirnov two-sample test, $p < 0.05$). Second, we approximated (see Materials and Methods) the effects of excitatory recurrent collaterals between place cells (Samsonovich and McNaughton, 1997; Redish and Touretzky, 1998; Tsodyks, 1999; Witter, 2007a). The resulting place-field distribution (Fig. 8C) shows both less coverage of the environment (85%) attributable to higher network sparsity (77.2%, 114 of 500 active place units) and substantially larger place fields (325 ± 13.1 cm²) than the simulation with recurrent inhibition alone (Table 1). The interquartile range of place-field areas (181 – 429 cm²) is greater than the median of the original distribution (150 cm²) (Fig. 2B, right). Furthermore, the “novel” and “familiar” (before and after these alterations of the model) representations overlapped quite well. Applying our pairwise measure, familiarization provoked remapping of 17.9% (the turnover measure does not apply because it assumes similar network sparsity), which was generated primarily by small shifts: 103 of 113 units shifted peak firing location by <10 cm.

These studies show that learning and recurrent excitation can extend our original model by constructing spatial maps with the

higher sparsity and larger place fields typical of hippocampal activity. The overall similarity between the place-field locations of the original and extended models suggests that our results for different grid-cell transformations and numbers of modules (Fig. 7) apply to familiarized as well as novel representations.

Discussion

We investigated a recurrent-inhibition model of initial place-cell activity in novel environments to assess the hypothesis of grid-cell modules as a basis of hippocampal remapping. Although simultaneously recorded grid cells in remapping experiments shift coherently (Fyhn et al., 2007), the tetrode recordings have typically been restricted to colocalized grid cells. These findings demonstrate that grid realignment is at least locally coherent. Two hypotheses were originally suggested by Fyhn et al. (2007) to reconcile the local coherence of spatial inputs to hippocampus with the randomization inherent to hippocampal remapping. First, they hypothesized independent grid-cell modules, as we have explicitly tested here. Second, they hypothesized that the spatial phases of grid cells may represent a position code for an infinitely large map of space. To elicit remapping, this position code is shifted to a new random location in the infinite map. The infinite-map hypothesis requires that grid cells of the same spatial frequency be displaced by the same amount. As a result, our simulations using frequency modules in the limit of large numbers of modules correspond to the infinite-map hypothesis. Our simulation results then show that both hypotheses can viably produce complete or nearly complete remapping. Simultaneous recordings of grid cells with significantly different spacing that nonetheless realign coherently during remapping would provide strong support for frequency-independent modules over the infinite-map hypothesis. Because we found that a small number of modules can be effective at remapping, the modularity hypothesis may require large distances between tetrodes along MEC to be proven or falsified. Such long-range recordings may be technically difficult, but the necessary experiments are being pursued (Stensland et al., 2010).

Although remapping based on grid-cell responses has been discussed (O’Keefe and Burgess, 2005; McNaughton et al., 2006) and demonstrated in modeling work (Fuhs and Touretzky, 2006; Hayman and Jeffery, 2008), the relative remapping effectiveness of various grid manipulations has not been systematically quantified. With data showing that ellipticity (Barry et al., 2007; Stensland et al., 2010) and rescaling (Barry et al., 2007, 2009) may be modes of grid-cell realignment, this sort of quantification has become necessary to understanding the relationship between the spatial activity patterns of entorhinal cortex and hippocampus. Specifically, changes in grid ellipticity are supported by observations of elliptical or compressed grids in both altered (Barry et al., 2007) and familiar (Stensland et al., 2010) environments. Preliminary observations have shown rescaling consisting of the uniform expansion of grid scale on introduction to a novel environment (Barry et al., 2009). Theoretically, these geometric transformations could result from changes to the synaptic weights in attractor network models (Fuhs and Touretzky, 2006; McNaughton et al., 2006; Burak and Fiete, 2009) or to the frequency modulation of theta-frequency oscillators in temporal interference models (O’Keefe and Burgess, 2005; Burgess et al., 2007; Blair et al., 2008; Hasselmo, 2009) of grid-cell activity. The reduced remapping capabilities of ellipticity and rescaling as putative forms of realignment could nevertheless be functional, producing for example the sort of partial or graded remapping

that has been observed in CA1 (Lee et al., 2004; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004).

Despite changes in grid orientation when an animal is moved to a new environment, we do not consider modularity of grid orientation to be an experimentally supported mode of realignment. If humans have grid cells, then the grid-like periodic signal from a recent functional imaging study would likely not have been apparent in the presence of any significant orientational incoherence (Doeller et al., 2010). In addition, directionality in rodent MEC is strongly coupled with the head-direction system (Hargreaves et al., 2007), which is itself internally coherent during remapping conditions (Yoganarasimha et al., 2006). However, small angular rotations may accompany other geometric modifications of grid-cell responses. To examine this possibility, we tested small modular rotations (Fig. 6A) and found that differential orientations can elicit strong remapping alone and complete remapping in combination with shifts.

Remapping may involve additional computational components such as pattern separation provided by dentate gyrus (Acádý and Káli, 2007; Leutgeb and Leutgeb, 2007) and the integration of changes in external sensory information represented by activity carried by the lateral entorhinal projection to the hippocampus (Knierim et al., 2006; Lisman, 2007; Witter, 2007b). Grid lattices are only part of a broader functional diversity of nonspatial and spatial activity in entorhinal cortex (Hargreaves et al., 2005; Sargolini et al., 2006; Savelli et al., 2008; Solstad et al., 2008). Although MEC and hippocampus respond concurrently during remapping, we do not address whether changes in grid-cell response result from direct environmental input or other mechanisms such as hippocampal feedback (O'Keefe and Burgess, 2005; McNaughton et al., 2006; Burgess et al., 2007; Hasselmo, 2008; Burak and Fiete, 2009). The model here approaches one particular spatial mode of hippocampal processing of its cortical inputs.

A number of grid-to-place models were posited after the discovery of grid cells (McNaughton et al., 2006; Rolls et al., 2006; Solstad et al., 2006; Franzius et al., 2007). Most rely on a combination of activity thresholds, grid inputs with correlated spatial phases, and associative or competitive learning rules. Associative learning with heterosynaptic depression generates place fields in both temporal and rate coding models (Molter and Yamaguchi, 2008; Savelli and Knierim, 2010). Fuhs and Touretzky (2006) demonstrated that place fields could be randomly displaced by progressively resetting the spatial phases of its grid inputs. Hayman and Jeffery (2008) showed that learning a dendritically organized spatial-phase partition of grids could provide the context dependence necessary for partial and complete remapping. de Almeida et al. (2009) proposed a model that, similarly to our model, uses randomly aligned grid inputs and does not require learning to produce place-like activity. However, their model depends on an adaptive activity threshold and a skewed weight distribution that prevents the recruitment of independent active subsets of place units in new environments. We use feedback inhibition rather than an adaptive threshold or synaptic modification to produce spatially selective outputs from the global spatial representation of grid-cell inputs. We showed that simple threshold mechanisms, representing feedforward inhibition, tend to follow overall input strength in novel environments and do not produce informative spatial representations.

Inhibition plays a key role in our place-cell network model. Hippocampal interneurons are diverse (Sik et al., 1997) and constitute up to 20% of hippocampal cells, approximately one-third of which directly innervate pyramidal cells (Buzsáki et al., 2007).

The dentate gyrus and subregion CA3 both have extensive recurrent and feedforward inhibitory microcircuits that are crucial to hippocampal computation (Sik et al., 1997; Acádý and Káli, 2007) and are supported by diverse mechanisms of synaptic plasticity (Pelletier and Lacaille, 2008). Inhibitory network dynamics are integral to many hippocampal functions, such as the competitive transformation of redundant inputs into more informative outputs (Leutgeb et al., 2007; Karlsson and Frank, 2008). We hypothesize that this general computation could produce place fields on the first pass through unfamiliar environments (Hill, 1978; Frank et al., 2004). Several studies have shown that interneurons in CA3 mediate disynaptic-latency feedback inhibition across long distances within the pyramidal cell layer (Glickfeld et al., 2009; Bazelot et al., 2010), which is consistent with the global inhibition modeled here. A detailed compartmental pyramid-interneuron model of fast feedback loops in CA3 demonstrated higher spike and burst rates with stronger inhibitory gains (Zeldenrust and Wadman, 2009). Our network-level model similarly uses strong feedback inhibition to enhance competition.

The activity patterns that we have modeled could form the basis of a “seed” representation that becomes refined with longer-timescale familiarization (Gerstner and Abbott, 1997; Lee et al., 2004; Leutgeb et al., 2004; Karlsson and Frank, 2008). Indeed, Savelli and Knierim (2010) used their spatial learning model to show that the details of initial activity may critically determine the spatial representations that are learned as a novel environment becomes more familiar. Our model predicts that nonspecific firing in the initial activity of place cells results from correlated grid inputs. Learning mechanisms may act to enhance spatial specificity while also broadening and shaping place fields, consistent with previous modeling and experimental findings (Mehta et al., 1997, 2000; Wallenstein and Hasselmo, 1997). We showed that *post hoc* application of associative synaptic modification (Fig. 8C) can both enhance sparsity and broaden the place fields of our model. These effects of familiarization may also contribute to remapping.

In conclusion, the effectiveness of a small number of shifting grid modules indicates that entorhinal realignment could be significantly coherent and still contribute substantially to hippocampal remapping. Although we did not test all combinations of realignment types, combinations of various grid transformations can provide enhanced remapping with minimal grid-cell modularity. Although shift-based realignment is more effective than ellipticity or scale changes, fully orthogonal remapping may not be necessary for spatial memory encoding in rodent hippocampus. These transformations could contribute to partial remapping or other graded response changes critical to hippocampal function.

Notes

Several supplemental figures as well as the model simulation software are available at <http://jdmnaco.com/remapping>. This material has not been peer reviewed.

References

- Acádý L, Káli S (2007) Models, structure, function: the transformation of cortical signals in the dentate gyrus. *Prog Brain Res* 163:577–599.
- de Almeida L, Idiart M, Lisman JE (2009) The input-output transformation of the hippocampal granule cells: from grid cells to place fields. *J Neurosci* 29:7504–7512.
- Amaral DG, Ishizuka N, Claiborne B (1990) Neurons, numbers and the hippocampal network. *Prog Brain Res* 83:1–11.

- Barry C, Hayman R, Burgess N, Jeffery KJ (2007) Experience-dependent rescaling of entorhinal grids. *Nat Neurosci* 10:682–684.
- Barry C, O'Keefe J, Burgess N (2009) Effect of novelty on grid cell firing. *Soc Neurosci Abstr* 35:101.24.
- Bazelot M, Dinocourt C, Cohen I, Miles R (2010) Unitary inhibitory field potentials in the CA3 region of rat hippocampus. *J Physiol* 588:2077–2090.
- Blair HT, Wolday AC, Zhang K (2007) Scale-invariant memory representations emerge from moiré interference between grid fields that produce theta oscillations: a computational model. *J Neurosci* 27:3211–3229.
- Blair HT, Gupta K, Zhang K (2008) Conversion of a phase- to a rate-coded position signal by a three-stage model of theta cells, grid cells, and place cells. *Hippocampus* 18:1239–1255.
- Bostock E, Muller RU, Kubie JL (1991) Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1:193–205.
- Burak Y, Fiete IR (2009) Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol* 5:e1000291.
- Burgess N, Barry C, O'Keefe J (2007) An oscillatory interference model of grid cell firing. *Hippocampus* 17:801–812.
- Buzsáki G, Kaila K, Raichle M (2007) Inhibition and brain work. *Neuron* 56:771–783.
- Doeller CF, Barry C, Burgess N (2010) Evidence for grid cells in a human memory network. *Nature* 463:657–661.
- Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. *J Neurosci* 24:7681–7689.
- Franzius M, Vollgraf R, Wiskott L (2007) From grids to places. *J Comput Neurosci* 22:297–299.
- Fuhs MC, Touretzky DS (2006) A spin glass model of path integration in rat medial entorhinal cortex. *J Neurosci* 26:4266–4276.
- Fyhn M, Hafting T, Treves A, Moser MB, Moser EI (2007) Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446:190–194.
- Fyhn M, Hafting T, Witter MP, Moser EI, Moser MB (2008) Grid cells in mice. *Hippocampus* 18:1230–1238.
- Gerstner W, Abbott LF (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. *J Comput Neurosci* 4:79–94.
- Giocomo LM, Hasselmo ME (2008) Computation by oscillations: implications of experimental data for theoretical models of grid cells. *Hippocampus* 18:1186–1199.
- Glickfeld LL, Roberts JD, Somogyi P, Scanziani M (2009) Interneurons hyperpolarize pyramidal cells along their entire somatodendritic axis. *Nat Neurosci* 12:21–23.
- Guzowski JF, McNaughton BL, Barnes CA, Worley PF (1999) Environment-specific expression of the immediate-early gene *Arc* in hippocampal neuronal ensembles. *Nat Neurosci* 2:1120–1124.
- Hafting T, Fyhn M, Molden S, Moser MB, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–806.
- Hargreaves EL, Rao G, Lee I, Knierim JJ (2005) Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science* 308:1792–1794.
- Hargreaves EL, Yoganarasimha D, Knierim JJ (2007) Cohesiveness of spatial and directional representations recorded from neural ensembles in the anterior thalamus, parasubiculum, medial entorhinal cortex, and hippocampus. *Hippocampus* 17:826–841.
- Hasselmo ME (2008) Grid cell mechanisms and function: contributions of entorhinal persistent spiking and phase resetting. *Hippocampus* 18:1213–1229.
- Hasselmo ME (2009) A model of episodic memory: mental time travel along encoded trajectories using grid cells. *Neurobiol Learn Mem* 92:559–573.
- Hayman RM, Jeffery KJ (2008) How heterogeneous place cell responding arises from homogeneous grids: a contextual gating hypothesis. *Hippocampus* 18:1301–1313.
- Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Exp Neurol* 62:282–297.
- Karlsson MP, Frank LM (2008) Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *J Neurosci* 28:14271–14281.
- Kjelstrup KB, Solstad T, Brun VH, Hafting T, Leutgeb S, Witter MP, Moser EI, Moser MB (2008) Finite scale of spatial representation in the hippocampus. *Science* 321:140–143.
- Knierim JJ, Lee I, Hargreaves EL (2006) Hippocampal place cells: parallel input streams, subregional processing, and implications for episodic memory. *Hippocampus* 16:755–764.
- Lee I, Yoganarasimha D, Rao G, Knierim JJ (2004) Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* 430:456–459.
- Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, McNaughton BL, Moser MB, Moser EI (2005) Progressive transformations of hippocampal neuronal representations in “morphed” environments. *Neuron* 48:345–358.
- Leutgeb JK, Leutgeb S, Moser MB, Moser EI (2007) Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315:961–966.
- Leutgeb S, Leutgeb JK (2007) Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. *Learn Mem* 14:745–757.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305:1295–1298.
- Lisman JE (2007) Role of the dual entorhinal inputs to hippocampus: a hypothesis based on cue/action (non-self/self) couplets. *Prog Brain Res* 163:615–625.
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB (2006) Path integration and the neural basis of the “cognitive map.” *Nat Rev Neurosci* 7:663–678.
- Mehta MR, Barnes CA, McNaughton BL (1997) Aug. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci U S A* 94:8918–8921.
- Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25:707–715.
- Molter C, Yamaguchi Y (2008) Entorhinal theta phase precession sculpts dentate gyrus place fields. *Hippocampus* 18:919–930.
- Muller RU, Kubie JL (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. *J Neurosci* 9:4101–4110.
- O'Keefe J, Burgess N (2005) Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15:853–866.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34:171–175.
- Pelletier JG, Lacaille JC (2008) Long-term synaptic plasticity in hippocampal feedback inhibitory networks. *Prog Brain Res* 169:241–250.
- Redish AD, Touretzky DS (1998) The role of the hippocampus in solving the Morris water maze. *Neural Comput* 10:73–111.
- Rolls ET, Stringer SM, Elliot T (2006) Entorhinal grid cells can map to hippocampal place cells by competitive learning. *Network* 17:447–465.
- Samsonovich A, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *J Neurosci* 17:5900–5920.
- Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser MB, Moser EI (2006) Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312:758–762.
- Savelli F, Knierim JJ (2010) Hebbian analysis of the transformation of medial entorhinal grid-cell inputs to hippocampal place fields. *J Neurophysiol* 103:3167–3183.
- Savelli F, Yoganarasimha D, Knierim JJ (2008) Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus* 18:1270–1282.
- Sik A, Penttonen M, Buzsáki G (1997) Interneurons in the hippocampal dentate gyrus: an in vivo intracellular study. *Eur J Neurosci* 9:573–588.
- Solstad T, Moser EI, Einevoll GT (2006) From grid cells to place cells: a mathematical model. *Hippocampus* 16:1026–1031.
- Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI (2008) Representation of geometric borders in the entorhinal cortex. *Science* 322:1865–1868.
- Stensland H, Kirkesola T, Moser MB, Moser E (2010) Orientational geometry of entorhinal grid cells. *Soc Neurosci Abstr* 36:101.14.
- Tsodyks M (1999) Attractor neural network models of spatial maps in hippocampus. *Hippocampus* 9:481–489.

- Vazdarjanova A, Guzowski JF (2004) 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.
- Wallenstein GV, Hasselmo ME (1997) GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. *J Neurophysiol* 78:393–408.
- Walling SG, Bromley K, Harley CW (2006) Glycogen phosphorylase reactivity in the entorhinal complex in familiar and novel environments: evidence for labile glycogenolytic modules in the rat. *J Chem Neuroanat* 31:108–113.
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Witter MP (2007a) Intrinsic and extrinsic wiring of CA3: indications for connectional heterogeneity. *Learn Mem* 14:705–713.
- Witter MP (2007b) The perforant path: projections from the entorhinal cortex to the dentate gyrus. *Prog Brain Res* 163:43–61.
- Witter MP, Moser EI (2006) Spatial representation and the architecture of the entorhinal cortex. *Trends Neurosci* 29:671–678.
- Yoganarasimha D, Yu X, Knierim JJ (2006) Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. *J Neurosci* 26:622–631.
- Zeldenrust F, Wadman WJ (2009) Two forms of feedback inhibition determine the dynamical state of a small hippocampal network. *Neural Netw* 22:1139–1158.