Neuron, Volume *83*

Supplemental Information

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

John Widloski and Ila R. Fiete

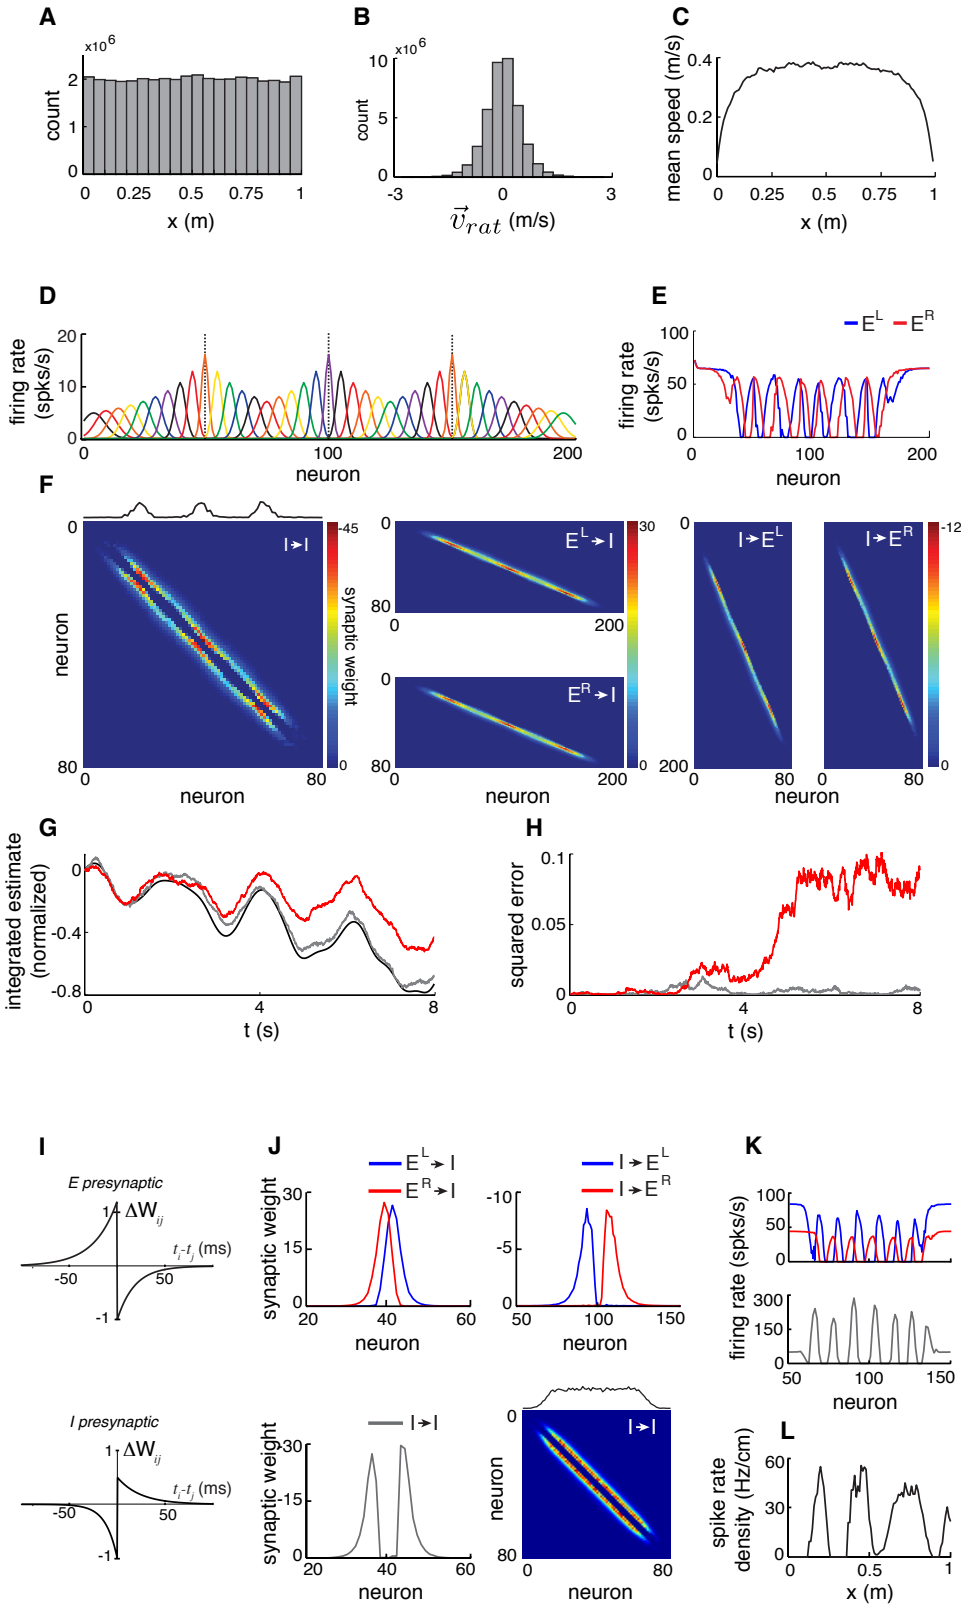

Figure S1: See caption on next page.

Figure S1: **Supplemental figures associated with Figure 1.** (A-C) *Statistics of the (quasi)random 1D trajectories generated for development.* (A) Histogram of visited locations in 1-meter enclosure for the entire training trajectory (4 hours). (B) Histogram of instantaneous animal velocities for training trajectory (mean speed, $\overline{v_{rat}} = 0.36$ m/s). (C) Mean speed as a function of position within enclosure, computed from training trajectory. (D-H) *Development with non-uniform location-specific inputs.* (D) Snapshots of the location-specific input during training. Both the width and amplitudes of the location-specific inputs are varied, across locations. The three vertical black lines mark landmark locations. The standard deviations of location-specific inputs are proportional to the distance from the nearest landmark. (E) Snapshot of the excitatory L and R population activities from the mature GCN, after training. (F) Synaptic weights of the mature GCN. The curve above the I-to-I matrix shows the off-diagonal, which reveals the non-uniformity in the weights. (G) Path integration performance: Red curve: neural representation of location over an 8-second animal trajectory (black curve), from the GCN trained on non-uniform location-specific inputs, as in (D). Gray curve: neural representation of location by mature GCN from main text (trained with uniform visual inputs). The red and gray curves represent integrated location in GCN coordinates; to find scale-factors to convert each curve into location coordinates, we simply performed a least-squares linear regression (with zero y-intercept) on the first second of each curve onto the actual location (black curve). (H) Squared error, measured point-by-point, between the normalized GCN trajectories in (G) and the true animal trajectory. (I-L) *GCN development with time-reversed plasticity windows: Hebbian STDP in I-to-E synapses and anti-Hebbian STDP in E-to-I synapses.* (I) Time-reversed plasticity kernels (relative to Figure 1D) used for development. All other parameters identical to 1D model parameters from main text. (J) Top: E-to-I (left) and I-to-E (right) synaptic weight profiles in the mature GCN (at T = 4 hours). Note reversal of weight asymmetries as compared to Figure 2A-B, row 4 in main text. Bottom: I-to-I synaptic weight profile (left) and full connectivity matrix (right). (K) Snapshot of the E (top) and I (bottom) population activity patterns in the mature GCN in the activation phase. (L) Sample spatial tuning curve from a single neuron in the E population, measured over a 10-second variable-velocity trajectory spanning the enclosure (see SI).

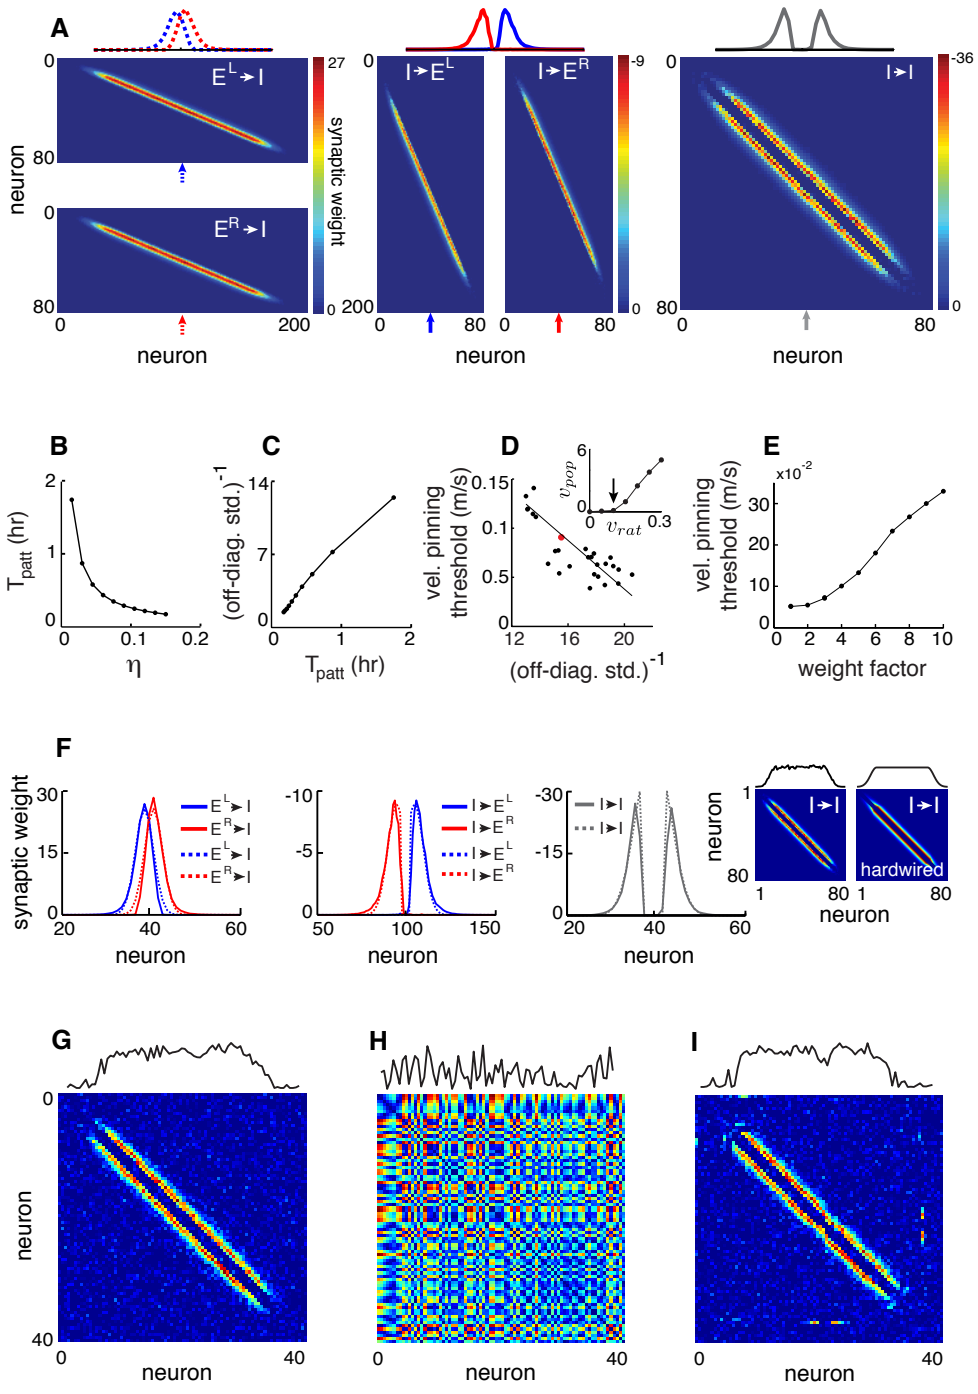

Figure S2: See caption on next page.

Figure S2: **Supplemental figures associated with Figure 2** (A) Synaptic weight matrices corresponding to the 1D GCN and figures in the main text, learned under conditions in which the visual cues are uniform (contrast with Figure S1). Top: Slices from the weight matrices, taken at the locations of the color-coded arrows. (B-F) *Tradeoffs in translation invariance and learning time.* (B) The population patterning time, T_{patt} (see Figure 5), increases as the learning rate parameter, η, decreases. (C) Translation invariance (assessed by the inverse standard deviation of the off-diagonal of the I-to-I synaptic weight matrix of the GCN at T_{patt}; see SI) increases as T_{patt} increases. (D) The pinning threshold (defined as the maximum animal speed that fails to elicit a translation of the population activity pattern, and given by the threshold in the GCN's velocity response as a function of input animal speed, arrow in inset; see SI), drops with increasing translation invariance. Black line: linear fit to results (Pearson's correlation: r = -0.75 ($p<0.01$)). (E) In a hard-wired network with weights based on the mature GCN but with hand-imposed translation invariance (at least down to the the discreteness of single neurons; see (F) and SI), the pinning threshold grows as the synaptic weights are scaled up in size. *Weight factor* is a multiplicative scaling; a 1 corresponds to weights of same strength as in the mature GCN of Figure 2. (F) Hardwired GCN with translation-invariant weights used in (E). Left three panels, dotted lines: Cross-sections of the E-to-I (first panel), I-to-E (second panel), and I-to-I (third panel) synaptic weight profiles in the hardwired network. Solid lines: the mature 1D GCN weights profiles from the main text. Rightmost two panels: weight matrices of the mature (left) and hardwired (right) GCNs. The hardwired network is translation-invariant (down to the resolution of single-neuron discreteness) by construction, but otherwise is matched to the mature GCN. (G-I) *Discovery of simple architectures in weight matrix data from a non-topographic GCN.* (G) I-to-I synaptic weight matrix organized according to inherited preferred location. To each element of the matrix, independent white noise samples (drawn from a normal distribution $\mathcal{N}(0, a^2)$, where a is 0.5 times the standard deviation of the distribution of weights in the I-to-I synaptic matrix) have been added. A cross section of the matrix is plotted above. (H) I-to-I synaptic weight matrix as it would appear if the inherited location preferences were unknown, as expected from a connectomics dataset generated from a network that is not topographically organized. (I) Matrix in (H) reordered via simulated annealing, sorted by synaptic weight size (see SI).

Interpretation of results in (B-E): The learning time T_{patt} can be shortened with larger learning rate (B), but at a fixed weight threshold the result is less translation invariance (C) and more pinning during velocity integration (D). If the network weights continue growing after the weight threshold for pattern formation is crossed, to increase translation invariance, this can also hurt integration even if the final weights are translation-invariant (E), simply because stronger weights than necessary tend to enhance pinning. These observations illustrate the tradeoff between a fast approach to network patterning and the final quality of integration in the network, and suggest that it is optimal for the learning rate to be low enough so that at threshold for pattern formation, enough traversals have been performed for translation invariance.

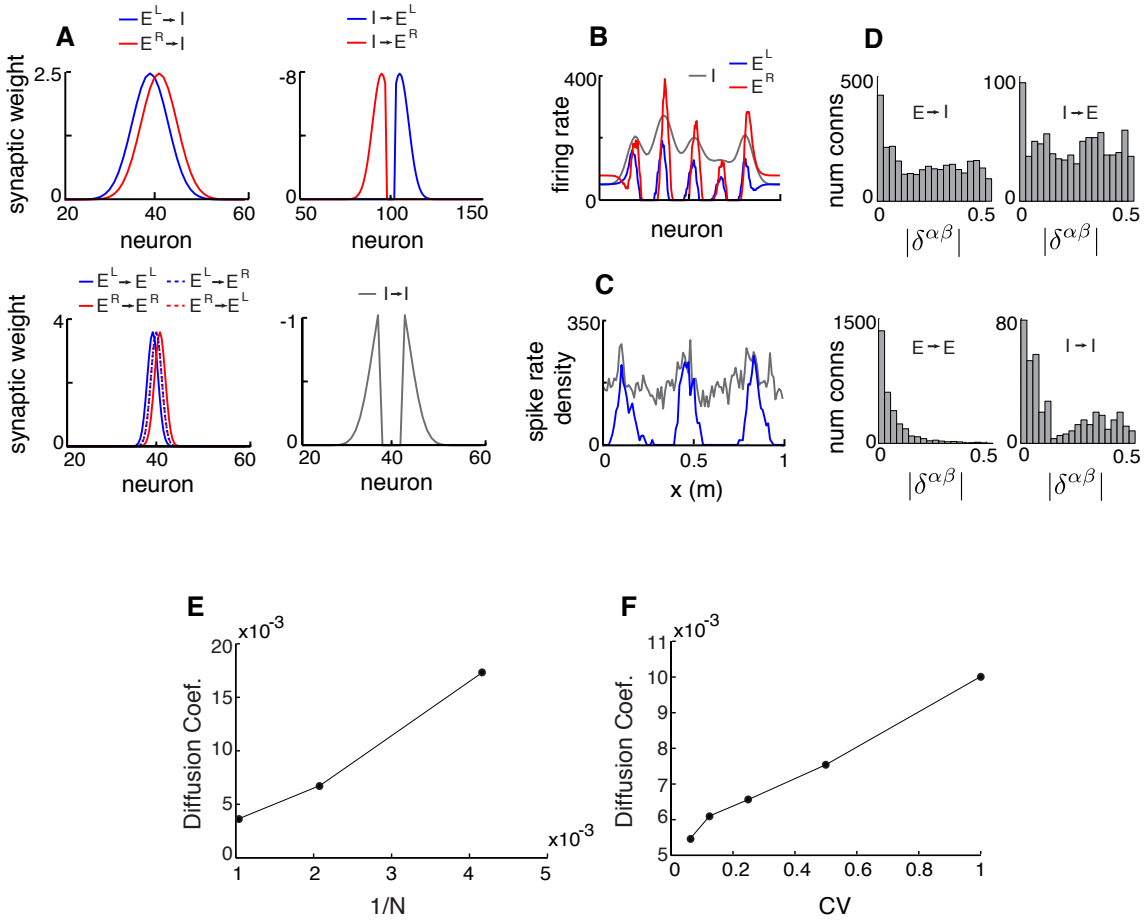

Figure S3: **Supplemental figures associated with Figure 3.** (A-D) *E-to-E connectivity enables grid-like spatial tuning in the E population with little tuning in the I population (compare with Figure 3).* (A) E-to-I (top left), I-to-E (top right), E-to-E (bottom left), and I-to-I (bottom right) synaptic weight profiles in a modified version of the mature GCN of the main manuscript (see SI). (B) Snapshot of the E (red and blue) and I (gray) population activity patterns. (C) Sample spatial tuning curves from non-edge neurons of the E (blue) and I (gray) populations, measured over a 10-second trajectory. (D) The number of connections between cell pairs as a function of the magnitude of the relative phase between the cells, $|\delta^{\alpha\beta}|$, for the E-to-I (top left), I-to-E (top right), E-to-E (bottom left), and I-to-I (bottom right) weights. A synapse whose strength is >= 5% of the largest strength for that synapse type counts as a connection. The distribution of spatial phases of the E cells that project to any I cell is broad and I cells display weak or no spatial tuning, consistent with the findings in Buetfering et al., 2014. (E-F) *Parametric quantification of integration performance in the mature GCN.* (E) Diffusivity or the diffusion coefficient (i.e. the slope of the mean squared displacement of the population pattern trajectory, measured as the center of one of the activity bumps over time, as a function of temporal displacement) of the population pattern at zero input velocity, as a function of GCN size (N = 480 is the size of the GCN in the main text). Diffusivity characterizes how much (squared) integration error is introduced by ongoing noise in the GCN per unit time. (F) Diffusion coefficient as a function of the coefficient of variance (CV) of the spiking process (see Experimental Procedures and SI).

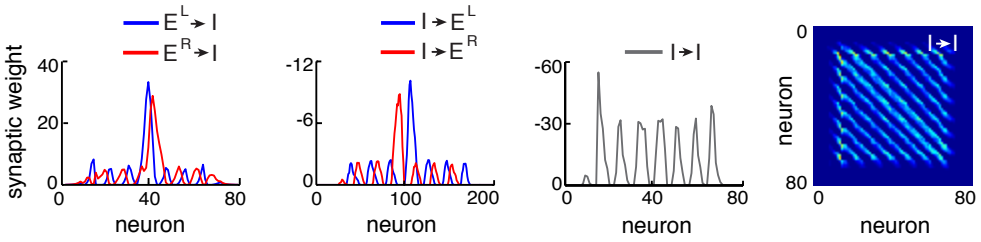

Figure S4: **Supplemental figures associated with Figure 4.** *Development of a GCN with periodic connectivity.* Beginning with the mature 1D GCN, we allow synaptic modification to occur through the same STDP rules by releasing the restriction on plasticity during the activation phase. During this period, the learning rate is sharply diminished, and we have added adaptation to the neural dynamics (see SI for details). These modifications are all applied at T = 4 hours in the original simulations. With these modifications, at T = 5 hours, cells with similar spatial phases all wire together, and the GCN becomes topologically periodic: E-to-I (first panel), I-to-E (second panel), and I-to-I (third panel) synaptic weight profiles, and the full I-to-I synaptic weight matrix (fourth panel).
Notes: This kind of learning is unstable because of strong positive feedback that results from the expression of recurrent weights during their activity-dependent plasticity. The instability creates biases in the previously symmetric or previously balanced weight profiles, driving pattern flow in the absence of velocity inputs. The added neural adaptation also works against faithful integration. Thus, it remains to be seen whether GCNs with periodic structure are plausible.

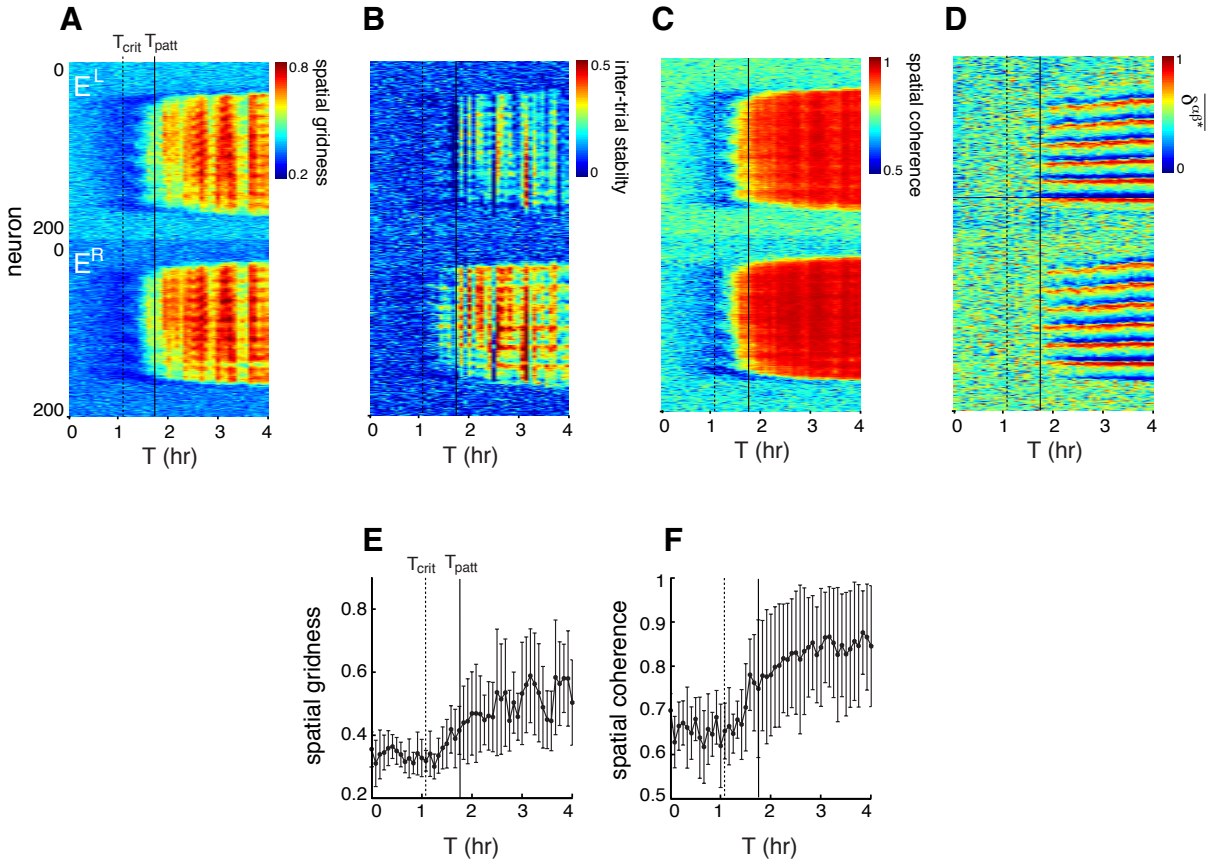

Figure S5: **Supplemental figures associated with Figure 5.** (A-D) *Developmental scores for the entire E population.* (A) Spatial gridness, (B) inter-trial stability, (C) spatial coherence, and (D) relative phase, $\delta^{\alpha\beta_*}$ (β^* indicates fixed reference cell, which is marked by the dark horizontal line in the figure, see SI), as a function of cell number in the E population (y-axis) and development time (x-axis). Each element in a matrix is the score for a particular cell at a particular development time point, averaged over 10 trials (trials are 10-second, random trajectories; the same set of random trajectories are used to test the GCN at each developmental time point). (E-F) *Developmental scores obtained from a small number of trials and neurons have poorer resolution in assessing the onset of pattern formation.* The development of (E) spatial gridness and (F) spatial coherence, averaged across randomly chosen cells in the E population (n = 15; drawn from the pool of cells with gridness score > 0.5) and over n = 4 trials (error bars are $\pm$ one standard deviation). Vertical lines at same times as in Figure 5A. Compare with Figure 5.

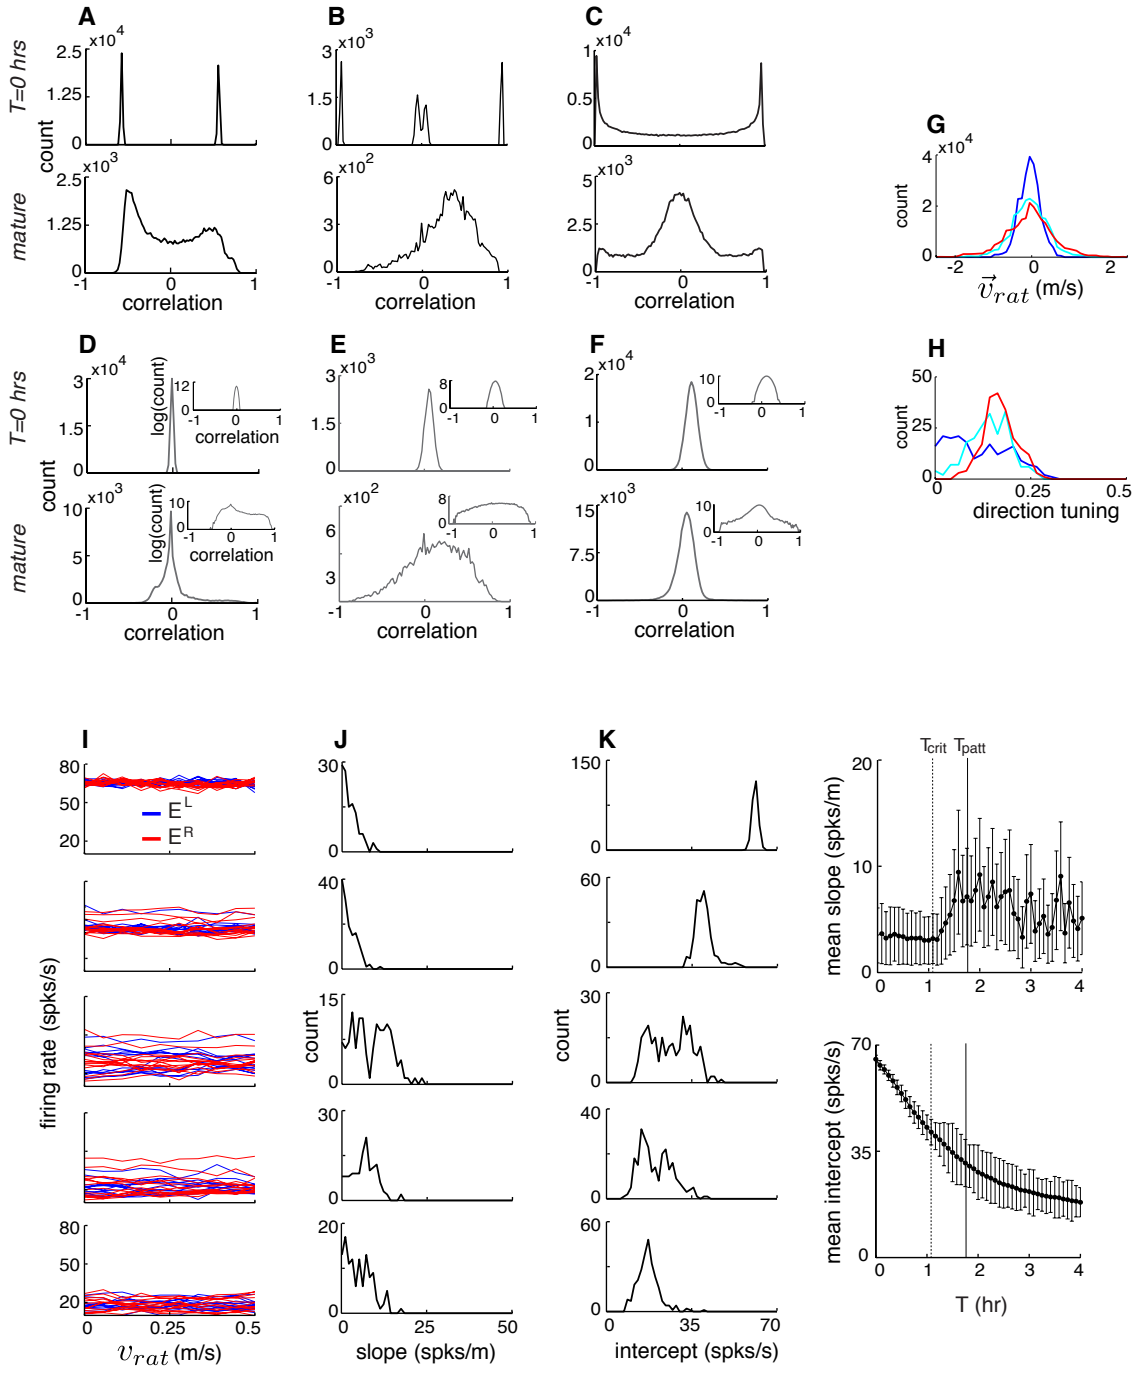

Figure S6: See caption on next page

Figure S6: **Supplemental figures associated with Figure 6.** (A-F) *The development of pairwise correlations between excitatory neurons in 1D and 2D GCNs.* The distribution of pairwise correlations in: left column: the 1D GCN; middle column: the 2D GCN with 4 different directionally tuned populations; right column: the 2D GCN with omnidirectional tuning preferences. (A-C) *In vivo* correlations, in the presence of velocity inputs. (A) The immature 1D GCN (top) exhibits strong correlations and anticorrelations in the form of two sharp peaks, because of the common, dominant velocity inputs to the two populations. The mature 1D GCN (bottom; T = 4 hours) has gained a clear uniform component. (B) The immature 2D GCN with 4 directional populations exhibits three peaks (top): strong velocity-driven correlations within-population, strong anticorrelations between opposing direction populations, and zero correlations for orthogonal direction populations. The mature GCN (bottom) has gained a uniform component. (C) The immature 2D GCN with arbitrary direction preference in individual neurons already exhibits a uniform component (top), because of the evenly distributed velocity-driven direction preferences. In the mature GCN(bottom; after completion of the training trajectory), the velocity-driven components have shrunk and the uniform component has grown, but the gain in the uniform component in this GCN is quantitative rather than qualitative. (D-F) The same correlations as in (A-C), but in a simulated *in vitro* condition, in which the velocity inputs are removed. Here, the velocity-induced correlations have vanished, and all immature GCNs simply display a peak at zero (top plots). Similarly, all the mature GCNs display a uniform component related to pattern formation, that is uncomplicated by velocity correlations (bottom plots). The uniform component does not extend as far toward ± 1 as in the *in vivo* conditions, because in the absence of velocity drive, the pattern can only translate based on random fluctuations, and does not uniformly translate over all phases. [Correlation distributions are computed from a 2-minute long trajectory in which the animal executes movements (*in vivo*) (in 2D, the trajectory is experimentally derived, see Experimental Procedures; in 1D, the trajectory is random, see SI), or over trajectory segments where speed is lower than a threshold (boxed insets), or when the velocity input is zero (*in vitro*). Correlations in (A,D) are measured between E cells whose 1D spatial gridness scores in the mature GCN is greater than 0.5 (see SI); correlations in (B,E) are measured between E cells whose 2D spatial gridness score (see SI) in the mature GCN exceeds 0; correlations in (C,F) are measured between cells drawn randomly from the E population.] (G-H) *Mean direction tuning score as a measure is not independent of the statistics of the trajectory used to assess it.* (G) Distribution of velocities for three different 2-minute 1D trajectories. (H) The distribution of direction tuning scores measured in the mature GCN for the three trajectories in (G) changes substantially with mean speed of the trajectory. Direction tuning scores should therefore only be compared across experiments with caution. (I-K) *Modulation of firing rate by animal speed, as computed in Wills* et al.*, 2012.* (I) Sample speed tuning curves, binned according the animal's speed (movement in the preferred and anti-preferred directions are conflated), for different cells in the E population, color-coded according to the cell's direction preference. Contrast with Figure 6C (see SI). Cells and trajectory used to generate spike trains same as in Figure 6C. (J) Snapshots of the distribution of the slopes (left column) and intercepts (right column) of the regression lines used to fit speed tuning curves for the cells in the E population. Computed from same cells and trajectory as in Figure 6D. (K) Development of the mean of the speed tuning slope (top) and intercepts (bottom) (error bars are $\pm$ one standard deviation), as in Figure 6G.

Interpretation of results in (K): In Wills et al., 2012, it is reported that the mean slopes and intercepts decrease with development. This is true of our data if considered from point T_{patt} onwards; after T_{patt}, spatial tuning is still evolving to become grid-like, and the mean slope and intercept do decrease with development (Figure 5B-D).

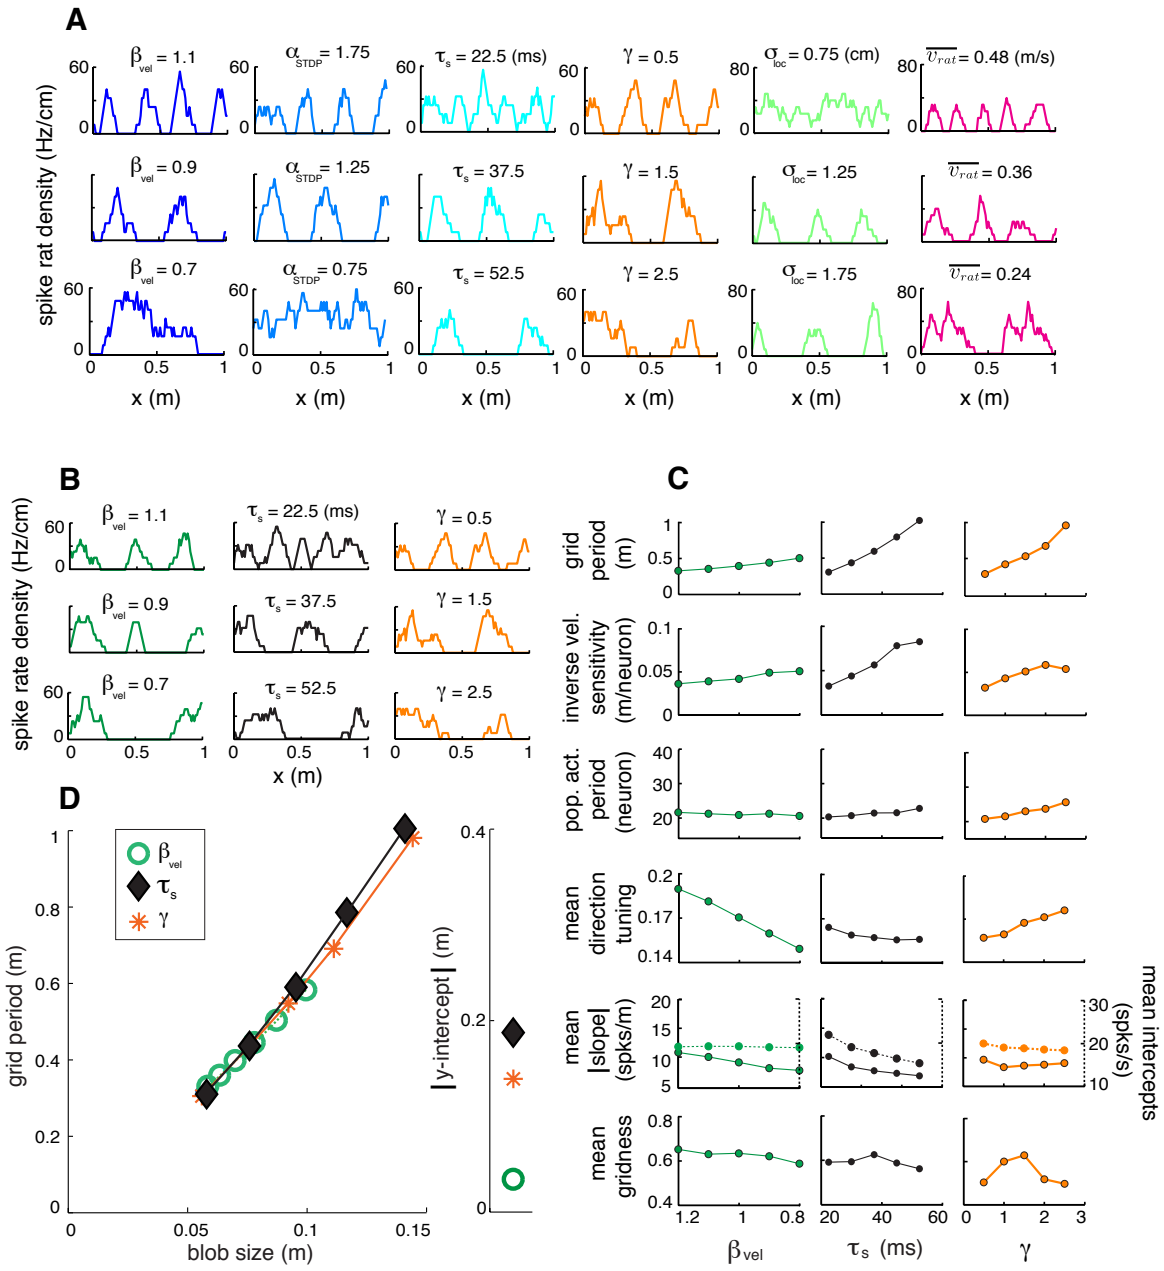

Figure S7: See caption on next page.

Figure S7: **Supplemental figures associated with Figure 7.** (A) Sample spatial tuning curves for GCNs developed with different parameters, color-coded according to Figure 3. (B-D) *Rapid rescaling of grid cell tuning: possible mechanisms.* (B) Columns: The effects of varying individual parameters in the mature GCN. The gain of the velocity input (β_{vel}; column 1), the biophysical time-constant of neurons (τ_s; column 2), and the strength of inhibition (γ; column 3) are varied individually (see Experimental Procedures). (C) Systematic variation in properties of the GCN in response to changing parameters in the mature GCN. Rows: Mean grid period; inverse GCN velocity sensitivity (see SI); population pattern period; average direction tuning score; average of the magnitudes of slopes (solid lines) and intercepts (dotted lines) of the regression lines that fit the speed tuning data; mean gridness score. (D) Left: Relationship between blob size and grid period, as different parameters are individually varied. Inset: Absolute value of the y-intercept of the linear fits in the left plot.

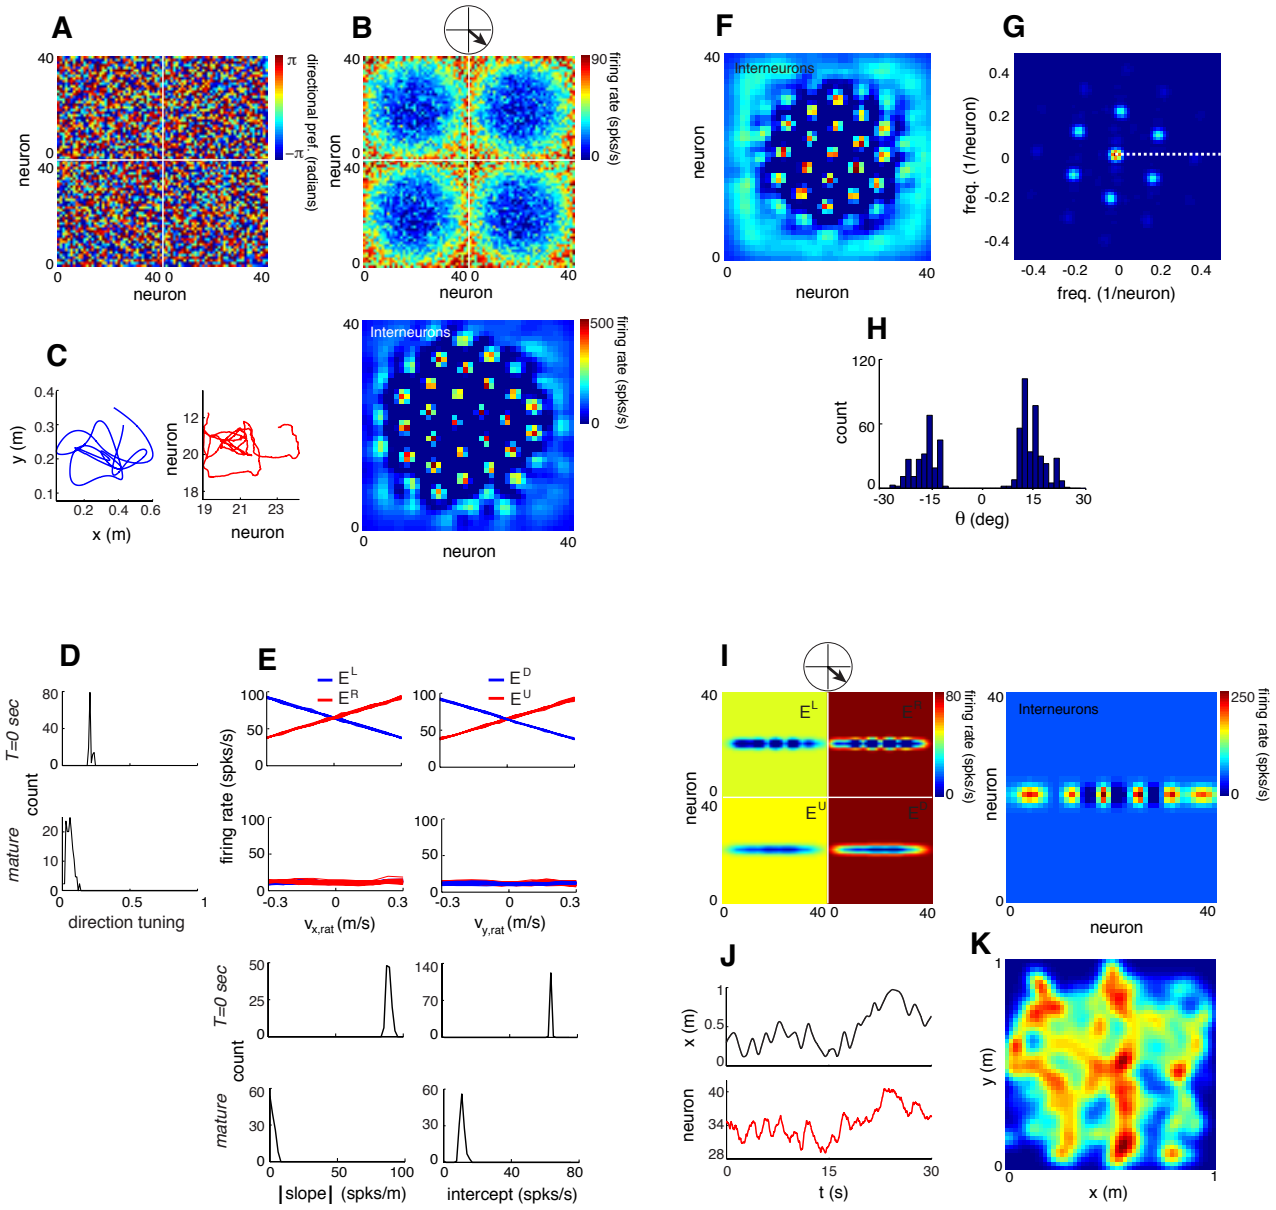

Figure S8: See caption on next page.

Figure S8: **Supplemental figures associated with Figure 8.** (A-C) *2D GCN with randomly assigned, omnidirectional direction preferences.* (A) Direction preferences of the four excitatory populations (a distinction that is not meaningful now that a population does not correspond to a specific direction preference; we preserve a partitioning of cells into four groups simply to keep figures maximally similar for comparison). Each cell in each population is assigned a direction preference randomly drawn from the uniform distribution over the interval $[-\pi, \pi)$. (B) Snapshot of the population activities of the excitatory (top) and inhibitory (bottom) GCNs after training. Training trajectory and parameters identical to 2D GCN simulation details in SI. (C) Left: 30-second animal trajectory. Right: GCN estimate of location. Same parameters and training trajectory as described in Experimental Procedures, except that $\beta_{vel} = 3.5$, $\gamma_{II} = 8.3$, $\gamma_{EI} = 0.25$, and $\gamma_{IE} = 20.8$. (D-E) *Development of direction and speed tuning statistics in the 2D grid cell GCN.* Statistics in the immature (at T = 0 s; top) and mature GCN (at the end of learning – see Experimental Procedures; bottom) of Figure 8 from the main manuscript. Scores obtained from a 2-minute long trajectory in which the animal executes movements according to a recorded trajectory from a randomly foraging rat (see Experimental Procedures). Data shown for cells with gridness score> 0 in the mature GCN (see SI). (D) Direction score distribution. (E) Speed tuning. Top: Speed tuning curves sampled from the population of cells in (D) (speed is binned along the axis of the cell's preferred direction); the distribution of slopes (bottom left) and intercepts (bottom right) include all of the cells in (D). (F-H) *Systematic orienting of the 2D population activity pattern in the mature GCN.* (F) Snapshot of the inhibitory population activity pattern. (G) Power spectrum of the population activity pattern in (F). (H) Distribution of angles (measured with respect to the horizontal) of the two non-central peaks closest to the positive x-axis in the power spectrum, over several trials (n = 10) and over time within each trial (n=10 per trial). In each trial, the GCN activity is seeded with random initial conditions; as the population activity quickly becomes patterned, the GCN is driven by a random 10-second trajectory segment from a recorded animal trajectory (as in Figure 8). Angles are sampled every 0.5 seconds for the last 5 seconds of each trial). The distribution shows that *de novo* patterns form with a narrow range of orientations, and the patterns, once formed, rotate at most within a narrow range. The square-like GCN edges induce a particular orientation in the population pattern, directed away from the horizontal ($\approx \pm 15^\circ$). Intriguingly, grid cell responses in square environments, reported in Stensola et al., 2013), tend to align to the square environment with approximately the same orientation ($\approx 7 - 18$ degrees). We wonder whether inhibitory border cells activated in the physical environment (Stewart et al., 2014) somehow impose constraints on the allowable orientations in the population pattern by associating environmental boundaries with network edges, and thus impose a systematic orienting of the grid fields. (I-K) *2D GCN in animal confined to a narrow (1D) corridor during development.* (I) Snapshot of the population activities of the 2D GCN after development in a confined 1D environment (for example, a long, narrow track in which the external 2D world is not visible, where the animal can only move along a single line on the track, with textures and landmarks to distinguish different parts of the track). The training trajectory consists of a sweep from left to right and back at speed $v_{rat} =$ 1 m/s. (J) During random motion in an open 2D environment (30-second trajectory), the GCN is able to accurately integrate and thus track displacements along only one dimension of the animal's 2D trajectory. This dimension corresponds to displacements of the population activity pattern along its one extended dimension. Black line (top): x-coordinate of animal displacement. Red line (bottom): center of one of the activity bumps in right panel of (I). (K) Smoothed rate map of one cell during random navigation of the full 2D space (5 minutes). Note the banded structure of the spatial response. Same parameters and training trajectory as described in Experimental Procedures, except that $\beta_{vel} =$ 3.5, $\eta = 0.12$, $\gamma_{II} = 4$, $\gamma_{EI} = 0.25$, and $\gamma_{IE} = 21.6$.

Interpretation of results in (I-K): This result shows that if a GCN or a subpopulation within a GCN receives only 1D velocity inputs (even if the animal actually explores 2D space), that subpopulation will form stripe cells (resembling those seen in Krupic et al., 2012). Raising an animal on non-Euclidean surfaces, for instance on a sphere (e.g. as done in Kruge et al., 2013), should distort the relationship between velocity, time, and distance travelled and should lead to corresponding conformal distortions in GCN wiring and population patterning. These conformal distortions, as well as others that might arise from other unusual conditions during development (e.g. other non-Euclidean surfaces and radial-arm mazes), can result in complex changes in spatial tuning on flat 2D surfaces, that can be directly queried in our developmental model.

Supplemental Experimental Procedures

Scores and measures

Network velocity response: The GCN velocity response quantifies the translation velocity of the population pattern, v_{pop} (neuron/s), as a function of input animal velocity, v_{rat} (m/s) (see Figure S2D, inset). Animal velocity is held constant for 10 seconds as the location of one of the peaks in the population pattern is tracked; animal velocity is incremented in steps (increments of 0.05 m/s). Within a step, the mean velocity of the population pattern is calculated over the 10-second window, after filtering the population pattern trajectory with a 2-second moving average filter. We fit a line to the plot of GCN velocity vs. animal velocity, for all data points with v_{pop} >0.1. *Velocity sensitivity* is the slope of this line; the *velocity pinning threshold* is the x-intercept of this line.

Estimation of translation invariance: Translation invariance is measured as the inverse of the normalized local standard deviation of an off-diagonal row of the weight matrix (in this case, the middle half of the 5th off-diagonal of the W_{II} matrix). Normalization is by the mean value.

Estimation of spatial tuning: The spatial tuning curve of a neuron is measured by first building a histogram of locations when the cell spikes (bin size = 1 cm). This histogram is then normalized (divided bin-by-bin) by the time spent in that spatial bin. The normalized histogram is smoothed by convolving with a boxcar filter (width = 5 bins).

Estimation of spatial tuning period and blob size: To measure *period*: For each cell, the power spectrum of the cell's spatial tuning curve for a single trial is computed. The period is taken to be the wavelength at which the power spectrum has the largest peak. To measure *blob size*: a Gaussian is fit to the central peak (i.e., the points between the troughs immediately to the left and right of the central peak) in the auto-correlogram of the cell's spatial tuning curve. The blob size is taken as the standard deviation of this Gaussian.

1D spatial gridness score: There is no clear 1D version of the commonly used 2D gridness score (see e.g. Sargolini et al., 2006). We developed a 1D gridness measure for spatial tuning. (Generalizing our 1D gridness score to 2D yields results that are similar to those derived from the common 2D gridness score, as in Sargolini et al., 2006 (data not shown)). For each cell, we compute the power spectrum of its normalized spatial tuning curve (normalization involves mean subtraction and division by the standard deviation). The power spectrum is rescaled by $2/L^2$, where L is the number of bins in the tuning curve (1 cm bin size yields $L = 100$). The spatial gridness is taken to be the power of the largest frequency component in the power spectrum. This normalization yields a gridness score of 1 if the the spatial tuning curve is a perfect sinusoid.

2D spatial gridness score: We use the same gridness score for 2D spatial tuning as used in Sargolini et al., 2006. Given the autocorrelogram of a cell's spatial tuning curve, we define an annular region with inner and outer radii, R_i and R_o (R_i and R_o are chosen by hand such that the annulus only contains the 6 peaks closest to the origin, excluding the peak at the center). Then, the Pearson correlation coefficient is computed between this map and the same map rotated by an angle $\phi = i$ degrees, where i is an integer that increments from 1 to 180; this defines a vector of correlation coefficients,

$\rho(\phi)$. The gridness score is defined as

$$gridness = \min\{\rho(60^\circ), \rho(120^\circ)\} - \max\{\rho(30^\circ), \rho(90^\circ), \rho(150^\circ)\}. \quad (1)$$

Inter-trial stability score: For each cell, the Pearson correlation coefficient is computed between the spatial tuning curves associated with two different trials. Trials begin with the same initial populaton phase, and involve two separate trajectories selected from a larger random trajectory (see below on the generation of a the random trajectory) to have similar starting locations (with a difference no greater than 1 cm).

If n_i is the spike rate density in the i^{th} bin of the cell's spatial tuning curve for a given trial, then the cell's inter-trial stability, ρ_{ITS}, measured between trials t_1 and t_2 is given by

$$\rho_{ITS} = \frac{\sum_{i=1}^{L}(n_i^{(t_1)} - \overline{n}^{(t_1)})(n_i^{(t_2)} - \overline{n}^{(t_2)})}{\sqrt{\sum_{i=1}^{L}(n_i^{(t_1)} - \overline{n}^{(t_1)})^2}\sqrt{\sum_{i=1}^{L}(n_i^{(t_2)} - \overline{n}^{(t_2)})^2}} \quad (2)$$

where L is the number of bins in the spatial tuning curve and $\overline{n}$ is the mean spike rate density in a given trial.

Spatial coherence: For each cell, the Pearson correlation coefficient is measured between the normalized spatial tuning curve and the same spatial tuning curve, but whose ith element is replaced by the mean of the 8 nearest elements (excluding itself). Formally, if n_i is the spike rate density in the i^{th} bin of the cell's spatial tuning curve, then coherence, ρ_{coh}, is given by

$$\rho_{coh} = \frac{\sum_{i=1}^{L}(n_i - \overline{r})(m_i - \overline{m})}{\sqrt{\sum_{i=1}^{L}(n_i - \overline{n})^2}\sqrt{\sum_{i=1}^{L}(m_i - \overline{m})^2}} \quad (3)$$

where L is the number of bins in the spatial tuning curve and

$$m_i = \frac{1}{M}\sum_{\substack{j=-\frac{M}{2} \\ j\neq 0}}^{\frac{M}{2}} n_{i+j}, \quad (4)$$

and M is the number of bins included in the average.

Population activity gridness and period: Similar to the spatial gridness, the population activity gridness is taken to be the power of the largest frequency component of the power spectrum measured from a normalized snapshot of the population activity (normalized = mean subtracted, followed by division by standard deviation) The power spectrum is rescaled by the factor $2/L^2$, where L is the number of bins in the population activity vector from which the power spectrum was computed. The population activity vector is shortened to include only the middle one-half of the population, so that for the E^L population, L is 100. From the power spectrum, the population activity period is taken to be the wavelength at which the power spectrum has the largest peak.

Spatial phase and relative *phase*: The phase of a cell is the normalized offset in the central peak from zero in the the cross-correlation of the spatial response of the cell of interest and a reference spatial

response (the reference is the first-trial spatial tuning of a particular cell in the same GCN; it is held fixed across trials). If d is the offset in the central peak of the cross-correlation, the spatial phase of the α^{th} cell in a given trial is defined as

$$\phi^{\alpha} = \frac{d}{\lambda} \bmod 1, \tag{5}$$

where λ is the cell's spatial tuning curve period (see above) measured from that trial, and $\phi^{\alpha} \in [0, 1)$. The *relative* spatial phase between cells α and β in a given trial is given by

$$\delta^{\alpha\beta} = (\phi^{\alpha} - \phi^{\beta}) \bmod 1. \tag{6}$$

For any phase ϕ, phase magnitude is defined as $|\phi| = \min\{\phi, 1 - \phi\}$, where $|\phi| \in [0, 0.5)$. If ϕ_{t_1} and ϕ_{t_2} are phases measured in two separate trials, then the change in phase is defined as

$$\Delta_t(\phi) = (\phi_{t_1} - \phi_{t_2}) \bmod 1. \tag{7}$$

To compute the development of relative spatial phase as in Figure 5E, at each developmental time point, and for each trial at that time point, the relative phase $\delta^{\alpha\beta_*}$ is computed for each cell α with respect to a fixed reference cell β_*, and then averaged over all trials.

Pairwise temporal correlation coefficient: The pairwise temporal correlation coefficient is the Pearson's correlation coefficient measured between spike trains for a pair of cells, after first convolving each spike train with with a Gaussian kernel ($\sigma = 15$ ms for the 1D GCN and $\sigma = 270$ ms for the 2D GCN).

Direction tuning: First, spikes are binned according to the animal's heading direction (in 1D, there are two directions and thus two bins, left and right; in 2D, the binning is over $[0°, 360°)$, with bin size $\Delta = 6°$) and then normalized (divided bin-by-bin) by the time spent in that bin. This gives a directional spike rate density. The mean direction vector is

$$\rho_{dir} = \frac{\sum_{\theta=0}^{360} r_{\theta} e^{\imath\theta \frac{\pi}{180}}}{\sum_{\theta=0}^{360} r_{\theta}}. \tag{8}$$

where r_{θ} is the spiking rate in the bin defined by $[\theta, \theta + \Delta]$. The mean vector length, i.e., the direction tuning strength, is given by $|\rho_{dir}|$.

Speed tuning: To compute the *speed tuning* of a given cell, spikes are binned by animal speed and running direction. Separate bins are assigned for the same speed if along different directions (the preferred and anti-preferred directions); bin = 0.1 m/s. Within each bin, we compute the mean number of spikes. We then compute the slopes and intercepts of the resulting mean spike rate per bin versus speed bin plots, for each direction. For the speed tuning of Figure S6, in the manner of Wills et al., 2012, firing rates are binned according to the absolute value of the animal's velocity, i.e., the preferred and anti-preferred directions are combined. The rest is the same as above.

Generating spike trains with CV < 1.

To generate stochastic spikes with a lower CV given by CV $= 1/\sqrt{M}$ (with M an integer greater than 1), we use the following procedure, unmodified, from Burak and Fiete, 2009: First subdivide

each time interval dt into M sub-intervals of length dt/M. In each of these finer time-steps $[t, t+dt)$, generate k spikes according to an inhomogeneous Poisson process with instantaneous rate $\lambda = Mr(t)$:

$$P(k|\lambda) = \frac{e^{-\lambda dt}(\lambda dt)^k}{k!}. \tag{9}$$

Finally, go through the spikes in order of emission, counting them and retaining only every Mth spike. This procedure generates a spike train with rate r and CV $= 1/\sqrt{M}$.

Generating a quasi-random 1D exploration trajectory.

The quasi-random trajectory is generated as follows. Starting from an initial location, a trajectory segment is drawn by picking a velocity and a time interval from uniform distributions on the intervals [-1,1] m/s and [0,0.02] seconds, respectively, and integrating them to obtain position coordinates over the interval. If any part of this segment would extend outside the enclosure (i.e., outside the [0,1]-meter interval), a different velocity and time interval are picked, until one is found that respects the enclosure confines. At the end of that segment, another is picked, and so on. A 4-hour duration trajectory is assembled in this way. The resulting trajectory is smoothed with a moving average filter of width 1s (2000 bins). Because smoothing leads to possible expansion of the trajectory outside the bounds of the enclosure, the smoothed trajectory is globally rescaled so that the minimum and maximum are contained within [0,1] m. This process determines animal location $\vec{x}(t)$, and by its derivative, the velocity $\vec{v}(t)$ (see Figure S1). The exploration speeds in this 1D trajectory are consistent with recorded animal exploration speeds in 2D.

Because of the steps taken to ensure that the trajectory remains within the enclosure, it is technically quasi-random rather than random. For simplicity, we refer to it as random in the main text and elsewhere.

The quasi-random 10-second trajectories used for GCN testing in Figures 4, 5, and 6, were extracted from the longer developmental trajectory such that, within 10 seconds, the animal begins within the same spatial interval, $[x_0, x_0 + \delta x]$, where $x_0 = 0.1$ and $\delta x = 0.01$, and touches both boundaries. Also, the trajectories are selected to be non-overlapping. The 60-second trajectories used in Figure 5 are just the extensions of the set of 10-second trajectories taken from the developmental trajectory. The 2-minute long trajectory used in Figure 6 is taken to be the first 2 minutes of the developmental training trajectory.

1D hardwired model with translation-invariant weights.

The weights going from population P' to P, and from cells i and j, are described as follows (see Figure S2):

$$W^{PP'}(i,j) = A^{PP'}(i,j)\Theta(\mu(i-\gamma j))\Theta(|i-\gamma j|-\delta)\alpha\left[\exp\left[\frac{-(i-\gamma j-\Delta)^2}{2\sigma^2}\right] + \beta\exp\left[\frac{-(i-\gamma j+\Delta)^2}{2\sigma^2}\right]\right], \tag{10}$$

where $\gamma = \frac{N_P}{N_{P'}}$ (N_P is the size of population P), Θ is the Heaviside function ($\Theta(x) = 0$ for $x < 0$ and is 1 otherwise), and A is an envelope function that tapers the weights, where $A^{PP'}(i,j) = A_i^P A_j^{P'}$,

$$A_i^P = \begin{cases} 1 & r_i^P < \eta N_P \\ \exp\left[-a_0 \left(\frac{r_i^P - \eta N_P}{(1-\eta)N_P}\right)^2\right] & \text{otherwise} \end{cases} \tag{11}$$

and $r_i^P = |i - \frac{N_P}{2}|$, $\eta = 0.28$, and $a_0 = 60$. The resulting weights are translation invariant (minus the tapering at the boundaries).

$E^L \to I$: $\alpha = 52$; $\beta = 0$; $\Delta = -1$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;
$E^R \to I$: $\alpha = 52$; $\beta = 0$; $\Delta = 1$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;
$I \to E^L$: $\alpha = -9$; $\beta = 0$; $\Delta = 4$; $\sigma = 5$; $\mu = 1$; $\delta = 3$;
$I \to E^R$: $\alpha = -9$; $\beta = 0$; $\Delta = -4$; $\sigma = 5$; $\mu = -1$; $\delta = 3$;
$I \to I$: $\alpha = -60$; $\beta = 1$; $\Delta = 2$; $\sigma = 3$; $\mu = 0$; $\delta = 3$;

In Figure S3, the weight parameters are as follows:
$E^L \to I$: $\alpha = 4.6$; $\beta = 0$; $\Delta = -1$; $\sigma = 4$; $\mu = 0$; $\delta = -1$;
$E^R \to I$: $\alpha = 4.6$; $\beta = 0$; $\Delta = 1$; $\sigma = 4$; $\mu = 0$; $\delta = -1$;
$I \to E^L$: $\alpha = -7.5$; $\beta = 0$; $\Delta = 4$; $\sigma = 5$; $\mu = 1$; $\delta = 3$;
$I \to E^R$: $\alpha = -7.5$; $\beta = 0$; $\Delta = -4$; $\sigma = 5$; $\mu = -1$; $\delta = 3$;
$I \to I$: $\alpha = -2.2$; $\beta = 1$; $\Delta = 2$; $\sigma = 3$; $\mu = 0$; $\delta = 3$;
$E^L \to E^L$: $\alpha = 13$; $\beta = 0$; $\Delta = -2$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;
$E^R \to E^R$: $\alpha = 13$; $\beta = 0$; $\Delta = 2$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;
$E^L \to E^R$: $\alpha = 13$; $\beta = 0$; $\Delta = 0$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;
$E^R \to E^L$: $\alpha = 13$; $\beta = 0$; $\Delta = 0$; $\sigma = 2$; $\mu = 0$; $\delta = -1$;

Rescaling the learning rate of inhibition or the inhibitory strength.

The scale of learning rate of inhibition, γ, (Figure 7) modifies the learning rates of the inhibitory learning rates: $\gamma_{II} \to \gamma\gamma_{II}$ and $\gamma_{EI} \to \gamma\gamma_{EI}$. Similarly, the strength of inhibition (Figure S7) modifies the inhibitory synaptic weights: $W^{II} \to \gamma W^{II}$ and $W_{EI} \to \gamma W^{EI}$.

Development of periodic GCN.

If plasticity is enabled during the activation phase post pattern-formation, then neurons with similar phases should become coupled. If neurons across the GCN with a similar spatial tuning phase are wired up, and if this happens for neurons of all spatial phases, then the GCN topology is equivalent to a single-bump network with periodic boundary conditions (a ring network in 1D or a single-bump twisted-torus network in 2D).

However, as discussed in the text, enabling plasticity during the activation phase is problematic because of the positive feedback loop between recurrent-driven activity and potentiation of the recurrent connections themselves. The tendency is that at whichever specific population pattern phase the recurrent activity-based plasticity is turned on, weights become strengthened and that specific

population pattern becomes a discrete fixed-point of the network dynamics; the tendency is toward a degradation of the translation-invariance across phases of the population pattern.

To minimize and overcome this effect, we assume a sharply reduced learning rate for the recurrent synapses during this phase and incorporate neural spike frequency adaptation dynamics, which force the network pattern to flow rather than remain pinned at specific phases. The total synaptic current $I_i^P(t)$ is modified as follows:

$$I_i^P(t) = A_i^P \alpha^{P,vel}(v,t)(g_i^{P,rec}(t) + g_i^{P,loc}(t) + g_i^{P,adapt}(t) + g^0) + A_i^P g^{0'}, \tag{12}$$

where $g_i^{P,adapt}(t) = \alpha_{adapt} a_i^P(t)$ and

$$\frac{da_i^P(t)}{dt} + \frac{a_i^P(t)}{\tau_{adapt}} = \sum_b \delta(t - t_{i,b}^P). \tag{13}$$

For the results shown in Figure S4, $\alpha_{adapt} = 4$ and $\tau_{adapt} = 120$ ms, and $\eta = 7.5 \times 10^{-5}$ s^{-1}. All other parameters and equations remain unchanged from those specified in the Experimental Procedures.

Revealing low-dimensional local connectivity in the weight matrix: Sorting by simulated annealing.

We define a cost function on neural indices and weights, that imposes a penalty when cells that share are strong connection are assigned distant indices. Minimizing the cost function with respect to neural indices provides a sorting order for cells. We use this ordering to visualize structure in the randomized connectivity matrix of Figure S2. The cost function is given by:

$$E = \sum_{i,j} W_{ij}^2 (i-j)^2. \tag{14}$$

To minimize the cost function, we employ simulated annealing (Kirkpatrick et al., 1983). First, index the N cells and construct a synaptic matrix from this ordering. In each step of the algorithm, the indices of two cells chosen at random are swapped, and the difference in the cost before and after the swap is evaluated. If $\Delta E \leq 0$, i.e., the cost has decreased, the swap is accepted; if $\Delta E > 0$, the swap is accepted with probability $P(\Delta E)$ given by the Boltzmann distribution:

$$P(\Delta E) = e^{-\Delta E/T}, \tag{15}$$

where T is a control parameter. T is initially large ($T_0 = 1 \times 10^5$) and is lowered in stages such that $T_n = 0.9^n T_0$ in the nth stage. Transition to the next stage is made after 1×10^4 flips have been accepted in the present stage. The process is terminated when the number of steps within any given stage reaches a maximum value, 1×10^5, before 1×10^4 flips have been accepted.

Estimating the learning time T_{2D} for a 2D GCN.

Let us assume that the main scale-up in time cost for learning a 2D GCN with noisy cells rather than a 1D GCN with noisy cells comes from the time taken to cover the 2D training enclosure with

a random trajectory, compared to the time taken to cover the 1D training enclosure with a random trajectory.

In the theory of random walks on lattices, the quantity *cover time* is the maximum over all x, of the expected time taken for the walk to visit all locations on the lattice starting from location x. If the lattice is the 1D finite chain of n nodes, $\mathbb{Z}_n$, then cover time scales as $\sim n^2$ (Guruswami and Kannan, 2012). For a random walk on the 2D regular lattice $\mathbb{Z}_n^2$ of $n \times n$ nodes, the cover time scales as $n^2 \log^2(n)$ (Dembo et al., 2004).

Here, n represents the number of discrete spatial sites. Space is not discrete in our 1D and 2D simulations; instead we obtain an effective n by dividing T_{1D} by τ_{auto}, which is the autocorrelation time over which an animal trajectory continues in a given direction before significantly changing direction. From Yoon et al., 2013, the trajectory autocorrelation time from recorded animal trajectories is computed to be between 1-10 seconds. Equating n in 1D and 2D, we obtain:

$$T_{2D} \sim T_{1D} \frac{1}{4} \left(\log \left(\frac{T_{1D}}{\tau_{auto}} \right) \right)^2 \tag{16}$$

The factor $\frac{1}{4} \left(\log \left(\frac{T_{1D}}{\tau_{auto}} \right) \right)^2$ determines how much longer a random walk would take to cover a 2D space compared to a 1D space, of equal length per linear dimension. It is clear that the factor is relatively insensitive to detailed values of T_{1D} and τ_{auto} because they appear in the argument of the logarithm. Using $T_{1D} = 7200$ seconds and $\tau_{auto} = 1$ second gives a factor of 4. Using $\tau_{auto} =$ 10 seconds gives a factor of 2. This analysis results in our estimate, stated in the main manuscript, that 2D GCN development with random trajectories in 2D space should take only 2-4 times as long as development of a 1D GCN (with noisy cells in both cases). Possibly an additional factor of 2 may be necessary to include approaches from different directions to each point in the training enclosure, which would bring 2D learning times to 4-8 times longer than the 1D learning times, but it is not clear the extra factor is necessary.

References

Buetfering, C., Allen, K. and Monyer, H. (2014). Parvalbumin interneurons provide grid cell-driven recurrent inhibition in the medial entorhinal cortex. Nat Neurosci *17*, 710–8.

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

Dembo, A., Peres, Y., Rosen, J. and Zeitouni, O. (2004). Cover times for brownian motion and random walks in two dimensions. Annals of Mathematics *160*, 433–464.

Guruswami, V. and Kannan, R. (2012). Markov chains. Lecture notes from Computer Science Theory for the Information Age.

Kirkpatrick, S., C. D. Gelatt, J. and Vecchi, M. P. (1983). Optimization by Simulated Annealing. Science *220*, 671–680.

Kruge, I. U., Wernle, T., Moser, E. I. and Moser, M. B. (2013). Grid cells of animals raised in spherical environments. In SFN Poster.

Krupic, J., Burgess, N. and O'Keefe, J. (2012). Neural representations of location composed of spatially periodic bands. Science *337*, 853–7.

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*, 758–62.

Stensola, T., Stensola, H., Moser, M. B. and Moser, E. I. (2013). The grip map anchors systematically to environment geometry. In SFN Poster.

Stewart, S., Jeewajee, A., Wills, T. J., Burgess, N. and Lever, C. (2014). Boundary coding in the rat subiculum. Philos Trans R Soc Lond B Biol Sci *369*.

Wills, T. J., Barry, C. and Cacucci, F. (2012). The abrupt development of adult-like grid cell firing in the medial entorhinal cortex. Front Neural Circuits *6*, 1–13.

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*, 1077–84.