

Optimizing Working Memory with Heterogeneity of Recurrent Cortical Excitation

Zachary P. Kilpatrick,¹ Bard Ermentrout,^{2,3} and Brent Doiron^{2,3}

¹Department of Mathematics, University of Houston, Houston, Texas 77204, ²Department of Mathematics, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, and ³Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania 15213

A neural correlate of parametric working memory is a stimulus-specific rise in neuron firing rate that persists long after the stimulus is removed. Network models with local excitation and broad inhibition support persistent neural activity, linking network architecture and parametric working memory. Cortical neurons receive noisy input fluctuations that cause persistent activity to diffusively wander about the network, degrading memory over time. We explore how cortical architecture that supports parametric working memory affects the diffusion of persistent neural activity. Studying both a spiking network and a simplified potential well model, we show that spatially heterogeneous excitatory coupling stabilizes a discrete number of persistent states, reducing the diffusion of persistent activity over the network. However, heterogeneous coupling also coarse-grains the stimulus representation space, limiting the storage capacity of parametric working memory. The storage errors due to coarse-graining and diffusion trade off so that information transfer between the initial and recalled stimulus is optimized at a fixed network heterogeneity. For sufficiently long delay times, the optimal number of attractors is less than the number of possible stimuli, suggesting that memory networks can under-represent stimulus space to optimize performance. Our results clearly demonstrate the combined effects of network architecture and stochastic fluctuations on parametric memory storage.

Introduction

Persistent neural activity occurs in prefrontal (Fuster, 1973; Funahashi et al., 1989; Romo et al., 1999) and parietal (Pesaran et al., 2002) cortex during the retention interval of parametric working memory tasks. Model networks of stimulus-tuned neurons that are connected with local slow excitation (Wang, 1999) and broadly tuned inhibitory feedback (Compte et al., 2000; Goldman-Rakic, 1995) exhibit localized and persistent high-rate spike train patterns called “bump” states (Compte et al., 2000; Renart et al., 2003). Bumps have initial locations that are stimulus-dependent, so population activity provides a code for the remembered stimulus (Durstewitz et al., 2000). These models relate cortical architecture to persistent neural activity, and are a popular framework for studying working memory (Wang, 2001; Brody et al., 2003).

Neural variability is present in all brain regions and limits neural coding in many sensory, motor, and cognitive tasks (Stein et al., 2005; Faisal et al., 2008; Laing and Lord, 2009). In parametric working memory networks, dynamic input fluctuations cause bump states to wander diffusively (Compte et al., 2000; Laing and Chow, 2001; Wu et al., 2008; Polk et al., 2012; Burak and Fiete,

2012; Kilpatrick and Ermentrout, 2013), degrading stimulus storage over time. Psychophysical data show that the spread of the recalled position increases with delay time (White et al., 1994; Ploner et al., 1998), consistent with diffusive wandering of a bump state. While several results examine how bump formation depends upon neural architecture, little is known about how cortical wiring affects the diffusion of persistent neural activity.

The response properties of cells are often heterogeneous (Ringach et al., 2002), a feature that can improve population-based codes (Chelaru and Dragoi, 2008; Shamir and Sompolinsky, 2006; Marsat and Maler, 2010; Osborne et al., 2008; Padmanabhan and Urban, 2010). In particular, there is a large degree of variation in synaptic plasticity and cortical wiring in prefrontal cortical networks involved in persistent activity during working memory tasks (Rao et al., 1999; Wang et al., 2006). Heterogeneity in excitatory coupling quantizes the neural space used to store inputs, reducing the network’s overall storage capacity (Renart et al., 2003; Itskov et al., 2011). On the other hand, stabilizing a discrete number of network states improves the robustness of working memory dynamics to parameter perturbation (Rosen, 1972; Koulakov et al., 2002; Brody et al., 2003; Goldman et al., 2003; Miller, 2006). In this study, we investigate how stabilization introduced by synaptic heterogeneity affects the temporal diffusion of persistent neural activity.

We show that spatial heterogeneities in the excitatory architecture of a spiking network model of working memory reduce the rate with which bumps diffuse away from their initial position. However, the same heterogeneities limit the number of stable network states used to store memories. A tradeoff between these consequences maximizes the transfer of stimulus informa-

Received April 17, 2013; revised Oct. 14, 2013; accepted Oct. 19, 2013.

Author contributions: Z.P.K., B.E., and B.D. designed research; Z.P.K. and B.D. performed research; Z.P.K. and B.D. analyzed data; Z.P.K., B.E., and B.D. wrote the paper.

This work was supported by National Science Foundation Grants NSF-DMS-1121784 (B.D.), NSF-DMS-1311755 (Z.P.K.), and NSF-DMS-1219753 (B.E.). We thank Robert Rosenbaum, Ashok Litwin-Kumar, and Krešimir Josić for useful discussions.

The authors declare no competing financial interests.

Correspondence should be addressed to Zachary P. Kilpatrick, Department of Mathematics, University of Houston, Phillip G. Hoffman 651, Houston, TX 77204-3008. E-mail: zpkipat@math.uh.edu.

DOI:10.1523/JNEUROSCI.1641-13.2013

Copyright © 2013 the authors 0270-6474/13/3318999-13\$15.00/0

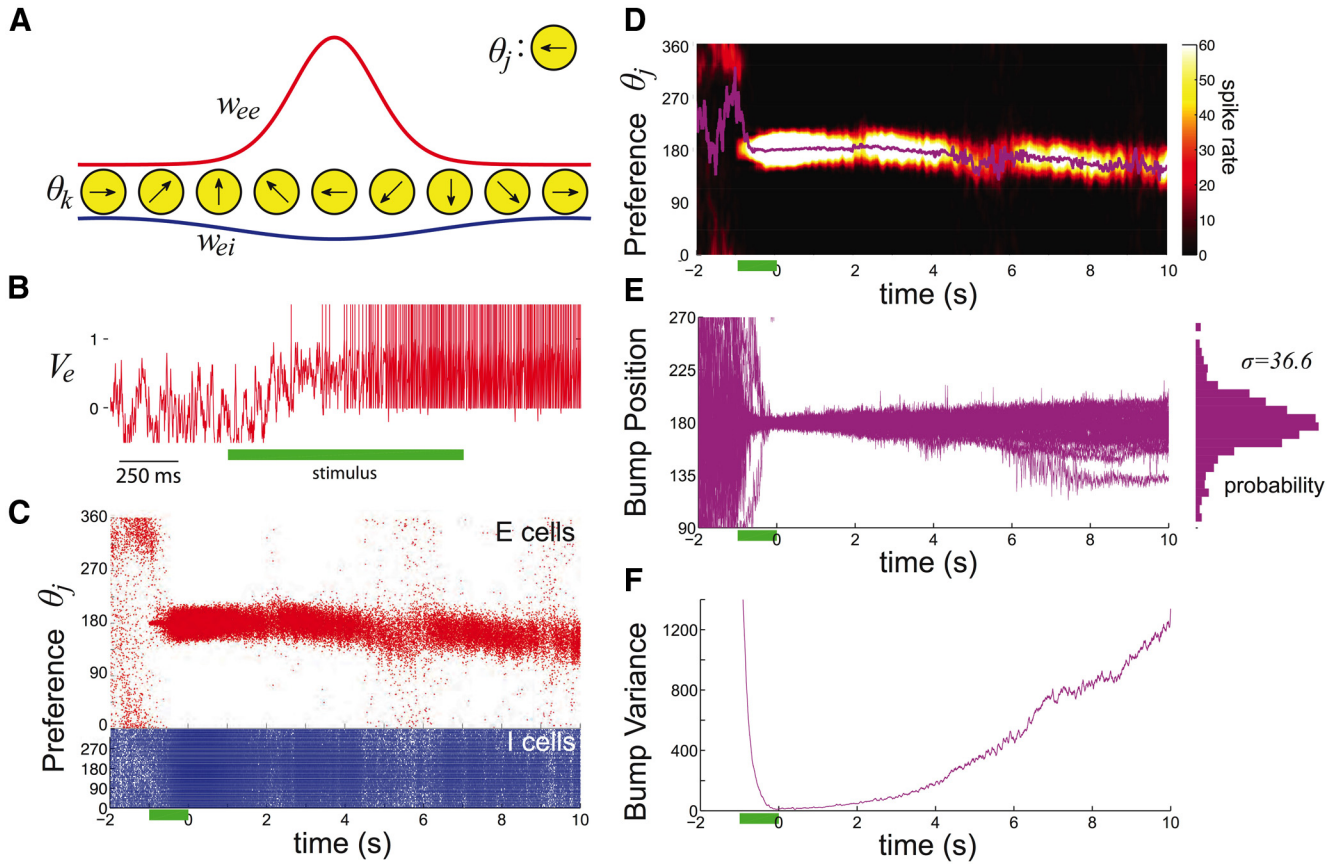


Figure 1. Spiking network model with spatially homogeneous synaptic connectivity. **A**, Strength of connections from pyramidal neurons to pyramidal neurons (red) and synapses from interneurons to pyramidal neurons (blue). A neuron of preferred stimulus angle θ_j receives synaptic inputs from all neurons spanning preferred stimulus angles indexed by θ_k (see Materials and Methods). **B**, Voltage of the pyramidal cell with stimulus preference 170° before, during (green bar), and after cue presentation. **C**, Formation of a bump of spiking activity in the pyramidal neurons (red) following cue presentation (green bar). **D**, Spike rate, locally averaged across space and time (see Materials and Methods). Plot shows that the position of the bump's peak ($\Delta(t)$, magenta) diffuses in space due to voltage and synapse noise. **E**, Bump position ($\Delta(t)$) plotted for 32 realizations. The resulting probability density of bump positions from 1000 realizations after 10 s is approximately Gaussian. **F**, Variance of the bump's position ($\langle \Delta(t)^2 \rangle$), across 1000 realizations, scales approximately linearly as a function of time.

tion at a specific degree of network heterogeneity. For a large number of stimulus locations and long retention times, we show that network architectures that under-represent stimulus space can optimize performance in working memory tasks.

Materials and Methods

Recurrent network architecture. We used for our network a ring architecture commonly used for generating persistent activity to represent direction between 0 and 360° (Ben-Yishai et al., 1995; Compte et al., 2000) with $N_E = 256$ pyramidal cells (E) and $N_I = 64$ interneurons (I). Each leaky integrate-and-fire neuron (Laing and Chow, 2001) was distinguished by its cue orientation preference θ_j , where $\theta_j(E) = \Delta_E \cdot j$ ($j = 1, \dots, N_E$) and $\theta_j(I) = \Delta_I \cdot j$ ($j = 1, \dots, N_I$) for $\Delta_E = 360/256$ and $\Delta_I = 360/64$, respectively. The subthreshold membrane potential of each neuron, $V_\alpha(\theta_j, t)$ ($\alpha = E, I$), obeyed the following equation:

$$\frac{dV_\alpha(\theta_j, t)}{dt} = -V_\alpha(\theta_j, t) + I_\alpha + I_{ext, \alpha}(\theta_j, t) + I_{syn, \alpha}(j, t) + I_{n, \alpha}(t),$$

where $I_E = 0.6$ and $I_I = 0.6$ are bias currents that determine the resting potential of E and I neurons. The external current, expressed in the following equation:

$$I_{ext, E}(\theta_j, t) = I_0 \exp \left[- \left(\frac{\theta_j - \theta_{ext}}{I_d} \right)^2 \right], t \in [T_{ON}, T_{OFF}];$$

$$I_{ext, E}(\theta_j, t) = 0, \quad \text{otherwise,}$$

represents sensory input received only by pyramidal neurons, where $I_0 = 2$ is the input amplitude, $I_d = 3$ determines input width, and θ_{ext} is the cue position. The stimulus was turned on at $T_{ON} = -1$ s and off at $T_{OFF} = 0$ s. Interneurons received no external input, so $I_{ext, I} = 0$. Voltage fluctuations were represented by the white noise process $I_{n, \alpha}(t)$ with variance $\sigma_{V, \alpha}^2$ ($\sigma_{V, E} = 0.5$ and $\sigma_{V, I} = 0.3$). We scaled and nondimensionalized voltage so the threshold potential $V_{th} = 1$ and the reset potential $V_{res} = 0$ for all neurons.

Synaptic currents were mediated by a sum of AMPA, NMDA, and GABA currents:

$$I_{syn, \alpha}(j, t) = I_{AMPA, \alpha}(j, t) + I_{NMDA, \alpha}(j, t) - I_{GABA, \alpha}(j, t)$$

each modeled as

$$I_{AMPA, \alpha}(j, t) = A_{AMPA, \alpha} \sum_{K=1}^{N_E} W_{AMPA, \alpha}(\theta_j, \theta_K) s_{AMPA, E}(\theta_K, t),$$

$$I_{NMDA, \alpha}(j, t) = A_{NMDA, \alpha} \sum_{K=1}^{N_E} W_{NMDA, \alpha}(\theta_j, \theta_K) s_{NMDA, E}(\theta_K, t),$$

$$I_{GABA, \alpha}(j, t) = A_{GABA, \alpha} \sum_{K=1}^{N_E} W_{GABA, \alpha}(\theta_j, \theta_K) s_{GABA, I}(\theta_K, t),$$

where $A_{AMPA, E} = 1$, $A_{NMDA, E} = 2$, $A_{GABA, E} = 0.81$, $A_{AMPA, I} = 1$, $A_{NMDA, I} = 1$, and $A_{GABA, I} = 0$. Orientation preference was introduced

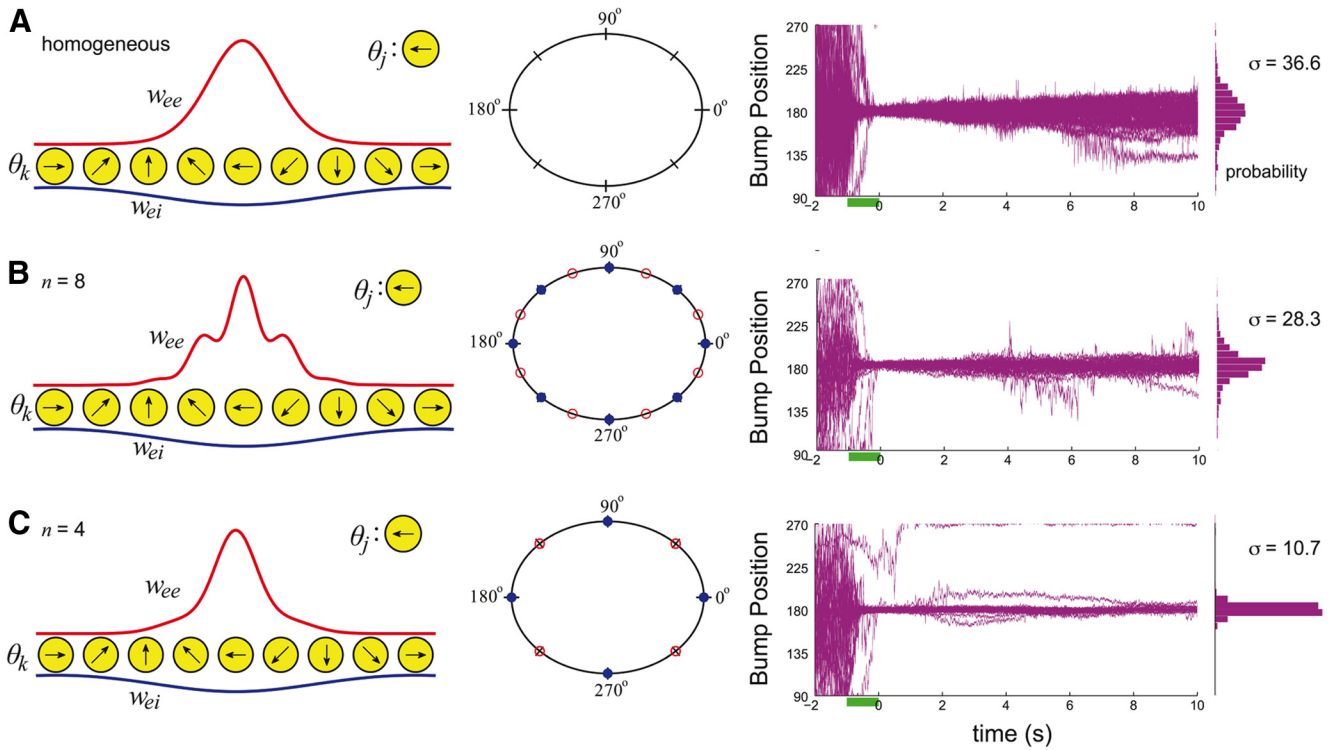


Figure 2. *A*, Distance-dependent synaptic connections lead to a spatially homogeneous system. Bump dynamics lie on a line attractor, so bumps diffuse with ease. *B*, Periodically breaking the spatial homogeneity of synaptic connections with $n = 8$ -fold heterogeneity leads to bump dynamics evolving on a chain of $n = 8$ attractors (red) each separated by repelling states (blue). Bumps do not wander away from their initial position as easily (position plot), which tightens the resulting probability density after 10 s. *C*, Effect of synaptic heterogeneity is more noticeable for a $n = 4$ -fold break in homogeneity. Bump position rarely strays from 180° , as shown by the very tight probability density.

into synaptic conductance by the spatially decaying functions (Fig. 1A) expressed:

$$W_{\beta,\alpha}(\theta_j, \theta_k) = \exp \left[\frac{\cos(\pi[\theta_j - \theta_k]/180) - 1}{d_{\beta,\alpha}} \right], \quad (1)$$

where $d_{AMPA,E} = 0.32$, $d_{NMDA,E} = 0.32$, $d_{GABA,E} = 5$, $d_{AMPA,I} = 5$, and $d_{NMDA,I} = 5$. Equation 1 was used for excitatory (AMPA and NMDA) synapses between pyramidal (E) cells in the case of spatially homogeneous connectivity. In the case of spatially heterogeneous synaptic strength (Fig. 2), the strength of AMPA and NMDA connections between pyramidal cells (E) was given by

$$W_{\beta,E}(\theta_j, \theta_k) = [1 + h \cos(\pi n \theta_k / 180)] \times \exp \left[\frac{\cos(\pi[\theta_j - \theta_k]/180) - 1}{d_{\beta,E}} \right],$$

where $h = 0.025$ represents the strength of the heterogeneity and n is the frequency of the heterogeneity, which must be integer valued.

Synaptic gating variables of type $\beta = AMPA, NMDA$, or $GABA$ associated with a neuron at location θ_j were instantaneously activating and exponentially decaying as described by

$$\frac{ds_\beta(\theta_j, t)}{dt} = -\frac{s_\beta(\theta_j, t)}{\tau_\beta} + \delta(V_\alpha(\theta_j, t) - 1) + I_{s,\beta}(t),$$

where $\alpha = E$ for $\beta = AMPA$ and $NMDA$ while $\alpha = I$ for $\beta = GABA$. Instantaneous activation is represented here using the delta function δ , so $s_\beta(\theta_j, t)$ increments by 1 when $V_\alpha(\theta_j, t)$ attains the spike threshold $V_{th} = 1$. Decay time constants for each synapse type are $\tau_{AMPA} = 5$ ms, $\tau_{NMDA} = 100$ ms, and $\tau_{GABA} = 20$ ms. Fluctuations in conductance were introduced into each synapse with the term $I_{s,\beta}(t)$, which is white noise with variance $\sigma_{s,\beta}^2$ ($\sigma_{s,AMPA} = 0.1$, $\sigma_{s,NMDA} = 0.45$, and $\sigma_{s,GABA} = 0.05$). We take the variance of noise to NMDA synapses to be high, $\sigma_{s,NMDA} = 0.45$, because it leads to high variances in the spike times, as commonly ob-

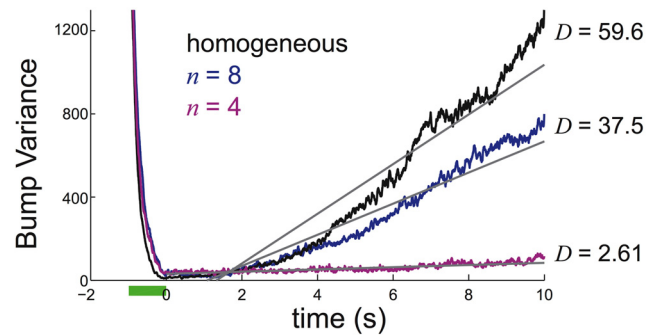


Figure 3. Variance of bump position as a function of time, averaged across 1000 realizations for spatially heterogeneous structure of pyramidal-to-pyramidal synapses with frequency $n = 4$ and $n = 8$, as well as spatially homogeneous structure. We fit each curve to straight lines to generate an approximation of the effective diffusion coefficient D (see Materials and Methods).

served in prefrontal cortical neurons during the delay period of working memory tasks (Compte et al., 2003). In addition, this generates an error of a few degrees in the recall of cue position for delay periods of 2–10 s, as observed in psychophysical experiments (White et al., 1994).

Numerical simulations were done using an Euler–Maruyama method with timestep $dt = 0.1$ ms and normally distributed random initial conditions. Spike time rastergrams were smoothed to generate population firing rates as a function of degree and time, whose maximum at each time were used to calculate the centroid of the bump (Figs. 1D, E, 2). Variances (Figs. 1F, 3) and probability densities (Figs. 1E, 2) were computed using 1000 values for the bump centroid across 10 s. Linear fits of variance in the case of spatially homogeneous synapses and spatially heterogeneous synapses with $n = 8$ and $n = 4$ (Fig. 3) were performed using linear regression.

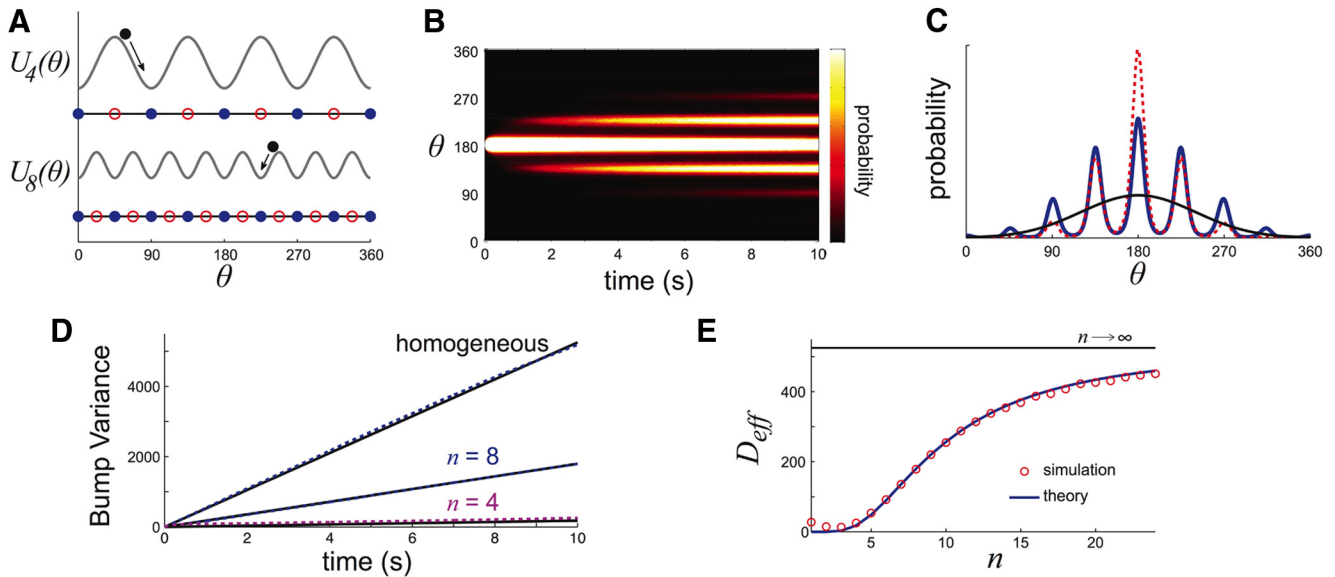


Figure 4. *A*, Particle diffusing in periodic potential well Equation 3 is an idealized model of the bump diffusing over the network with spatially heterogeneous synapses, given by the stochastic process Equation 2. *B*, Probability density $p(\theta, t)$ of particle position θ spreads diffusively in time ($n = 8; h = 1$; noise variance $\sigma^2 = 0.16$). *C*, Profile of $p(\theta, t)$ computed from 10,000 realizations (red); effective diffusion theory Equation 6 (black); and effective diffusion theory with periodic correction Equation 7 (blue). *D*, Effective diffusion theory using D_{eff} in Equation 8 matches variance scaling from simulations of stochastic equation. Variance increases monotonically with the well frequency n . *E*, An effective diffusion coefficient D_{eff} can be computed by treating well hopping as a jump Markov process (Lifson and Jackson, 1962; Lindner et al., 2001), yielding formula Equation 8.

Diffusion in the potential well model. To analyze the diffusive dynamics of the bump, we studied an idealized model of bump motion. In this model, the bump position $\phi(t)$ obeys the stochastic differential equation, (Lifson and Jackson, 1962; Lindner et al., 2001):

$$d\phi(t) = -h \sin [n\phi(t)]dt + \sigma dW(t). \quad (2)$$

Here $\phi(t)$ was restricted to the periodic domain $\phi \in (-\pi, \pi]$ and dW was a standard white noise process. The first term in Equation 2 models the periodic spatial heterogeneity that is responsible for attractor dynamics. Heterogeneity is parametrized by its strength h and spatial frequency n . In this framework, the dynamics of $\phi(t)$ is a diffusive process occurring on an energy landscape defined by the periodic potential

$$U_n(\phi) = \frac{h}{n} \cos(n\phi), \quad (3)$$

producing n attractors (Fig. 4*A*). These attractors occur at the minima of Equation 3, given by $\phi = 2j\pi/n$ where $j = 1, \dots, n$. They are separated from one another by repellers or saddles at the maxima of Equation 3.

To analyze the model, we reformulated Equation 2 as an equivalent Fokker–Planck equation (Risken, 1996),

$$\frac{\partial p(\phi, t)}{\partial t} = \frac{\partial}{\partial \phi} [h \sin(n\phi)p(\phi, t)] + \frac{\sigma^2 \partial^2 p(\phi, t)}{\partial \phi^2}, \quad (4)$$

where $p(\phi, t)$ is the probability density of finding the bump at a given value ϕ at time t . For ease of analytic calculations, we let ϕ evolve on an infinite domain. Since we worked in parameter regimes where the resulting spread of probability densities was relatively narrow, this did not considerably alter the results. Also, experimentally measured errors in cue recall are typically not large enough to span more than a quarter of the possible stimulus space (Ploner et al., 1998).

For large times and sufficiently high-frequency n , the variance of the stochastic process Equation 2 can be quantified using an effective diffusion coefficient (Lindner et al., 2001)

$$D_{\text{eff}} = \frac{1}{2} \lim_{t \rightarrow \infty} \frac{\langle \Delta \phi(t)^2 \rangle}{t} = \frac{1}{2} \lim_{t \rightarrow \infty} \frac{\langle [\phi(t) - \langle \phi(t) \rangle]^2 \rangle}{t}. \quad (5)$$

The density $\rho(\phi, t)$ tends asymptotically:

$$p_{\text{asy}}(\phi, t) = \frac{p_0(\phi)}{\sqrt{4\pi D_{\text{eff}} t}} \exp \left[-\frac{\phi^2}{4D_{\text{eff}} t} \right], \quad (6)$$

where $p_0(\phi)$ is the stationary, periodic solution, Equation 4, given by

$$p_0(\phi) = \chi \exp \left[\frac{2h \cos(n\phi)}{\sigma^2 n} \right],$$

where χ is a normalization factor. The approximation, Equation 6, matches realizations of the full stochastic process Equation 2 very well (Fig. 4*C*). Clearly, the frequency of the probability distribution’s microscale is commensurate with that of the periodic potential Equation 3. We were mainly concerned with computing the effective diffusion of the stochastic process defined by Equation 2. Remarkably, second-order statistics are still well approximated by ignoring the microperiodicity of the density in Equation 6, just using:

$$p_{\text{gauss}}(\phi, t) = \frac{1}{\sqrt{4\pi D_{\text{eff}} t}} \exp \left[-\frac{\phi^2}{4D_{\text{eff}} t} \right], \quad (7)$$

(Fig. 4*C*). Previous authors have used asymptotic methods for computing the associated effective diffusion coefficient D_{eff} inherent in the formula Equation 6 (Lifson and Jackson, 1962; Lindner et al., 2001). The long-standing result is (Lifson and Jackson, 1962)

$$D_{\text{eff}} = \frac{\sigma^2/2}{\int_0^{2\pi/n} \int_0^{2\pi/n} e^{2h[\cos(n\phi) - \cos(n\psi)]/(n\sigma^2)} d\psi d\phi}$$

and we can compute the integrals in the denominator to find:

$$D_{\text{eff}} = \frac{\sigma^2}{2I_0\left(\frac{2h}{n\sigma^2}\right)}, \quad (8)$$

where $I_0(x)$ is the modified Bessel function of the zeroth kind. Equation 8 demonstrates the monotone increasing dependence of the effective diffusion upon the number of potential wells n .

To calculate the probability density $p(\phi, t)$, we simulated 10,000 realizations of Equation 2 using an Euler–Maruyama integration scheme with a timestep $dt = 0.001$ from $t = 0$ to $t = 10$ s (Fig. 4B,C). The effective diffusion coefficient was calculated as the gradient of the variance across the time window and converted to degrees with the change of variables $\theta = 180(\phi + \pi)/\pi$ yielding:

$$D_{\text{eff}}(\theta) = \frac{180^2 D_{\text{eff}}(\phi)}{\pi^2} = \left[\frac{180^2}{\pi^2} \right] \left[\frac{\sigma^2}{2I_0 \left(\frac{2h}{n\sigma^2} \right)} \right]. \quad (9)$$

Adding unstructured heterogeneity. Effects of unstructured heterogeneity were studied by perturbing the potential in Equation 2 with a combination of random periodic functions,

$$d\phi(t) = \left(-h \sin[n\phi(t)] - \eta \frac{dU_p(\phi)}{d\phi} \right) dt + \sigma dW(t). \quad (10)$$

The unstructured heterogeneity was given by the random potential

$$U_p(\phi) = \sum_{j=1}^{N_h} [a_j \cos(j\phi) + b_j \sin(j\phi)], \quad (11)$$

where a_j, b_j ($j = 1, \dots, N_h$) are randomly drawn from normal distributions. Here η scales the amplitude of the random potential and the maximal frequency of the unstructured components is given by N_h . We take the rounded integer $N_h = [0.05(1 + 0.1\xi)/(\eta + 0.0005) + 1]$, ξ is a normally distributed random variable, so larger η values reduce the number of modes added to the potential, decreasing the maximal number of attractors in the system, as in other studies of unstructured heterogeneity in bump attractor networks (Zhang, 1996; Renart et al., 2003; Itskov et al., 2011; Hansel and Mato, 2013). To calculate an effective diffusion coefficient D_h , we initialized 10,000 simulations of Equation 10 at $\phi(0) = 0$ and computed $D_h = \langle \phi(T)^2 \rangle / T$ for $T = 10$ s.

Information measures. To measure the performance of the network on a working memory task, we used a Shannon measure of mutual information for a noisy channel (Cover and Thomas, 2006). We considered a channel receiving one of m possible stimuli (X), storing an input as one of $n \leq m$ possible states ($Y(t)$), and reading out the remembered stimulus as one of the original m possible values (Z). The stored variable $Y(t)$ evolves during storage time $t \in [0, T]$ due to degradation of the initially loaded signal $Y(0)$ by dynamic noise. The stimuli were presented with equal probability $p_j = 1/m$ ($j = 1, \dots, m$), so that the stimulus entropy was as follows:

$$H(X) = -\sum_{j=1}^m p_j \log_2 p_j = \log_2 m.$$

The network represented a stimulus as the bump position at one of the system's n attractors. If m was a multiple of n ($m = qn$ with q an integer) then the mapping from stimulus to loaded representation was straightforward with $Y(0) = \text{ceil}(X/q)$. When m was not a multiple of n , we allowed the potential well structure of the system to guide the loaded state to the nearest attractor. This led to slightly nonuniform distributions of loaded stimuli. However, the effects of diffusion made this slight nonuniformity insignificant, especially as the length of storage time T was increased. In our theoretical calculations, we assumed that the loading algorithm maximized the entropy of the neural representation $Y(0)$; this sometimes involved random assignments from X to $Y(0)$. In fact, we found our numerical results did not stray too far from this approximation (see Figs. 7, 8).

If $n = m$, then $Y(0) = X$ and the loaded representation had the same entropy as the stimulus $H(Y(0)) = H(X)$. If $n < m$, then the representation entropy was smaller than the stimulus entropy $H(Y(0)) < H(X)$, since the space into which the stimulus is represented is smaller than the original stimulus space. Since the stimulus X has a uniform distribution, then so does Y , with the probability that attractor j loaded as $p_j = 1/n$ ($j = 1, \dots, n$) and the entropy of the representation as $H(Y(t)) = \log_2 n$ (this

holds for all t). The readout of the neural representation required an expansion from $Y(T)$ to Z . If m was a multiple of n , then the expansion was straightforward with Z having probability $1/q$ for $Z = q(Y(T) - 1) + 1, \dots, qY(T)$ and zero otherwise. In this case $H(Z) = \log_2 m$. If m was not a multiple of q , we subdivided the domain into m evenly spaced subdomains and assigned Z accordingly, and for theoretical calculations we again assumed $H(Z) = \log_2 m$.

While $H(Y) \leq H(Z)$, $H(Y)$ nevertheless set an upper limit for the mutual information between the stimulus X and readout Z , expressed as:

$$I(X;Z) = I(Y;Z) = H(Y) - H(Y|Z), \quad (12)$$

To compute the conditional entropy $H(Y|Z)$ we first calculated the probability of transition between one state and another during the diffusion phase ($0 \leq t \leq T$). A direct estimate of the transition probabilities was obtained by numerically simulating many realizations of the model and estimating $p(\phi(0)|Z)$ where $\phi(0)$ is the center of mass of the bump at time $t = T$. We subdivided ϕ into n subdomains of equal width, and the area of each subdomain is $p_{j \rightarrow k}$, the transition probability from the loaded state $X = j$ to another state ($k \neq j$) or itself ($k = j$), where $k = 1, \dots, n$. Due to discrete translation symmetry in both systems, we expected $p_{j \rightarrow k} = p_{j+t \rightarrow k+t}$. The conditional entropy can then be computed as follows (Cover and Thomas, 2006):

$$H(Y|Z = j) = -\sum_{k=1}^n p_{j \rightarrow k} \log_2 p_{j \rightarrow k}. \quad (13)$$

Our second method of computing the conditional entropy employed the effective diffusion coefficient D_{eff} associated with the probability density for locations of the bump or particle. D_{eff} depends upon n , ultimately introducing this further n dependence into the mutual information (Eq. 12). Using the associated pure Gaussian probability (Eq. 7), we computed transition probabilities analytically for the case $j = 1$, so, as Equation 14 states:

$$\begin{aligned} p_{1 \rightarrow k} &= \frac{1}{\sqrt{4\pi D_{\text{eff}} T}} \int_{a(k)}^{a(k)+2\pi/n} \exp \left[-\frac{x^2}{4\pi D_{\text{eff}} T} \right] dx \\ &= \frac{1}{2} \left[\text{erf} \frac{a(k) + 2\pi/n}{\sqrt{4D_{\text{eff}} T}} - \text{erf} \frac{a(k)}{\sqrt{4D_{\text{eff}} T}} \right], \end{aligned} \quad (14)$$

for a set delay time T , where $a(k) = (2\pi k - 1)/n$ is the lower boundary of the k th subdomain and $k = 1, \dots, \text{ceiling}(n/2)$. Due to reflection symmetry of the Gaussian, we expected $p_{1 \rightarrow k} = p_{1 \rightarrow n-k}$. For even n , the $(n/2 + 1)$ th subdomain is $(-\pi, \pi/n - \pi) \cup (\pi - \pi/n, \pi)$, which would lead to two integrals in Equation 14. As mentioned, the transition probabilities $p_{j \rightarrow k}$ for $j = 2, \dots, n$ were easily computed as $p_{j \rightarrow k} = p_{1 \rightarrow k-j+1}$ (see Fig. 6B,C, insets). We then plugged each $p_{j \rightarrow k}$ value into our formula for conditional entropy, Equation 13. The analytic and numeric calculation of $p_{j \rightarrow k}$ led to similar results for $I(X;Z)$, the calculated value of mutual information (Eq. 12; see Fig. 7).

Results

Diffusion of bumps in a spatially homogeneous network

The neural mechanics of parametric working memory has a long history of theoretical investigation (Amari, 1977; Camperi and Wang, 1998; Compte et al., 2000; Laing and Chow, 2001; Wang, 2001; Brody et al., 2003; Renart et al., 2003). Motivated by the working memory of visual cue orientations, we consider a network of spiking model neurons where each neuron has a preferred orientation in its feedforward input. Persistent neural spiking within the network is due to a combination of assumptions about synaptic connectivity (Goldman-Rakic, 1995; Rao et al., 1999; Lewis and Gonzalez-Burgos, 2000). First, the strength of pyramidal-to-pyramidal connectivity decreases as the distance between the tuning peak of each neuron increases (Fig. 1A, red line). Second, excitatory synaptic currents involve

both fast-acting AMPA and slow-acting NMDA components (see Materials and Methods). Third, feedback connections from interneurons are broadly tuned (Fig. 1A, blue line). With these architectural features, neurons in the network respond to a transient stimulus (Fig. 1, green bar) with an elevated rate of spiking that persists long after the stimulation ceases (Fig. 1B). Short-range excitation leads to high-rate pyramidal spiking across a short range of orientations, while wide-range inhibition localizes this spiking (Fig. 1C); we refer to this pattern of activity as a “bump.” The position of the bump encodes the initial stimulus position in working memory (Compte et al., 2000; Wang, 2001; Brody et al., 2003).

We model the inherent trial-to-trial variability of neural response with an orientation-independent fluctuating input to each neuron, as well as a stochastic component of the recurrent synaptic feedback (see Materials and Methods). These fluctuations degrade the storage of the orientation cue by causing the bump to stochastically wander away from its initial position (Fig. 1C,D). Spatiotemporal averaging of the spike time raster plots identifies the maximal firing rate at each time point, and visualizes the bump wandering across the network (Fig. 1D, magenta line). We fix the stimulus orientation and perform many trials of the network simulation, with the only difference between trials being the realization of the stochastic forces in the network. The bump’s position after a delay period of 10 s can be described by a probability density having an overall Gaussian profile (Fig. 1E), and the variance in bump position increases linearly as a function of time (Fig. 1F). These last two properties suggest the bump position behaves as a diffusion process (Risken, 1996).

Diffusive dynamics in working memory networks have been studied in several different frameworks (Compte et al., 2000; Miller, 2006; Wu et al., 2008; Burak and Fiete, 2012; Polk et al., 2012; Kilpatrick and Ermentrout, 2013). The intuition for the diffusive character of these networks is best gained from an analysis of the deterministic network. A bump can be formed with its center of mass located at any orientation, allowing for the storage of a continuum of stimuli (Amari, 1977; Camperi and Wang, 1998). However, perturbations that change the bump’s position will be integrated and stored as if they were another input. Stochastic inputs lead to a continuous and random displacement of the bump, without the bump relaxing back to its original location. Over time, the position of the bump effectively obeys Brownian motion and recall error increases with the delay period. This diffusion-based error is consistent with psychophysical studies that show the spread of recalled continuous variables scales sublinearly with time (White et al., 1994; Ploner et al., 1998).

Reduced diffusion in a spatially heterogeneous network

Previous models of working memory have considered networks that use neuronal units with bistable properties (Rosen, 1972; Koulakov et al., 2002; Brody et al., 2003; Goldman et al., 2003; Miller, 2006). These networks lack the homogeneity required for a continuum of neutrally stable stimulus representations, and rather have a discrete number of stable states. One advantage of this network heterogeneity is a “robustness” of representation with respect to parameter perturbation, a feature that is absent in homogeneous networks (Brody et al., 2003; Goldman et al., 2003). We consider spatially periodic modulation of excitatory coupling (Fig. 2, left), where the period of the modulation is $360/n$ degrees, so that n cycles cover orientation space. We assume such an architecture would not be biased to favor one particular cue location because errors reported in recalling cues are

approximately the same for each cue location (White et al., 1994). Such an architecture may develop from Hebbian plasticity rules during training in working memory tasks, since orientation cues are typically chosen at fixed and evenly spaced locations around the circle (Funahashi et al., 1989; White et al., 1994; Goldman-Rakic, 1995; Meyer et al., 2011). In this situation, some neuron pairs are activated more than other pairs, leading to relative strengthening of their recurrent connections (Clopath et al., 2010; Ko et al., 2011). Alternatively, reward-based plasticity mechanisms could also set up spatial heterogeneity in synapses if it improved a subject’s performance during a task (Schultz, 1998; Wang, 2008; Klingberg, 2010).

Spatial biases introduced to network architecture shift the smooth continuum of stable states to a chain of discrete attractors, each separated by a repeller (Fig. 2, compare A, B, C, middle column). This discrete attractor structure occurs because some pyramidal neurons receive stronger excitatory projections than others (Zhang, 1996; Itskov et al., 2011; Hansel and Mato, 2013). Spatial heterogeneity in the strength of excitatory connections (decreasing n) stabilizes bump positions to perturbations by noise (Fig. 2, compare A, B, C, right column). For all n tested, the probability density of bump positions retains an approximately Gaussian shape with periodic modulation, so the variance of bump positions still grows nearly linearly with time (Fig. 3), and it is well approximated by Dt where D is the diffusion coefficient (see Materials and Methods). The coefficient D drops considerably for a network with spatially heterogeneous synapses, compared with the bump diffusion measured with the homogenous network. In total, spatial heterogeneity of excitatory coupling helps stabilize bump position in models of working memory with fluctuating stochastic inputs.

Potential well model for bump diffusion

To analyze the relationship between network heterogeneity and bump diffusion more deeply, we now study an idealized model for parametric working memory. Briefly, a noise-driven particle on a periodic potential landscape retains the essential effects of noise and spatial heterogeneity in our spiking network model (see Materials and Methods). Our simplified model treats the bump position as a particle moving in a landscape of peaks, from which it is repelled (Fig. 4A, maxima), and wells, to which it is attracted (Fig. 4A, minima). In the potential well model, the memory of the stimulus location is tracked by the particle’s position $\theta(t)$ obeying the following stochastic differential equation,

$$d\theta(t) = \left(-h \sin \left[\frac{180}{\pi} n\theta(t) \right] \right) dt + dW(t). \quad (15)$$

Here h is the amplitude of the periodic potential and $W(t)$ is a Wiener process (Risken, 1996). The sine function determines the θ -dependent drift of the particle, which ultimately affects its diffusive motion. The positive integer n determines the number of stable attractors. Similar reduced neural models have been explored (Renart et al., 2003; Itskov et al., 2011), and the general problem of noise-induced behavior in periodic potentials has been well studied (Lifson and Jackson, 1962; Risken, 1996; Lindner et al., 2001).

Diffusive behavior occurs in periodic potentials, yet the mechanics are different than diffusion on a free potential landscape ($h = 0$). In periodic potentials, a particle typically undergoes small variance dynamics confined within a well. However, rare but large noise kicks eventually push the particle to a neighboring well. These noise-induced well transitions continue indefinitely,

and diffusion over the potential landscape occurs in a punctuated fashion. Across many trials, the probability $p(\theta, t)$ of finding the particle at position θ at time t evolves like a Gaussian kernel modulated by a periodic function (Fig. 4B; see Materials and Methods). The maxima of $p(\theta, t)$ are centered at the minima of the potential, indicating a higher likelihood of finding the particle at the bottom of a well than in transition between wells.

Treating the transitions between wells as a jump process, we can approximate the diffusion of the sine potential model, Equation 15, with the following Brownian walk:

$$d\theta(t) = \sqrt{2D_{\text{eff}}(n)}dW(t). \quad (16)$$

The effective diffusion coefficient $D_{\text{eff}}(n)$ is derived using standard approaches (see Materials and Methods, Eq. 9). Under this approximation, $\theta(t)$ obeys a Gaussian distribution with variance $D_{\text{eff}}(n)t$, and $p(\theta, t)$ does not possess the periodic microstructure of the actual probability density (Fig. 4C, compare blue and black curves). Despite these differences, the approximation agrees very accurately with the variance in the particle's position in the periodic potential (Fig. 4D). In this framework we can directly relate the frequency of heterogeneity n to the overall diffusivity of the particle through $D_{\text{eff}}(n)$. As with bumps in the spiking network, increasing the frequency n of the spatial heterogeneity increases the effective diffusion D_{eff} (Fig. 4D). This is true across the entire range of frequencies n , and as $n \rightarrow \infty$ the particle's variance saturates to that of a system with a flat potential (Fig. 4E). Thus, despite the simplicity of the potential well framework, it can qualitatively explain the diffusivity observed in the spiking network. In both Figure 3 and Figure 4D, the variance in bump position is fit well by a linear function whose slope decreases with the number of attractors n .

Impact of unstructured heterogeneity

As we have shown, a structured heterogeneity of cortical architecture that generates evenly spaced attractors (Figs. 2, 4A) curtails the rate of diffusion (Figs. 3, 4D). However, synaptic architecture may also have random components that are neither related to task specifics nor optimized to any specific computation (Wang et al., 2006). In principle, such perturbations in architecture could degrade the performance of working memory networks that require a fine-tuned architecture. Renart et al. (2003) demonstrated that the deleterious effects of such unstructured heterogeneity in bump attractor networks can be mitigated by homeostatic plasticity, which spatially homogenizes network excitability. In our model, considering such a process would bar the system from establishing spatially structured heterogeneity, which we have shown improves storage accuracy. Thus, we next study how a combination of structured and unstructured spatial heterogeneity affects the diffusive dynamics in working memory networks. We show that the system is still robust to noise, even when the potential is altered in this way.

Specifically, we modify the shape of the potential in Equation 15, so that,

$$d\theta(t) = \left(-h \sin \left[\frac{180}{\pi} n\theta(t) \right] - \eta \frac{dU_p(\theta)}{d\theta} \right) dt + dW(t), \quad (17)$$

where $U_p(\theta)$ is an unstructured perturbation of the underlying cosine potential function (see Materials and Methods). Briefly, $U_p(\theta)$ is a component of the potential that is randomized from trial to trial (two realizations of the full potential are shown in Fig.

5A, B). Adding a small component of unstructured heterogeneity ($\eta > 0$) to a network with an initially flat potential function ($h = 0$) substantially alters the attractor structure (Fig. 5A). The network shifts from having a continuum of preferred locations to having a small number of preferred locations at disordered positions. This is analogous to the drastic collapse in the number of possible stable bump locations observed in bump attractor models whose synaptic structure is randomly perturbed (Zhang, 1996; Renart et al., 2003; Itskov et al., 2011; Hansel and Mato, 2013). On the other hand, a network that possesses structured heterogeneity ($h > 0$) retains the original positions of its stable attractors after unstructured heterogeneity is added, even though the profile of the potential is distorted (Fig. 5B). When the severity of heterogeneity is increased (larger η), the number of attractors is considerably reduced in the network without structured heterogeneity, while remaining the same in the network with structured heterogeneity (Fig. 5C). Last, the effective diffusion D_h (see Materials and Methods) of the network containing structured heterogeneity increases only gradually as the degree of unstructured heterogeneity is increased (Fig. 5D), contrasting the distinct rise in the effective diffusion D_h in the network without structured heterogeneity.

Therefore, the spatial organization of attractors in the network with structured heterogeneity is robust to random perturbations of the underlying potential landscape. Recent studies have shown that parametrically perturbed spiking network models of bumps retain dynamics whose spatial profile is bump-shaped (Brody et al., 2003; Itskov et al., 2011; Hansel and Mato, 2013; Kilpatrick and Ermentrout, 2013). Brody et al. (2003) showed the effective dynamics of the resulting system can then be numerically approximated by a potential well model like Equation 17. Thus, the low-dimensional dynamics of the spiking network model can still be described by the potential well model, so we believe the spiking network will also be robust to unstructured perturbations in its spatial architecture. This robustness allows the reduction of diffusion due to structured heterogeneity to be relatively unaffected by sources of unstructured heterogeneity that undoubtedly exist in most cortical networks (Wang et al., 2006).

Memory storage as a noisy channel

Structured spatial heterogeneity in recurrent excitatory coupling has two distinct influences on response fidelity in working memory networks. First, it produces a finite set of attractors with which to store stimuli. Second, as we have shown, heterogeneity reduces the diffusion of persistent bump states across the network. These two influences have consequences for the overall storage performance by the network. We next characterize the working memory network as a noisy information channel (Cover and Thomas, 2006) and show how spatial heterogeneity of excitatory coupling mitigates a tradeoff between errors due to these two influences.

Consider a stimulus that chosen from m equally likely values and is to be stored by a working memory network. The network has n attractors and must store the stimulus value for T seconds before being read out. From a coding perspective, we have a chain where random input $X \in [1, m]$ is loaded into attractor $Y(0) \in [1, n]$ and remains in storage until $Y(T) \in [1, n]$, after which it is finally read out as response $Z \in [1, m]$ (Fig. 6A). If $n < m$ then the transition $X \rightarrow Z$ involves a compression ($X \rightarrow Y(0)$) and expansion ($Y(T) \rightarrow Z$) of data, causing errors in transmission due to the quantization of the neural representation (Materials and Methods).

The transition $Y(0) \rightarrow Y(T)$ involves diffusion across the network, which also degrades storage. To compute the probability of transitioning from one attractor to another during the storage phase, we need only integrate the Gaussian approximation with variance $D_{\text{eff}}(n)T$ over the appropriate domain (Fig. 6B,C). In this way, we calculate the matrix of transition probabilities from one attractor to another during the retention interval. With this matrix, we can calculate the information lost due to diffusion (Materials and Methods). Naturally, as the product $D_{\text{eff}}(n)T$ increases, diffusion becomes more prominent (Fig. 6B,C), and information loss due to diffusion increases.

In total, an increase in n has the dual effect of reducing quantization error, yet increasing diffusion error. Thus, we predict that spatial heterogeneity (measured by n) causes a tradeoff between quantization and diffusion-based error, and an optimal heterogeneity will maximize the overall information flow across the channel. We explore this prediction in the next section.

Optimizing information flow with heterogeneous network coupling

Delay time T and the number of possible stimuli m can be easily controlled in working memory experiments (Funahashi et al., 1989). By fixing the protocol in working memory tasks, it has been shown animals can improve their performance through extensive training, and boost their average reward rate (Meyer et al., 2011). Performance improvements are likely caused by modifications to the structure of networks underlying working memory, so we presume the spatial heterogeneity of the network, parametrized by n , evolves internally through reward-based plasticity mechanisms (Schultz, 1998; Wang, 2008; Klingberg, 2010). To measure the overall success of storage, we consider the mutual information I between X and Z for both the potential well model and the full spiking network (Materials and Methods). Mutual information measures the reduction in uncertainty in stimulus X when response Z is known. Furthermore, mutual information allows for a clean dissection of the information loss due to quantization-based and diffusion-based errors.

The stimulus space compression involved in $X \rightarrow Y(0)$ involves a loss of $\log_2(m/n)$ bits, a quantity that decreases with n . Calculating information loss due to diffusion ($Y(0) \rightarrow Y(T)$) requires that we compute the effective diffusion coefficient $D_{\text{eff}}(n)$. To obtain $D_{\text{eff}}(n)$ for the potential well model, we use our analytical approach (see Materials and Methods, Eq. 9), while for the spiking model we use a numeric fit to $D_{\text{eff}}(n)$ (Fig. 3). Information loss due to diffusion increases with $D_{\text{eff}}(n)$, which in turn

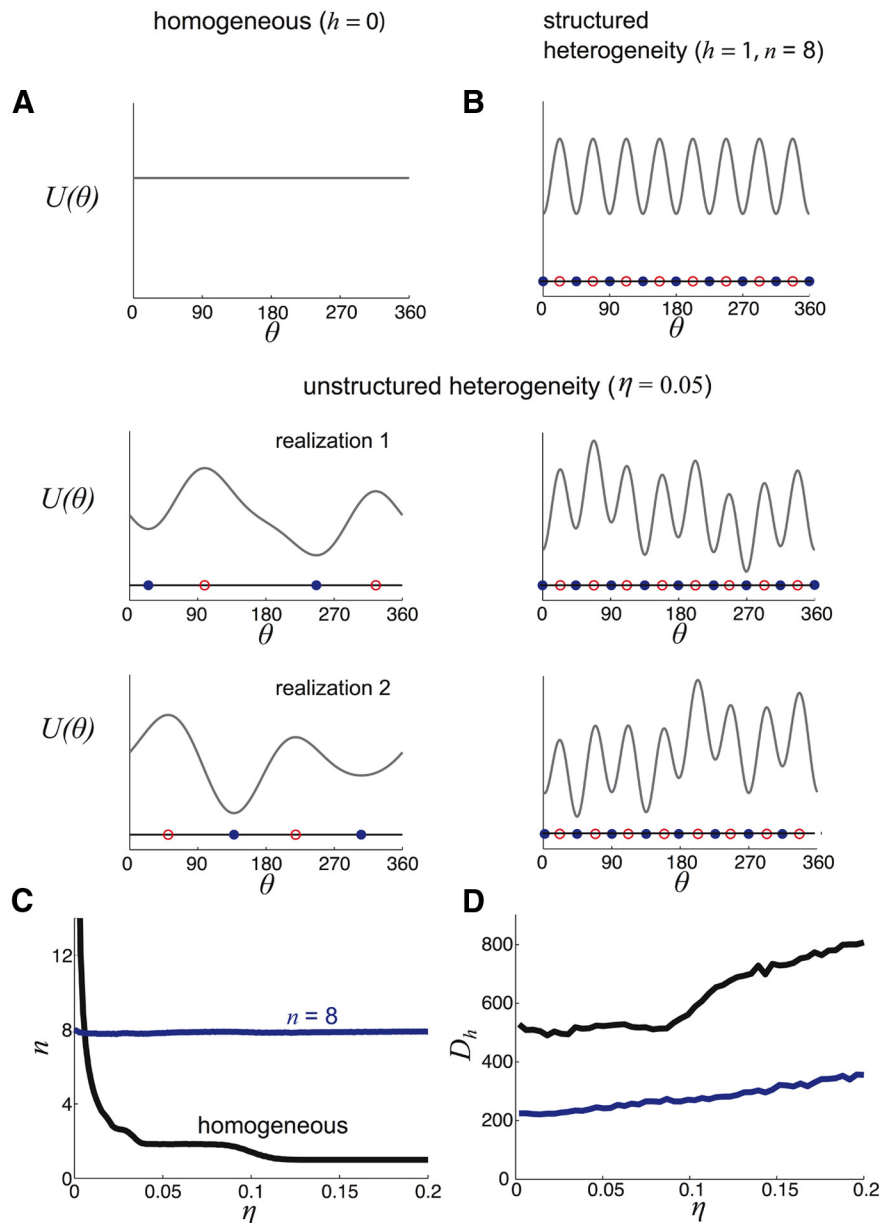


Figure 5. Storage in potentials with structured heterogeneity resists degradation from unstructured heterogeneity. **A**, Adding unstructured heterogeneity to a flat potential function drastically alters the state space of attractors. **B**, When structured heterogeneity is already present, adding unstructured heterogeneity does not change the number of attractors or their positions. **C**, The number of attractors n is strongly influenced by the severity of unstructured heterogeneity η in the homogeneous potential. Starting with $n = 8$ attractors, adding unstructured heterogeneity does not alter n . **D**, Effective diffusion increases more for the homogeneous potential as a function of η than for the potential containing structured heterogeneity.

increases with n (Figs. 3, 4E). Given both sources of information loss, we compared the channel theory prediction for $I(X, Z)$ to direct estimates of $I(X, Z)$ based on the joint density $p(X, Z)$; Materials and Methods), for both the spiking and potential well models.

For a fixed m and very short delay time T the information $I(X, Z)$ monotonically increases with n and is maximized when $n = m$ (Fig. 7A,B; $T = 0.1$ s). This is expected, since recall is near immediate, so that diffusion-based error is negligible and quantization error dominates the information loss. This error vanishes when $n = m$ and $I(X, Z)$ approaches the stimulus entropy ($\log_2(m)$ bits). As the delay time is increased, information loss due to quantization error does not change, but errors due to diffusion increase. For sufficiently long T , information peaks at a value of

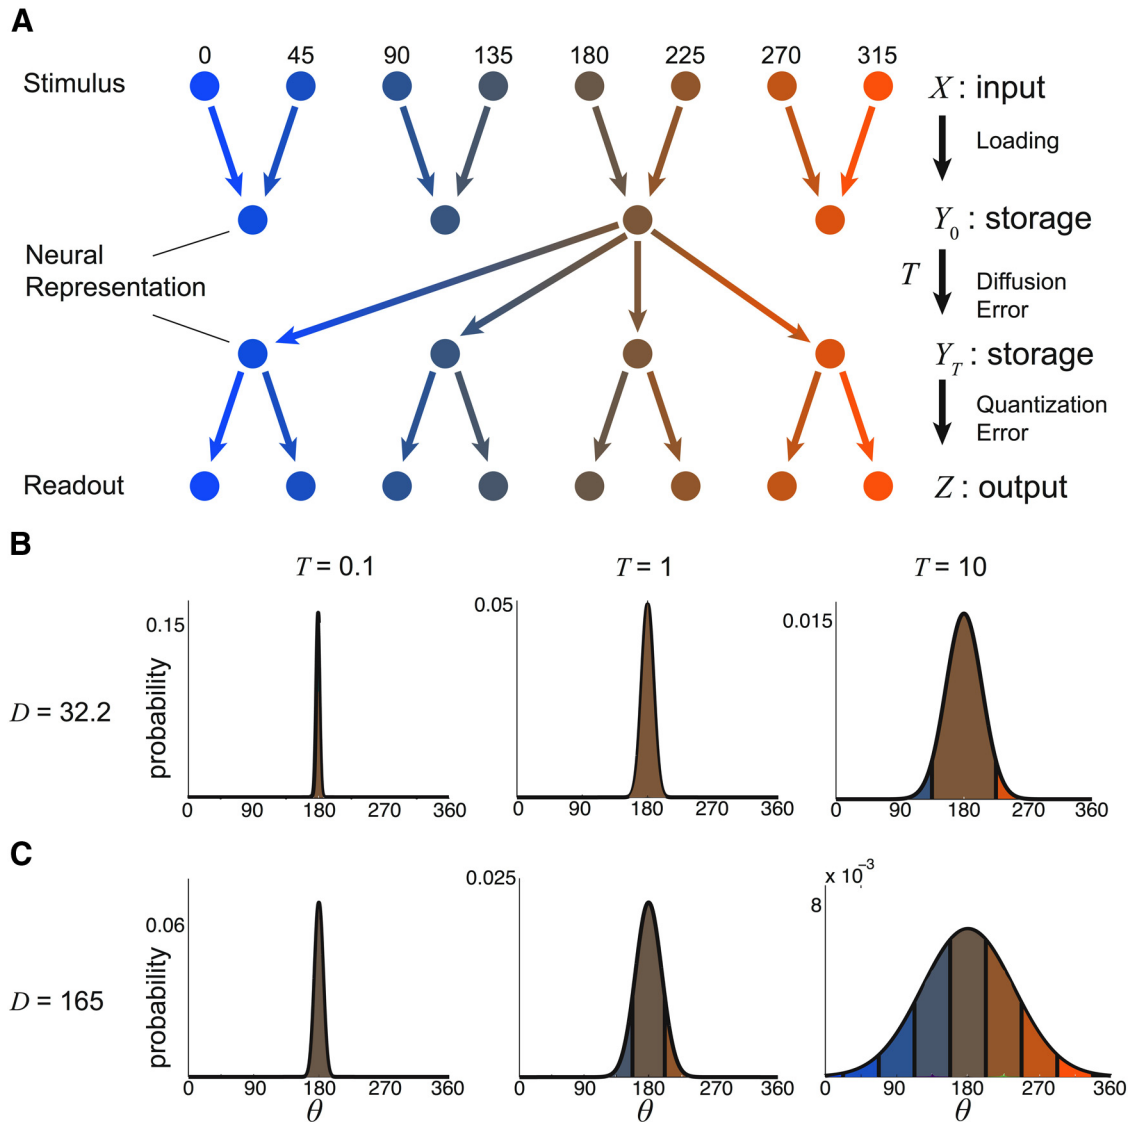


Figure 6. Noisy channel description of memory storage. **A**, Loading m possible initial conditions into n possible wells initially reduces information. After the storage period (T), information may have been lost due to hops between wells. **B**, Purely Gaussian probability density with the effective diffusion coefficient D_{eff} calculated when $n = 4$ and $\sigma^2 = 0.16$. The area of each filled portion represents the probability of recalling the cue angle associated with that color. Each area corresponds to the probability of transitioning from the original state to that state $p_{j \rightarrow k}$. **C**, For $n = 8$, the effective diffusion coefficient D_{eff} is larger, leading to faster spreading.

$n_{\text{max}} < m$ in both the potential well model and the spiking model (Fig. 7A,B; $T = 10$ s). The value n_{max} marks a compromise between quantization and diffusion errors. For the potential well model, the optimal heterogeneity n_{max} decreases as the delay time T increases (Fig. 7C), since diffusion error grows as T increases. In total, we find that for sufficiently long delay times the degree of heterogeneity n should be less than the stimulus size m to optimize information transfer.

For a fixed delay time T , varying the number of possible inputs m also shifts n_{max} . Diffusion error is independent of the number of possible inputs m ; however, the total possible information increases with the number of inputs m . For small m , we have that $n_{\text{max}} = m$ since when $n \geq m$, the quantization error is always zero and network diffusion increases with n (Fig. 8A,B; $m = 4$). However, for larger m , we find $n_{\text{max}} < m$, due to a compromise between quantization error and diffusion (Fig. 8A,B; $m = 16$). These results hold for the potential well model over a wide range of m (Fig. 8C). Overall, we highlight that a combination of varying T and m uncovers the effect of diffusion and quantization

error on mutual information in a working memory network. In particular, for many combinations of T and m , an optimal spatial heterogeneity for information transfer can be found.

The information $I(X, Z)$ measures the general relation between X and Z , one that is decoder independent. However, in psychophysical experiments, a reward is only given when the recall is correct, i.e., $X = Z$. Thus, it is important to consider how the probability of correct recall depends on the spatial heterogeneity of excitatory connections. In both the potential well and spiking network models, the probability of correct recall is maximized for a fixed $n < m$ when T is sufficiently long (Fig. 9A,B), consistent with observations of $I(X, Z)$ (Fig. 7A,B). The n that maximizes the probability of correct recall decreases as storage time increases (Fig. 9C), also in agreement with results using $I(X, Z)$ (Fig. 7C). The probability of correct recall provides a quantification of error that weights all incorrect responses the same. To use knowledge of the spatial organization of the cue set in determining error, we also measure the impact of the number of attractors on the angular difference between the recalled and cued

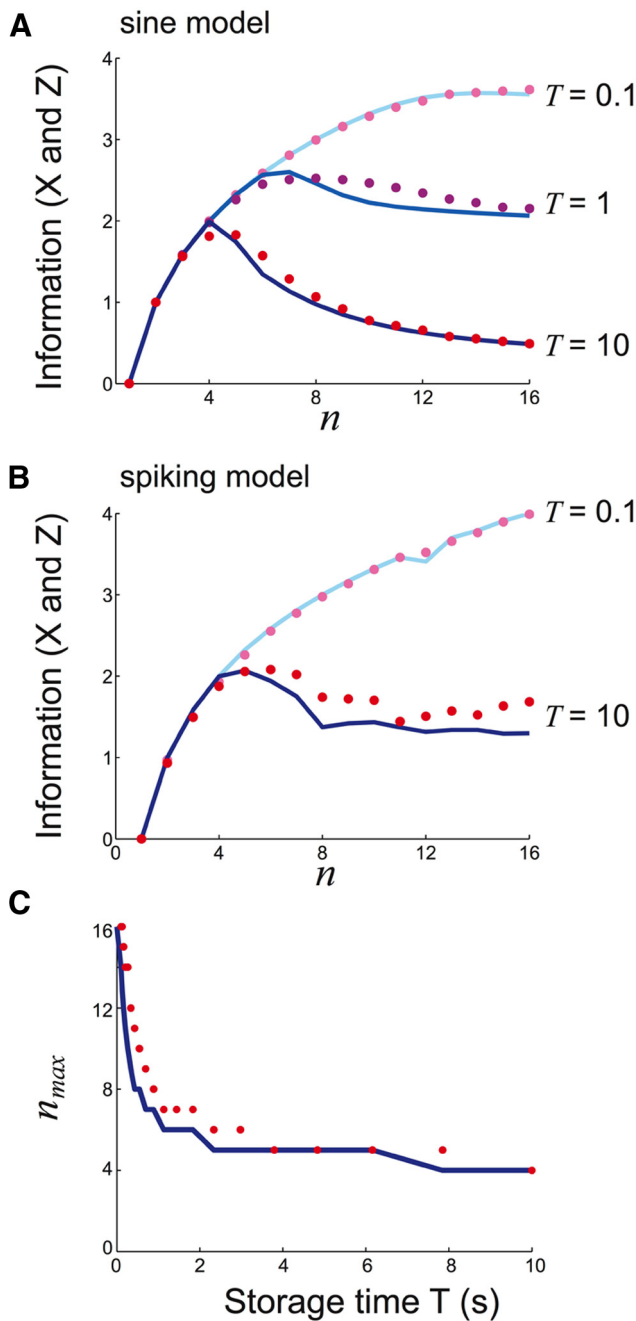


Figure 7. An optimal number of attractors $n_{max} < m$ emerges as the delay time T grows with m fixed. **A**, Mutual information $I(X; Z)$ between input particle position X and output recalled position Z varies with n in the sine model for delay (storage) times $T = 0.1$, $T = 1$, and $T = 10$; input number $m = 16$; well height $h = 1$; and noise variance $\sigma^2 = 0.16$. **B**, Mutual information $I(X; Z)$ calculated between initial X and final Z position of the bump in the spiking network for delay times $T = 0.1$ and $T = 10$ with $m = 16$ inputs. **C**, Keeping the number of possible initial positions $m = 16$ fixed reveals that n_{max} decreases monotonically with delay time T .

stimulus position. Specifically, we compute the variance of the difference between the recalled and input cue location $X - Z$. The magnitude of the recall error is minimized for a fixed $n < m$ (Fig. 9D,E), corroborating our findings for $I(X, Z)$ and proportion correct. Thus, our core finding that information transfer across the memory network is maximized for a specific degree of spatial heterogeneity also holds for a measure of task performance.

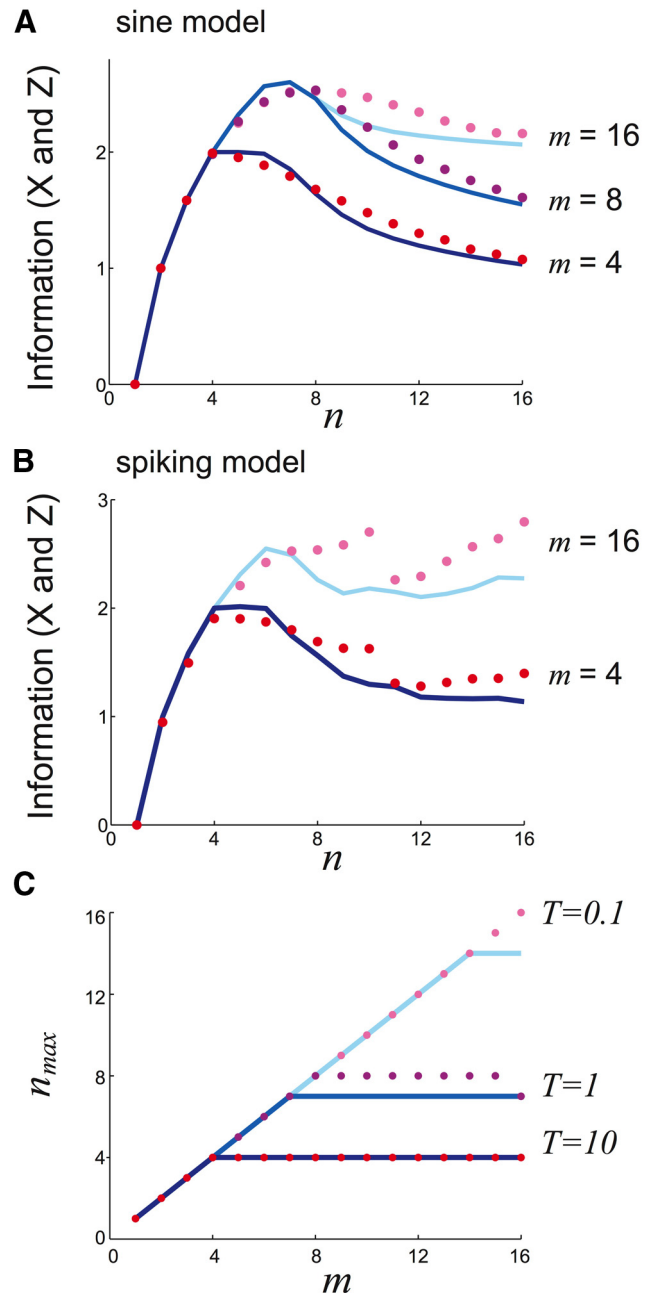


Figure 8. An optimal number of attractors $n_{max} < m$ emerges as m is increased with T fixed. **A**, $I(X; Z)$ in sine model for input numbers $m = 4$, $m = 8$, and $m = 16$ with delay time $T = 1$ fixed. **B**, $I(X; Z)$ in spiking model for input numbers $m = 4$ and $m = 16$ with delay time $T = 5$ s fixed. **C**, As the delay time T increases, n_{max} reaches an optimum at smaller cue numbers m . Noise variance $\sigma^2 = 0.16$.

Discussion

We have outlined how both neural architecture and noisy fluctuations determine error in working memory codes. In working memory networks, the position of a bump in spiking activity encodes the memory of a stimulus, and input fluctuations cause diffusion of the bump position, which degrades the memory. Spatially heterogeneous recurrent excitation reduces the diffusion of bumps by stabilizing a discrete set of bump positions. However, this also introduces memory quantization, limiting the capacity of information transfer. By analyzing the information loss incurred by both error sources, we can maximize the transfer

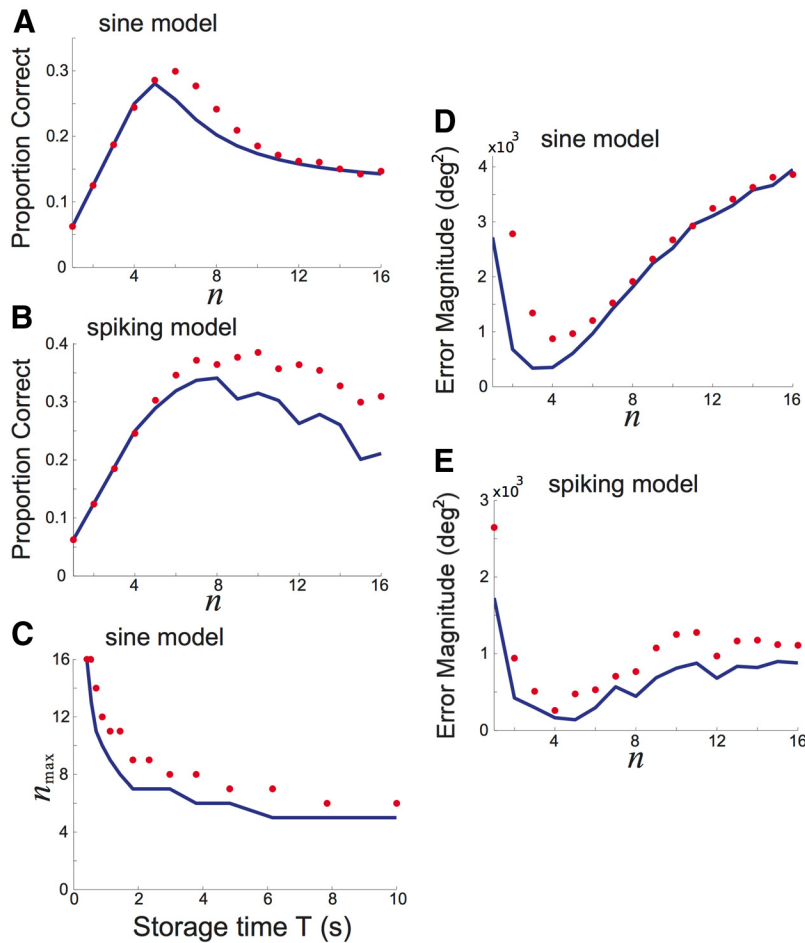


Figure 9. The proportion of correct responses depends on the number of possible outputs, n . Blue curves are computed using theoretical probability densities and red dots employ numerically computed probability densities (see Materials and Methods). **A**, Proportion $p(Z = X)$ of correct responses Z varies with n in the sine model for delay (storage) time $T = 10$ s. **B**, Proportion of correct responses $p(Z = X)$ as a function of n in the spiking model for delay (storage) time $T = 10$ s and input number $m = 16$. **C**, The number of outputs n_{max} , that maximizes the proportion of correct responses as the storage time T is varied in the sine model. **D**, Error magnitude (in degrees squared) varies with n in the sine model for delay (storage) time $T = 10$ s. **E**, Error magnitude (in degrees squared) varies with n in spiking model for delay (storage) time $T = 10$. Input number $m = 16$. In the sine model, well height $h = 1$ and noise variance $\sigma^2 = 0.16$.

of information between the stimulus and the memory output by tuning the spatial heterogeneity of recurrent excitation. We found that the ideal heterogeneity gives a number of attractors in the network n_{max} , which can be less than the number of possible inputs to the network m .

Robust bump dynamics through quantization

Networks whose dynamics lie on a continuum attractor have steady-state activity that can be altered by arbitrarily weak noise and input (Bogacz et al., 2006). The advantage of this feature is that two stimuli with an arbitrarily fine distinction can be reliably stored and distinguished upon recall. However, this structure requires fine-tuning of network architecture, since any parametric jitter will destroy a continuum attractor. Previous work has shown how spatial heterogeneity in recurrent excitatory coupling quantizes the continuum attractor and stabilizes persistent network firing rates to perturbations in model parameters (Koulikov et al., 2002; Brody et al., 2003; Goldman et al., 2003; Cain and Shea-Brown, 2012) and fluctuations (Fransén et al., 2006). We believe our results apply to these models and have extended this previous work in two major ways.

First, we have shown that quantizing the state space of a spatially structured network into a finite number of attractors stabilizes bump position to dynamic noise. Second, we have shown that there is an optimal number of attractor positions for storing stimuli when dynamic noise is present. The optimal number can be lower than the actual quantization of possible stimuli, so that under-representing stimulus space can lead to more reliable coding. Studies of networks encoding the memory of eye position show individual neurons exhibit bistability in their firing rates (Aksay et al., 2003), which motivated modeling their firing rate to input relations as quantized, staircase-shaped functions (Goldman et al., 2003). This provides an example of a working memory network thought to provide a discrete delineation of a continuous variable. Our results also suggest parametric working memory networks should coarsen the stored signal to guard against diffusion error.

The advantage of spatial heterogeneity

Past work has suggested spatial heterogeneity in working memory networks is a barrier to reliable memory storage (Zhang, 1996; Renart et al., 2003; Itskov et al., 2011; Hansel and Mato, 2013). In these studies, parameters of single neurons (Renart et al., 2003) or synaptic architecture (Itskov et al., 2011; Hansel and Mato, 2013) change throughout the network in a spatially aperiodic way. Substantial quantization error results since the bump drifts toward one of a finite number of attractor positions that may not be evenly spread over representation space. Renart et al. (2003) show this effect can be overcome

by considering homeostatic mechanisms that balance excitatory drive to each neuron in the network, halting drift of bumps altogether. On the other hand, both Itskov et al. (2011) and Hansel and Mato (2013) show drift can be slowed by including short-term facilitation in the network. Rather than exploring ways to remove the effective drift introduced by spatial heterogeneity, we have shown spatial heterogeneity can improve network coding. As long as quantization error is outweighed by a reduction in diffusion error, heterogeneous networks make less overall error in recall tasks than spatially homogeneous networks.

Our model represents the space of possible oriented cues as evenly distributed in space with uniform probability of presentation, a protocol often used in experiments (Funahashi et al., 1989; White et al., 1994; Goldman-Rakic, 1995; Meyer et al., 2011). Thus, we reason that the ideal covering of stimulus space by the network will have a uniform distribution. This translates into network spatial heterogeneity that is exactly periodic. The periodicity allows for a compact derivation of the effective diffusion coefficient $D_{eff}(n)$. Were the stimulus set to have an asymmetric probability distribution, we would expect the ideal spatial heterogeneity would not produce evenly spaced attractors. In this case, approximating bump

position is possible, but motion between attractors will depend on θ and will be difficult to interpret as a simple diffusion process. Nevertheless, we expect that the specifics of spatially uneven heterogeneous coupling will significantly impact both the attractor quantization and stochastic drift across the network, and control the information transfer from stimulus to recall.

Mechanisms that produce structured heterogeneity

We conceive of two main biophysical processes that could produce structured spatial heterogeneity in a working memory network. First, Hebbian plasticity may operate at locations in the network that are driven by common external cues. Such cues will consistently activate neurons of similar orientation preference, so clusters of similarly tuned cells will tend to strengthen recurrent excitation between each other (Goldman-Rakic, 1995; Clopath et al., 2010; Ko et al., 2013). Recent experiments show training does increase the delay period firing rates of neurons with a preference for the encoded cue (Meyer et al., 2011), which may occur due to reinforcement of recurrent excitation. This mechanism would create attractors only at locations in the network that consistently receive feedforward input during training. Neurons that are never directly stimulated by cues would be deprived of continual reinforcement of their excitatory inputs, allowing broadly tuned inhibition to decrease their delay period firing (Wang, 2001). In the framework of our models, a network trained on n cue locations would form n attractors. Depending on the length of delay, this might not be the optimal number of attractors, but it would improve coding compared with the network without quantization.

Second, reward-based plasticity mechanisms signaled by dopamine may supervise the reinforcement of synaptic excitation to form a network with the optimal number of attractors. Many studies have verified that dopamine can carry reward signals back to the network responsible for a correct action (Schultz, 1998). Selectively acting on specific subsets of neurons, dopamine can prompt plasticity in network architecture to improve future chances of rewards (McNab et al., 2009; Klingberg, 2010). Such supervisory mechanisms could seek an optimal architecture in the network to maximize reward yields for a fixed retention time and number of possible cues. As we have demonstrated, this resulting structured heterogeneity will improve coding, even if there is unstructured heterogeneity present (Wang et al., 2006).

Relating diffusion of neural activity to behavior

Our results (and those of many other past studies) assume that neural activity has a diffusive component. However, how exactly neural variability drives behavioral variability is largely unknown (Britten et al., 1996; Churchland et al., 2011; Brunton et al., 2013; Haefner et al., 2013). Psychophysical studies of spatial working memory tasks reveal that subjects typically respond with nonzero error. In particular, Ploner et al. (1998) show that the midspread of memory-guided saccades in humans scales sublinearly with delay time over 0.5–20 s. This scaling is consistent with a diffusion of neural activity involved in the storage of memory, giving support to our modeling assumptions.

In contrast to these data, recent psychophysiological work in rodents and humans performing decision-making tasks lasting 0.5–2 s suggests that models with sensory noise, rather than internal diffusion, best capture these behavioral data (Brunton et al., 2013). However, this study ultimately considers a two-alternative forced-choice task, and does not consider the storage and recall of inputs over a large stimulus space. When there are only two attractors in our network, the diffusion coefficient is

near zero, consistent with Brunton et al. (2013). In addition, the timescale of tasks studied by Brunton et al. (2013) may not be long enough to substantially reveal the effects of internal diffusion. These differences between the diffusive nature of working memory and decision integrator networks suggest that more work needs to be done to link variability of neural and behavioral response.

Implications for multiple object working memory

We emphasize that we did not study network encoding of multiple object working memory. Added complications arise when several items must be remembered at once (Luck and Vogel, 1997). For instance, the error made in recalling the value of a set of multiple continuous variables increases with the set size (Wilken and Ma, 2004). Recently, it has been shown that a spiking network model can recapitulate many of these set size effects (Wei et al., 2012). Interestingly, there is an optimal spread of pyramidal synapses that minimizes errors due to set size. However, reduction of the effects of dynamic noise on the accuracy of memories has yet to be studied. Our ideas could be extended to analyze how networks that encode multiple object memory could be made more robust, applying network quantization to the storage of multiple bump attractors.

References

- Aksay E, Major G, Goldman MS, Baker R, Seung HS, Tank DW (2003) History dependence of rate covariation between neurons during persistent activity in an oculomotor integrator. *Cereb Cortex* 13:1173–1184. CrossRef Medline
- Amari S (1977) Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern* 27:77–87. CrossRef Medline
- Ben-Yishai R, Bar-Or RL, Sompolinsky H (1995) Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci U S A* 92:3844–3848. CrossRef Medline
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765. CrossRef Medline
- Britten KH, Newsome WT, Shadlen MN, Celebri S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100. CrossRef Medline
- Brody CD, Romo R, Kepecs A (2003) Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol* 13:204–211. CrossRef Medline
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340:95–98. CrossRef Medline
- Burak Y, Fiete IR (2012) Fundamental limits on persistent activity in networks of noisy neurons. *Proc Natl Acad Sci U S A* 109:17645–17650. CrossRef Medline
- Cain N, Shea-Brown E (2012) Computational models of decision making: integration, stability, and noise. *Curr Opin Neurobiol* 22:1047–1053. CrossRef Medline
- Camperi M, Wang XJ (1998) A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *J Comput Neurosci* 5:383–405. CrossRef Medline
- Chelaru MI, Dragoi V (2008) Efficient coding in heterogeneous neuronal populations. *Proc Natl Acad Sci U S A* 105:16344–16349. CrossRef Medline
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision making. *Neuron* 69:818–831. CrossRef Medline
- Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nat Neurosci* 13:344–352. CrossRef Medline
- Compte A, Brunel N, Goldman-Rakic PS, Wang XJ (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex* 10:910–923. CrossRef Medline
- Compte A, Constantinidis C, Tegner J, Raghavachari S, Chafee MV, Goldman-Rakic PS, Wang XJ (2003) Temporally irregular mnemonic

- persistent activity in prefrontal neurons of monkeys during a delayed response task. *J Neurophysiol* 90:3441–3454. [CrossRef Medline](#)
- Cover T, Thomas J (2006) *Elements of information theory*. New York: Wiley-Interscience.
- Durstewitz D, Seamans JK, Sejnowski TJ (2000) Neurocomputational models of working memory. *Nat Neurosci* 3 [Suppl]:1184–1191. [Medline](#)
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303. [CrossRef Medline](#)
- Fransén E, Tahvildari B, Egorov AV, Hasselmo ME, Alonso AA (2006) Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. *Neuron* 49:735–746. [CrossRef Medline](#)
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349. [Medline](#)
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol* 36:61–78. [Medline](#)
- Goldman MS, Levine JH, Major G, Tank DW, Seung HS (2003) Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. *Cereb Cortex* 13:1185–1195. [CrossRef Medline](#)
- Goldman-Rakic PS (1995) Cellular basis of working memory. *Neuron* 14:477–485. [CrossRef Medline](#)
- Haefner RM, Gerwinn S, Macke JH, Bethge M (2013) Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat Neurosci* 16:235–242. [CrossRef Medline](#)
- Hansel D, Mato G (2013) Short-term plasticity explains irregular persistent activity in working memory tasks. *J Neurosci* 33:133–149. [CrossRef Medline](#)
- Itskov V, Hansel D, Tsodyks M (2011) Short-term facilitation may stabilize parametric working memory trace. *Front Comput Neurosci* 5:40. [CrossRef Medline](#)
- Kilpatrick ZP, Ermentrout B (2013) Wandering bumps in stochastic neural fields. *SIAM J Appl Dyn Syst* 12:61–94. [CrossRef](#)
- Klingberg T (2010) Training and plasticity of working memory. *Trends Cogn Sci* 14:317–324. [CrossRef Medline](#)
- Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrsic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473:87–91. [CrossRef Medline](#)
- Ko H, Cossell L, Baragli C, Antolik J, Clopath C, Hofer SB, Mrsic-Flogel TD (2013) The emergence of functional microcircuits in visual cortex. *Nature* 496:96–100. [CrossRef Medline](#)
- Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. *Nat Neurosci* 5:775–782. [CrossRef Medline](#)
- Laing C, Lord GJ (2009) *Stochastic methods in neuroscience*. New York: Oxford UP.
- Laing CR, Chow CC (2001) Stationary bumps in networks of spiking neurons. *Neural Comput* 13:1473–1494. [CrossRef Medline](#)
- Lewis DA, Gonzalez-Burgos G (2000) Intrinsic excitatory connections in the prefrontal cortex and the pathophysiology of schizophrenia. *Brain Res Bull* 52:309–317. [CrossRef Medline](#)
- Lifson S, Jackson JL (1962) On self-diffusion of ions in a polyelectrolyte solution. *J Chem Phys* 36:2410.
- Lindner B, Kostur M, Schimansky-Geier L (2001) Optimal diffusive transport in a tilted periodic potential. *Fluct Noise Lett* 1:R25–R39. [CrossRef](#)
- Luck SJ, Vogel EK (1997) The capacity of visual working memory for features and conjunctions. *Nature* 390:279–281. [CrossRef Medline](#)
- Marsat G, Maler L (2010) Neural heterogeneity and efficient population codes for communication signals. *J Neurophysiol* 104:2543–2555. [CrossRef Medline](#)
- McNab F, Varrone A, Farde L, Jucaite A, Bystritsky P, Forssberg H, Klingberg T (2009) Changes in cortical dopamine D1 receptor binding associated with cognitive training. *Science* 323:800–802. [CrossRef Medline](#)
- Meyer T, Qi XL, Stanford TR, Constantinidis C (2011) Stimulus selectivity in dorsal and ventral prefrontal cortex after training in working memory tasks. *J Neurosci* 31:6266–6276. [CrossRef Medline](#)
- Miller P (2006) Analysis of spike statistics in neuronal systems with continuous attractors or multiple, discrete attractor states. *Neural Comput* 18:1268–1317. [CrossRef Medline](#)
- Osborne LC, Palmer SE, Lisberger SG, Bialek W (2008) The neural basis for combinatorial coding in a cortical population response. *J Neurosci* 28:13522–13531. [CrossRef Medline](#)
- Padmanabhan K, Urban NN (2010) Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. *Nat Neurosci* 13:1276–1282. [CrossRef Medline](#)
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811. [CrossRef Medline](#)
- Ploner CJ, Gaymard B, Rivaud S, Agid Y, Pierrot-Deseilligny C (1998) Temporal limits of spatial working memory in humans. *Eur J Neurosci* 10:794–797. [CrossRef Medline](#)
- Polk A, Litwin-Kumar A, Doiron B (2012) Correlated neural variability in persistent state networks. *Proc Natl Acad Sci U S A* 109:6295–6300. [CrossRef Medline](#)
- Rao SG, Williams GV, Goldman-Rakic PS (1999) Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: evidence for microcolumnar organization in PFC. *J Neurophysiol* 81:1903–1916. [Medline](#)
- Renart A, Song P, Wang XJ (2003) Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* 38:473–485. [CrossRef Medline](#)
- Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22:5639–5651. [Medline](#)
- Risken H (1996) *The Fokker-Planck equation: methods of solution and applications*, volume 18. Berlin: Springer.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473. [CrossRef Medline](#)
- Rosen MJ (1972) A theoretical neural integrator. *IEEE Trans Biomed Eng* 19:362–367. [Medline](#)
- Schultz W (1998) Predictive reward signal of dopamine neurons. *J Neurophysiol* 80:1–27. [Medline](#)
- Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. *Neural Comput* 18:1951–1986. [CrossRef Medline](#)
- Stein RB, Gossen ER, Jones KE (2005) Neuronal variability: noise or part of the signal? *Nat Rev Neurosci* 6:389–397. [CrossRef Medline](#)
- Wang XJ (1999) Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci* 19:9587–9603. [Medline](#)
- Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* 24:455–463. [CrossRef Medline](#)
- Wang XJ (2008) Decision making in recurrent neuronal circuits. *Neuron* 60:215–234. [CrossRef Medline](#)
- Wang Y, Markram H, Goodman PH, Berger TK, Ma J, Goldman-Rakic PS (2006) Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat Neurosci* 9:534–542. [CrossRef Medline](#)
- Wei Z, Wang XJ, Wang DH (2012) From distributed resources to limited slots in multiple-item working memory: a spiking network model with normalization. *J Neurosci* 32:11228–11240. [CrossRef Medline](#)
- White JM, Sparks DL, Stanford TR (1994) Saccades to remembered target locations: an analysis of systematic and variable errors. *Vision Res* 34:79–92. [CrossRef Medline](#)
- Wilken P, Ma WJ (2004) A detection theory account of change detection. *J Vis* 4(12):1120–1135. [Medline](#)
- Wu S, Hamaguchi K, Amari S (2008) Dynamics and computation of continuous attractors. *Neural Comput* 20:994–1025. [CrossRef Medline](#)
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci* 16:2112–2126. [Medline](#)