Supplementary Figure 1

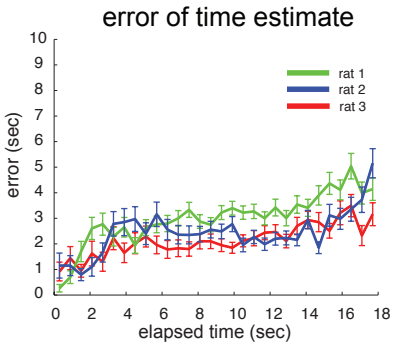

Supplementary Figure 1. Error of time estimate using phase prediction. Same as in Figure 1D, using the time prediction model via phase rather than rate.

Supplementary Figure 2

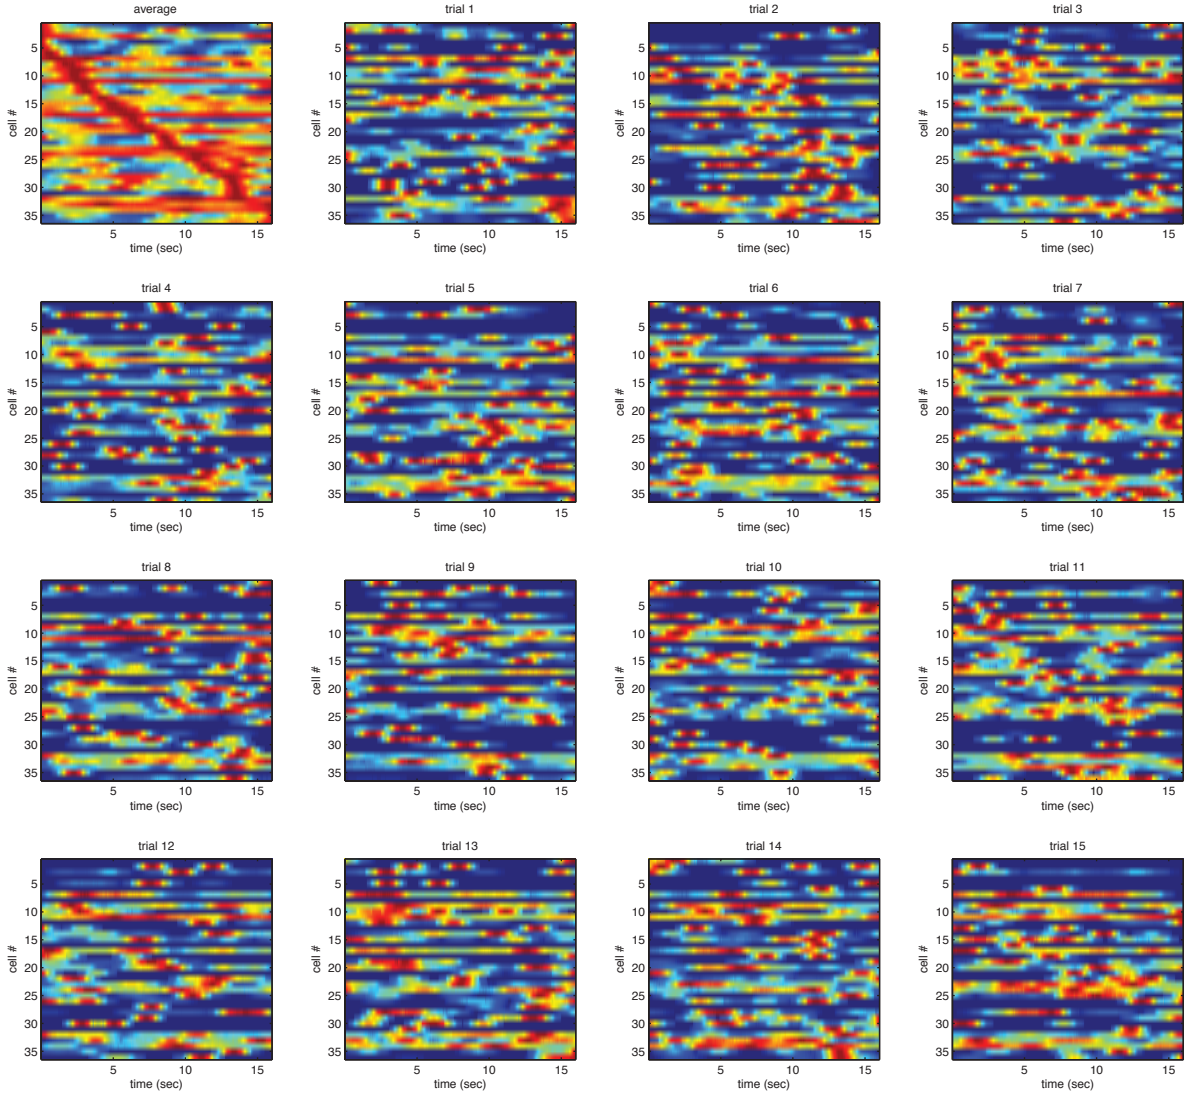

Supplementary Figure 2. Single trials from the same session of the homecage recordings, all sorted according to the best ordering for the average across trials (top left). Sequential activity is not evident on individual trials.

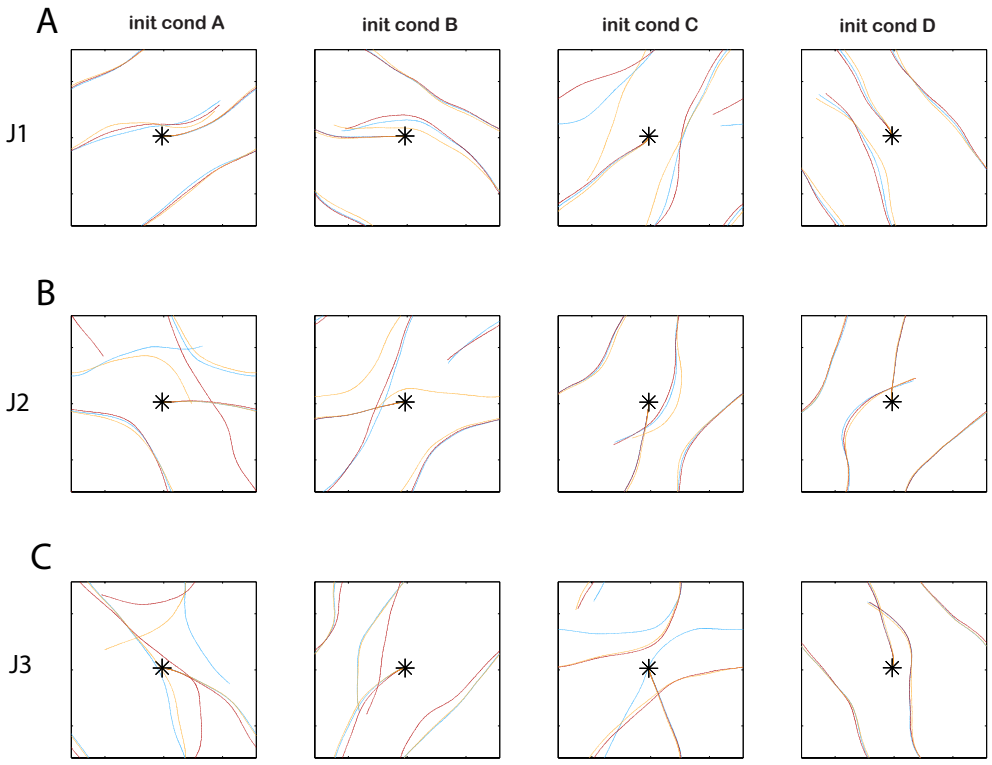

Supplementary Figure 3. Trajectory dependence on initial conditions. (A) For each of the four initial conditions, A, B, C and D, three trials were chosen at random. The trajectories of the center of the bump for each trial (red, blue and yellow curves), resulting from 10 seconds of simulated data, are reliable for each initial condition but distinct across initial conditions. **(B, C)** Same as in (A), but using matrices J^2 and J^3 produced by using different instances of the random heterogeneity matrix J^{het}, where $J^i = J^0 + J^{het}$ for $i = 1, 2, 3$. Note that for identical initial conditions, the particular trajectory differs for different synaptic connectivity matrices J^i.

Supplementary Figure 4

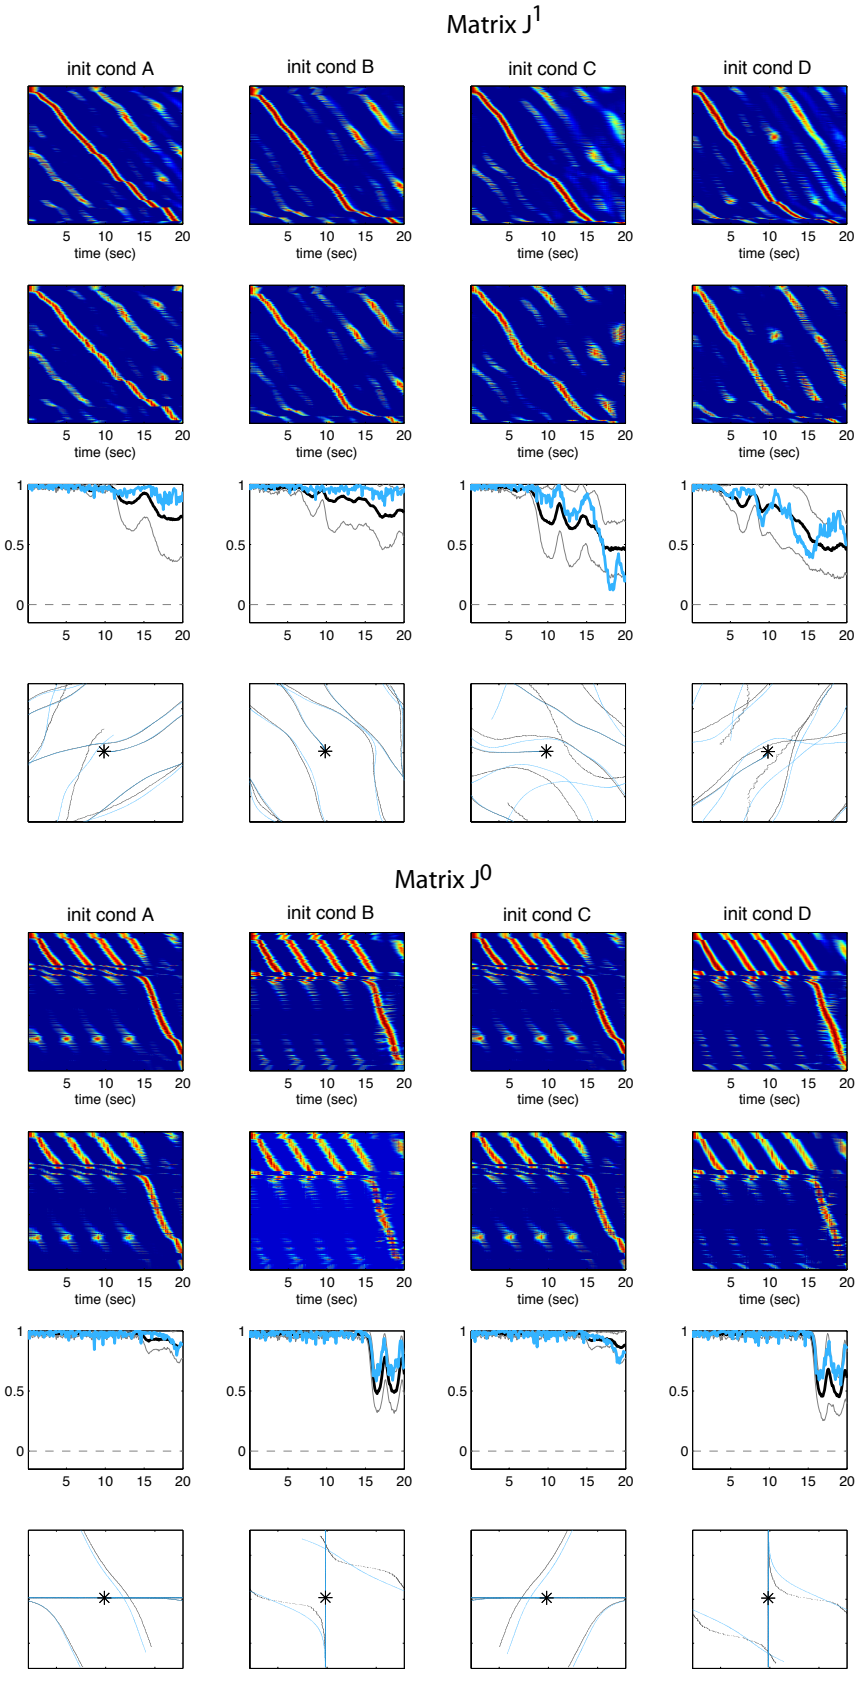

Supplementary Figure 4. Simulation results for the synaptic connectivity matrices J^1 and J^0. **(J1)** Same as Figure 3, but with all 4 initial conditions shown (A,B,C and D) and no homecage condition. The first two columns are identical to those in Figure 3. The heterogeneous synaptic weight matrix used in all cases is $J^1 = J^0 + J^{het}$. Note that the initial conditions C and D produced sequences that are less reliable at later times than those for initial conditions A and B. **(J0)** Same as (J1), but with *homogeneous* synaptic connectivity matrix J^0. While reliability is very good using the perfectly symmetric torus architecture J^0, the sequences are all short-lived, repeating after 5 seconds. This makes them unsuitable for timekeeping as multiple time points during a 20 second trial yield similar activity patterns, introducing ambiguity in the representation of time.

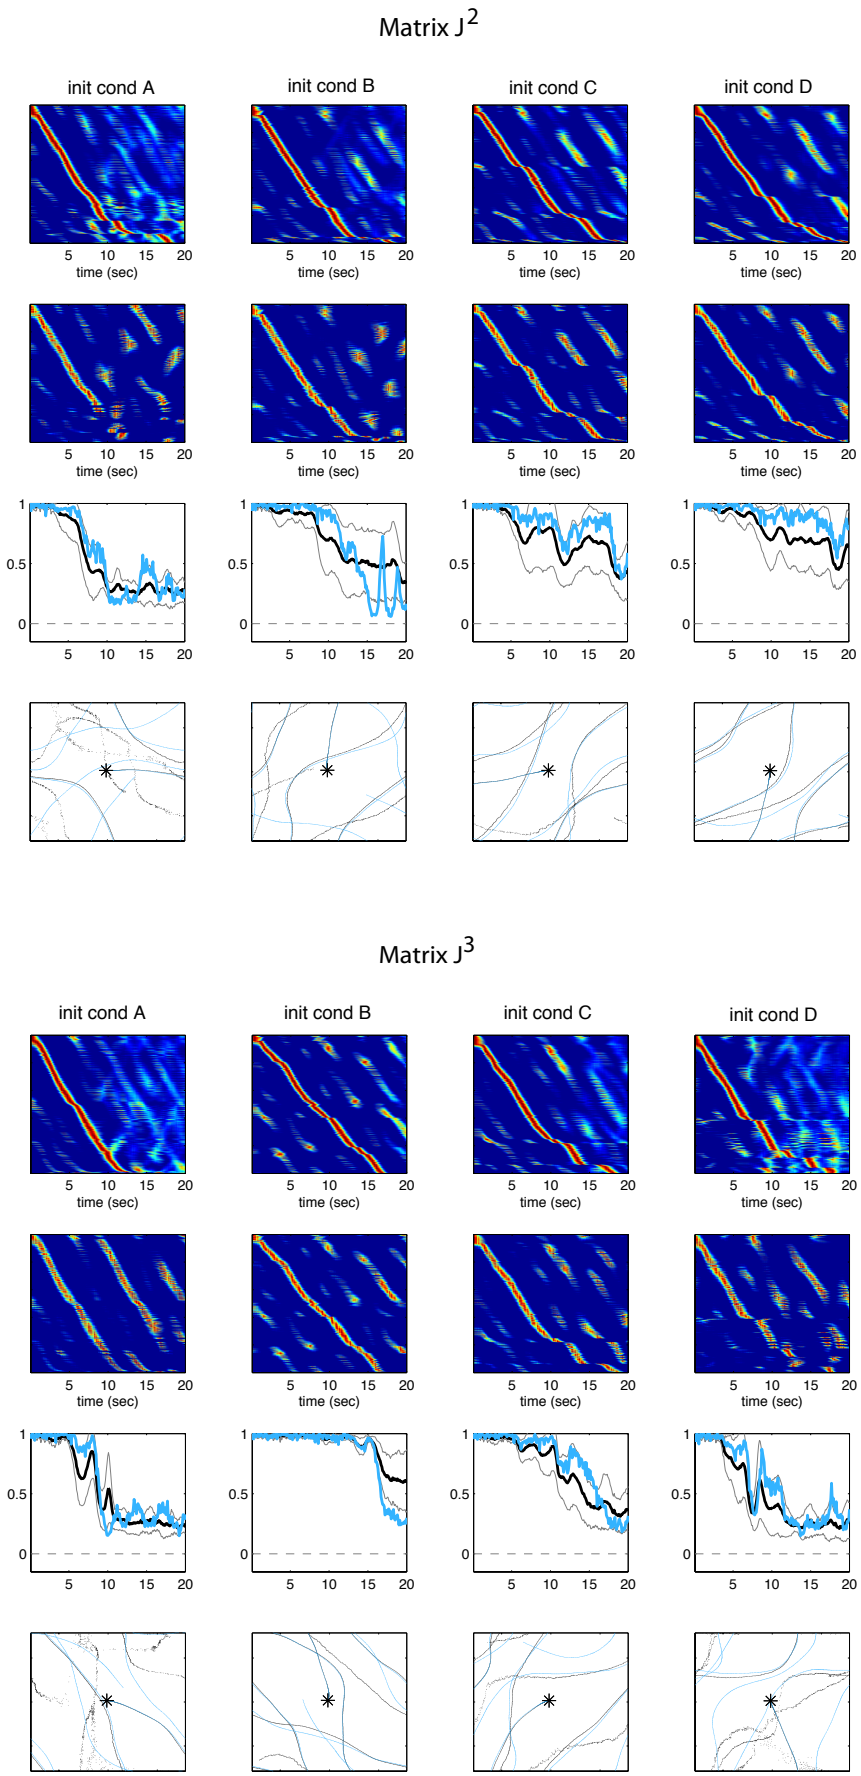

Supplementary Figure 5. Simulation results using matrices J^2 and J^3. Same as Supp Fig 4, this time with heterogeneous synaptic connectivity matrices J^2 (top) and J^3 (bottom). A variety of sequence lengths are produced, depending on the combination of matrix and initial condition.

Supplementary Figure 6

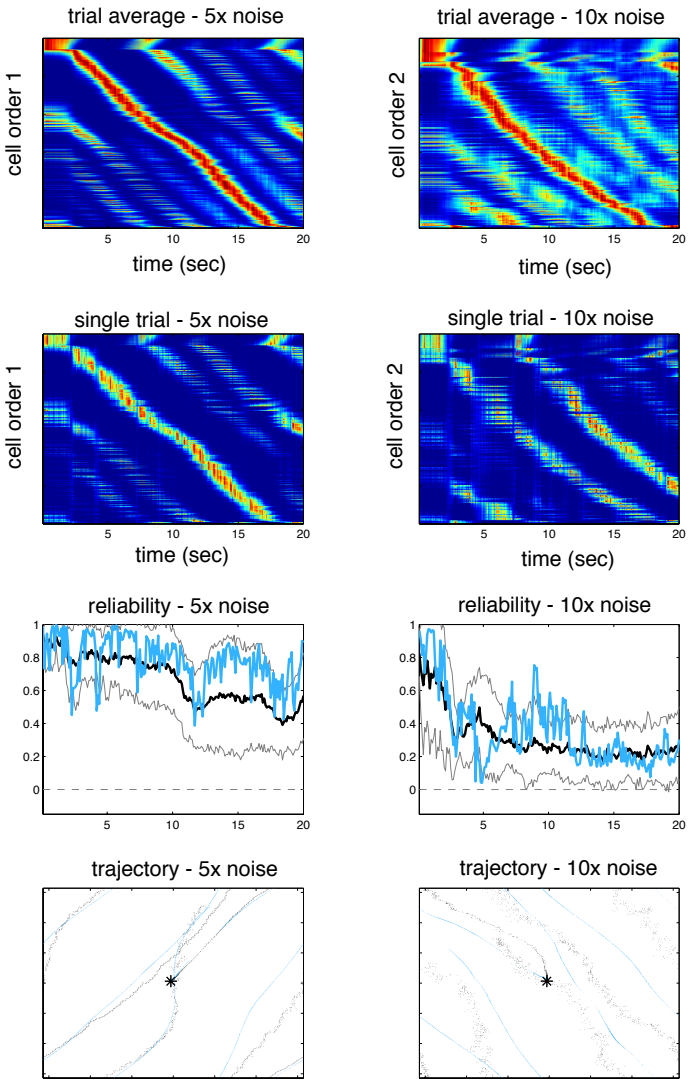

Supp Fig 6. Simulation results using greater input noise strength. Same as the left column of Figure 3 (panels A,D,G, and J, with connectivity matrix J^1 and initial condition 'A'), repeated using 5-fold (left column) and 10-fold (right column) increases in the input noise to each neuron. Sequence is still present, but reliability is lower and drops off more quickly as a function of time.

Supplementary Figure 7

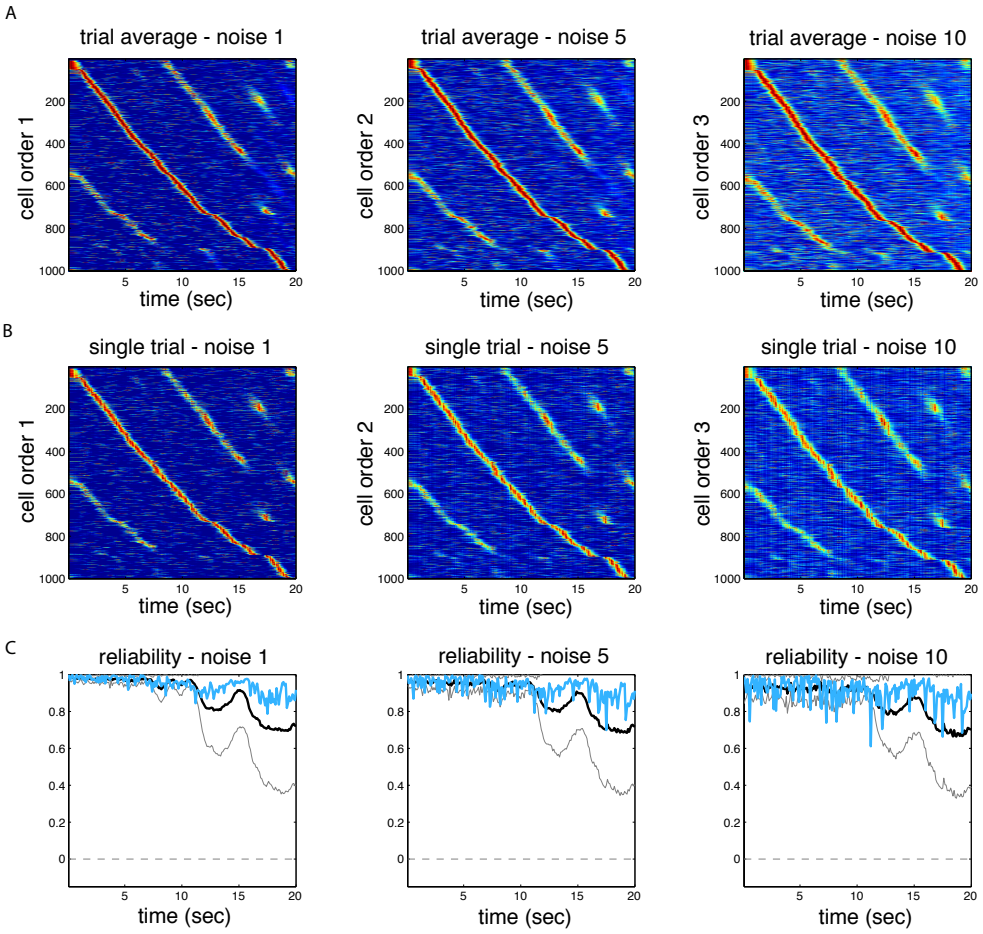

Supp Fig 7. Sequence reliability is preserved for different levels of input noise to layer 2. Same as Figure 4 (left column), but including simulations with 5-fold and 10-fold increases in input noise to layer 2 (middle and right columns, respectively). Although increasing levels of input noise to the first layer, where the sequence is generated, decreases sequence reliability, this is not the case for increased noise levels to layer 2, where the sequence is inherited.

Supplementary Movie.

A movie of the neural activity as a function of time. (**Top**) The activity pattern on the torus resembles a ``bump'' at all times. Neurons are organized in a 50x50 grid, similar to Figure 2A. White pixels correspond to neurons with firing rate exactly 0, while dark blue corresponds to very small, but nonzero, firing rates. (**Bottom**) Tracing the center of the bump (red dot) one can see the trajectory of activity (black curve). The trial shown here is from simulations with the matrix J1, using initial condition A.

Supplementary text for "Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus"

V.Itskov, C. Curto, E. Pastalkova, G. Buzsáki

Here we describe the models used in the main text in more detail. We have used two different types of models. The first are *time prediction models*, which are *phenomenological* models fit to data and used to infer time from experimentally recorded neural activity. The purpose of these models is to show that the sequences observed in the hippocampus have a possible functional role - i.e., they could be used for the animal to keep track of elapsed time. These models do not reflect a cellular or network mechanism for either generating the sequences or for extracting temporal information from them in the brain, rather they show that reliable temporal information is present in the sequential cell assembly activity on a trial-to-trial basis.

The second type of model is a *mechanistic* model, showing a simple yet biologically plausible network mechanism for generating long-lasting, context-dependent, temporally reliable sequences in a recurrent network such as hippocampus in the absence of structured inputs. We call it the *threshold adaptation model* (TAM).

1 The time prediction models

For time prediction only one kind of correct trials was used for each session: 'right' or 'left' trials were chosen depending on which kind of trial had the animal making the fewest number of mistakes in the alternation task. All putative pyramidal cells with average firing rates in the range $0.1Hz \leq r \leq 15Hz$, were used in the time prediction models. The models are distinguished by whether they use firing rate (*rate-only* model) or theta phase (*phase-only* model) as the feature of neural activity used to estimate elapsed time.

1. Each individual spike train was smoothed with a Gaussian of width $\sigma = 0.25s$ and then binned into time bins of size $0.5s$. This resulted in a time series of rates $r_k(t)$ for each k-th cell.

2. For each cell, each time bin was assigned a discrete activity *state* using one of the following two rules:

   a) *Rate-only.* For each k-th cell an average maximal firing rate (across trials), $r_k^{\max}$, was computed. Then, for each time-bin t, the cell was assigned one of 3 states $u_k(t)$ based on the relative firing rate $\tilde{r}_k(t) = r_k(t)/r_k^{\max}$ as follows: $u = 1$ if $\tilde{r}_k(t) \leq 0.05$, $u = 2$ if $0.05 < \tilde{r}_k(t) \leq 0.25$, and $u = 3$ if $\tilde{r}_k(t) > 0.25$. The values $u_k(t) = 1, 2$ or 3 represent 'low', 'medium', of 'high' firing rates, respectively.

   b) *Phase-only.* For each k-th cell and each time bin t the average phase $\theta_k(t)$ of spike times with respect to the theta oscillation in that time bin was computed. Note that each time bin contains several periods of theta. For each time bin and each cell one of 4 states $u_k(t)$ was assigned. One state was reserved for having no spikes in that time-bin ($u = 0$), and the other

three states ($u = 1, 2$, or 3) were computed based on the phase $\theta_k(t)$ by dividing the unit circle into three equal parts.

3. For each individual neuron, k, the following state-based probabilistic model was used for both the *rate-only* and *phase-only* versions of time prediction. The population model assumed that the states of different neurons are independent from each other, and thus the probability of observing a state vector $\mathbf{u}(t) = (u_1, .., u_N)$ at a time-bin t is

$$P_t(\mathbf{u}) = \prod_{k=1}^{N} P_t^k(u_k), \quad \text{where}$$
$$P_t^k(u) = \text{Prob}(u_k(t) = u), \text{ with } u = 0, 1, 2, \text{or } 3.$$

   The independence is of course an incorrect assumption. However, the prediction based on this assumption can be no better than that of a model taking correlations into account, and is less susceptible to overfitting.

4. For each trial, we obtained a time prediction for each time bin as follows. The trial was selected as a "test set," with all the other trials serving the role of "training set." The probabilistic model $P_t(\mathbf{u})$ was fit on the training set by computing the frequency of each state for each neuron in each time bin, and using the above product formula to generate from this the probability of observing a given state vector $\mathbf{u}$. The model was then used on the selected "test set" trial with population vectors $\mathbf{u}(t)$ to obtain a time prediction for each time bin via the maximum-likelihood estimate

$$\hat{t}(t) = \arg\max_{t'} P_{t'}(\mathbf{u}(t)),$$

   which selects the time $\hat{t}(t)$ at which the population pattern $\mathbf{u}(t)$ was most likely to occur. Notice that the estimates at different times are computed independently of each other since they used only instantaneous neuronal activity; therefore the inference method we used does not require any kind of working memory or integration. The time-estimate errors were computed for each time bin as the average error across trials.

We have chosen to use a model with a few discrete states, rather than analog firing rate values, in order to reduce overfitting of the probabilistic model. Overfitting could also be avoided using other techniques such as smoothing in continuous models. We have found that this simple model is sufficient to show that elapsed time can in principle be inferred by brain structures downstream from the CA1 area of hippocampus (on behaviorally relevant time scales), but a different model may have served this purpose equally well.

2 The threshold-adaptation model (TAM)

Model equations

We model network dynamics using a standard firing rate model, with threshold nonlinearity. At any point in time, the vector $x(t) = (x_1(t), ..., x_N(t))$ represents a population vector of firing rates for each of N neurons. Another ingredient is the activity-dependent adaptation of the thresholds for each cell, represented by the dynamic variables $h(t) = (h_1(t), ..., h_N(t))$. The model equations are thus:

$$\tau_m \frac{dx_i}{dt} = -x_i + \left[\sum_{j=1}^{N} J_{ij} x_j + I_i - h_i \right]_+, \qquad i = 1, \ldots, N \tag{1}$$

$$\tau_a \frac{dh_i}{dt} = -h_i + c x_i, \qquad i = 1, \ldots, N \tag{2}$$

where τ_m and τ_a are membrane and threshold-adaptation time constants, respectively, J is the matrix of synaptic weights for the recurrent network, and $I = (I_1, .., I_N)$ is a (time-dependent) vector of external inputs to the considered recurrent network. Here the brackets denote the threshold-linear function

$$[y]_+ = \begin{cases} y & \text{if} \quad y > 0, \\ 0 & \text{if} \quad y \leq 0. \end{cases}$$

The constant c controls the strength of the activity-dependent adaptation, whereas τ_a determines the timescale with which the thresholds recover in the absence of the cell's firing.

The synaptic matrix J: torus architecture with heterogeneity

The neurons in the recurrent network were organized on a two-dimensional sheet, with periodic boundary conditions yielding a torus-like grid of neurons ($N = 50^2 = 2500$ in simulations). Each neuron was indexed by a location (x_i, y_i) on this sheet, with the $\{x_i\}$ and $\{y_i\}$ values evenly partitioning the interval $[0, 2\pi]$ with periodic boundary condition ($0 \equiv 2\pi$.)

The matrix $J = J^0 + J^{\text{het}}$ is the sum of a matrix J^0, with perfect torus-like topography, and a matrix J^{het}, which represents heterogeneity of the synaptic weights. By the "perfect torus-like topography" we mean that the strength of the (symmetric) connection between neurons i and j is purely a function of their distance on the torus:

$$J^0_{ij} = \frac{4\pi^2}{N} \varphi \left(\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \right),$$

where $\varphi(d) = j_0 + j_1 \cos(d/\sqrt{2})$ was chosen similar[1] to the "ring model" [1], however a wide class of Mexican-hat-type connectivity functions $\varphi(d)$ may be chosen to yield bump attractor solutions (c.f. [2]). The parameters j_0 and j_1 ($j_0 = -.6, j_1 = 1$ in simulations) were chosen to to yield recurrent excitation for nearby neurons, and inhibition for far away neurons on the torus, so that the resulting dynamics of equation (1) with spatially and temporally constant inputs and thresholds ($I_i - h_i = \text{const}$) results in "bump attractor" dynamics.

The matrix of heterogeneities is constant in time and is drawn from the normal distribution:

$$J^{\text{het}}_{ij} = \frac{\varepsilon}{\sqrt{N}} n_{ij},$$

[1] One difference is the $\sqrt{2}$, which is added here to compensate for the fact that the longest distance on the torus is $\sqrt{2}\pi$, rather than π as it is on the circle.

where $n_{ij} \sim N(0,1)$ are independent gaussian-distributed variables with mean 0 and variance 1. The scaling of $N^{-\frac{1}{2}}$ is chosen so that the spectral radius ($= \varepsilon$) of this matrix does not change with finer and finer discretizations of the torus sheet. Note that due to the scaling by N, the individual matrix elements J^{het}_{ij} are much larger than J^0_{ij} in the limit of large N; nevertheless, the correlated structure of the matrix J^0 determines the "bump attractor" dynamics of equation (1). The heterogeneity matrix J^{het} is important not just to avoid fine-tuning of the model, but also for the reliability. By breaking the symmetry in the two-dimensional sheet of bump attractors, J^{het} has the effect of "carving out" out an irregular trajectory for the movement of the bump; this makes the sequences last longer without repeating the same neurons.

Simulations with TAM

Here we describe in detail all of the particular choices made in simulations (parameters, noise, synaptic weight matrix, etc.). In all, we generated 500 simulated data sets, each consisting of a single 20-second 'trial' with 2500 neurons. We used 4 different synaptic-weight matrices J^0, J^1, J^2, and J^3, where J^0 had perfect torus-like symmetry in the connections, and the others had heterogeneities introduced by the addition of different 'noise' matrices J^{het}_{ij}. For each matrix $J^l, l = 0, ..., 4$, we ran 25 trials for each of 5 different trial-types: 4 trial types were 'task' trials with different adaptation initial conditions (tasks A, B, C, and D), and the fifth trial-type represented 'homecage' trials (hc).

Parameters used in simulations

$$
\begin{aligned}
N &= 50^2 = 2500 \text{ (number of neurons, in a 50 by 50 grid)} \\
T &= 20\text{s (time length for simulations)} \\
\tau_m &= 30\text{ms (membrane time constant)} \\
\tau_a &= 2\text{s (timescale of the adaptation)} \\
c &= 0.5 \text{ (adaptation constant, controls 'ergodicity' of the trajectory)} \\
J^0_{ij} &= \frac{4\pi^2}{N}\left(j_0 + j_1 \cos\sqrt{\frac{1}{2}(x_i - x_j)^2 + \frac{1}{2}(y_i - y_j)^2}\right), \text{ and} \\
J^l_{ij} &= J^0_{ij} + \frac{\varepsilon}{\sqrt{N}} n_{ij}, \text{ where} \\
\varepsilon &= 0.5,\ j_0 = -0.6,\ j_1 = 1, \text{ and } n_{ij} \sim N(0,1). \\
I_i(t) &= I_0 + I^{\text{noise}}_i(t) \text{ (input noise varies over neurons, time bins, and trials)} \\
I_0 &= 1 \text{ (constant across neurons, time bins, and trials)} \\
I^{\text{noise}}_i(t) &\sim N(0,1)
\end{aligned}
$$

What changed between trials and trial-types in the 'task' conditions

For each choice of matrix J^l in the 'task' conditions, we simulated trials with initial conditions 'A', 'B', 'C', and 'D'. The differences in initial conditions are for the *adaptation* variables only. The initial conditions 'A', 'B', 'C', and 'D' have neurons with adapted thresholds to the left, bottom, right and top of the initial activity bump, respectively.

Within a fixed matrix choice J^l and fixed initial condition, the only aspect of the simulation that changed from trial to trial was the white noise $I^{\text{noise}}_i(t)$. Each neuron was driven by a random input,

normally distributed with mean 1 and variance 1. The drive was independent across neurons and from one 1ms-time bin to the next. Each noise instance was re-sampled from the distribution for each trial.

Simulations for the second layer.

In order to investigate if the reliable sequences can be "inherited" by a second layer via random projections, we have considered a second layer of $N_2 = 1000$ cells that were receiving downstream projections from the layer with torus-like mexican-hat connectivity having $N = 2500$ cells (as described above). Each cell in the second layer was receiving an input from the first layer as well as a noisy unstructured input. The second layer also had non-specific recurrent inhibition. The dynamics of the second layer are thus described by the following equations:

$$\tau_m \frac{dy_k}{dt} = -y_k + \left[\sum_{i=1}^{N} M_{ki} x_i - \sum_{l=1}^{N_2} y_l + \tilde{I}_k - g_k \right]_+, \qquad k = 1, \ldots, N_2 \tag{3}$$

$$\tau_a \frac{dg_k}{dt} = -g_k + c y_k, \qquad k = 1, \ldots, N_2. \tag{4}$$

Here $y_k(t)$ are the firing rates of the cells in the second layer, $x_i(t)$ are the firing rates in the first layer, τ_m and τ_a are the membrane time-constant and the timescale of the synaptic adaptation respectively (the same as in the first layer). The synaptic weights M_{kl} of the projection from the first to the second layer were chosen to be nonzero with probability $p = 0.1$, and the nonzero elements were drawn from a uniform distribution on the unit interval (i.e. $0 \leq M_{kl} \leq 1$). The input noise $\tilde{I}_k = \tilde{I}_k(t)$ in the feedforward synapses was drawn from a Gaussian distribution with zero mean and the same variance as the input noise $I_i^{\text{noise}}(t)$ in the first layer. The second layer was assumed to have the same dynamics of threshold adaptation (the variables $g_k(t)$) as in the first layer.

The 'homecage' condition

There are two differences between our 'homecage' condition simulations and the above simulations for the 'task' conditions: (1) On each trial, we randomly selected one of the initial conditions 'A', 'B', 'C', or 'D', to reflect the fact that – unlike in the task conditions – the specific initial condition on adaptation variables in the homecage will vary unpredictably from trial to trial; (2) Instead of having noise inputs $I(t)$ that are independently sampled on each 1ms time step of the simulation, the noise input to each neuron is constant over a duration of 125ms (approx. one theta-cycle). This reflects the fact that in inattentive states, cortical activity often shows higher temporal correlations than in attentive states [3]. These two changes destroyed the sequences produced by the model.

References

[1] R Ben-Yishai, R L Bar-Or, and H Sompolinsky. Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci U S A*, 92(9):3844–3848, Apr 1995.

[2] K. Kishimoto and S. Amari. Existence and stability of local excitations in homogeneous neural fields. *J. Math. Biol.*, 7(4):303–318, 1979.

[3] J Poulet and C Petersen. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*, Jan 2008.