Neuron
Article

A Model of Grid Cell Development through Spatial Exploration and Spike Time-Dependent Plasticity

John Widloski[1] and Ila R. Fiete[1,*]
[1]Center for Learning and Memory and Department of Neuroscience, The University of Texas at Austin, Austin, TX 78712, USA
*Correspondence: ilafiete@mail.clm.utexas.edu
http://dx.doi.org/10.1016/j.neuron.2014.06.018

SUMMARY

Grid cell responses develop gradually after eye opening, but little is known about the rules that govern this process. We present a biologically plausible model for the formation of a grid cell network. An asymmetric spike time-dependent plasticity rule acts upon an initially unstructured network of spiking neurons that receive inputs encoding animal velocity and location. Neurons develop an organized recurrent architecture based on the similarity of their inputs, interacting through inhibitory interneurons. The mature network can convert velocity inputs into estimates of animal location, showing that spatially periodic responses and the capacity of path integration can arise through synaptic plasticity, acting on inputs that display neither. The model provides numerous predictions about the necessity of spatial exploration for grid cell development, network topography, the maturation of velocity tuning and neural correlations, the abrupt transition to stable patterned responses, and possible mechanisms to set grid period across grid modules.

INTRODUCTION

The unusual tuning properties of mammalian grid cells (Hafting et al., 2005; Fyhn et al., 2008; Doeller et al., 2010; Yartsev et al., 2011; Killian et al., 2012) have spurred a number of theoretical and experimental efforts to dissect their function and mechanisms. The periodic response of grid cells to 2D animal location is relatively independent of nonspatial variables. Thus, grid cells are widely conjectured to be responsible for computing spatial displacements, by integrating self-motion cues (Hafting et al., 2005; Fuhs and Touretzky, 2006; Fiete et al., 2008; Burak and Fiete, 2009). This hypothesized function is consistent with the observation that the grid code has several unique features for the fine-grained representation of space (Fuhs and Touretzky, 2006; Fiete et al., 2008; Gorchetchnikov and Grossberg, 2007; Sreenivasan and Fiete, 2011; Mathis et al., 2012) including the capacity to represent variables over exponentially large ranges at fixed resolution, in contrast to the polynomial scaling of range over resolution permitted by other known population codes (Sreenivasan and Fiete, 2011).

On the question of mechanism, grid cells from a module (defined as the set of all cells with similar spatial period and orientation) (Stensola et al., 2012) appear to collectively exhibit 2D continuous attractor dynamics (Yoon et al., 2013). Consistent with this finding, network input in the form of slow depolarizing current ramps drives the spatial firing rate patterns of grid cells (Domnisoru et al., 2013; Schmidt-Hieber and Häusser, 2013). Several hardwired models demonstrate how recurrent circuits that exhibit low-dimensional, continuous attractor dynamics can give rise to grid-cell-like activity (Fuhs and Touretzky, 2006; McNaughton et al., 2006; Guanella et al., 2007; Burak and Fiete, 2009; Pastoll et al., 2013; Blair et al., 2008; Mhatre et al., 2012). Other models of individual grid cells are based on interfering temporal oscillations (Burgess et al., 2007; Hasselmo et al., 2007) that are mapped into spatially periodic responses.

Nevertheless, key mechanistic questions remain unanswered. It is poorly understood how the system might form. In rats, grid cell responses emerge at 3–4 weeks of age, after eye opening, (Wills et al., 2010; Langston et al., 2010), suggesting that the system is not hardwired at birth. Published experiments do not determine, however, which kinds of neural plasticity are necessary for development of the grid cell system or even whether development is experience dependent.

The question of grid cell development poses several conceptual challenges. It is unclear how a periodic response to space can arise, given that neither space nor animal location within it are periodic. Second, it is common to model the emergence of tuning curves by assuming that a supervisory input imposes the desired tuning on the network and that plasticity rules consolidate the imposed patterns (Stringer et al., 2002; Hahnloser, 2003). But it is implausible that fully functional supervisory grid-patterned inputs exist before grid cell maturation. Third, local inhibition is instrumental in recurrent models and can be easily realized by local sprouting of arbors, but such a developmental rule would predict a topographic arrangement of cells within a module based on their spatial tuning, which may or may not be consistent with existing data. Fourth, if grid cells are path integrators, the plasticity rules must generate grid responses even in the absence of visual or spatially informative inputs (contrast with Kropff and Treves, 2008). Across model classes, velocity integration requires elaborate circuitry between cells that depends on their

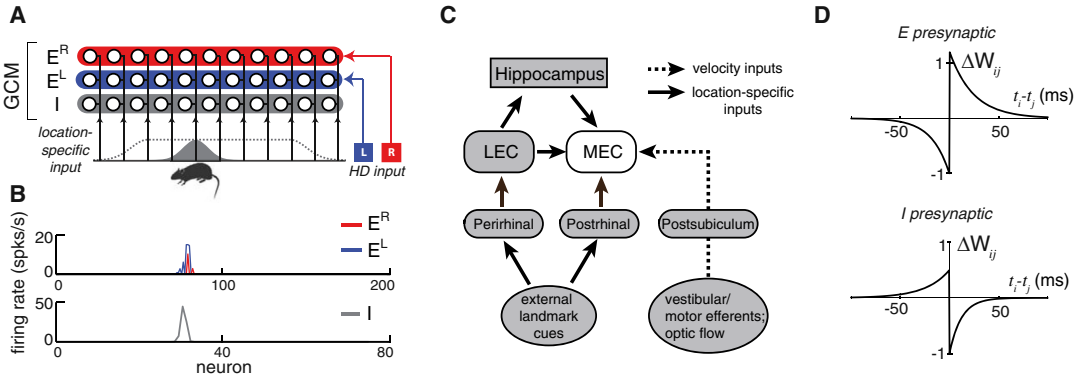

Figure 1. Initial Architecture and Learning Rule

(A) The 1D GCN consists of inhibitory and velocity-sensitive excitatory cells. All cells are assigned location-specific inputs (gray bell-shaped curve: schematic of a location-specific input; dotted gray envelope suppresses location-specific inputs near the environment boundaries.)

(B) Snapshot of population activity during the plasticity phase.

(C) Possible pathways for velocity and location-specific inputs to entorhinal grid cells.

(D) The STDP windows (kernels) used in this work. i and j are the indices of the postsynaptic and presynaptic neurons, respectively.

velocity tuning, suggesting that activity-based mechanisms are important.

Here, we describe a neural network model for the emergence of grid cells based on spatial experience. Despite its relative simplicity, the model overcomes the conceptual hurdles described above. The mature network exhibits grid-cell-like activity patterns and is capable of path integration. It is consistent with the existing (but limited) data on connectivity in layer II of the medial entorhinal cortex (MEC) (Dhillon and Jones, 2000; Pastoll et al., 2013; Couey et al., 2013), and like the grid cell data, it exhibits approximate 2D continuous attractor dynamics (Yoon et al., 2013).

RESULTS

Initial Architecture and Learning Rules

Our model relies on active exploration of a spatially cue-rich environment. This work focuses on the assembly of an individual grid cell network (GCN) whose model neurons correspond to grid cells in one experimentally observed module. The GCN consists of excitatory (E) and inhibitory (I) neurons in a 5:1 ratio, each modeled as a linear-nonlinear Poisson neuro n. The main reason to use stochastic model neurons is to establish an estimate of the required learning time, something that is not possible with rate-based deterministic dynamics in which one sweep through the environment is sufficient to generate translation-invariant weight profiles.

During exploration (see Figure S1, available online, for trajectory statistics), the E and I populations receive location-specific inputs with unimodal tuning to animal location (Figures 1A and 1B). The location-specific inputs are assumed to be derived from external sensory cues (Figure 1C) and to uniformly cover the environment. (When the uniformity constraint is relaxed, the qualitative results are unchanged, but the accuracy of path integration in the mature GCN is limited; see Figure S1.) As we will see, once the GCN is mature, the location-specific inputs become ineffective in driving neural activation and setting the network state.

The E cells are divided into two populations, distinguished by their velocity inputs: one population (E^R) receives biased excitation when the animal moves rightward, the other (E^L) when the animal moves leftward (Figure 1A). The I cells receive no velocity input. The recurrent synapses (E-to-I, I-to-E, and I-to-I; to be consistent with experimental data [Dhillon and Jones, 2000; Pastoll et al., 2013; Couey et al., 2013; Beed et al., 2013], there are no E-to-E connections), initially random and weak, are subject to change according to spike-time-dependent plasticity (STDP) rules (conventional Hebbian STDP for E-to-I synapses [Lu et al., 2007; Huang et al., 2013; Fino and Venance, 2010; Feldman, 2012]; anti-Hebbian STDP for I-to-E and I-to-I synapses, in which plasticity is generally poorly characterized [Holmgren and Zilberter, 2001]). Simultaneously reversing the time axes of all STDP windows used here, which would correspond to anti-Hebbian STDP in E-to-I synapses (Haas et al., 2006; Fino and Venance, 2010), results in quantitatively the same outcome (up to a sign-flip during integration; see Figure S1). During plasticity, neural activity is driven by the location- and velocity-specific feedforward inputs (the recurrent inputs and global tonic feedforward excitation are suppressed).

Development of Connectivity

Emergence of "Local" Connectivity

With the feedforward-driven activity patterns shown in Figure 1B, cells with similar location inputs are activated within a short latency of each other. These short-latency activations rapidly result in stronger synaptic coupling that is "local" in the functional sense of preferred location, if not in the actual positions of neurons in the GCN (Figures 2A–2C, row 1).

GCN neurons also spontaneously develop weight asymmetries based on direction preference (Figures 2A and 2B). During rightward traversals, the positive STDP lobe strengthens weights

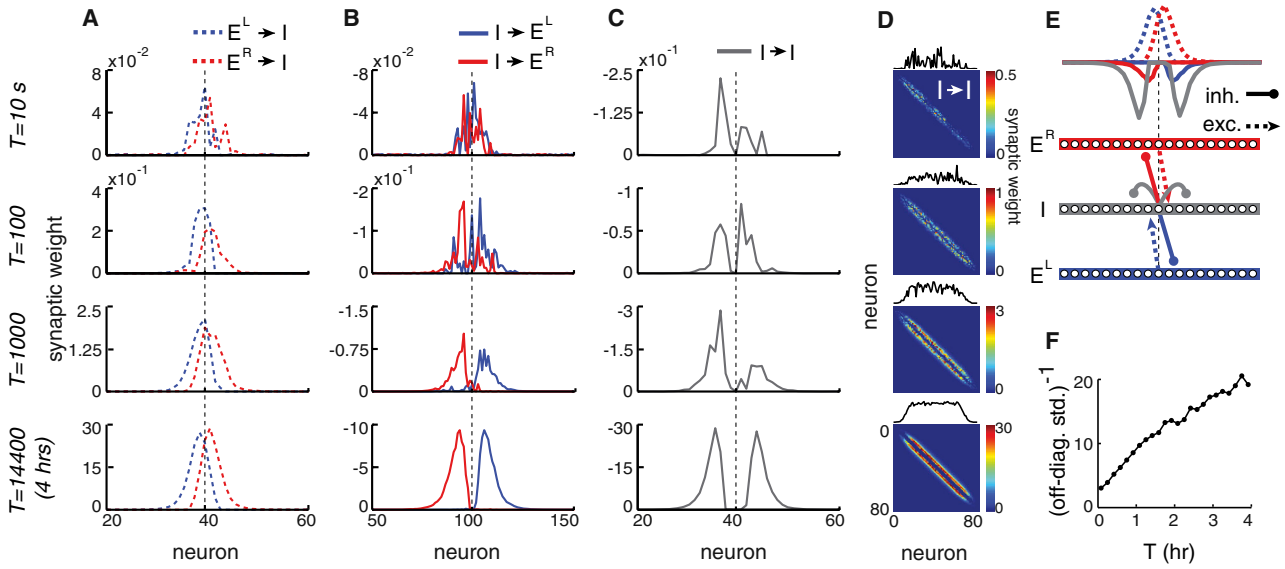

Figure 2. Evolution of Network Connectivity

(A–C) Recurrent weights across development. Each curve represents the output projections of a single cell (located at the black dashed line) to its postsynaptic targets. Note the change in scale across rows.

(D) Development of the I-to-I synaptic weight matrix, with the off-diagonal plotted above.

(E) Top: weight profiles in the mature GCN (from [A]–[C], row 4). Bottom: schematic of mature connectivity between populations.

(F) Measure of translation invariance (inverse SD of the off-diagonal shown in [D]) as a function of time during development.

from E^R cells to I cells with slightly more rightward inputs, which fire a short time later. The weight gain is not fully cancelled by the opposite lobe of the STDP kernel during leftward trajectories, because during such runs, the net activation of the E^R neurons is lower. Thus, E^R cells project to I cells with slightly right-shifted location inputs (Figure 2A, row 4). The opposite happens for the E^L population.

By a similar argument, I-to-E projections acquire slight shifts in the opposite direction, because the STDP window is flipped when the presynaptic cell is inhibitory (Figure 2B, rows 3 and 4). The I-to-I connections remain symmetric and unbiased because I cells receive no velocity input (Figures 2C and 2D, rows 3 and 4). The symmetries and asymmetries of these network projections are summarized in Figure 2E (full mature weights are in Figure S2).

Learning Time

The basic architecture of the GCN is established within the first few minutes of exploration and plasticity (Figures 2A–2C, rows 1–3). However, two key developments unfold over hours. The first is a strengthening of the weight profiles: until the weights reach a threshold in size, they cannot drive pattern formation (next subsection). The second is a progressive increase in translation invariance of the outgoing weight profiles across cells. Translation invariance, which is important for the formation of a continuum of fixed points for analog integration and memory (Zhang, 1996), is quantified by the increase in smoothness along the off-diagonals of the weight matrices (Figures 2D and 2F). There is a tradeoff between rapidity of weight growth and translation invariance in the mature GCN, so that major changes in learning rate either cause learning to be too slow or cause the mature network to exhibit too little translation invariance (Figure S2). Thus, the estimate of learning time is not susceptible to order-of-magnitude changes through corresponding adjustments of the rate parameter.

We have not explored the dependence of learning time on neuron number. However, larger networks allow for more averaging, so we expect that larger networks may reach a comparable level of performance sooner. Finally, note that the accrual of ~4 hr of plasticity time, as taken for maturation by the GCN in Figure 2, may span half a day or several days to weeks in animals if plasticity occurs only part of the time during exploration and exploration itself occurs only in small bouts over the day.

Emergence of Patterned Activity

Next, we examine how the growing recurrent weights shape neural activity and spatial tuning. In all that follows, neural responses are probed in the activation phase, with all recurrent, tonic and location- and velocity-specific inputs active.

Population Activity and Path Integration

Early in development, population activity is roughly uniform across cells (Figures 3A and 3B, rows 1 and 2). In the presence of the global tonic excitatory input, the relative contribution of the location-specific input is small, and the largest source of activity modulation is the velocity input (Figures 3A and 3B, rows 1 and 2).

The growing recurrent weights eventually destabilize uniform activity states within the GCN and force a periodic patterning. The pattern becomes apparent if cells are ordered based on their location-specific inputs (Figures 3A and 3B, rows 3–5). Cells in

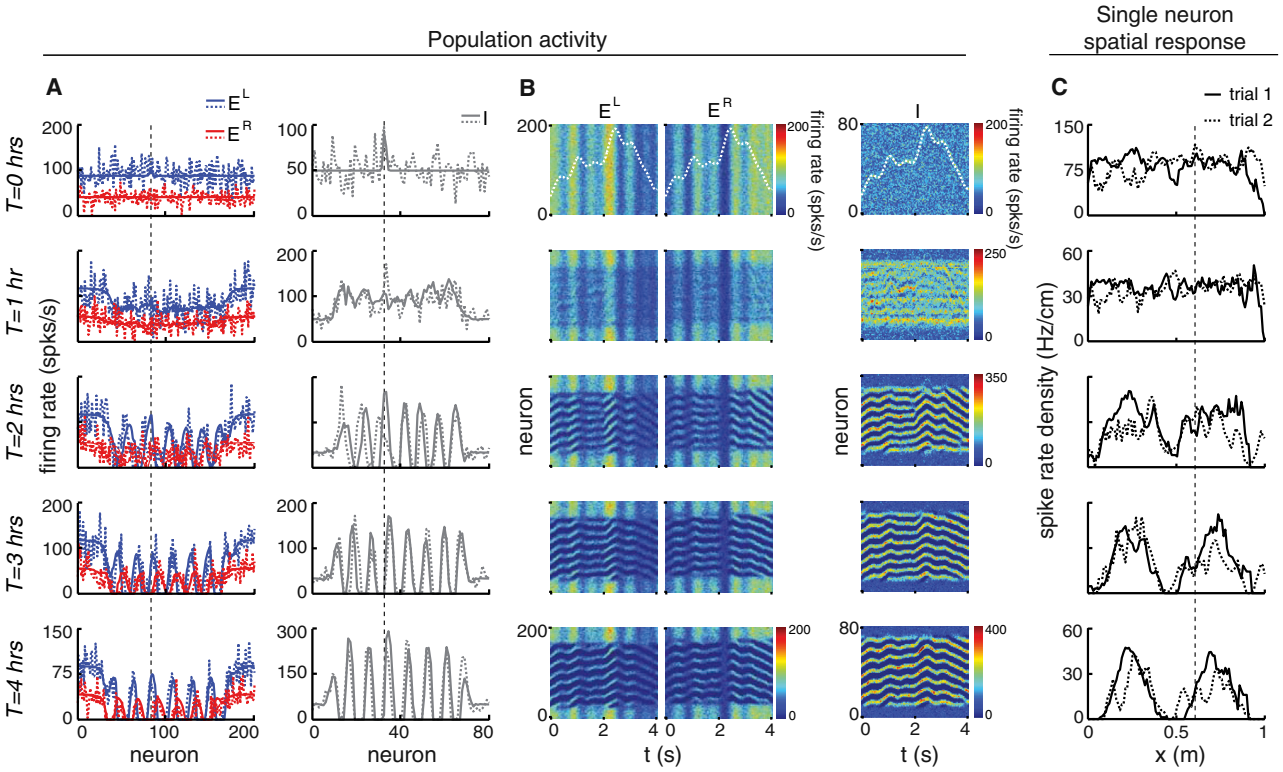

Figure 3. Evolution of Population Activity, Path Integration, and Spatial Tuning

(A) Snapshot of population activities in the activation phase during a leftward run (v_{rat} = 0.4 m/s), across development. Solid lines: mean rate used to generate Poisson spikes; dashed lines: filtered spiking activity (spike trains convolved with Gaussian, σ = 15 ms). The location-specific bump in the population activity (location marked by the vertical dashed line) is obscured by the larger-amplitude spike variability.

(B) Population activities over a 4 s interval (spiking activity filtered as in [A]). The white dashed line (top row) tracks the location-specific input. Note the difference in movement gain between the white line and the population phase.

(C) Spatial tuning (see Supplemental Information) of a cell across two trials, measured across development (the two trials consist of different 10 s variable-velocity runs that span the enclosure, beginning at the same initial population pattern phase and similar initial location of the animal within some tolerance; see Supplemental Information). The location-specific input of this cell drives (insignificant) activation at the dashed vertical line.

the E^R population drive the activity pattern rightward by exciting I cells with a rightward bias, which in turn inhibit the left flanks of the corresponding E bumps (Figure 2E). When the animal is still, the E^L population exerts an equal and opposite influence, and the pattern remains stationary. When the animal moves, one of the E populations receives biased excitation (Figure 3A, row 5) and succeeds in driving the activity pattern along its preferred direction.

Note that both the population pattern phase and the location-specific input track animal location (compare pattern phase with the white dashed line in Figure 3B); however, the movement gain of the location-specific input and the population pattern phase across the GCN are mismatched. Therefore, the GCN pattern phase is based on integrating the velocity input rather than being dictated by the location-specific input.

In this model, both E and I populations become patterned. This result is a consequence of prohibiting direct E-to-E coupling. When E-to-E coupling is permitted, it is possible for E cells to be patterned while I cells remain largely unpatterned and exhibit minimal spatial tuning (Figure S3).

Integration accuracy, and thus, by extension, the fidelity of spatial tuning, improves with the size of the GCN and decreases with the variability of neural spiking (Figure S3).

Spatial Tuning of Individual Cells

During early development, GCN neurons are not spatially tuned (Figure 3C, rows 1 and 2), consistent with the uniform, unpatterned population response at the same stage. Stable spatial tuning appears at a lag relative to the appearance of population patterning (compare the emergence of a population pattern in rows 2 and 3 of Figure 3C with spatial tuning in rows 3 and 4; quantification to follow).

The firing locations of a GCN neuron at this stage are unrelated to its location-specific input (Figure 3C). This follows because firing locations are set by the speed with which the population pattern moves through the network, which depends on the recurrent weight asymmetry and the gain of the feedforward velocity input in the GCN, and thus in general differs from the speed of the location-specific input. Thus, the GCN's internal location estimate is decoupled from the location-specific input's estimate, and the latter has become irrelevant to GCN performance.

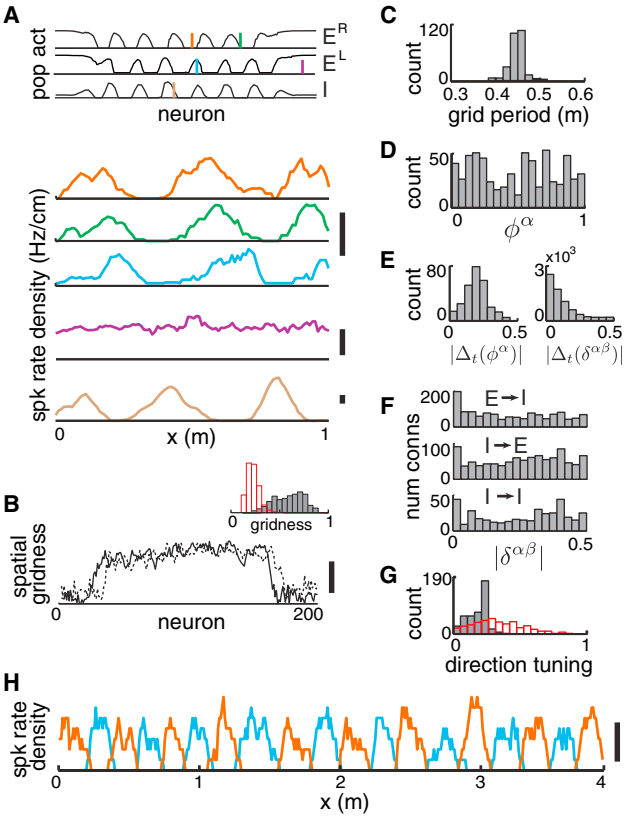

Figure 4. Spatial Tuning in the Mature Network

(A) Top: population activity snapshot; cells probed for spatial tuning (shown below) are indicated by colored bars. Bottom: spatial tuning of a sample of cells from the mature GCN, obtained from a 10 s variable-speed run across the environment. Scale bars: mean spike rate of 60 Hz (averaged over 1 cm).

(B–D) (B) Spatial gridness scores across the mature E (solid line) and I (dashed line) populations (scale bar: gridness of 0.5). Inset: distribution of E population gridness scores (gray bars), as well as when the inputs are "lesioned" (red bars; see Experimental Procedures). For cells in the E population with spatial gridness $\geq$0.5, grid periods are narrowly distributed (C), and spatial phases (ϕ^{α}) are approximately uniformly distributed (D).

(E) Left: distribution of the change across trials in the spatial phase of individual E cells (shown is the distribution of the magnitude of these changes, $|\Delta_t(\phi^{\alpha})|$) assuming the same initial conditions in the two trials, pooled over all cells α that appear in (C) and (D). Right: change across trials in the relative phase between E cell pairs (shown is the distribution of the magnitude of changes in relative phase, $|\Delta_t(\delta^{\alpha\beta})|$).

(F) The number of connections between cell pairs as a function of the relative phase difference between the cells, $|\delta^{\alpha\beta}|$, for the E-to-I (top), I-to-E (middle), and I-to-I (bottom) weights. Any synapse whose strength exceeds 5% of the strongest synapse of that type counts as a connection.

(G) Direction tuning scores in the mature GCN (gray bars) and after "lesion" (red bars) (see Supplemental Information; plotted for E cells in [B] and [C]).

(H) Spatial tuning in a 4 m space (four times larger than the training environment) for two cells from (A). The trajectory is a single, constant-speed unidirectional sweep (v_{rat} = 0.4 m/s) across the space.

Properties of the Mature Network

The mature GCN has aperiodic boundaries (cells at one edge of the neural sheet do not connect to neurons at the opposite, and neurons of a given preferred spatial phase are not connected with all others of the same phase). Edge cells display poor spatial tuning, but those in the bulk—a majority of the total—have high gridness scores (Figure 4A). If weak synaptic plasticity is permitted in the activation phase of the mature GCN, the same STDP rules will wire together all cells with common preferred phases, producing a network that is topologically equivalent to a single-bump patterned network with periodic boundary conditions (Guanella et al., 2007; Burak and Fiete, 2006, 2009) (Figure S4). However, plasticity in the activation phase is a strong positive feedback process and generically leads to instability and bias in the GCN dynamics (Figure S4). Therefore, it remains an open question whether there is a stable way to developmentally obtain a GCN with periodic boundaries.

The mature GCN exhibits the key signatures of (approximate) continuous attractor dynamics (Fuhs and Touretzky, 2006; Burak and Fiete, 2006, 2009; Yoon et al., 2013). For instance, (non-edge) neurons in the mature GCN exhibit very similar spatial tuning periods (Figures 4A and 4C). The distribution of spatial phases is uniform (Figure 4D), and relative phases between cells are stably preserved even when the spatial tuning of individual cells drifts over time (Figure 4E).

When the uniform tonic feedforward drive is removed, the population pattern is lost and gridness scores drop precipitously (Figure 4B). The GCN neurons instead display strongly direction-tuned, head-direction-cell-like responses (Figure 4F), consistent with results from experimental studies in which inputs to MEC were lesioned (Brandon et al., 2011; Bonnevie et al., 2013).

Cells in the mature GCN receive lateral input connections from other cells with disparate spatial phases (Figure 4G). This is the case despite the highly structured, "local" inhibitory weight profile of the mature GCN.

Spontaneous Periodic Tuning in Large, New Environments

Cells in the mature GCN generate periodic spatial responses on the very first run-through in much larger environments than the training environment (Figure 4H). This is possible because mature GCN neuron responses are generated by integration of velocity inputs, independent of external location-specific inputs whose spatial tunings would have had to be learned in new places.

The Emergence of Patterning Is Abrupt

To quantify how grid-cell-like features emerge over development, we examine several metrics of patterning and stability as a function of time during development (Figures 5 and S5).

As weights gradually reach and then exceed a threshold strength, population patterning emerges abruptly (around T_{crit} in Figure 5A) because of a weight-driven linear instability in the GCN dynamics. The population pattern is fully formed, in terms of reaching a maximal gridness score (computed on the periodic population pattern just as it is for the periodic spatial tuning of individual cells; see Supplemental Information), by about T_{patt}. To display spatial tuning, the GCN must be capable of sufficiently good path integration over the trajectory on which tuning is assessed, in addition to displaying a population pattern. Thus, the emergence of grid-like spatial tuning in cells consistently lags, and is smoother in onset, than population patterning (Figure 5B). Spatial tuning gridness scores for shorter trajectories are sharper and better probes of population patterning than

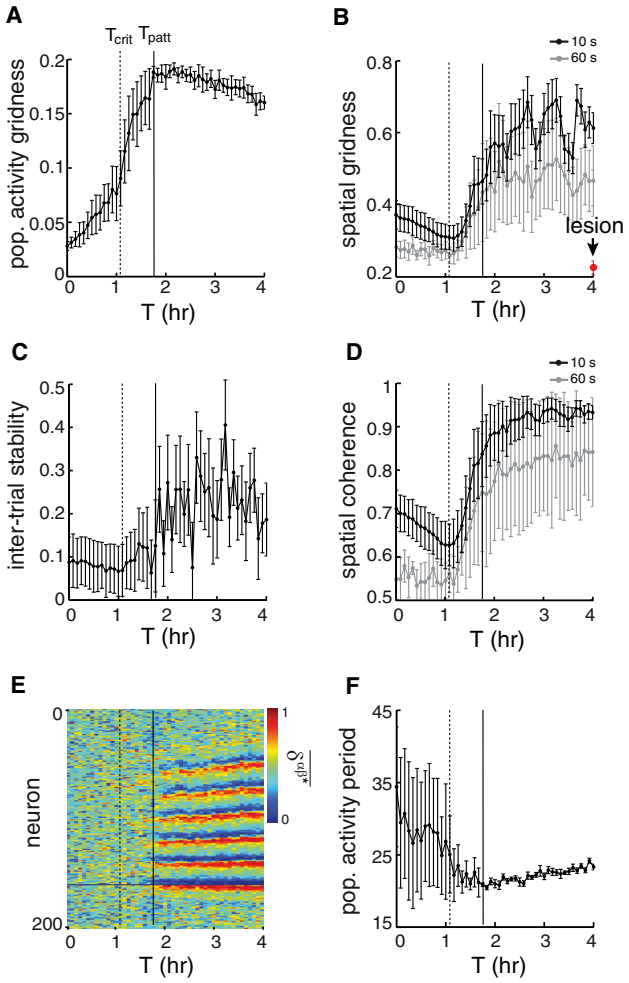

Figure 5. Abrupt Onset of Patterning

(A) Development of population activity gridness in the GCN (only E^L neurons are used in this analysis). At several points during development, the GCN is probed in the activation phase. At one such point, we collect a set of population activity snapshots (one every 5 ms, for 1 s); from this set, an average gridness score (error bars are ±1 SD here and in rest of figure) is obtained. Long vertical lines in (A)–(F) designate $T_{crit} \approx 1.08$ hr (dotted) and $T_{patt} \approx 1.75$ hr (solid), defined as the times at which the population pattern gridness score reaches its half-maximum and maximum value, respectively.

(B–D) Development of gridness (B), intertrial stability (C), and spatial coherence in the spatial tuning of cells (see Supplemental Information). Black dots ([B]–[D]) and gray dots ([B] and [D]) are average scores from a set of 10 s and 60 s trajectories (see Supplemental Information), respectively. (Average computed across n = 263 E cells and n = 10 trials. The same ten trials are used at each point in development. Cells are included if their spatial tuning scores exceed 0.5 in the mature [at 4 hr] GCN.) Red dot: average gridness of spatial tuning in the GCN with "lesioned" feedforward input.

(E) Relative phases ($\delta^{\alpha\beta^*}$; see Supplemental Information) of E^L cells (labeled by α), with respect to one reference cell (β^*; dark horizontal line marks the reference cell), averaged across trials (same 10 s trajectories as in [B]–[D]).

(F) Development of the period of the population activity pattern (in neurons, see Supplemental Information), estimated using the same data as in (A).

longer trajectories because they involve less integration (Figure 5B, black versus gray). The evolution of across-trial tuning stability and spatial coherence (how well spiking in a spatial bin is predicted by spiking in neighboring bins; see Supplemental Information) (Figures 5C and 5D) closely resembles the evolution of spatial tuning. Stability over long trajectories is evident later in development than stability over short trajectories (Figure 5D).

Early in development ($T < T_{crit}$), cells exhibit an artifactual spatial tuning over short trajectories. This is because of their strong velocity modulation coupled with the fact that variations in velocity are not averaged over space in short trajectories. As recurrent weights develop and the velocity modulation diminishes, the artifactual spatial tuning and spatial coherence decline slightly before T_{crit} (black curves in Figures 5B and 5D). This effect is weaker in longer trajectories because of averaging (gray curves). However, longer trajectories obscure the emergence of grid-like tuning because of the accumulation of path integration errors. The emergence of spatial tuning can also be obscured in data sets with fewer cells and trials (Figure S5).

Finally, the relative phase of spatial tuning between cell pairs (Figure 5E) is a good measure of population patterning, because patterning in the relative phases is arguably more abrupt and possibly emerges sooner than spatial tuning gridness or spatial coherence. The relative phases between cells, though clearly themselves patterned (in the sorted population), are not constant over development because of a slight gradual expansion of the population period after T_{patt} (Figure 5F). This expansion is partly responsible for the oscillatory variation in the spatial gridness and coherence measures late in development.

Signatures of Development in Cell-Cell Correlations and Velocity Tuning

By definition, population activity patterning involves the stable coactivation of some cells and counteractivation of others. Thus, hallmarks of population patterning should be visible in the emergence of stable cell-cell correlations. The key prediction associated with the emergence of a canonical population pattern and all its translations is that the pairwise correlation distribution should develop a uniform component.

In vitro, the velocity inputs are absent, and the GCN exhibits only very weak correlations in early development (Figure 6A, rows 1 and 2; in vitro conditions are simulated by removing velocity inputs but providing sufficient tonic drive, corresponding to pharmacological treatments that promote activity in slices). A uniform platform of correlations emerges with pattern formation (Figure 6A, row 3; especially see inset). The uniform component is not large and not fully flat because, in the absence of velocity inputs, the population pattern does not efficiently flow; as a result, distant pattern phases, and thus the larger anticorrelations, are simply not well-sampled (cf. Figure S6). The SD of the in vitro correlation distribution grows slowly (Figure 6B) even though population patterning is abrupt and the uniform component in the correlation distribution appears suddenly (insets, Figure 6A) because the amplitude of the uniform component is small.

The predicted emergence of a uniform component in the in vitro correlation distribution is consistent with MEC slice

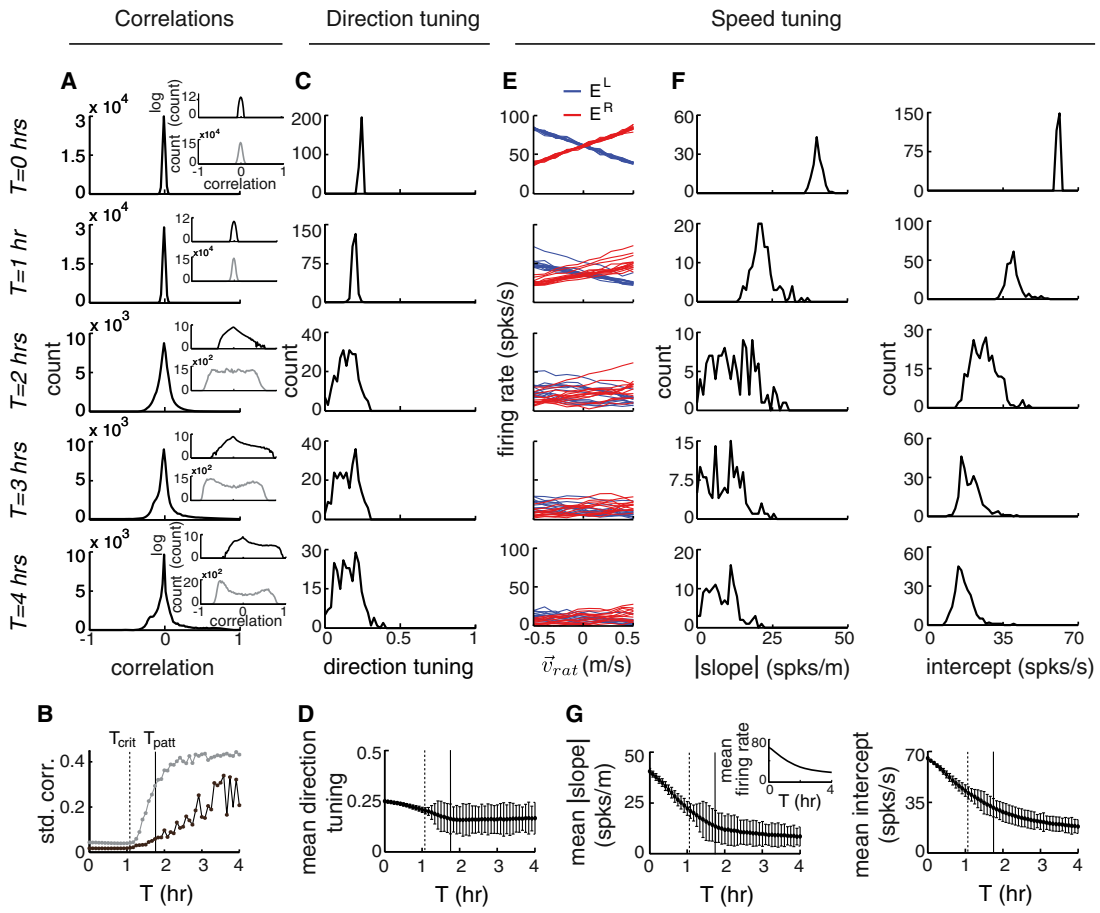

Figure 6. Windows into Development: Correlations, Direction Tuning, and Speed Tuning

(A) Main plots: development of in vitro pairwise activity correlations (Pearson's) between E cells. Pairs drawn from cells used in Figures 5B–5D; correlations assessed over a 120 s trial (see Supplemental Information; same trial and same cells used for remaining panels). Inset, Top: Semi-log version of main plots. Inset, Bottom: in vivo correlations computed only for low-velocity (v_{rat} < 0.1 m/s) trajectory segments.

(B) Finely sampled evolution of the SD of the pairwise correlation distribution for the in vitro (black) and low-velocity in vivo (gray) cases. Vertical lines as in Figure 5A.

(C) Histogram of direction tuning scores (which measure the strength of direction tuning and range between 0 and 1; see Supplemental Information) for cells in the E population.

(D) Finely sampled evolution of mean direction tuning (error bars are ±1 SD, here and in rest of figure).

(E) Examples of speed tuning in E cells; color-coded according to preferred direction.

(F) Absolute value of the slopes (left column) and intercepts (right column) of the regression lines used to fit speed tuning curves.

(G) Finely sampled evolution of the mean slopes and intercepts of speed tuning. Inset: evolution of mean firing rate for cells in E population.

data showing that (anti)correlations grow with maturation (Langston et al., 2010).

When correlations are measured in vivo, the velocity input induces strong (anti)correlations early in development that can obscure the emergence of pattern-based correlations (Figure S6). However, in vivo data can provide a reasonable proxy for in vitro correlations, and show more clearly the abrupt emergence of a uniform component in the pairwise correlations, if the correlations are based only on segments of the trajectory when the animal is moving slowly (so that the velocity input is small, thus facilitating the sampling of diverse pattern phases) (Figures 6A and 6B, gray curves).

Over GCN development, the mean strength of direction tuning decreases slightly, while the variance grows (Figures 6C and 6D). Note that inferred direction tuning (through the mean vector length statistic) depends on exploration speed: in a fixed network, faster speeds produce larger vectors (Figure S6), so comparisons across development must be made carefully, with statistically matched trajectories.

Speed tuning, or the ability of animal speed to modulate the firing rate of a cell, broken down by speed along its preferred and antipreferred directions (see Supplemental Information; Figure 6E, red and blue), decreases in strength over development, as quantified by the absolute value of the slopes and the intercepts (Figures 6F and 6G). This decrease can be attributed to the increasing influence of recurrent collaterals and diminishing influence of velocity input on cell firing over development.

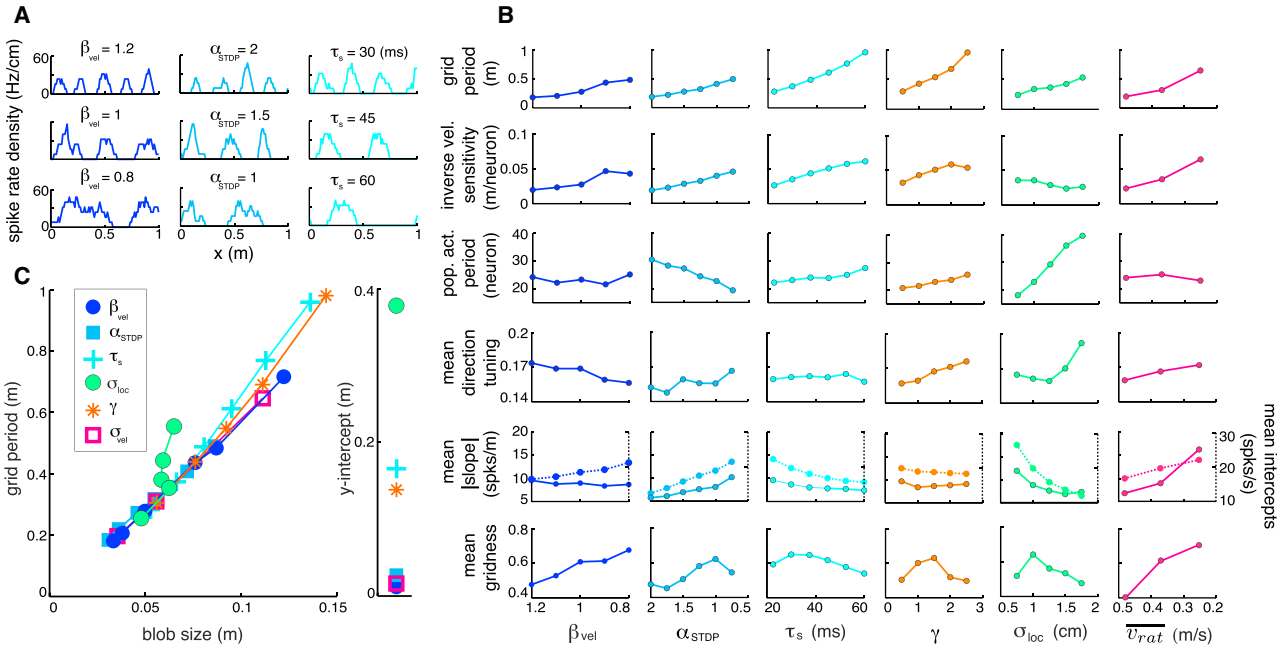

Figure 7. Mechanisms for Controlling Spatial Period across Modules

(A) Differences in the velocity input gain, β_{vel}; the scaling of the STDP time constant, α_{STDP}; and biophysical time constant, τ_s (columns 1–3, respectively; all other parameters held fixed), result in systematic variations in the spatial tuning period (assessed over a 10 s, constant-speed [0.4 m/s] sweep through the environment).

(B) Metrics of the mature GCN (rows) and how they vary as a function of parameter settings over development. Note that many scales on the abscissa are inverted. GCN metrics (in order): average spatial tuning period, inverse velocity sensitivity (see Supplemental Information), population activity period, average strength of direction tuning, average strength (slope) of absolute values of speed tuning curves, and average intercepts of the speed tuning curves. Parameters (in addition to those mentioned in [A]): the scale of learning of inhibition (γ; see Supplemental Information), the width of the location-specific developmental input (σ_{loc}), and the mean speed of exploration during development ($\overline{v_{rat}}$).

(C) Left: relationship between blob size (the widths of the activity bumps in the spatial tuning curves; see Supplemental Information) and grid period for different spatial tunings that result from varying different single parameters. Right: absolute value of the y intercept of the linear fits in the left plot. Each color marks the effects of variations in a single parameter.

If speed tuning is instead computed by averaging together changes in firing rate as a function of speed without taking into account motion direction (and thus without taking the absolute values of the firing rate versus input speed curves for different directions before averaging), the results are different (Figure S6) but largely consistent with experimental results that use this definition of speed tuning.

Multiperiod GCNs

The present model allows us to parametrically explore possible mechanisms for the differences in grid period of GCNs along the dorsoventral (DV) axis of the MEC.

If put in place over development, a smaller gain in feedforward velocity input, a smaller temporal width of the STDP windows, a longer intrinsic biophysical time constant (τ_s) in neurons or synapses, or a larger learning rate at the inhibitory synapses all result in larger-period spatial tuning (Figure 7A; Figure 7B, row 1, columns 1–4, respectively; Figure S7).

Decreasing the feedforward velocity gain leaves the population period unchanged but reduces the ability of animal velocity to translate the population pattern, thus increasing the spatial tuning period (Figure 7B, column 1, rows 1–3). In general, changing developmental parameters can affect two or more distinct properties of the GCN that influence the spatial tuning period, sometimes in opposite directions. For instance, as the width of the STDP windows decreases, the period of the population pattern shrinks, but the GCN's velocity sensitivity increases (pattern translates faster for a given animal velocity). In total, the enhanced velocity sensitivity wins out over population pattern expansion, and grid periods actually increase (Figure 7B, column 2, rows 1–3).

Although variations in any of several parameters can in theory produce a range of spatial tuning periods (Figures 7B, row 1, and S7), they do so in different and experimentally distinguishable ways. A mechanism based on velocity gain alone would predict weaker direction tuning more ventrally (Figure 7B, column 1, rows 1–3), consistent with recent reports (Giocomo et al., 2014), and unchanged lateral connectivity across GCNs. By contrast, increasing the spatial tuning period ventrally by increasing the STDP window width would predict narrower lateral inhibition (and a smaller population period), together with a reduction in the intercepts and slopes of speed tuning (Figure 7B, column 3). Increasing τ_s is predicted to decrease the intercepts and slopes of speed tuning more ventrally

(Figure 7B, column 3, rows 4 and 5) while not significantly affecting the strength of direction tuning; these effects are in contrast to the former two mechanisms. A 3-fold variation in τ_s can, in our simulations, explain most of the (roughly 10-fold) variation in grid period along the DV axis (data not shown), consistent with the experimentally estimated DV variation in the membrane and synaptic integration time constant of layer II stellate cells (Garden et al., 2008).

These and other mechanisms for varying GCN period—scaling the learning rate (or equivalently, the overall strength) of lateral inhibition, scaling the width of the location-specific inputs, or differentially scaling the velocity input across GCNs— produce additional experimentally distinguishable predictions for DV variation across modules (Figure 7B, columns 4–6). A corollary of the predicted effect of exploration speed on grid period is that if animals are forced to move more slowly during development, all grid periods would be larger than in control animals (Figure 7B, column 6).

Grid cells in MEC with different periods look like globally rescaled versions of one another; therefore, a plot of grid period against blob size should be linear, with zero intercept. Figure 7C illustrates which parameter scalings produce an appropriate relationship between grid period and blob size. We find that the velocity gain, STDP window width, and mean exploration speed parameters produce the smallest offsets from a zero-intercept line.

In contrast to the relationship between inhibition strength and grid period in the GCN, inhibition is found to be stronger more dorsally in MEC layer II stellate cells (Beed et al., 2013). In single-population grid cell models (with neurons that do not obey Dale's law; e.g., Burak and Fiete, 2009), the spatial tuning period does shrink with increasing inhibition amplitude, but in the present model, scaling up all inhibitory weights (or only the I-to-E weights; data not shown) results in larger periods both in the population pattern and in the spatial tuning (Figure 7B, column 4, rows 1 and 3). In conductance-based neuron models, inhibitory inputs can effectively shorten the biophysical time constant; if this effect is strong enough, more inhibition might result in smaller-period spatial tuning (Figure 7B, column 3, row 1). Alternatively, other variables explored above also influence grid period and may dominate over the effects of inhibition strength in MEC.

Finally, under certain experimental conditions, spatial responses can rescale essentially instantly (Barry et al., 2007, 2012). Several of the parameters considered above are related to plasticity and, because of their slower timescales, cannot be responsible for rapid rescaling. The parameters capable of driving fast change (e.g., through neuromodulation) include the feedforward velocity gain, the strength of recurrent inhibition (which developmentally was equated with the learning rate of inhibitory synapses), and the biophysical time constant of neurons.

Analysis of grid cell data in rapid rescaling experiments indicates that the underlying population pattern remains unchanged (Yoon et al., 2013). In contradiction with this empirical result, varying the strength of inhibition in the mature GCN changes the population pattern (Figure S7), as does changing τ_s (Figure S7). This lends support to the possibility that, regardless of the mechanisms underlying the gradient in spatial tuning period along the DV axis, fast rescaling is driven by a change in the feedforward gain of the velocity inputs to the mature grid cell system (Figures 7B and S7). The possibility that rapid rescaling is caused by a gain change in the feedforward velocity input can be tested by looking for variations in direction and speed tuning that are predicted to accompany such a change (Figure S7).

2D Network

The principles illustrated above for 1D GCN development carry directly over to 2D (Figure 8). There are now four sets of E cells, defined by whether the cells receive cosine-tuned inputs for north (E^U), south (E^D), east (E^R), or west (E^L) motion. The I population does not receive velocity inputs, as in the 1D GCN. Each E set contains $40 \times 40 = 1{,}600$ cells, as does the I population, so that the E:I cell ratio is 4:1. The reason to choose four distinct sets of E cells with discrete direction tuning is to illustrate the systematic weight shifts of each population. We obtain qualitatively similar functionality if each E cell is independently and randomly assigned a preferred direction from a continuous and uniform distribution over all possible angles (Figure S8). We already explored the effects of noise and variability (stochastic GCN neurons and random trajectories) on development in 1D; thus, for tractability, we employed deterministic cells and a simple trajectory to train the 2D system (see Experimental Procedures). For testing, however, we reverted to fully stochastic dynamics, as in 1D. The STDP kernels are identical to those used in the 1D GCN, and other parameters are similar (Experimental Procedures; Supplemental Information).

Cells in the mature 2D GCN display grid-like spatial tuning in response to velocity inputs (Figure 8A). Note that the trajectory is 5 min long, with no corrective mechanisms from outside the GCN to reduce the accumulation of path integration errors over this long interval. As expected, the population activity underlying spatial tuning is itself patterned (Figure 8B). The population flows in proportion to and in the direction of animal displacement (Figure 8C); therefore, the GCN performs path integration.

The evolution of in vitro correlations, direction tuning, and speed tuning in the 2D GCN qualitatively track the trends from 1D development (Figure S8; for a description and comparison of in vivo correlations, see Figure S6). In particular, the in vitro cell-cell correlation distribution is initially narrow and centered around low correlation values, but it evolves to display long tails that reflect strong pattern-related correlations (Figure S8). Also consistent with 1D is the reduction in direction and speed tuning with development (Figure S8).

To estimate the learning time for the 2D GCN with neural noise and a more natural random trajectory, we use mathematical results on the cover time of random walks in different dimensions (see Supplemental Information) and use as a starting point our learning time results for the development of a noisy 1D GCN with a random 1D trajectory (Figure S1). If T_{1D} is the learning time in 1D, then T_{2D} is estimated to be $\sim$2–4 times T_{1D}. Assuming from above that $T_{1D} \approx 4$ hr, this implies that for our parameter settings (with noisy cells and a random trajectory, as in 1D), 2D GCN development will require about 8–16 hr of plasticity during exploration.

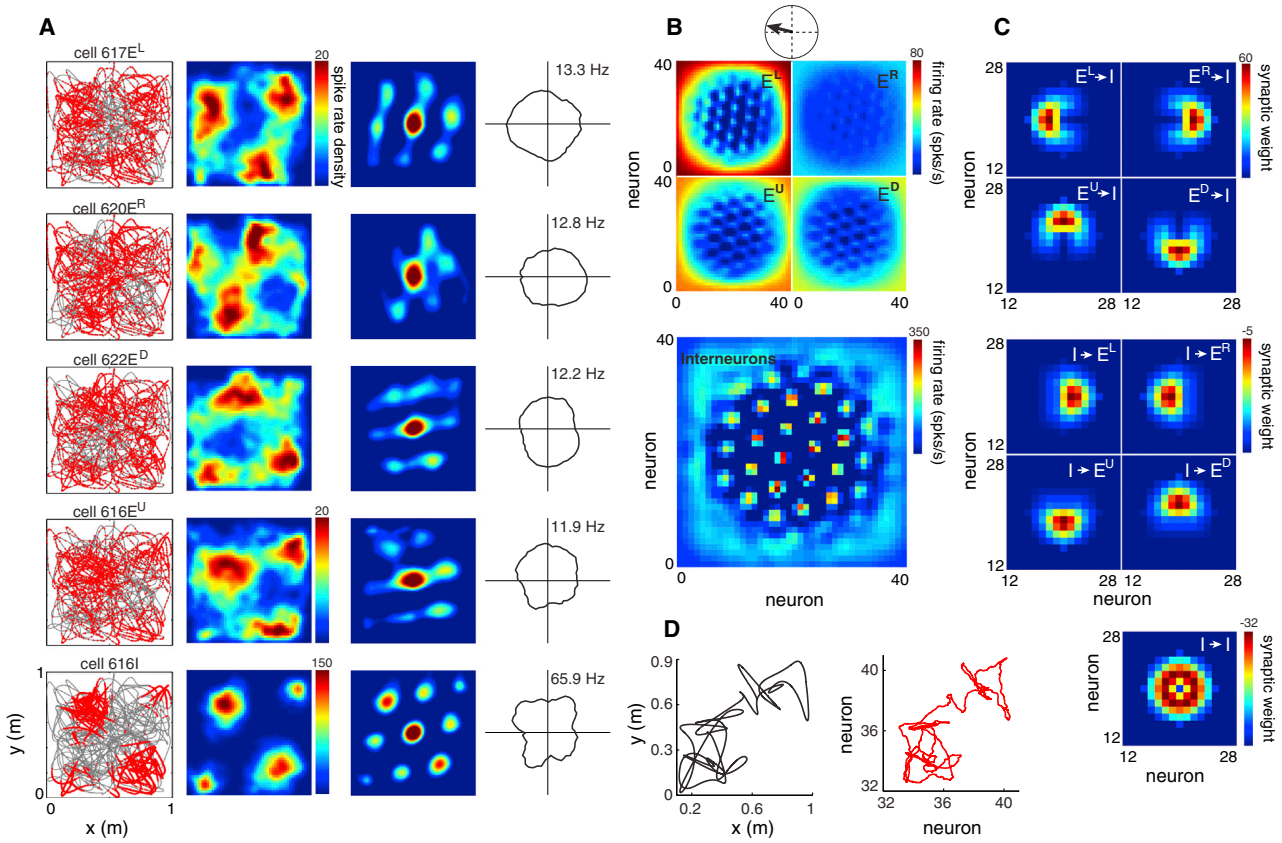

Figure 8. 2D Grid Cell Network

(A) Spatial tuning of five cells in a mature 2D GCN, following development. The trajectory is a 5 min segment taken from actual animal motion recorded in a square environment. The location-specific input used during training was removed for this trajectory. Column 1: spiking response of a cell (each red dot = 1 spike). Column 2: spatially filtered response (bin = 2 cm; convolved with Gaussian, $\sigma_x = \sigma_y = 3$ cm). Column 3: autocorrelogram of column 2. Column 4: direction tuning curve, with peak firing rates indicated.

(B) Snapshot of population activity when the animal is moving in the direction of the black vector shown above, with cells ordered according to their location-specific training inputs (E cells above; I cells at bottom).

(C) Recurrent weights in the mature GCN. Each matrix shows output projections from one cell located at the GCN center.

(D) Left: 30 s recorded animal trajectory. Right: path-integrated estimate of the same trajectory by the GCN (GCN estimate is defined as the location of one of the population pattern bumps in the neural sheet [B] as the pattern flows with input velocity).

DISCUSSION

Summary

We have presented a model for the development of GCNs with continuous attractor dynamics, a recently substantiated property of grid cells in animals (Yoon et al., 2013). This experience-dependent model is a proof-of-principle demonstration of how grid tuning and path integration functionality might arise through synaptic plasticity, with no assumptions about topographical order in the GCN. The mature cells in our model resemble layer II grid cells: Mature cells have strong grid-like spatial tuning and weak direction tuning (Sargolini et al., 2006) but lose their gridness and become strongly directional when the feedforward tonic inputs are removed (Bonnevie et al., 2013); consistent with existing data on connectivity in MEC layer II (Dhillon and Jones, 2000; Pastoll et al., 2013; Couey et al., 2013), the principal (E) cells interact only through inhibitory interneurons; over development, gridness emerges abruptly and speed tuning decreases, and the in vitro correlation strengths increase (Langston et al., 2010).

Assumptions of Model

If the role of grid cells is to estimate location during navigation, it may seem like we are putting the cart before the horse by assuming that location-specific inputs exist during GCN development. However, in our model, GCN development is a process of bootstrapping on spatially informative input derived from external cues in highly familiar cue-rich environments, so that the mature GCN becomes capable of autonomous spatial estimation in novel and cue-poor environments through integration of internal self-motion cues. This is an important gain in functionality, because a major computational challenge of navigation involves self-localization in novel spaces and across relatively featureless stretches of familiar environments.

Location-specific input to the GCN is assumed to derive from multisensory constellations of proximal and distal external cues during exploration around the familiar home nest area, possibly via other spatially tuned cell types of the hippocampal formation, including place cells (Langston et al., 2010; Wills et al., 2010), border/boundary cells (Solstad et al., 2008; Lever et al., 2009; Bjerknes et al., 2014), and landmark-specific LEC cells (Deshmukh and Knierim, 2013), or via bottom-up pathways including the visuo-spatial stream through the postrhinal cortex (Kerr et al., 2007). We found that pattern formation can proceed with sparser or less uniform location-specific inputs, but the development of translation invariance, and thus path integration, suffers.

The suppression of location-specific inputs at the boundaries of the environment during plasticity may be performed by inhibitory border cells, similar to those found in the subiculum (Stewart et al., 2014). The network ''edge'' need not be its topographic edge; it merely consists of cells that receive location-specific input from the boundaries of the training environment. These cells are edge cells, topologically speaking, because they are missing half of their potential lateral partners: the GCN contains no cells whose input preferred location is directly adjacent to the edge cells', because that location would fall outside the environment. If border cells were to simply globally inhibit the GCN at the training environment boundary, the result would be a weaker activation of the then-active GCN cells (the edge cells). Attenuated activation of cells at the GCN edge during plasticity results in a mature GCN with less pinning and more accurate path integration (even though the GCN is tested in the activation phase without attenuated input at the borders). Boundary-evoked activity suppression might also be linked to the systematic orienting of grid fields in square environments (Figure S8), as seen in experiments (Stensola et al., 2013).

The self-motion-based velocity inputs required in our model are likely derived from a combination of vestibular, optic flow, motor efference, and proprioceptive cues (Taube, 2007; Canto et al., 2012; Coogan and Burkhalter, 1993; Muir and Bilkey, 2003). The main requirements for the velocity inputs are that the direction of movement be encoded by unimodal tuning curves and that speed inputs (whether they arise from the same or a different pathway than the direction input) modulate the overall activity level of the GCN subpopulation corresponding to the present motion direction.

GCN dynamics are divided into plasticity and activity phases (in which recurrent inputs or plasticity are suppressed, respectively) to avoid the deleterious effects of positive feedback associated with STDP, in which elevated weights lead to elevated activity that further enhances the weights. Positive feedback tends to create discrete, fixed points at the cost of translation invariance. Restricting recurrent input during plasticity (or partitioning plasticity into phases of potentiation with feedforward drive and depotentiation during purely recurrent activation) breaks the positive feedback loop; such procedures are widely espoused in the learning of continuous attractors or when learning the statistics of the external world (Stringer et al., 2002; Hahnloser, 2003; Hasselmo et al., 2002; Hinton et al., 1995; Hinton, 2002). Activation may occur during some runs in the home environment and when the animal explores test environments away from home (Wills et al., 2010; Langston et al., 2010). Alternatively, the network may spend a fraction of each theta-cycle in an activation phase and another fraction in a plasticity phase, if modulators can control the alternation of recurrent synaptic transmission and plasticity on that timescale (Hasselmo et al., 2002). In either case, the neural response collected in a given environment would sample from both plasticity and activity phases and thus would reflect the contribution of recurrent weights.

Future Questions and Relationship to Existing Work

In Hahnloser (2003), a continuous attractor network is organized by supervised learning—the desired patterns are imposed on the network, and an error-driven rule makes these patterns permanent. In Stringer et al., 2002, the desired patterns are imposed on the network, and are consolidated by associative learning rules. However, to train a GCN with the desired patterns would require grid-like population pattern inputs that translate with animal motion, just like mature grid cells. Plasticity in Stringer et al. (2002) is governed by firing rates, not spikes. Modeling noisy spiking allows us to make an estimate of learning time; but other differences between spiking versus rate models are less significant because the key time-dependence of spikes in our model arises from the underlying time-varying rates, and the temporal asymmetries of our STDP rules play a functionally similar role to the asymmetric way in which presynaptic and postsynaptic neural firing rates drive plasticity in Stringer et al. (2002).

In Kropff and Treves (2008) and Mhatre et al. (2012), competitive learning rules act on location-specific inputs to produce stripe-like or grid-like spatial tuning. However, in these models, velocity inputs do not influence the network's spatial response, and the network is unable to path integrate; spatial tuning remains entirely dependent on the continued presence of location-specific inputs. In Fuhs and Touretzky (2006), associative plasticity rules acting on traveling activity waves in the neural sheet generate grid-like population patterning. However, there is no mechanism to associate activity patterns with animal location or velocity, so the model does not produce spatial representations.

The present model overcomes several conceptual hurdles to show that simple associative rules can result in the development of periodic, path-integrating neural representations with the help of inputs that do not possess such features. However, fundamental questions about how the brain could form continuous attractors remain unanswered by our and all other works on the topic: are translation invariant training inputs necessary for building a translation-invariant network? Are infinitesimally small weight changes (and thus a very large dynamic range in synaptic weights) necessary to overcome noise and nonuniformity in the inputs? Is either the suppression of weights during plasticity or an alternation of learn-and-erase phases necessary to suppress positive feedback instability? Our next steps are to attempt to understand how the brain might solve these problems.

Predictions

Our model is robust—but therefore also nonspecific—in the sense that various parameters may be varied substantially and

yet produce a functional GCN. The model is insensitive to certain modifications of the STDP windows, for which there are many suitable combinations (for one example, see Figure S1; various other examples: data not shown), although a temporal asymmetry is required. STDP windows for the three types of synapses between the E and I populations are underconstrained by the experimental data, even though there is support for the windows we used (Lu et al., 2007; Huang et al., 2013; Fino and Venance, 2010; Feldman, 2012; Holmgren and Zilberter, 2001). The GCN can be modified to model MEC layer III rather than layer II, by adding E-E connectivity and a larger feedforward velocity gain (thus increasing the strength of direction tuning; data not shown), and is robust to adding direct velocity inputs to the inhibitory cells. When velocity inputs drive I cells, these cells also develop asymmetries in their outgoing weight profiles and become direction selective. We can modify the sign and shapes of the STDP windows to generate antigrid spatial tuning (constant background firing with inactivity at every vertex of a triangular lattice) in inhibitory neurons while the E population response remains grid like (Pastoll et al., 2013). For all these reasons, the key predictions of our model are not a specific set of STDP windows or weight profiles, but certain qualitative features and parametric trends, described next.

The model predicts that spatial experience is necessary for grid cell development. While it does not specify a postnatal date for the appearance of grid cells, our results require a minimum of ~8–16 hr of spatial experience in a familiar environment (for 2D GCN formation). Restriction of spatial exploration during the normal period of grid cell development should delay or—if development occurs in a critical period that is not extended by experiential deprivation—prevent development. Changing the trajectory statistics (e.g., restricting the animal to a radial or linear track) (Figure S8) or changing the spatial metric of the environment (e.g., raising the animal on a non-Euclidean surface, as in Kruge et al., 2013) is predicted to lead to distortions in network wiring and population patterning and thus to qualitatively different single neuron spatial responses.

The prediction that velocity inputs are necessary for development is consistent with the experimental observation that head direction responses are stable before grid cells (Langston et al., 2010; Wills et al., 2010). It is unclear whether the predicted location-specific inputs arrive in MEC during development, though both place cells and border cells, which might supply such inputs, do display spatial tuning before grid cells (Langston et al., 2010; Wills et al., 2010; Bjerknes et al., 2014). Moreover, while the location-specific inputs implicitly carry all motion information (the time derivative of location is velocity), such inputs would fail, in our model, to induce the requisite asymmetries in the weights needed for path integration. Thus, an explicit velocity input is necessary for development. It follows that the replay of trajectories (e.g., in sleep), if unaccompanied by the proper velocity inputs (Brandon et al., 2012), is itself insufficient for training the GCN.

The mature GCN exhibits simple and local connectivity (the simplicity and locality can be inferred from the set of pairwise weights, even if neuron locations in the network are not topographically organized by spatial tuning) (Figure S2). Despite the local connectivity, GCN cells are predicted to project with equal frequency to cells tuned to similar and orthogonal spatial phases (Figure 4E). All neurons with direct velocity input and synapses subject to asymmetric STDP are predicted to exhibit asymmetries in their outgoing weight profiles.

The model predicts that population patterning arises abruptly. The abruptness of patterning may be assessed by spatial tuning over short trajectories. The abruptness of patterning is likely to be masked or smoothed when assessed by spatial tuning over longer trajectories or when assessed based on a limited number of neurons and trials. Thus, grid cell responses on short linear tracks might more readily reveal pattern formation.

The mean direction tuning strength and the strength of speed tuning are predicted to decline (Figures 6F and 6G) because of the growth in recurrent weights over development. If pairwise neural correlations are computed in the absence of velocity inputs to the GCN, as when the animal is at rest or in slice preparations, the model predicts a considerable broadening of the correlation distribution over development, with the emergence of a uniform platform-like component. The in vivo system can, in the presence of velocity inputs, display strong velocity-driven (anti)correlations in early development that are difficult to tell apart from population pattern-induced (anti)correlations (Figure S6).

The model also makes a number of parametric predictions about how spatial tuning in the mature entorhinal cortex depends on the biophysical time constant of cells and synapses, the STDP window widths, the strength of the feedforward velocity input, the relative gain in excitatory and inhibitory synaptic strengths, and the statistics of spatial exploration. Our results show how to experimentally discriminate between possible mechanisms underlying the variation in grid period along the DV axis of the entorhinal cortex, based not only on spatial tuning period but also on other aspects of neural response, including direction and speed tuning, population pattern period, and the scaling of grid period to blob width in different modules.

EXPERIMENTAL PROCEDURES

Roman subscripts (e.g., i, j) refer to individual cells within population P. The population index P can take the values {I, E^R, E^L} in the 1D GCN and {I, E^R, E^L, E^U, E^D} in the 2D GCN. Integration in all simulations is by the Euler method with time-step dt.

In the 1D GCN, unless otherwise noted, the trajectories used during the plasticity phase and for probes of GCN development consist of random paths across a 1D environment (see Supplemental Information). In the 2D GCN, to reduce simulation time during development, the exploration trajectory consists of fixed-speed (v_{rat} = 1 m/s) sweeps vertically and horizontally across the environment. Successive parallel sweeps are staggered (separated by a distance of 0.0125 m) to provide approximately uniform coverage of the environment. The mature 2D GCN is evaluated with velocity inputs derived from rat trajectories recorded in the open field (Barry et al., 2007).

Neural and Synaptic Dynamics

Given a summed current input $I_i^P(t)$ to the (P, i)th cell, the instantaneous firing rate of the cell is

$$r_i^P(t) = f\left(I_i^P(t)\right), \tag{1}$$

with the neural transfer function f given by

$$f(x) = \begin{cases} 0 & x \le 0 \\ x & x > 0. \end{cases} \tag{2}$$

Based on this time-varying firing rate, neurons fire spikes according to an inhomogeneous (sub-Poisson) point process with a coefficient of variance of CV = 0.5 (see Burak and Fiete, 2009 and Supplemental Information for details on generating a sub-Poisson point process).

The activation $s_i^P(t)$ of synapses from the (P, i)th cell is given by

$$\frac{ds_i^P(t)}{dt}+\frac{s_i^P(t)}{\tau_s}=\sigma_i^P(t), \tag{3}$$

where

$$\sigma_i^P(t)=\begin{cases} r_i^P(t) & \text{(deterministic dynamics)} \\ \sum_b \delta\left(t-t_{i,b}^P\right) & \text{(stochastic dynamics)}, \end{cases} \tag{4}$$

where $t_{i,b}^P$ specifies the time of the bth spike of the cell and the sum is over all spikes of the cell. For both the 1D and 2D GCN, the neural dynamics are stochastic during testing in the activation phase. During the plasticity phase, dynamics are stochastic for the 1D GCN and deterministic for the 2D GCN.

The total synaptic current $I_i^P(t)$ into the (P, i)th cell is given by

$$I_i^P(t)=A_i^P\alpha^{P,vel}(\vec{v}(t))\left(g_i^{P,rec}(t)+g_i^{P,loc}(t)+g^0\right)+A_i^P g^{0'}, \tag{5}$$

where g^0 and $g^{0'}$ are small, positive, constant bias terms common to all cells; $g_i^{P,rec}$ are the recurrent inputs; $g_i^{P,loc}$ are the location-specific inputs; $\alpha^{P,vel}$ are the velocity inputs; and A_i^P is the suppressive envelope function. The recurrent input is

$$g_i^{P,rec}(t)=\sum_{P'}\sum_{j=1}^{N}W_{ij}^{PP'}s_j^{P'}(t), \tag{6}$$

where $W_{ij}^{PP'}$ are the recurrent weights. The location-specific input is a Gaussian bump of height $W^{P,loc}$ and width σ_{loc}:

$$g_i^{P,loc}(\vec{x}(t))=W^{P,loc}\exp\left[\frac{-\left|\left|\vec{x}(t)-\vec{x}_i^P\right|\right|^2}{2\sigma_{loc}^2}\right], \tag{7}$$

where $\vec{x}(t)$ is the location variable and $\vec{x}_i^P$ is the preferred location of the input to the (i, P)th cell. (In the 1D GCN, location and preferred location preference are scalars.) The preferred locations are evenly distributed over the unit interval in R^1 and R^2 for the 1D and 2D GCNs, respectively.

All cells in the Pth population (with preferred direction given by the unit vector $\hat{e}^P$) receive a common velocity input:

$$\alpha^{P,vel}(\vec{v}(t))=1+\beta^{vel}\vec{v}(t)\cdot\hat{e}^P \tag{8}$$

where $\vec{v}(t)$ is velocity of the animal and β^{vel} sets the gain of the velocity input; $\hat{e}^P=(0,0),(0,1),(0,-1),(1,0),(-1,0)$ for the I, E^U, E^D, E^R, E^L populations, respectively. The multiplicative influence of velocity inputs on the rest of the inputs to each cell may be viewed as a shunting effect (Heeger, 1992) from inputs that arrive simultaneously at different parts of the dendritic tree and combine nonlinearly.

The envelope function, which is only applied in the familiar home environment (and not during testing, although performance would improve during testing if the envelope were also applied then), is given by Burak and Fiete (2009):

$$A_i^P=\begin{cases} 1 & X_i^P<1-\Delta X \\ \exp\left[-a_0\left(\frac{X_i^P-1+\Delta X}{\Delta X}\right)^2\right] & \text{otherwise} \end{cases} \tag{9}$$

where $X_i^P=\left|\left|\vec{x}_i^P-(0.5,0.5)\right|\right|$ (in 1D, $X_i^P=\left|x_i^P-0.5\right|$), ΔX determines the range over which tapering occurs, and a_0 controls the steepness of the tapering.

Plasticity Rule and Development

The recurrent weights $W_{ij}^{PP'}$ are drawn initially from a uniform distribution on the interval $[0,\pm w_0]$ (− when the presynaptic cell is inhibitory, and + when it is excitatory). We assume no direct E-to-E connections; thus, these weights are 0. During plasticity in the home environment, g^0, $g^{0'}$, and $g_i^{P,rec}$ are set to zero; neural activity is based only on the feedforward inputs (which are tapered at the edges according to the envelope function A_i^P). Weights are incremented during exploration according to the following:

$$\Delta W_{ij}^{PP'}(t)=\eta\gamma_{PP'}\left[\int_0^\infty\sigma_i^P(t)k_+^{P'}(\tau)\sigma_j^{P'}(t-\tau)d\tau+\int_0^\infty\sigma_i^P(t-\tau)k_-^{P'}(\tau)\sigma_j^{P'}(t)d\tau\right], \tag{10}$$

where η is the learning rate, $\gamma_{PP'}$ is a term of order 1 that controls the relative speed of learning for the different types of recurrent weights, and σ_i^P is given by (4). There are two learning kernels, k^E and k^I, depending on whether the presynaptic cell is excitatory or inhibitory, respectively (Figure 1). The causal (+) and acausal (−) sides of these two learning kernels are given by:

$$k_+^E(t)=Ae^{-t/2\alpha_{STDP}\tau_{STDP}} \tag{11}$$

$$k_-^E(t)=-e^{-t/1.5\alpha_{STDP}\tau_{STDP}}; \tag{12}$$

$$k_+^I(t)=Be^{-t/2\alpha_{STDP}\tau_{STDP}} \tag{13}$$

$$k_-^I(t)=-e^{-t/\alpha_{STDP}\tau_{STDP}}. \tag{14}$$

The coefficients A and B control the relative magnitudes of the two sides of the learning kernels, and the time constant τ_{STDP}, with a scale factor α_{STDP} of order 1, controls the widths.

In the testing condition (activation phase), η is set to zero and the envelope is removed; all figures probing GCN dynamics, Figures 3–8, are generated in the activation phase.

1D Simulation Parameters

$N_I = 80$; $N_E = 400$ (200 per E population); CV = 0.5; dt = 0.5 ms; τ_s = 30 ms; σ_{loc} = 1 cm; $W^{I,loc} = 50$; $W^{E,loc} = 10$; $g^0 = 50$ (= 1 for "lesioned" feedforward input); $g^{0'} = 15$ (= 0 for "lesioned" feedforward input); $\beta^{vel} = 0.9$; $w_0 = 0.001$; $\eta = 0.015\ s^{-1}$; $\gamma_{II} = 7$; $\gamma_{EI} = 2$; $\gamma_{IE} = 1$; $\tau_{STDP} = 12$ ms; $\alpha_{STDP} = 1$; A = 1.2; B = 0.5; $\Delta r = 0.72$; $a_0 = 60$.

2D Simulation Parameters

$N_I = 1{,}600$; $N_E = 6{,}400$ (40 × 40 = 1,600 per E population); $\beta^{vel} = 2$; $w_0 = 0$; $\eta = 0.012\ s^{-1}$; $\gamma_{II} = 5$; $\gamma_{EI} = 0.25$. $\gamma_{IE} = 16.7$; $\Delta r = 36$; $a_0 = 10$. All other parameters identical to 1D parameters.

SUPPLEMENTAL INFORMATION

Supplemental Information includes eight figures and Supplemental Experimental Procedures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2014.06.018.

ACKNOWLEDGMENTS

We are grateful to Yoram Burak, Michael Buice, Dori Derdikmann, Mike Mauk, Matt Nolan, Haim Sompolinsky, and Michael Marder for helpful comments and to Tom Wills for discussions about the home nest environment. I.R.F. is a McKnight Scholar and an ONR Young Investigator (ONR YIP N00014-13-1-0529) and acknowledges support from the NSF through NSF-EAGER IIS-1148973.

Accepted: June 6, 2014
Published: July 16, 2014

REFERENCES

Barry, C., Hayman, R., Burgess, N., and Jeffery, K.J. (2007). Experience-dependent rescaling of entorhinal grids. Nat. Neurosci. *10*, 682–684.

Barry, C., Ginzberg, L.L., O'Keefe, J., and Burgess, N. (2012). Grid cell firing patterns signal environmental novelty by expansion. Proc. Natl. Acad. Sci. USA *109*, 17687–17692.

Beed, P., Gundlfinger, A., Schneiderbauer, S., Song, J., Böhm, C., Burgalossi, A., Brecht, M., Vida, I., and Schmitz, D. (2013). Inhibitory gradient along the dorsoventral axis in the medial entorhinal cortex. Neuron *79*, 1197–1207.

Bjerknes, T.L., Moser, E.I., and Moser, M.B. (2014). Representation of geometric borders in the developing rat. Neuron *82*, 71–78.

Blair, H.T., Gupta, K., and 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.

Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J.L., Roudi, Y., Moser, E.I., and Moser, M.B. (2013). Grid cells require excitatory drive from the hippocampus. Nat. Neurosci. *16*, 309–317.

Brandon, M.P., Bogaard, A.R., Libby, C.P., Connerney, M.A., Gupta, K., and Hasselmo, M.E. (2011). Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. Science *332*, 595–599.

Brandon, M.P., Bogaard, A.R., Andrews, C.M., and Hasselmo, M.E. (2012). Head direction cells in the postsubiculum do not show replay of prior waking sequences during sleep. Hippocampus *22*, 604–618.

Burak, Y., and Fiete, I. (2006). Do we understand the emergent dynamics of grid cell activity? J. Neurosci. *26*, 9352–9354, discussion 9354.

Burak, Y., and Fiete, I.R. (2009). Accurate path integration in continuous attractor network models of grid cells. PLoS Comput. Biol. *5*, e1000291.

Burgess, N., Barry, C., and O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. Hippocampus *17*, 801–812.

Canto, C.B., Koganezawa, N., Beed, P., Moser, E.I., and Witter, M.P. (2012). All layers of medial entorhinal cortex receive presubicular and parasubicular inputs. J. Neurosci. *32*, 17620–17631.

Coogan, T.A., and Burkhalter, A. (1993). Hierarchical organization of areas in rat visual cortex. J. Neurosci. *13*, 3749–3772.

Couey, J.J., Witoelar, A., Zhang, S.J., Zheng, K., Ye, J., Dunn, B., Czajkowski, R., Moser, M.B., Moser, E.I., Roudi, Y., and Witter, M.P. (2013). Recurrent inhibitory circuitry as a mechanism for grid formation. Nat. Neurosci. *16*, 318–324.

Deshmukh, S.S., and Knierim, J.J. (2013). Influence of local objects on hippocampal representations: Landmark vectors and memory. Hippocampus *23*, 253–267.

Dhillon, A., and Jones, R.S. (2000). Laminar differences in recurrent excitatory transmission in the rat entorhinal cortex in vitro. Neuroscience *99*, 413–422.

Doeller, C.F., Barry, C., and Burgess, N. (2010). Evidence for grid cells in a human memory network. Nature *463*, 657–661.

Domnisoru, C., Kinkhabwala, A.A., and Tank, D.W. (2013). Membrane potential dynamics of grid cells. Nature *495*, 199–204.

Feldman, D.E. (2012). The spike-timing dependence of plasticity. Neuron *75*, 556–571.

Fiete, I.R., Burak, Y., and Brookings, T. (2008). What grid cells convey about rat location. J. Neurosci. *28*, 6858–6871.

Fino, E., and Venance, L. (2010). Spike-timing dependent plasticity in the striatum. Front. Synaptic Neurosci *2*, http://dx.doi.org/10.3389/fnsyn.2010.00006.

Fuhs, M.C., and Touretzky, D.S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. J. Neurosci. *26*, 4266–4276.

Fyhn, M., Hafting, T., Witter, M.P., Moser, E.I., and Moser, M.-B. (2008). Grid cells in mice. Hippocampus *18*, 1230–1238.

Garden, D.L.F., Dodson, P.D., O'Donnell, C., White, M.D., and Nolan, M.F. (2008). Tuning of synaptic integration in the medial entorhinal cortex to the organization of grid cell firing fields. Neuron *60*, 875–889.

Giocomo, L.M., Stensola, T., Bonnevie, T., Van Cauter, T., Moser, M.B., and Moser, E.I. (2014). Topography of head direction cells in medial entorhinal cortex. Curr. Biol. *24*, 252–262.

Gorchetchnikov, A., and Grossberg, S. (2007). Neural Netw. *21*, 182–193.

Guanella, A., Kiper, D., and Verschure, P. (2007). A model of grid cells based on a twisted torus topology. Int. J. Neural Syst. *17*, 231–240.

Haas, J.S., Nowotny, T., and Abarbanel, H.D.I. (2006). Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. J. Neurophysiol. *96*, 3305–3313.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801–806.

Hahnloser, R.H.R. (2003). Emergence of neural integration in the head-direction system by visual supervision. Neuroscience *120*, 877–891.

Hasselmo, M.E., Bodelón, C., and Wyble, B.P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput. *14*, 793–817.

Hasselmo, M.E., Giocomo, L.M., and Zilli, E.A. (2007). Grid cell firing may arise from interference of theta frequency membrane potential oscillations in single neurons. Hippocampus *17*, 1252–1271.

Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. Vis. Neurosci. *9*, 181–197.

Hinton, G.E. (2002). Neural Comp. *14*.

Hinton, G.E., Dayan, P., Frey, B.J., and Neal, R.M. (1995). The "wake-sleep" algorithm for unsupervised neural networks. Science *268*, 1158–1161.

Holmgren, C.D., and Zilberter, Y. (2001). Coincident spiking activity induces long-term changes in inhibition of neocortical pyramidal cells. J. Neurosci. *21*, 8270–8277.

Huang, S., Huganir, R.L., and Kirkwood, A. (2013). Adrenergic gating of Hebbian spike-timing-dependent plasticity in cortical interneurons. J. Neurosci. *33*, 13171–13178.

Kerr, K.M., Agster, K.L., Furtak, S.C., and Burwell, R.D. (2007). Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. Hippocampus *17*, 697–708.

Killian, N.J., Jutras, M.J., and Buffalo, E.A. (2012). A map of visual space in the primate entorhinal cortex. Nature *491*, 761–764.

Kropff, E., and Treves, A. (2008). The emergence of grid cells: Intelligent design or just adaptation? Hippocampus *18*, 1256–1269.

Kruge, I.U., Wernle, T., Moser, E.I., and Moser, M.B. (2013). Grid cells of animals raised in spherical environments. Society for Neuroscience.

Langston, R.F., Ainge, J.A., Couey, J.J., Canto, C.B., Bjerknes, T.L., Witter, M.P., Moser, E.I., and Moser, M.-B. (2010). Development of the spatial representation system in the rat. Science *328*, 1576–1580.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. J. Neurosci. *29*, 9771–9777.

Lu, J.T., Li, C.Y., Zhao, J.-P., Poo, M.M., and Zhang, X.H. (2007). Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type. J. Neurosci. *27*, 9711–9720.

Mathis, A., Herz, A.V., and Stemmler, M. (2012). Optimal population codes for space: grid cells outperform place cells. Neural Comput. *24*, 2280–2317.

McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.-B. (2006). Path integration and the neural basis of the 'cognitive map'. Nat. Rev. Neurosci. *7*, 663–678.

Mhatre, H., Gorchetchnikov, A., and Grossberg, S. (2012). Grid cell hexagonal patterns formed by fast self-organized learning within entorhinal cortex. Hippocampus *22*, 320–334.

Muir, G.M., and Bilkey, D.K. (2003). Theta- and movement velocity-related firing of hippocampal neurons is disrupted by lesions centered on the perirhinal cortex. Hippocampus *13*, 93–108.

Pastoll, H., Solanka, L., van Rossum, M.C.W., and Nolan, M.F. (2013). Feedback inhibition enables θ-nested γ oscillations and grid firing fields. Neuron *77*, 141–154.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.-B., and Moser, E.I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science *312*, 758–762.

Schmidt-Hieber, C., and Häusser, M. (2013). Cellular mechanisms of spatial navigation in the medial entorhinal cortex. Nat. Neurosci. *16*, 325–331.

Solstad, T., Boccara, C.N., Kropff, E., Moser, M.-B., and Moser, E.I. (2008). Representation of geometric borders in the entorhinal cortex. Science *322*, 1865–1868.

Sreenivasan, S., and Fiete, I. (2011). Grid cells generate an analog error-correcting code for singularly precise neural computation. Nat. Neurosci. *14*, 1330–1337.

Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.B., and Moser, E.I. (2012). The entorhinal grid map is discretized. Nature *492*, 72–78.

Stensola, T., Stensola, H., Moser, M.B., and Moser, E.I. (2013). Environmental constraints on grid cell orientation. Society for Neuroscience.

Stewart, S., Jeewajee, A., Wills, T.J., Burgess, N., and Lever, C. (2014). Philos. Trans. R. Soc. Lond. B Biol. Sci. *369*, http://dx.doi.org/10.1098/rstb.2012.0514.

Stringer, S.M., Trappenberg, T.P., Rolls, E.T., and Araujo, I.E.T. (2002). Network. Comput in Neural Sys *13*, 217–242.

Taube, J.S. (2007). The head direction signal: origins and sensory-motor integration. Annu. Rev. Neurosci. *30*, 181–207.

Wills, T.J., Cacucci, F., Burgess, N., and O'Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. Science *328*, 1573–1576.

Yartsev, M.M., Witter, M.P., and Ulanovsky, N. (2011). Grid cells without theta oscillations in the entorhinal cortex of bats. Nature *479*, 103–107.

Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., and Fiete, I.R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. Nat. Neurosci. *16*, 1077–1084.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. J. Neurosci. *16*, 2112–2126.