A Tradeoff Between Accuracy and Flexibility in a Working Memory Circuit Endowed with Slow Feedback Mechanisms

Jacinto Pereira[1,2] and Xiao-Jing Wang[1,2,3]

[1]Department of Neurobiology and Kavli Institute for Neuroscience, Yale University School of Medicine, New Haven, CT 06510, USA, [2]Center for Neural Science, New York University, New York, NY 10003, USA and [3]NYU-ECNU Institute of Brain and Cognitive Science, NYU Shanghai, Shanghai, China

Address correspondence to Xiao-Jing Wang. Email: xjwang@nyu.edu

Recent studies have shown that reverberation underlying mnemonic persistent activity must be slow, to ensure the stability of a working memory system and to give rise to long neural transients capable of accumulation of information over time. Is the slower the underlying process, the better? To address this question, we investigated 3 slow biophysical mechanisms that are activity-dependent and prominently present in the prefrontal cortex: Depolarization-induced suppression of inhibition (DSI), calcium-dependent nonspecific cationic current (I_{CAN}), and short-term facilitation. Using a spiking network model for spatial working memory, we found that these processes enhance the memory accuracy by counteracting noise-induced drifts, heterogeneity-induced biases, and distractors. Furthermore, the incorporation of DSI and I_{CAN} enlarges the range of network's parameter values required for working memory function. However, when a progressively slower process dominates the network, it becomes increasingly more difficult to erase a memory trace. We demonstrate this accuracy–flexibility tradeoff quantitatively and interpret it using a state-space analysis. Our results supports the scenario where *N*-methyl-ᴅ-aspartate receptor-dependent recurrent excitation is the workhorse for the maintenance of persistent activity, whereas slow synaptic or cellular processes contribute to the robustness of mnemonic function in a tradeoff that potentially can be adjusted according to behavioral demands.

Keywords: DSI, I_{CAN}, persistent activity, prefrontal cortex, synaptic plasticity

Introduction

Working memory is thought to be represented by persistent activity (Fuster and Alexander 1971; Gnadt and Andersen 1988; Funahashi et al. 1989; Amit 1995; Goldman-Rakic 1995; Miller et al. 1996; Romo et al. 1999; Wang 2001; Major and Tank 2004). Such activity patterns are likely sustained by positive feedback processes in a neural circuit, but the precise mechanisms remain unresolved. Computational models stressed the role of recurrent synaptic excitation (Amit 1995; Amit and Brunel 1997; Camperi and Wang 1998; Durstewitz et al. 2000; Brunel and Wang 2001) that depends on the *N*-methyl-ᴅ-aspartate (NMDA) receptors (Wang 1999; Compte et al. 2000; Lim and Goldman 2013), a prediction supported by findings from a recent experiment (Wang et al. 2013).

Other synaptic and cellular process, present in the prefrontal cortex (PFC), are likely involved in mnemonic persistent activity, including short-term facilitation (STF; Hempel et al. 2000; Wang et al. 2006; Mongillo et al. 2008; Szatmary and Izhikevich 2010; Hansel and Mato 2013), depolarization-induced suppression of inhibition (DSI; Carter and Wang 2007), and calcium-dependent nonspecific cationic current (I_{CAN}; Egorov et al. 2002; Tegnér et al. 2002; Fransén et al. 2006; Yoshida and

Hasselmo 2009; Kulkarni et al. 2011; Kalmbach et al. 2013). STF and I_{CAN} provide feedback excitation, whereas DSI is a disinhibition process. All are activity-dependent, thus become selective for neurons that show elevated persistent activity. Furthermore, these mechanisms operate with biophysical time constants much slower than the NMDA receptor-mediated synaptic excitation. Therefore, the long-standing question (Major and Tank 2004) has gained urgency: What may be the relative contributions to working memory function of these slow synaptic and cellular processes versus the recurrent network mechanism?

We analyzed the role of slow biophysical processes in mnemonic persistent activity, using a biologically based continuous spiking circuit model for spatial working memory. This model system is endowed with a resting state and a continuum of spatially tuned persistent activity patterns ("bump attractors") for memory storage of an analog quantity such as spatial location (Camperi and Wang 1998; Compte et al. 2000; Gutkin et al. 2001; Laing and Chow 2001; Renart et al. 2003; Carter and Wang 2007; Wei et al. 2012; Murray et al. 2014). During a mnemonic delay period, a bump attractor drifts over time (Compte et al. 2000; Carter and Wang 2007; Murray et al. 2014), resulting in random deviations of the memory away from the to-be-remembered sensory cue. Additionally, heterogeneity in single neurons disrupts the continuous family of attractors (Ben-Yishai et al. 1995; Tsodyks and Sejnowski 1995; Zhang 1996), leading to systematic drifts of memory trace (Renart et al. 2003; Itskov et al. 2011). Furthermore, the system may be perturbed by external distractor stimuli. Interestingly, we found that while STF, DSI, and I_{CAN} enhance the accuracy of a memory trace, they hinder rapid memory erasure and network reset. The latter is functionally desirable, since behavior demands that brief transient inputs should be sufficient to switch a working memory system from its resting state to a memory state or vice versa (Compte et al. 2000; Gutkin et al. 2001). Therefore, our study reveals a fundamental tradeoff between robustness and flexibility of working memory function instantiated by slow neurobiological mechanisms in a recurrent network.

Materials and Methods

In an oculomotor delayed response (ODR) task, the sensory stimulus is a visual cue and the motor response is a saccade to the cued location. A subject is briefly shown a visual cue that must be remembered during a delay period of a few seconds. This memory is subsequently used to perform a memory-guided behavioral response (the saccade). During the delay period, many neurons in the dorsolateral PFC show high persistent activity that is spatially selective (Funahashi et al. 1989). The present work uses a spiking network model for the ODR task that has been tested thoroughly (Compte et al. 2000; Carter and Wang 2007; Wei et al. 2012; Murray et al. 2014). The parameters were

modified starting with the original "control parameter set" in Compte et al. (2000). The model consists of a population of excitatory pyramidal cells and a population of inhibitory interneurons. Pyramidal cells are arranged in a ring-like fashion and labeled by their preferred cue direction, from 0 to 360°. A schematic of the network structure is shown in Figure 1A.

Single Neuron Model

Both pyramidal cells and interneurons are modeled as leaky integrate-and-fire units (Tuckwell 1988). Each type of cell is characterized by total capacitance C_m, total leak conductance g_L, leak reversal potential V_L, threshold potential V_{th}, reset potential V_{res}, and refractory time τ_{ref}. The values that we use in the simulations are $C_m = 0.5$ nF, $g_L = 25$ nS, $V_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 2$ ms for pyramidal cells; and $C_m = 0.2$ nF, $g_L = 20$ nS, $V_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 1$ ms for interneurons. The subthreshold membrane potential, $V(t)$, follows:

$$C_m \frac{dV(t)}{dt} = -g_L(V(t) - V_L) - I_{syn}(t)$$

where $I_{syn}(t)$ is the total synaptic current to the cell.

Synaptic Interactions

The network consists of $N_E = 2048$ pyramidal cells and $N_I = 512$ inhibitory interneurons. Neurons receive recurrent, background, and external inputs. Excitatory synaptic currents are mediated by 2-amino-3-(3-hydroxy-5-methyl-isoxazol-4-yl) propanoic acid receptors (AMPARs) and NMDARs, and inhibitory synaptic currents are mediated by γ-aminobutyric acid type A receptors (GABA_ARs). The total synaptic current to each neuron is

$$I_{syn} = I_{NMDA} + I_{AMPA} + I_{GABA} + I_{ext}$$

where I_{ext} delivers stimulus input to pyramidal cells. The dynamics of synaptic currents for neuron i follow:

$$I_{i,AMPA} = (V_i - V_E) \sum_j g_{ji,AMPA} s_{j,AMPA}$$

$$I_{i,NMDA} = (V_i - V_E) \frac{\sum_j g_{ji,NMDA} s_{j,NMDA}}{1 + [Mg^{2+}]\exp(-0.062\, V_i/mV)/3.57}$$

$$I_{i,GABA} = (V_i - V_I) \sum_j g_{ji,GABA} s_{j,GABA}$$

where $V_E = 0$ mV and $V_I = -70$ mV and $g_{ji,syn}$ denotes the synaptic conductance strength on neuron i from neuron j. NMDAR-mediated currents exhibit voltage dependence controlled by the extracellular magnesium concentration $[Mg^{2+}] = 1$ mM (Jahr and Stevens 1990).

Given a spike train $\{t_k\}$ in the presynaptic neuron j, the gating variables $s_{j,AMPA}$ and $s_{j,GABA}$ for AMPAR- and GABAR-mediated currents, respectively, are modeled as:

$$\frac{ds}{dt} = \sum_k \delta(t - t_k) - \frac{s}{\tau_s}$$

The gating variable $s_{j,NMDA}$ for NMDAR-mediated current is modeled as:

$$\frac{dx}{dt} = \alpha_x \sum_k \delta(t - t_k) - \frac{x}{\tau_x}$$

$$\frac{ds}{dt} = \alpha_s x(1 - s) - \frac{s}{\tau_s}$$

with $\alpha_x = 1$ (dimensionless), $\tau_x = 2$ ms, and $\alpha_s = 0.5$ kHz. The decay time constant τ_s is 2 ms for AMPA, 10 ms for GABA, and 100 ms for NMDA. For simplicity, background inputs are mediated entirely by AMPARs, and recurrent excitatory inputs are mediated entirely by NMDARs, as they are critical for the stability of persistent activity (Wang 1999;

Compte et al. 2000; Wang et al. 2013). All cells receive background excitatory inputs from other cortical areas. This overall external input is modeled as uncorrelated Poisson spike trains to each neuron at a rate of $v_{ext} = 1800$ Hz per cell, with AMPAR maximal conductances of 3.1 nS on pyramidal cells and 2.38 nS on interneurons.

Network Connectivity

As stated above, pyramidal cells are organized in a ring architecture and are tuned to the angular location on a circle (0–360°, Fig. 1A), with uniform distribution of their preferred angles. The network structure follows a columnar architecture, such that pyramidal cells with similar stimulus selectivity are preferentially connected to each other. The synaptic conductance on neuron i from neuron j, $g_{ji,syn} = W(\theta_j - \theta_i)G_{syn}$, where θ_i is the preferred angle of neuron i, and $W(\theta_j - \theta_i)$ is the connectivity profile normalized such that:

$$\frac{1}{360°} \int_{0°}^{360°} W(\theta) d\theta = 1$$

For pyramidal-to-pyramidal connections, $W(\theta_j - \theta_i) = J^- + (J^+ - J^-) \exp[-(\theta_j - \theta_i)^2/2\sigma^2]$. We use $J^+ = 1.62$ and $\sigma = 14.4°$. J^- is determined using the normalization condition of W. All other synaptic connection profiles are unstructured. Synaptic conductance strengths are given by $G_{EE} = 0.381$ nS, $G_{EI} = 0.292$ nS, $G_{IE} = 1.336$ nS, $G_{II} = 1.024$ nS.

Stimulus

Inputs are modeled as an injected current with a Gaussian profile, $I(\theta) = I_0 \exp[-(\theta - \theta_c)^2/2\sigma_I^2]$, where the maximum current $I_0 = 200$ pA, except otherwise noted. θ_c is the stimulus location, and the width parameter $\sigma_I = 18°$.

Slow Calcium-Dependent Nonspecific Cationic Current

I_{CAN} can trigger a sustained depolarization outlasting the stimulus for several seconds (Haj-Dahmane and Andrade 1998; Strübing et al. 2001; Egorov et al. 2002; Tegnér et al. 2002). The activation of this current requires a rise in intracellular calcium. In some simulations (results in Figs 2, 4, and 5), I_{CAN} was added to the network model (described above) according to the following equation (Tegnér et al. 2002):

$$I_{CAN} = -g_{CAN} m_{CAN}^2 (V - E_{CAN})$$

$$\frac{dm_{CAN}}{dt} = \phi_{CAN} \times \left(\frac{m_\infty([Ca^{2+}]) - m_{CAN}}{\tau_{CAN}([Ca^{2+}])}\right)$$

$$m_\infty([Ca^{2+}]) = \frac{\alpha[Ca^{2+}]^2}{\alpha[Ca^{2+}]^2 + \beta}$$

$$\tau_{CAN}([Ca^{2+}]) = \frac{1}{\alpha[Ca^{2+}]^2 + \beta}$$

with $g_{CAN} = 1.5$ nS, $E_{CAN} = -20$ mV, $\beta = 0.002$ ms^{-1}, $\alpha = 0.0056$ ms^{-1} μM^{-2}. ϕ_{CAN} is used to adjust the effective time constant of I_{CAN}, without changing the steady-state levels of activity.

Calcium influx to pyramidal cells is triggered by spikes and obeys first-order kinetics as follows (Liu and Wang 2001):

$$\frac{d[Ca^{2+}]}{dt} = \alpha_{Ca} \sum_i \delta(t - t_i) - \frac{[Ca^{2+}]}{\tau_{Ca}}$$

When an action potential fires (at time t_i), $[Ca^{2+}]$ is incremented by α_{Ca} (0.2 μM). The calcium concentration decays back to zero exponentially, with a time constant τ_{Ca} (240 ms).

Depolarization-Induced Suppression of Inhibition

DSI is detected in various regions of the brain (Llano et al. 1991; Pitler and Alger 1992; Trettel and Levine 2003). DSI is dependent on endocannabinoids that are released by active pyramidal cells, triggered by calcium influx (Ohno-Shosaku et al. 2001; Wilson and Nicoll 2001; Wilson et al. 2001). These endogenous cannabinoids retrogradely

activate type 1 cannabinoid receptors (CB1R) located on the axon terminals of interneurons that coexpress GABA and cholecystokinin (Katona et al. 1999; Marsicano and Lutz 1999). The activation of CB1R results in the suppression of transmitter release to postsynaptic pyramidal cells.

DSI was added to the network model (Figs 3–5) as previously described in Carter and Wang (2007) and the same parameters were used, unless noted otherwise. Briefly, the inhibitory synaptic conductance g_{GABA} to a pyramidal cell is multiplied by a factor D, which is proportional to the fraction of inhibitory synapses that are sensitive to cannabinoid and their presynaptic release probability. D varies between 0 and 1. There is no DSI effect if D is set to 1. DSI is the fractional reduction in inhibitory event size or frequency. The dynamics of D are described by the following equation (Carter and Wang 2007):

$$\frac{\mathrm{d}D}{\mathrm{d}t} = \phi_D \times \left(\frac{1 - D}{\tau_D} - \beta_D \times [\text{Ca}^{2+}] \times (D - D_{min}) \right)$$

where $[\text{Ca}^{2+}]$ represents the intracellular calcium concentration in the pyramidal cell and has the same kinetics as I_{CAN}. When $[\text{Ca}^{2+}]$ accumulates, D decreases with a rate controlled by β_D ($1.66 \times 10^{-5}\,\mu\text{M}^{-1}\,\text{ms}^{-1}$, leading to disinhibition. D is bounded below at D_{min}, which determines the maximum disinhibition and biophysically corresponds to the maximum number of synapses that are cannabinoid-sensitive multiplied by the maximal reduction in release probability at each synapse due to DSI. Unless stated otherwise, D_{min} was set to 0.96, corresponding to a maximum DSI of 4%. When the pyramidal cell ceases to be active, D recovers back to a maximal value of 1 with a time constant τ_D (16.7 s). The factor ϕ_D accounts for temperature sensitivity and was used to adjust the effective time constant of DSI without changing the steady-state levels of activity.

Short-Term Facilitation

In simulations where we incorporated STF (results in Figs 6–8), only the recurrent excitatory synapses are facilitatory. To implement STF, the parameter α_x is multiplied by F, which is the facilitation factor and obeys the following dynamical equation (Matveev and Wang 2000):

$$\frac{\mathrm{d}F}{\mathrm{d}t} = \alpha_F \sum_i \delta(t - t_i)(1 - F) - \frac{F}{\tau_F}$$

The parameter α_F controls the facilitation potency and was set at 0.6. The facilitation factor F changes smoothly during a spike, but undergoes a discrete jump in the limit of approximating the spike by a delta function. In the numerical simulation, F is updated at each spike time as: $F^+ = 1 - (1 - F^-)\,e^{-\alpha F}$, where F^- and F^+ correspond to the values just before and after the spike, respectively.

Parameter Change

A key manipulation in our study is to gradually change the timescale of a biophysical process. For I_{CAN}, we varied the parameter ϕ_{CAN}, which scales the speed of the channel kinetics without affecting the averaged steady-state level of the activity variable m_{CAN}. Similarly, we varied the parameter ϕ_D to systematically change the time constant of DSI while preserving the average level of the activity variable D. Unlike I_{CAN} or DSI, for STF the activity variable F undergoes discrete jumps in time and what matters is its value immediately after each jump due to a presynaptic spike, rather than the temporal average. For this reason, we varied τ_F directly (see Results for more details).

When a slow mechanism is added to a network model, the overall level of activity of the excitatory population changes significantly, to a degree correlated to the nature and strength of the mechanism. This changes the shape of a population activity pattern and may even disrupt its stability. For this reason, when I_{CAN}, DSI, or STF were present in the model, G_{EE} was adjusted from 0.381 to 0.378, 0.379, or 0.383 nS, respectively. This way, the network maintained consistently a fixed steady-state activity across all simulations, allowing a fair comparison between scenarios.

Analysis of Simulation Data

To determine the remembered cue location at any given time, we used the population vector, which is a simple readout of the peak location of a spatially tuned persistent activity pattern (Georgopoulos et al. 1982).

The minimum time to shutdown ($t_{SHUT,MIN}$), in Figures 2E, 3E, and 6A, was determined as follows. For each time constant (τ), a range of shutdown pulse durations (t_{SHUT}) was considered. For each τ and t_{SHUT}, a set of model simulations was run, where an inhibitory input current lasting for t_{SHUT} was applied when the network was in a bump attractor state. At the end of each simulated trial (seconds after pulse offset), whether the bump state was still present or not was judged through the maximum of the firing rate profile. If >95% of simulations of a set yielded successful shutdowns, the corresponding pulse duration was accepted. Finally, for each τ, $t_{SHUT,MIN}$ was chosen as the minimum of those accepted pulse durations.

Bistability Analysis and Bifurcation Diagrams

To plot the bifurcation diagrams in Figures 4 and 7B, we ran simulations across a range of values for the varied parameter (G_{EE} and F profile, respectively) with and without cue input and measured the firing rate during the delay. The maximum firing rate across the network indicated whether the system had evolved to the memory state (typically >20 Hz) or remained at the baseline state (<5 Hz).

The simulations presented in Figure 7 were obtained with a modified model where the facilitation factor F is not a variable but is treated as a parameter with a particular spatial profile, obtained as follows. For $\tau_F = 1$ s, we determined the F profile at the onset of a shutdown input, averaged over a number of trials from previous simulations where F was a variable ($\tau_F = 1$ s, $N = 400$ trials, Fig. 7A, black curve). During the shutdown period, there is no spiking activity and F simply decays exponentially with τ_F. To reproduce this process, new F profiles were mathematically determined by decay of the original profile for periods of 4–328 ms, in 4 ms steps (gray profiles in Fig. 7A, only 16 examples are shown). Similar procedures were applied to $\tau_F = 0.5, 2, 3$, and 4 s. The profiles for longer τ_F are broader, as seen in Figure 6B. These facilitation profiles were supplied to the population of excitatory cells. Note that in both versions of the model relevant for these simulations (F as a variable or as a parameter), the stimulus was presented at the same location (180°).

Simulation Method

The model was implemented in python in the Brian simulator (Goodman and Brette 2009). The equations were integrated using a second-order Runge-Kutta algorithm (time step = 0.02 ms). The simulations were carried out in the cluster facilities of the Yale University Biomedical High Performance Computing Center.

Results

Our working memory model was designed for an ODR task, which proceeds from cue (angle) presentation, to a delay period and memory-guided behavioral response. The cue stimulus activates a group of pyramidal neurons with preferred directions around the sensory cue (first step current, lower panel of Fig. 1B). If the firing rate of this subpopulation of neurons is sufficiently elevated and mutual excitation among them is strong enough, reverberation can give rise to self-sustained persistent activity after the stimulus offset (plateau in Fig. 1B, upper panel; Wang 2001). At the end of the delay, a negative input is applied to all excitatory neurons in the network (Fig. 1B, lower panel, second step current). This shutdown pulse should be sufficient long to switch the network back to the baseline resting state.

The spatiotemporal activity pattern of the network model is shown in Figure 1C (left panel). The memory trace is encoded as a population activity pattern that persists during the delay
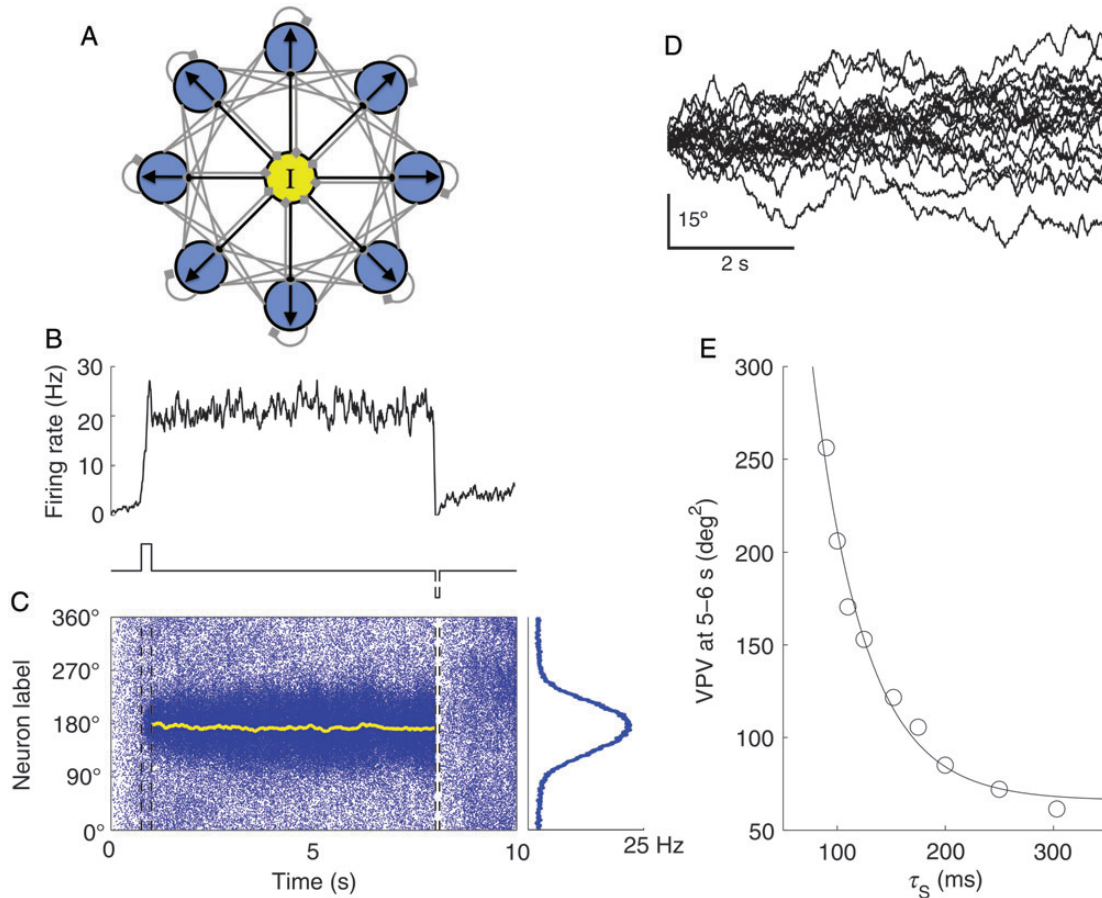
Figure 1. Persistent activity and random drifts of a memory trace in a spiking network model for spatial working memory. (*A*) Schematic of the network connectivity (all-to-all) between the excitatory (blue circles) and inhibitory (yellow circle) neurons. Light gray and black connectors indicate, respectively, excitatory and inhibitory synapses. Each excitatory cell is selective for a direction (black arrows), and the strength of connection between 2 excitatory cells is a decreasing function of the difference in their preferred directions. (*B*) Lower panel: applied current to excitatory cells. The first positive step current corresponds to cue presentation. The second negative current represents a shutdown signal. Upper panel: average firing rate of a group of 200 neurons (with preferred directions around cue location) during a trial. The activity ramps up during cue presentation, persists during delay, and is reset to a spontaneous baseline by the shutdown pulse. (*C*) Left panel: spatiotemporal pattern of excitatory cells of the same simulation as in (*A*) (cue presented at 180°). Each dot represents a spike. The yellow line is the population vector, which traces the peak of the bell-shaped persistent activity pattern (bump attractor) as the internal representation of the cue location. Right panel: Population firing profile, averaged over the delay period. (*D*) Remembered cue as measured by the population vector from 20 sample trials with the same cue location. The memory traces drift away from the initial cue during the delay, the VPV across trials quantifies this deviation so that the smaller is the VPV, and the more accurate is the memory readout. (*E*) Drift magnitude at 5–6 s of the delay period, as measured by the VPV ($N = 500$ trials), is plotted as a function of the time constant of the NMDAR-mediated synaptic excitation τ_S. The VPV decreases steeply with increasing τ_S; the fitting line is an exponential function for ease of eye inspection.

period. The spatial profile of the bump state, corresponding to the activity during the delay period, has a typical Gaussian shape (Fig. 1*C*, right panel). The population vector (shown in yellow) quantifies the peak location of the bump attractor as the internal representation of the sensory cue at any instant. In this example, the remembered cue location fluctuates slightly around the initial cue (180°) and remains reasonably close to it at the end of the delay period. Consequently, in this trial, the PFC circuit model successfully encodes and maintains a spatial memory trace, leading to an accurate readout.

Dominant Time Constant Determines Memory Accuracy

The analysis of simulations across trials reveals that the remembered cue (the population vector) as encoded by the network activity pattern displays random drifts over time (Fig. 1*D*). This is because the system is endowed with a continuous family of bump attractors, each for a directional angle as an analog quantity. During a delay period, irregular neural activity leads to random shifting of the network state among those bump

states. At the end of a trial, if the drifts have grown over time greatly, the remembered cue location could be located significantly away from the sensory cue angle. This is shown in some trials of Figure 1*D*, with deviations of >20°. These simulations therefore show a relatively low accuracy of memory representation, which implies poor performance. Note that, across trials, the average of random drifts is zero (i.e., there is no systematic drift), whereas the variance increases roughly linearly over time (Camperi and Wang 1998; Compte et al. 2000; Renart et al. 2003; Carter and Wang 2007). This variance of population vector (VPV) quantifies the magnitude of random drifts, which we used as a measure to assess the network's function: the smaller is the VPV, the more accurate is the representation of a memory trace, and the better is the behavioral performance.

A key ingredient in our working memory model is that persistent activity is stabilized by slow reverberation mediated by the NMDARs at the recurrent excitatory synapses (Wang 1999). The NMDAR-dependent synaptic current has a time constant τ_S on the order of 50–100 ms. We hypothesized that, the longer is τ_S, the more robust will be the memory trace. To test this
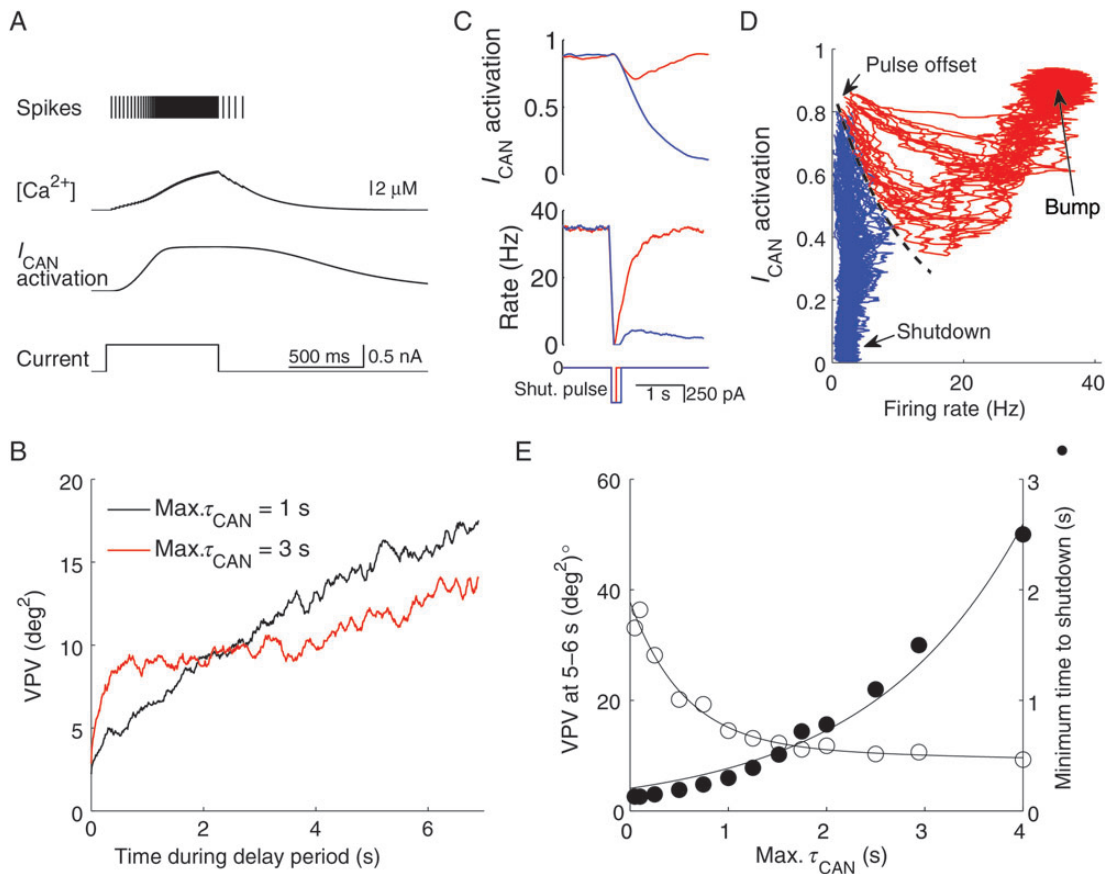
Figure 2. Tradeoff between memory accuracy and flexibility with I_{CAN}. (A) An integrate-and-fire neuron model endowed with I_{CAN}. A step current (bottom panel) induces initial firing activity (upper panel). Each spike triggers a small calcium influx (middle upper panel), which leads to a slow activation of I_{CAN} (middle lower panel). When the applied current stops, the high level of I_{CAN} activation is sufficient to induce afterdischarge of spikes. (B) Variance of the remembered cue location (VPV) during the delay period with max τ_{CAN} of 1 (black trace) and 3 (red trace) s (N = 500 trials). A longer time constant leads to smaller random drifts after an initial time needed for the mechanism to take effect. (C) With max τ_{CAN} = 500 ms, a negative pulse of 200 ms to excitatory cells is required in order to shutdown the bump state at the end of delay. Lower panel shows applied current with 2 negative pulses of lasting 100 (red) and 200 (blue) ms. Middle and upper panels: the average population firing rates and I_{CAN} activation, respectively, of 200 cells in the bump state around the initial cue location, under the 2 conditions (the same color scheme, N = 10 trials). With 100 ms, I_{CAN} activation decays by a small amount but immediately increases after the shutdown input is over, providing the necessary positive feedback for the return of the high-firing memory state. After a longer shutdown pulse (200 ms), the activation decays to such an extent that ultimately leads to the resting state. (D) State space analysis with the population rate and the I_{CAN} activation shown in (C) plotted against each other in phase space. Each trajectory corresponds to a trial and starts immediately at the shutdown pulse offset. Red trajectories evolve to the bump attractor; blue proceed to shutdown (resting state). There is a clear diagonal boundary that separates the 2 attractors (dashed black curve), suggesting the presence of an unstable manifold. (E) Tradeoff between decrease in variance of remembered cue location (VPV) and minimum time to shutdown ($t_{SHUT,MIN}$), with increasing max τ_{CAN}. Open circles were determined as in Figure 1E, with max τ_{CAN} between 50 ms and 4 s. Filled circles express $t_{SHUT,MIN}$ (see Materials and Methods) (N = 500 trials). The 2 data sets are fitted as a sum of 2 exponentials (VPV) or as a simple exponential ($t_{SHUT,MIN}$). A compromise corresponds to an optimal value of max $\tau_{CAN} \approx 1.5$ s.

possibility, we gradually varied the value of the NMDAR decay time constant, and measured the variance of the remembered cue location during a delay interval across hundreds of trials. The VPV decreases with increasing τ_S (Fig. 1E). The VPV is 206.2 deg^2 with τ_S equal to 100 ms. A substantial reduction in the VPV is observed when τ_S is increased 3-fold (300 ms, $\sigma^2 = 61.5$ deg^2). This result serves as a proof-of-principle of the idea that extending the dominant time constants decreases random drifts of persistent activity and improves the accuracy of memory representation. In the following, we will consider 3 slow, biophysically plausible mechanisms that are present in the PFC and may improve working memory function.

I_{CAN} Increases Memory Stability but Decreases System Flexibility

Figure 2A shows the spiking activity of an integrate-and-fire single neuron model endowed with the slow inward current

I_{CAN}. An external current results in action potentials that induce calcium influx, which in turn activates I_{CAN}. After the stimulus offset, the activation of I_{CAN} decays slowly, which allows it to provide positive feedback that is enough to trigger a few additional spikes (afterdischarges). It is worth noting that we assumed that I_{CAN} is not sufficiently strong to produce stable persistent activity in an isolated neuron (Fig. 2A), and we were interested in examining the contribution of the activity-dependent I_{CAN} in single neurons to the maintenance of a persistent firing pattern in a recurrent working memory circuit.

We ran simulations with I_{CAN} present in excitatory cells and measured the VPV of the delay period memory trace across trials. We tested 2 different values of max τ_{CAN} that lie within the experimentally measured range (Partridge and Valenzuela 1999; Faber et al. 2006; Gross et al. 2009; Sidiropoulou et al. 2009). With a shorter max τ_{CAN} (1 s), the VPV increases quasi-linearly with time (Fig. 2B, black curve). In contrast, with max τ_{CAN} = 3 s,
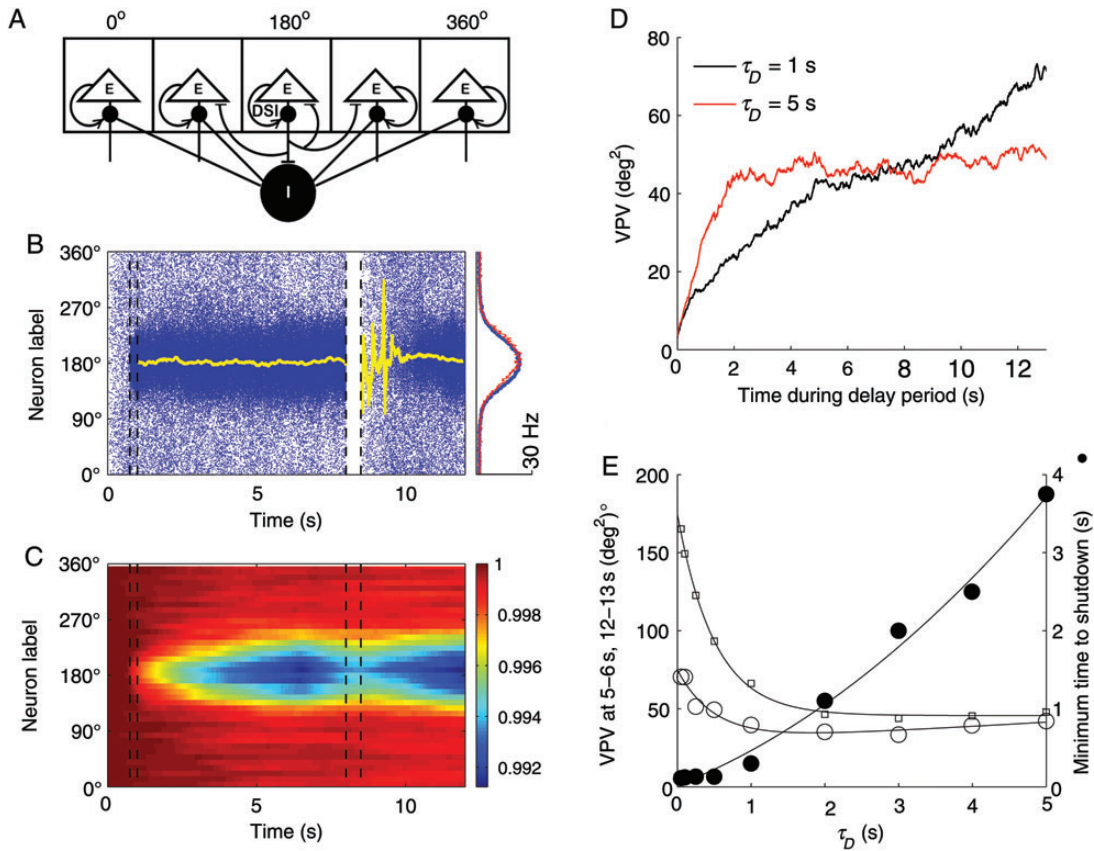
Figure 3. Tradeoff between memory accuracy and flexibility with DSI. (*A*) Schematic of network model of spatial working memory endowed with DSI. This mechanism is implemented as a cell-specific reduction in inhibitory input conductance. Adapted from Carter and Wang (2007). (*B*) Left panel: spatiotemporal pattern of excitatory cells endowed with DSI ($\tau_D = 5$ s). Cue was presented at 180° during the 0.75–1 s interval. A shutdown pulse of 500 ms was applied at 8 s. The yellow lines represent the remembered cue location during delay and after shutdown pulse. Right panel: population firing profiles, averaged over the delay period (blue) or over the last second of the simulation (red), showing that the bump state survives the shutdown input and the memory trace is not erased. (*C*) Spatiotemporal representation of the activation variable (*D*) of DSI (inverted scale, 1 means no DSI) of the same trial. Only the (*D*) value of 41 cells (recorded equidistantly in the network) is plotted. The lingering DSI trace, visible after the shutdown pulse, is sufficient to induce the re-emergence of the bump state (in *B*). (*D*) The accuracy–flexibility tradeoff with DSI. The variance of the remembered cue location (VPV) during the delay period with effective τ_D of 1 (black trace) and 5 (red trace) s (*N* = 500 trials). In the former scenario, the VPV keeps increasing almost linearly. In contrast, in the latter, it stabilizes after an initial period of 2 s. (*E*) Tradeoff between decrease in the VPV (open symbols) and $t_{SHUT,MIN}$ (closed circles), as τ_D is increased from 50 ms to 5 s (*N* = 500 trials). The VPV was determined during 2 intervals of the delay period: 5–6 s (open circles, same as Figs 1*E* and 2*E*) or 12–13 s (open squares). The data sets were fitted by solid curves for eye inspection.

the VPV shows a pronounced increase during the first second of the delay period and then plateaus in the range 10–15 deg² (Fig. 2*B*, red curve). A possible explanation for the initial rise in drifts (which is not visible for max $\tau_{CAN} = 1$ s) is that, with a slower time constant, the I_{CAN} takes longer to be activated and does not provide robustness against drifts as promptly. The crossover between the 2 time courses shows that shorter τ_{CAN} is more advantageous for shorter delay periods, whereas slower τ_{CAN} increases memory accuracy in longer delays.

The increase in memory robustness provided by I_{CAN}, however, is just one of the effects this current has in the working memory model; the incorporation of a slow mechanism also makes it harder to erase memory. At the end of a delay period, memory erasure was simulated using a negative current input to all excitatory cells, which completely silences the network. If this pulse is not sufficiently long, the network returns to the memory state, with high I_{CAN} activation and elevated neural firing (Fig. 2*C*, red traces, 100 ms pulse). With a longer shutdown pulse, in contrast, I_{CAN} deactivates to a sufficiently low level that does not allow the return of the high spiking activity and the network is switched off from a bump attractor state (Fig. 2*C*, blue traces, 200 ms pulse).

To further demonstrate the role of I_{CAN} in the memory erasure process, we studied the dependence between this current's activation and the activity level of the network. We recorded simultaneously the activation variable of I_{CAN} (m_{CAN}) and the firing rate of the network and plotted them in a state space, for several trials (Fig. 2*D*). We only recorded neurons around the cue location and in simulations that successfully maintained a memory during the delay. All trajectories initiate immediately after the shutdown ("pulse offset"). There is a clear divergence between 2 kinds of traces: In a given trial, the system's trajectory either revert back to the memory state (red traces, "bump") or decays to the resting state (blue traces, "shutdown"). A boundary (dashed line) separates the regions of attraction of the 2 states. This result shows that even though a relative weak I_{CAN} [which by itself does not yield persistent activity in a single neuron (Fig. 2*A*)] does not determine whether a network generates persistent activity per se, it can have a remarkably significant impact on the network's behavior.

Therefore, I_{CAN} stabilizes the memory trace by reducing memory drifts over time; at the same time it renders the network less flexible, that is, it may be harder to load inputs and discard old memories. This accuracy–flexibility tradeoff
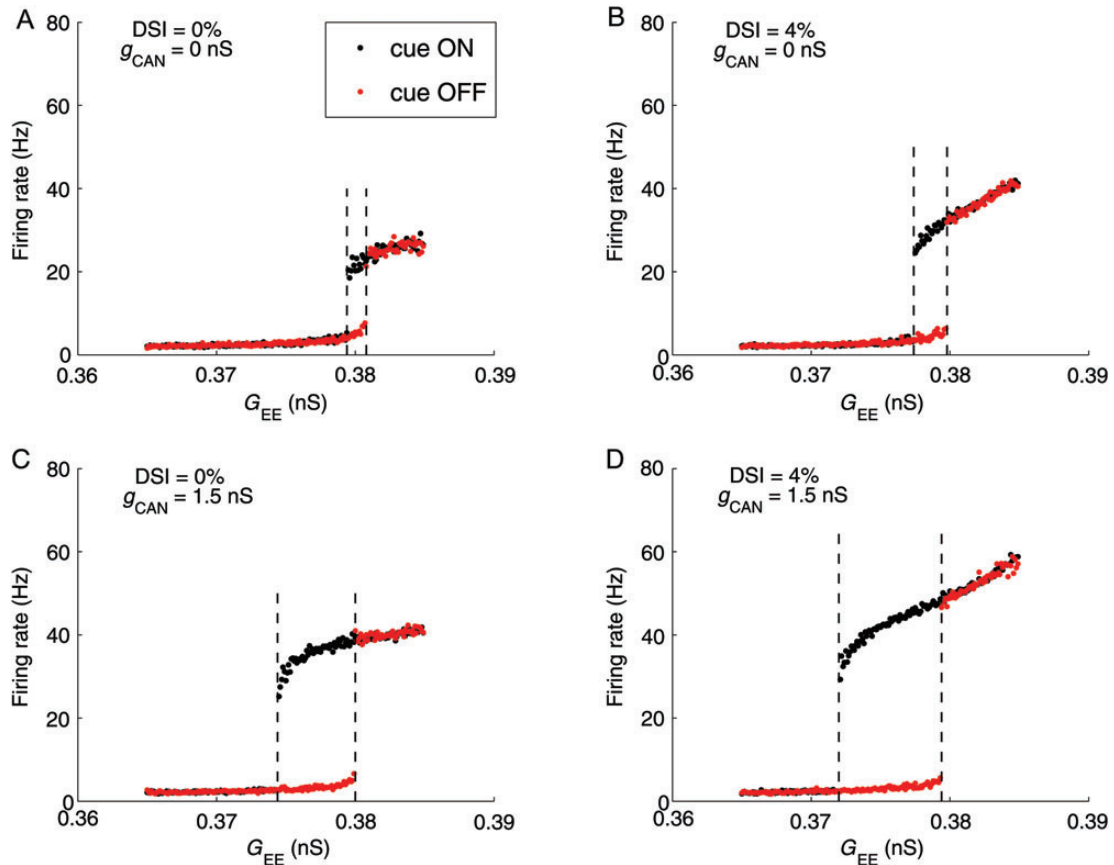
Figure 4. Multistability analysis of the working memory model as a dynamical system reveals that I_{CAN} and DSI increase the robustness of memory function. Simulations were ran with (black dots) or without (red dots) cue presentation, for a range of recurrent excitatory conductance (G_{EE}) values. The maximum firing rate among all excitatory cells, at the end of the delay period, is either low (2–6 Hz) corresponding to the resting state or higher than 20 Hz corresponding to a memory sate. The resulting state diagram is shown for the control network without slow mechanisms (*A*), with only DSI (*B*) or I_{CAN} (*C*) or both (*D*). The range of G_{EE} values for multistability are delimited by 2 vertical dashed lines. The presence of DSI (*B*) and I_{CAN} (*C*) alone increased the multistability range and also the firing rate separation between memory and resting states. These effects are larger when both mechanisms are combined (*D*).

was demonstrated more explicitly when we varied max τ_{CAN} parametrically (Fig. 2*E*). The increase in max τ_{CAN} decreases the variance of the remembered cue location (the VPV, open circles), but increases the minimum time required to shutdown the network (filled circles). A "sweet spot" corresponds to the crossover point of the 2 curves (max τ_{CAN} = 1–2 s), where the VPV is close to the minimum while $t_{SHUT,MIN}$ is reasonably short (a few hundreds of milliseconds). However, an optimal compromise for a working memory circuit could be different depending on the functional demand that may emphasize either accuracy or flexibility.

### DSI Also Shows Tradeoff Between Accuracy and Flexibility

DSI is a cannabinoid-dependent process through which synaptic inhibition to excitatory neurons is reduced by the magnitude of DSI, which in turn is controlled by the activity of the same *E* cells (Fig. 3*A*). Thus, for each neuron, a higher level of excitation leads to a weaker inhibition, resulting in an effective positive feedback. In the control network employed in this study, the recurrent excitation mostly mediated by the NMDAR is balanced by lateral inhibition, which prevents the outburst in activity. The incorporation of DSI only reduces inhibition by a modest degree (controlled by the parameter D_{min}) and does

not significantly alter the balance *E/I*. As a result, the firing activity remains at reasonable levels without diverging.

The cells that are most active during an ODR task are those around the peak of the bump activity pattern (Fig. 3*B*, cue location at 180°). Therefore, due to its activity-dependence, DSI is the strongest in this group as well. This is depicted in the blue region of the spatiotemporal activity pattern in Figure 3*C* (note the inverted scale, with hotter colors representing less DSI activation). This creates a favorable bias for the network at the location of the sensory cue, thereby reducing spontaneous drift and stabilizing the neuronal representation of the remembered cue (Carter and Wang 2007).

To quantify this DSI-induced effect, we determined the variance of the remembered cue location (the VPV). We proceeded in a similar way as described above, and the results are remarkably similar. When DSI is controlled by a long time constant (5 s), there is an initial period of rise in drifts (Fig. 3*D*, red trace, first 2 s of delay), similar to a network without DSI. However, once the mechanism is fully activated (with a longer delay), the VPV does not grow any longer, reaching a plateau instead. For the shorter time constant (1 s), the variance increases almost monotonically (Fig. 3*D*, black trace).

Another notable feature in the particular sample trial of Figure 3*B*,*C* is the persistence of the inhibition suppression. Given the slow nature of its decay (τ_D = 5 s), DSI does not
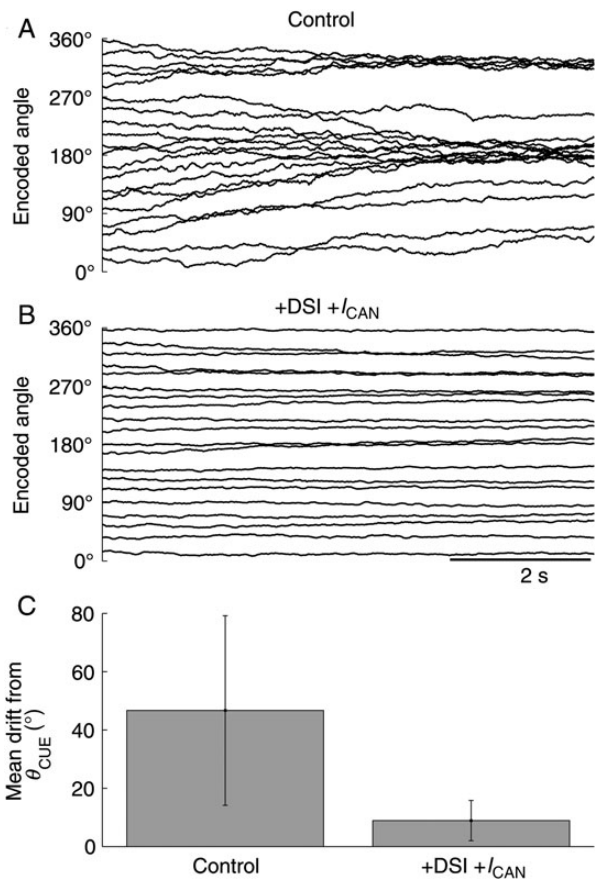
Figure 5. DSI and I_{CAN} stabilize the memory trace in the presence of heterogeneity across neurons in the network. Simulations were carried out where the cue was applied at 20 evenly spaced locations along the 360° space. The maintenance and retrieval of memory require that the remembered location at any given point in time should closely match that of the to-be-remembered cue. (*A*) The remembered cue locations of the simulations with the control parameter set systematically drift to a few privileged locations. (*B*) When DSI (4% maximum effect) and I_{CAN} (g_{CAN} = 1.5 nS) were incorporated in the network, the internal representation of the cue location becomes much better (the population vector is nearly stable across time). (*C*) The mean drift from the original cue location (at the end of a 9-s delay) is greatly reduced with DSI and I_{CAN} compared with the control (N = 500 trials). The time constant for DSI and I_{CAN} were, respectively, 2 and 0.5 s.

have sufficient time to fade away during a negative pulse lasting 0.5 s (compare with Fig. 1*C*). The remaining trace of disinhibition is strong enough to restart the memory bump at approximately the same angle, without a new cue presentation (Fig. 3*B*, right panel, red profile).

As shown in Figure 3*E* (open symbols), the duration of a step current required to reset the network increases dramatically with τ_D (0.5 s: $t_{SHUT,MIN}$ = 130 ms; 5 s: $t_{SHUT,MIN}$ = 3.75 s). On the other hand, the variance of the remembered cue location, the VPV, is larger in simulations with short τ_D and decreases for progressively longer τ_D, reaching a low plateau for τ_D larger than 1 s. Compared with the control (Fig. 1*E* with τ_S = 100 ms, VPV = 206 deg^2) with the same delay period duration of 5–6 s, a circuit endowed with DSI displays a smaller variance of drifts overall (VPV = 70.5 deg^2 with τ_D = 50 ms and 42.1 deg^2 with τ_D = 5 s; Fig. 3*E*, open circles). With larger delays (12–13 s), the simulations show higher variance due to the accumulation of drifts over a longer time (Fig. 3*E*, open squares). However, in agreement with the traces in Figure 3*D*, this relative increase of

VPV due to longer delays is mostly observed for shorter τ_D and is minimal for longer ones. Therefore, our analysis shows a tradeoff between the ease of shutdown and memory accuracy, which is the same with DSI as that observed with I_{CAN}.

I_{CAN} and DSI Enhance the Robustness of Working Memory

We next examined the network behavior when the model system is endowed with a combination of both DSI and I_{CAN}. To quantify the robustness of the working memory system, we determined the range of the parameter space where there is co-existence of a resting state (low-firing rates) with memory states (high rates). This multistability range corresponds to the regime where the system remains in the resting state in the absence of stimulation, but encodes a memory after a cue presentation, which is the desirable behavior of a working memory system. We determined the regime boundaries for a range of the recurrent connectivity strength between excitatory neurons (G_{EE}) in the form of bifurcation diagrams (Fig. 4). When DSI and I_{CAN} are not present, the multistability range (bounded by 2 dashed lines) is restricted to a narrow range around G_{EE} = 0.38 nS (Fig. 4*A*). If either DSI or I_{CAN} is incorporated (same as in previous simulations: 4% DSI or g_{CAN} = 1.5 nS), the lower boundary of the range is extended to smaller G_{EE} values (Fig. 4*B,C*). The maximum broadening effect occurs when both slow mechanisms are present (Fig. 4*D*). This is readily understood: with the help of DSI and I_{CAN}, less recurrent excitation is required to generate persistent activity.

A second noteworthy feature of Figure 4 is that the slow biophysical mechanisms increase the firing rate of memory states while that of the resting state remains roughly the same. This is because DSI and I_{CAN} are activity-dependent, therefore minimal in the low-firing spontaneous activity, but significant in the high-rate memory states. This leads to a larger separation between the resting and memory states. Consequently, a random fluctuation in spontaneous spiking activity will be less prone to give rise to a "false" memory, and the network function is more reliable.

To conclude, I_{CAN} and DSI are beneficiary to the system by making it less sensitive to variations of the network properties (such as G_{EE}), and less prone to spontaneous transitions by noise between the resting state and memory states. Both effects enhance the robustness of working memory behavior.

I_{CAN} and DSI Counteract Heterogeneity

A continuum of attractor states requires that the network is homogeneous and neurons are homogeneous, so that the system is translationally invariant (Ben-Yishai et al. 1995). Under this condition, if a localized pattern of activity is spatially displaced, it will lead to another identical pattern centered at the new location. However, any neural network shows a certain degree of variability across cells (Marder and Goaillard 2006). A homeostatic mechanism that equalizes the long-term firing rates of all cells to a predetermined level was shown to recover the accuracy of the memory trace (Renart et al. 2003). Alternatively, can DSI and I_{CAN} remedy the system's vulnerability to heterogeneity, by virtue of reinforcing a privileged location in the network in an activity-dependent manner? To investigate this question, we implemented a modest amount of heterogeneity, by assuming that the leak potential V_L varies from cell to cell according to a Gaussian distribution [mean V_L = −70 mV and standard deviation SD(V_L) = 1 mV].
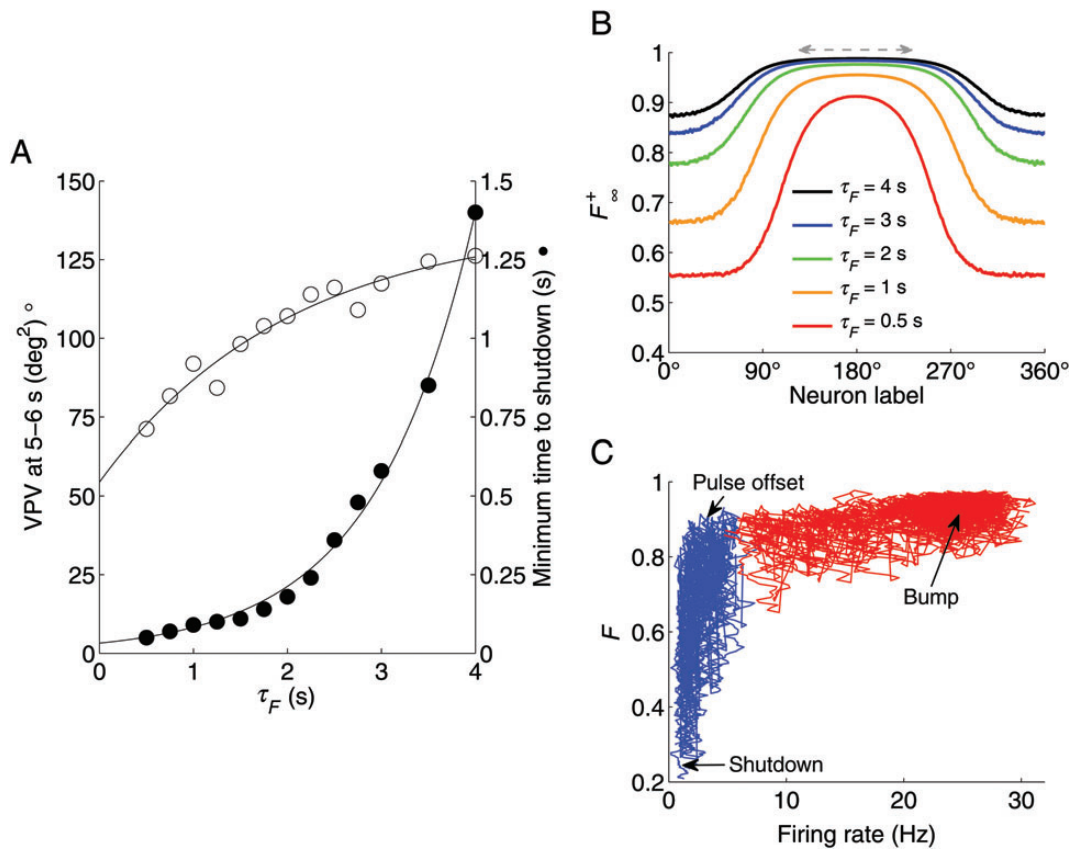
Figure 6. STF of recurrent excitatory synapses reduces random drifts. (*A*) $t_{\text{SHUT,MIN}}$ (filled circles) increases with τ_F (fitted with an exponential equation). Likewise, the variance of the remembered cue location (VPV) also increases with slower STF (exponential fit), but remains much smaller than that in the absence of STF (VPV = 206 deg^2 in Fig. 1*E*, τ_S = 100 ms; N = 500). (*B*) Steady-state profiles of F^+ (the facilitation variable, *F*, after a spike) for 5 different τ_F (7 s after delay start, N = 400). For longer time constants, the peak of the profile broadens (dashed gray double arrow), resulting in a region effectively without facilitation. This explains increased drifts with longer τ_F (*C*) Phase space plot of *F* and the population firing rate. Each trajectory corresponds to a trial and starts immediately at the shutdown pulse offset. The network either revert back to the mnemonic bump state (trials in red) or rest to the resting state (trials in blue), depending on the stochastic network dynamics. The *F* variable fluctuates from trial to trial and is significantly larger in red trajectories than blue ones (see Results). Note that, at the pulse offset, the population of excitatory cells was silent. However, due to the temporal sliding window (50 ms) used to calculate firing rates, the trajectories depicted start at >0 Hz.

Across a large number of trials, the input cues are presented at 20 angle locations equally distributed along the 360° of a circle. When both mechanisms are absent, the remembered cue locations display systematic drifts and, as previously reported (Renart et al. 2003), tend to converge to a few privileged locations (Fig. 5*A*, θ = 180 and 320°). These locations are determined by the heterogeneous distribution of the cellular excitability across the network, which disrupts the continuous family of bump attractors. The mean drift from the cue location is minimal in networks with DSI and I_{CAN} (8.9 ± 6.9°) and significantly different (2-sample *t*-test, $P = 5 \times 10^{-110}$) from that of the control network (46.7 ± 32.5°; Fig. 5*C*).

Intuitively, when DSI and I_{CAN} are included, the remembered cue locations show much smaller drifts (Fig. 5*B*). Both mechanisms add a second layer of activity-dependency besides NMDAR and operate on a slow timescale. Therefore, their presence "anchors" the original location of the bump attractor that encodes the sensory cue. These slow mechanisms are powerful enough to overcome the disrupting effect of heterogeneity.

Short-Term Facilitation Increases Memory Accuracy
Finally, we considered the effect of STF in our working memory model. STF shares similar features with I_{CAN} and DSI,

namely activity-dependence, positive feedback, and slow time course of activation (Zucker 1989; Fisher et al. 1997; Tsodyks and Markram 1997; Abbott and Regehr 2004). It is especially prevalent in excitatory synapses between pyramidal cells in the frontal cortex (Hempel et al. 2000; Wang et al. 2006).

The implementation of STF in the model reduced random drifts of the memory trace during the delay. Compared with the control network (Fig. 1*E*, τ_S = 100 ms, VPV = 206 deg^2), the variance of the remembered cue location was lower for any τ_F (Fig. 6*A*, open circles, VPV ranges 71–126 deg^2). However, contrary to DSI and I_{CAN}, the VPV increased rather than decreases with longer τ_F. This unexpected result is elucidated by the analysis of a profile of the peak value of the facilitation variable *F*, for a bump attractor. For each cell in the network, every time there is a spike, *F* is increased by a discrete jump and the resulting value F^+ is used to update the synaptic conductance. Between spikes, *F* decayed until the next spike takes place. Thus, neurons in the bump that had elevated firing rates also show higher F^+ (profile in Fig. 6*B*). For longer τ_F, the decay is very slow, resulting in more temporal summation and, eventually, in a saturation of F^+ (Fig. 6*B*, gray dashed double arrow). A wide steady-state F^+ profile effectively removed the facilitation effect in that spatial region and selective enhancement created by the activity-dependent-positive feedback. For
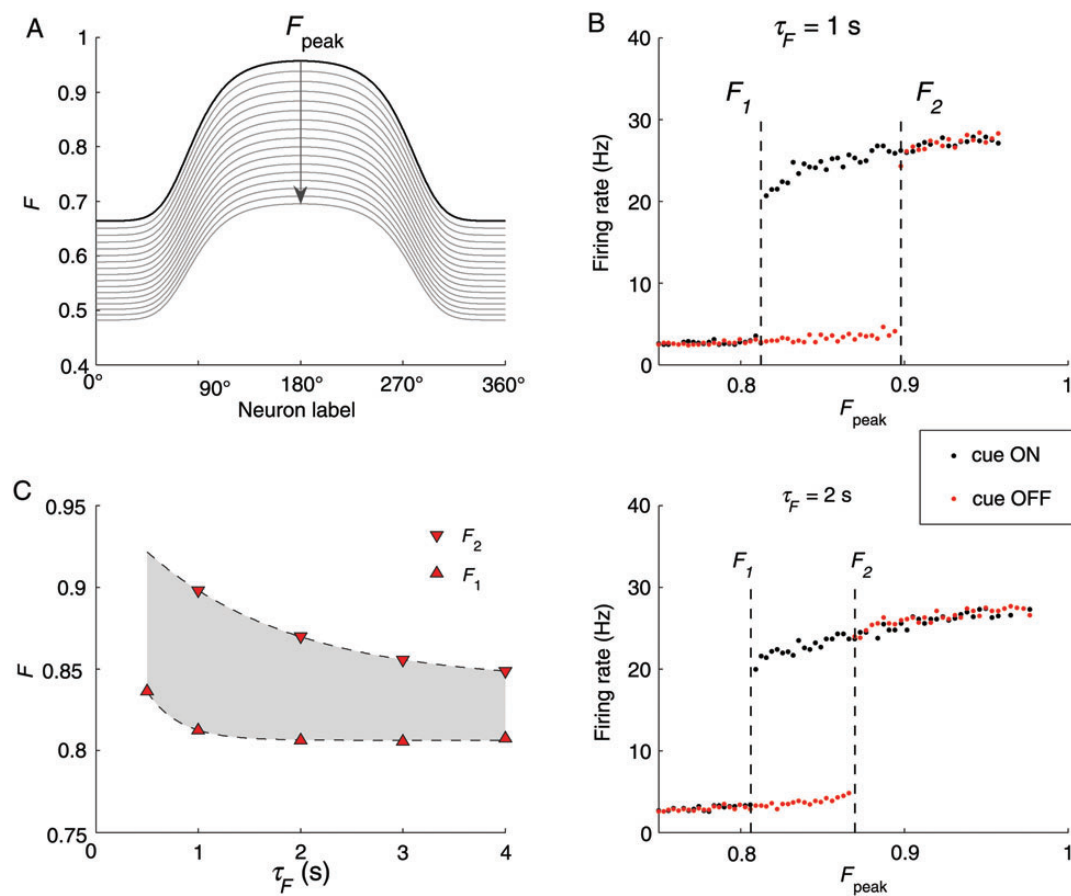
Figure 7. A simplified model with fixed F profile shows that the network is multistable within a range of STF values. (A) The black curve corresponds to the orange profile ($\tau_F = 1$ s) in Figure 6B, and the other curves were obtained by assuming an exponential decay in time of the black profile, during different temporal intervals (see Results). (B) Bifurcation diagram for $\tau_F = 1$ and 2 s (upper and lower panels, respectively). Simulations were run with (black dots) or without (red dots) cue presentation, and plotted is the maximum firing rate among all excitatory cells, at the end of the delay period. In these simulations, F did not change dynamically but was set as a parameter and given spatial profiles as those shown in (A). The peaks of the corresponding F profiles are shown in the abscissa. Below F_1, the network was always in the resting state. Above F_2, no cue was necessary to initiate a bump. (C) F_1 and F_2 as a function of $\tau_F = 0.5, 1, 2, 3, 4$ s (fit with single exponentials). The shaded area represents the presence of multistability.

this reason, augmenting the STF time constant increased memory drifts and, consequently, increased the variance of the remembered cue location (Fig. 6A, open circles).

This saturating feature was not observed with the other 2 slow mechanisms because of the following differences between the biological processes. The magnitudes of I_{CAN} and DSI vary quasi-continuously over time through their dependence on intracellular calcium, which accumulates and declines slowly. Furthermore, they influence the excitability of the cell at almost any point in time. Therefore, the spatial profile of the activity variable (m_{CAN} or D, respectively) can be fixed and remains not saturating, when the time constant is varied through a scaling factor (ϕ_{CAN} or ϕ_D). On the other hand, STF is not a continuous process, but acts only during synaptic events. This means that the value of the variable F is only used at times of spikes (F^+) and ignored when it decays away between spikes. For this reason, a scaling method is not appropriate, because it would only preserve the time-averaged steady state of F but not the steady state of F^+.

Similarly to what is observed for I_{CAN} and DSI, prolonged STF time constant makes it more difficult to reset the network (Fig. 6A, filled circles). When τ_F is 0.5 s, the required time to shutdown is just 50 ms. At the other end of the tested range, a

τ_F of 4 s requires a negative input pulse lasting for at least 1.4 s to erase a memory trace.

The minimum shutdown time is determined by the decay for F during the inhibitory input. At the end of the shutdown phase, the magnitude of F for neurons in the bump attractor reaches a level that depends on the pulse duration and τ_F. This level fluctuates from trial to trial, and has a large influence on whether the bump reappears or not in any given trial. In Figure 6C, are shown the system trajectories (state space) of F versus firing rate, for 40 trials with $\tau_F = 1$ s. This facilitation time constant corresponds to a minimum time to shutdown of 90 ms (Fig. 6A), which means that shorter pulses should not be able to shutdown the network. The red traces ($t_{SHUT,MIN} = 50$ ms, leading to return of the bump state) start at an average of $F = 0.89 \pm 0.02$, whereas the blue ones ($t_{SHUT,MIN} = 100$ ms, leading to shutdown) begin at $F = 0.84 \pm 0.04$—a significant difference (2-sample t-test, $P = 8.58 \times 10^{-6}$). Longer τ_F requires longer shutdown pulses in order for F to decay to a low enough level, so that the recurrent excitation is too weak to enable the bump to reemerge.

To shed further insights into how the degree of facilitation decay over the entire neuronal network determines the success of the memory shutdown, we ran new simulations in a
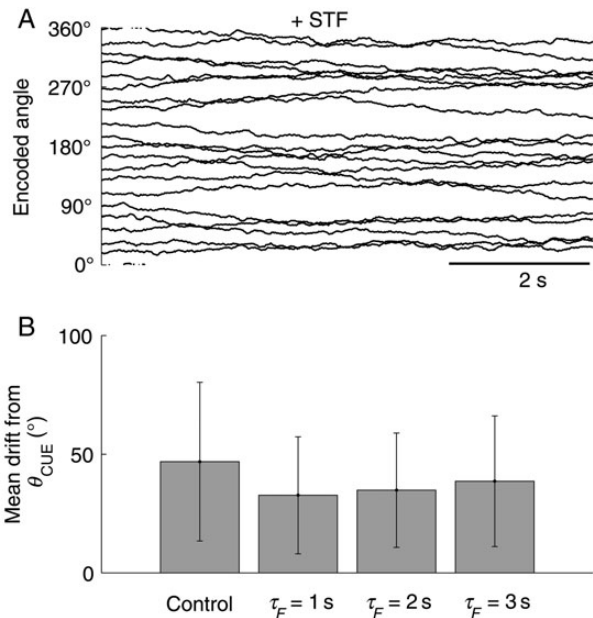
Figure 8. STF stabilizes the remembered cue locations in the presence of heterogeneity across neurons in the network. In stimulations, the cue location was applied at 20 evenly spaced locations along the 360° space. (*A*) The remembered cue locations with STF (τ_F = 1 s) show visibly less drifts than the control (Fig. 5*A*). (*B*) The mean heterogeneity-induced systematic drifts (at the end of a 9-s delay) for the network model without STF (control) or with STF operating at 3 different time constants (N = 400).

modified model where the facilitation factor F is no longer a variable but is treated as a parameter. To each simulation, we assigned a fixed F spatial profile. For realistic purposes, these are all mathematically derived (time decay) from the average facilitation profile at the end of delay of previous simulations, where F was a variable (Fig. 7*A*, see Materials and Methods). This replicates the decay of F during the shutdown phase, with different pulse durations. In these simulations, F_{peak} is a parameter that characterizes the profile of facilitation expected after a shutdown pulse with a given duration—lower F_{peaks} correspond to longer pulses. Roughly, the profiles that result from decays longer than $t_{SHUT,MIN}$ should not be able to sustain a bump without cue. The emergence of either a resting (low rates) or a memory state (high rates) as a function of F_{peak} was presented in bifurcation diagrams (Fig. 7*B*, τ_F = 1, 2 s). The thresholds F_1 and F_2 delimit the bistability regime (Fig. 7*B*,*C*; Fig. *C*, shaded area). With longer STF time constants, it is necessary to reach a profile with lower F_{peak} such that simulations without cue input remain in the resting state (below F_2). This behavior roughly corresponds to the increase in $t_{SHUT,MIN}$ for longer τ_F, in simulations where F changes dynamically.

As was shown in a recent study using a firing-rate model, systematic drifts of memory trace due to heterogeneity could be dramatically reduced by STF (Itskov et al. 2011). We checked the effect of STF in our spiking network model in the presence of cellular heterogeneity [$\langle V_L \rangle$ = −70 mV and SD(V_L) = 1 mV]. It is evident by visual inspection that, with STF, the memory storage of the sensory cue is more stable over time (Fig. 8*A*), compared with those under the control condition without STF (Fig. 5*A*). This impression is confirmed statistically (Fig. 8*B*). There is a significant (2-sample *t*-test) decrease in the mean drifts between the control network (without STF,

46.9 ± 33.4°) and each of the 3 scenarios with a different τ_F (1 s, 32.7 ± 24.7°, P = 2 × 10^{-11}; 2 s, 34.8 ± 24.1°, P = 8 × 10^{-9}; 3 s, 38.6 ± 27.5°, P = 1 × 10^{-4}).

In summary, like I_{CAN} and DSI, STF reduces noise-induced random drifts or heterogeneity-induced systematic drifts of memory traces, thereby rendering working memory function more robust. In contrast to the other 2 slow mechanisms, a longer time constant of STF leads to larger drifts of a memory trace, but drifts remain smaller than those in the control network without STF.

Slow Mechanisms Protect Memory Against Distractors

A cortical circuit assigned to store a particular stimulus in memory may receive, at any point in time, additional external signals with the potential to alter its network state and output. Depending on the nature of the new, distractor signal, the circuit may respond within a range of possible behaviors. It can erase the previous memory and encode the new one. The second stimulus can also modify quantitatively the established memory. Finally, the circuit may filter out the distractor completely. Considering the influence that slow mechanisms have on memory robustness and flexibility, they may also play a crucial role in this process.

In our network model, if we apply a particular distractor stimulus during the delay period, the new remembered cue location will shift toward its location and away from the cue stimulus angle. The deviation of the bump peak location is clearly visible for a network with the control parameter set (Fig. 9*A*, upper panel, θ_1 = 178°, θ_2 = 244°). The magnitude of this deviation ($\theta_2 - \theta_1$) depends on the angular difference between the cue stimulus (θ_S) and the distractor (θ_D). The distraction increases with $\theta_D - \theta_S$ before reaching a maximum. Beyond this point, the influence of the distractor decreases abruptly and the final location of the bump is much closer to the cue angle. Longer distractor durations result in significantly larger deviations of the final memory trace (Fig. 9*B*, 500 ms).

When DSI is incorporated in the network, the deviation of the bump induced by a distractor is visibly smaller (Fig. 9*A*, lower panel, θ_1 = 185°, θ_2 = 196°) than in the control network. This outcome is observed across the range $\theta_D - \theta_S$ and for all distractor durations (Fig. 9*C*). Consequently, the maximum distraction (Fig. 9*D*, upper panel) with DSI is smaller and grows slower with distractor duration (150 ms, 11.5 ± 3.7°; 500 ms, 29.0 ± 4.8°) than in control conditions (150 ms, 41.6 ± 12.3°; 500 ms, 135.0 ± 25.6°). The angle difference between distractor and cue that originated those maximum distractions [($\theta_D - \theta_S)_{max}$] corresponds to the distractor location that has maximum influence on the memory bump. This indicator was higher with control parameters than with DSI for all distractor durations. Remarkably, the presence of the slow mechanism resulted in a more stable ($\theta_D - \theta_S)_{max}$ (150 ms, 90°; 500 ms, 110°) than in the control network (150 ms, 110°; 500 ms, 155°). Similar results were obtained with I_{CAN} and STF (Supplementary Fig. 1).

Taken together, these results suggest that DSI decreases the influence of all distractors regardless of their location. Moreover, it reduces the range of distractor locations that significantly deviate the memory bump. Finally, the protection against farther distractors is almost independent of their duration.
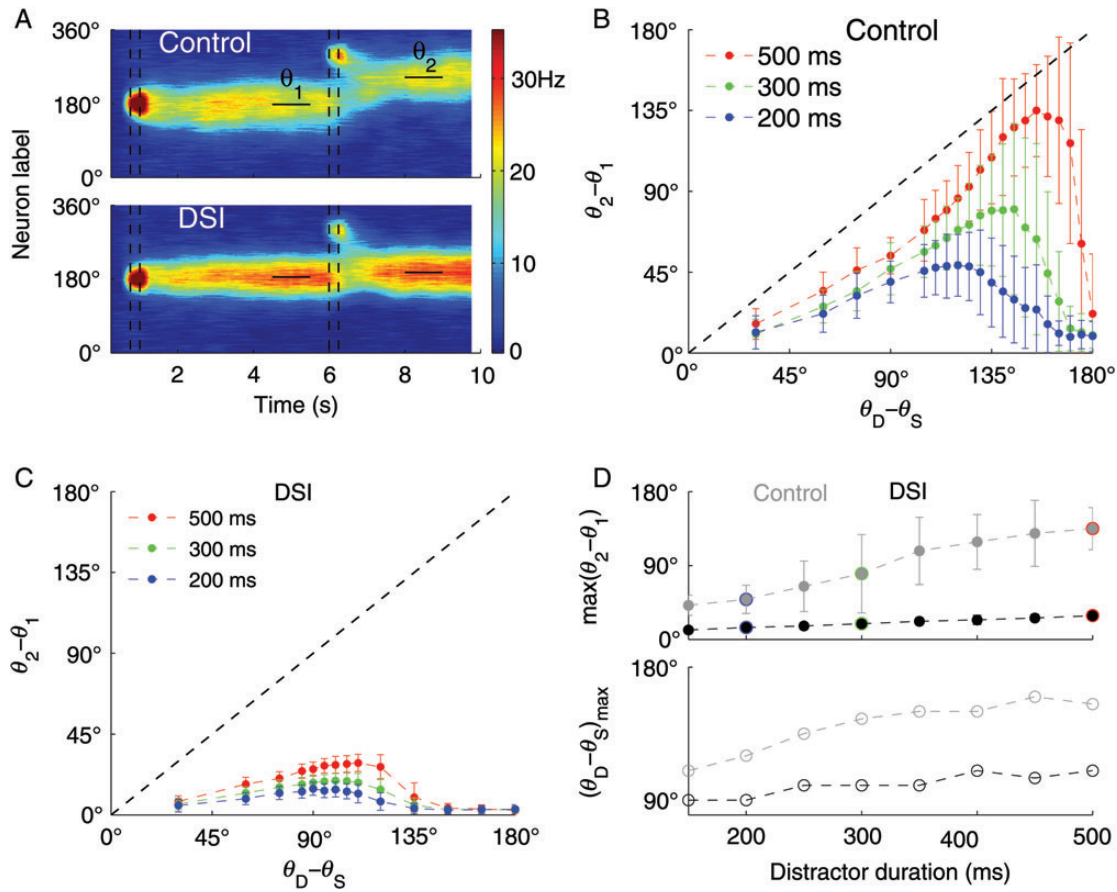
Figure 9. Slow mechanisms preserve cue representation and decrease the influence of long distractor stimuli. (*A*) Smoothed spatiotemporal activity pattern of the network's excitatory cells under control conditions (upper panel) or with DSI (lower panel), in the presence of a distractor. An initial cue stimulus (peak angle $\theta_S = 180°$, 750 ms–1 s, first pair of vertical dashed lines) drives the network to the memory state. The application of a distractor during the delay period (peak angle $\theta_D = 300°$, 100 pA, 6–6.25 s, second pair of dashed lines) pulls the location of the bump closer to it. In these 2 example trials, the deviation of the bump, measured as the difference between the remembered cue location after the distractor (θ_2, 8 – 9 s) and before (θ_1, 4.5–5.5 s), is larger in the control network than with DSI. (*B*) The average difference between θ_2 and θ_1 as a function of the difference in peak angles of distractor (θ_D) and cue stimulus (θ_S), for 3 distractor durations ($N = 150$). The deviation increases and approaches the perfect distraction (diagonal dashed line) before declining for more distant distractors. Longer durations produce generally larger deviations that have a maximum at larger distractor angles. (*C*) Same as in (*B*) but for network with DSI. The differences in remembered cue locations are visibly smaller than with the control network for all 3 distractor durations ($N = 150$). (*D*) Distraction indicators for sets of trials with different distractor durations, under control network (grey symbols) or with DSI (black symbols). Upper panel: the maximum distraction is small and increases almost linearly in a network with DSI. Under control conditions, this measure is larger throughout the whole range and has a more prominent increase. The edge-colored data points were taken from (*B*) and (*C*) with the same color scheme. Lower panel: similarly, the distraction angle ($\theta_D - \theta_S$) at which the maximum deviation of the bump is observed is wide and increases with duration in the control network, but is narrower and almost stable when DSI is present. This slow mechanism limits the effects of closer distractors and protects the memory against farther ones almost independently of their duration.

Discussion

Until now it is commonly recognized that a working memory circuit should not be conceptualized in terms of rapid switches between attractor states. Instead, reverberation underlying persistent activity must be slow, likely involving the NMDARs at recurrent excitatory synapses (Wang 1999; Wang et al. 2013). Slow network dynamics enables a single microcircuit mechanism to subserve working memory and decision-making. The latter requires accumulation of information over time by virtue of slow neural transients such as quasi-linear ramping activity (Wang 2002, 2008). It is noteworthy that persistent activity during a mnemonic delay period often displays slow temporal variations, as well as a rich heterogeneity across neurons (Batuev et al. 1979; Baeg et al. 2003; Miller et al. 2003; Goldman 2009; Machens et al. 2010; Barak et al. 2013; Stokes et al. 2013).

Is the slower the underlying mechanism, the better? In the present work, we investigated 3 biophysical mechanisms in a network model of spatial working memory. I_{CAN}, DSI, and STF are present in frontal neurons and are activity-dependent. They provide positive feedback to active excitatory cells and operate on a slow timescale. Our main finding was that slow timescale has a tradeoff effect. I_{CAN}, DSI, and STF render working memory representation more robust. However, their slow decay leaves a lingering memory trace even after the termination of persistent firing activity, which makes it difficult to reset the circuit by brief inputs, a fundamental requirement for normal function of a working memory system. These findings suggest that recurrent attractor dynamics are the "working horse" for mechanisms that sustain delay activity, and that very slow processes contribute to the accuracy of memory maintenance.

Random Drifts

Our study started with the premise that very slow processes are not necessary for the generation of persistent activity per se,

but may play a role in determining the accuracy and robustness of a working memory circuit's behavior. It has been previously reported that, in continuous attractor networks, a mnemonic activity pattern (bump attractor) exhibits random drifts that accumulate over time and deviate the stored spatial information away from the cue location (Camperi and Wang 1998; Compte et al. 2000; Carter and Wang 2007). This decreases the accuracy of the memory readout. We found that the incorporation of DSI, I_{CAN}, or STF helps to reduce random drifts of a memory trace. Moreover, we demonstrated that this stabilizing effect is dependent on the effective time constant of the mechanism considered. For I_{CAN} and DSI, the increase in τ is associated with a decrease in random drifts before reaching a plateau (~1–2 s). In contrast, with STF, although it also increases accuracy when compared with the control network, random drifts actually increase with τ_F. We showed that this happens because, for longer time constants, the facilitation variable F saturates around the bump, effectively removing facilitation in that part of the network.

Heterogeneity-Induced Drifts

A different type of drifts of memory trace arises from heterogeneity across neurons, which is detrimental to the realization of a continuous family of attractors. A homeostatic mechanism that scales the excitatory synapses was shown to recover the accuracy of the remembered cue location under those conditions (Renart et al. 2003). This activity-dependent mechanism scales the excitatory synaptic weights of each cell, so that the long-term average firing rate is similar for all and equal to a predetermined level. Recently, it was demonstrated using a firing-rate model that STF slows down the velocity of drifts in the presence of synaptic heterogeneity (Itskov et al. 2011). Building upon these insights and following our successful stabilization of random drifts, we tested the effect of the 3 slow mechanisms on our spiking model in the presence of cellular heterogeneity. A combination of I_{CAN} and DSI effectively counteracted the tendency of a bump to drift to privileged locations, during a relatively long delay period. Likewise, STF significantly reduced systematic drifts due to heterogeneity, strengthening the previous results with a biophysically realistic spiking network model.

Memory Flexibility

Whereas slow biophysical mechanisms increased the accuracy of memory representation, they have the opposite implications on the flexibility to switch between dynamical states. The precise mechanism used by the brain to erase a working memory trace remains poorly understood. We used a negative current of sufficiently strong amplitude to bring neurons below the firing threshold, which is the most efficiently way to switch off from a memory state. If we used a different input to only reduce the firing rate rather than silencing neurons, the duration of that input would need to be longer than $t_{SHUT,MIN}$. This analysis leads us to very similar conclusions regarding I_{CAN}, DSI, and STF. The minimum duration of the pulse required to shutdown the network dramatically increases with the effective time constant of the mechanism.

Our study approaches a relevant debate regarding the basis of neural circuit mechanism of working memory. At one extreme, a working memory system is viewed in terms of fast switches between multiple steady-state attractors, with

virtually no transient. At the other extreme, there are no multiple attractors but a single resting state. In this scenario, a transient cue simply perturbs the system to another location in the state space, and the network has a short-term memory merely in terms of very slow decay (returning to the resting state) after the stimulus offset. This typically requires really slow biophysical time constants, such as those provided by I_{CAN}, DSI, and STF. Our model is not yet at this extreme, but it should follow from our results that the scenario without multiple attractors would have the same problem of shutting down.

We also studied how a slow mechanism may help preserve the location of the memory bump against distractor stimuli. DSI limits the effects of closer distractors and protects the memory against farther ones almost independently of their duration. Whereas this effect is desirable to discard unwanted stimuli, it also uncovers the potential for intertrial persistency. When a new cue is presented at a different location, the trace of disinhibition from the previous trial will act as a distractor and pull the location of the new bump toward the old one. Taken together, these results establish a tradeoff between memory accuracy and flexibility.

An accuracy–flexibility tradeoff can also be obtained by the modulation of the strength of the slow mechanisms instead of their time constant. The strengthening of DSI or I_{CAN} increases the memory trace stability but, on the other hand, increases the ease to shutdown (Supplementary Fig. 2). However, in these simulations, the overall firing rates increase when the slow mechanism became stronger, which presents a confounding factor that hinders the comparison between different strengths.

Slow Mechanisms Modulate Dynamics of a Working Memory System

The bifurcation analysis of the network model with G_{EE} as a control parameter revealed that, in the presence of DSI and I_{CAN}, the model system shows a wider multistability range and a larger separation between the firing rate of persistent activity and resting states. A wider multistability range implies a higher degree of robustness because normal function is less sensitive to variations of network properties. A larger separation of firing rates implies that it is harder for spontaneous transitions between states to occur merely by noise. We also showed that the realistic range of the facilitation factor F contained a multistability range that shifts with the time constant τ_F. These raise the possibility that, in a working memory circuit such as PFC, some modulatory mechanisms could flexibly tune slow biophysical processes for optimal behavior.

Figure 10 offers a conceptual understanding common for all 3 slow mechanisms. This schematic depiction is partly deduced from the phase space plots in Figures 2D and 6C. It shows, in a state space, how activity of neurons engaged in working memory storage and the activation variable (X) of any of the 3 slow mechanisms interact with each other dynamically. Just before the shutdown pulse (Fig. 10A), the phase plane consists of a stable manifold, which separates the resting and memory attractors, and an unstable manifold. The intersection between the 2 lines creates an unstable saddle point. During the shutdown pulse presentation (Fig. 10B), only the resting state exists. The negative input pulse immediately inhibits all firing activity and X decays exponentially. At the pulse offset (Fig. 10C), the system regains its previous landscape with both
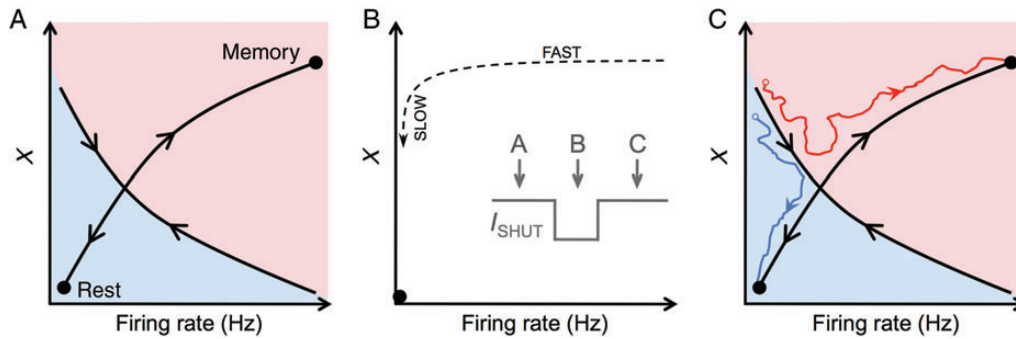
Figure 10. Schematic phase-plane diagram of our working memory model, during 3 stages of a shutdown process. This scheme applies to all 3 slow biophysical mechanisms considered in this paper, with *X* representing the activation variable of I_{CAN}, DSI, or STF. The inset in (*B*) displays the timing of the 3 stages according to the presentation of the negative shutdown input. (*A*) The state space displays a stable manifold (line with converging arrows) and an unstable manifold (line with diverging arrows), and their intersection creates a saddle point. There are 2 stable steady states (filled circle) representing a memory state and a rest state. At the end of delay, the system is in the memory state. (*B*) During the application of the negative pulse, there is only one steady state (filled circle), with a low-firing rate and low *X* magnitude. After the quick suppression of all firing activity ("FAST"), the system moves along the direction of the exponential decay of *X* ("SLOW") over the duration of the pulse. (*C*) The attractor landscape (*A*) is restored after the pulse offset. Depending on whether the state of the system at the offset of the shutdown input is on the left or the right side of the stable manifold, the system will revert back to the memory state (red trajectory) or reset to the resting state (blue trajectory, successful shutdown).

attractors. Afterwards, the network trajectory depends on the extent of *X* decay, which is an exponential function of the duration of the pulse. If *X* is below the stable manifold, the system will progress to the resting state, resulting in a successful shutdown (blue trace). Otherwise, if *X* decayed less and remains above the stable manifold at the pulse offset, the system will revert back to the mnemonic attractor state (red trace).

Accuracy–Flexibility Tradeoff

The circuits of the PFC that encode working memory, like all systems in the nervous system, have a rich variety of processes that modulate their performance. In this study, we considered a group of mechanisms that may be involved in the dynamical stabilization of the memory trace. The apparent conflict resulting from a tradeoff between accuracy and flexibility of the memory trace may turn out to be significant for neuronal modulation. According to environmental conditions and behavioral task demands, the network may be instructed to tilt the balance in favor of increased accuracy at the expense of flexibility. Under these circumstances, I_{CAN}, DSI, or STF may be strongly activated so that the memory is encoded as precisely as possible. On the other hand, when the task requires faster response to cue stimulation, the network may be tuned to decrease the activation of the slow mechanisms or shorten their time constants. This prevents the previous memory from interfering with the encoding of the new stimulus. Interestingly, we reveal the location of "sweet spots" for the models with I_{CAN} and DSI. In these time constant ranges (1–2 s), these slow mechanisms stabilize the memory to a great degree without significantly hampering the shutdown process. This observation raises the question of whether a working memory system in the brain can be tuned to that optimal configuration and what might be a neurobiological mechanism for achieving such optimality. Future experiments and theory are worth pursuing in this direction.

It has been proposed that an emphasis on robust online representation of information versus rapid switching could be adjusted by dopamine signaling, with D1 (respectively D2) receptors acting in favor of robustness (respectively flexibility; [Durstewitz and Seamans 2008](); [Rolls et al. 2008]()). Additionally, several other pathways modulate the 3 slow mechanisms. The

channels that mediate I_{CAN} are highly sensitive to muscarinic and metabotropic receptor activation ([Haj-Dahmane and Andrade 1998](); [Sidiropoulou et al. 2009]()), DSI is obviously dependent on endocannabinoids ([Ohno-Shosaku et al. 2001](); [Wilson and Nicoll 2001](); [Wilson et al. 2001]()) and STF is controlled by synaptic vesicles release ([Hempel et al. 2000]()). Our results suggest that slow processes, including those studied here, are potentially effective targets of action by dopamine or other neuromodulators, which can optimally adjust the tradeoff between robustness of memory storage and cognitive flexibility. This prediction may be tested through experiments to understand precisely how modulation occurs and to determine under which circumstances each slow mechanism is predominant in the encoding of working memory in the PFC.

At the present time, there still exists a large gap between neural circuits and behavior ([Carandini 2012]()); this wide gap must be bridged in order to achieve our goal of understanding the brain mechanisms of cognitive functions and their impairments associated with mental disorders. The present work illustrates how biophysically based computational modeling, in interplay with experimentation, can help make progress in this direction, through elucidation of how specific cellular and synaptic processes shape network activity patterns (persistent activity) and contribute to a key functional requirement (accuracy–flexibility tradeoff) in a cognitive process.

Supplementary Material

Supplementary material can be found at: [http://www.cercor.oxford journals.org/.]()

Notes

Conflict of Interest: None declared.

References

Abbott LF, Regehr WG. 2004. Synaptic computation. Nature. 431:796–803.

Amit DJ. 1995. The Hebbian paradigm reintegrated: Local reverberations as internal representations. Behav Brain Sci. 18:617–626.

Amit DJ, Brunel N. 1997. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb Cortex. 7:237–252.

Baeg E, Kim Y, Huh K, Mook-Jung I, Kim H, Jung M. 2003. Dynamics of population code for working memory in the prefrontal cortex. Neuron. 40:177–188.

Barak O, Sussillo D, Romo R, Tsodyks M, Abbott L. 2013. From fixed points to chaos: three models of delayed discrimination. Prog Neurobiol. 103:214–222.

Batuev AS, Pirogov AA, Orlov AA. 1979. Unit activity of the prefrontal cortex during delayed alternation performance in monkey. Acta Physiol Acad Sci Hung. 53:345–353.

Ben-Yishai R, Bar-Or RL, Sompolinsky H. 1995. Theory of orientation tuning in visual cortex. Proc Natl Acad Sci USA. 92:3844–3848.

Brunel N, Wang XJ. 2001. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. J Comput Neurosci. 11:63–85.

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.

Carandini M. 2012. From circuits to behavior: a bridge too far? Nat Neurosci. 15:507–509.

Carter E, Wang XJ. 2007. Cannabinoid-mediated disinhibition and working memory: dynamical interplay of multiple feedback mechanisms in a continuous attractor model of prefrontal cortex. Cereb Cortex. 17:i16–i26.

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.

Durstewitz D, Seamans JK. 2008. The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-O-methyltransferase genotypes and schizophrenia. Biol Psychiatry. 64:739–749.

Durstewitz D, Seamans JK, Sejnowski TJ. 2000. Neurocomputational models of working memory. Nat Neurosci. 3:1184–1191.

Egorov AV, Hamam BN, Fransen E, Hasselmo ME, Alonso AA. 2002. Graded persistent activity in entorhinal cortex neurons. Nature. 420:173–178.

Faber E, Sedlak P, Vidovic M, Sah P. 2006. Synaptic activation of transient receptor potential channels by metabotropic glutamate receptors in the lateral amygdala. Neuroscience. 137:781–794.

Fisher SA, Fischer TM, Carew TJ. 1997. Multiple overlapping processes underlying short-term synaptic enhancement. Trends Neurosci. 20:170–177.

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.

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.

Fuster JM, Alexander GE. 1971. Neuron activity related to short-term memory. Science. 173:652–654.

Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci. 2:1527–1537.

Gnadt J, Andersen R. 1988. Memory related motor planning activity in posterior parietal cortex of macaque. Exp Brain Res. 70:216–220.

Goldman MS. 2009. Memory without feedback in a neural network. Neuron. 61:621–634.

Goldman-Rakic P. 1995. Cellular basis of working memory. Neuron. 14:477–485.

Goodman DFM, Brette R. 2009. The brian simulator. Front Neurosci. 3:192–197.

Gross SA, Guzmán GA, Wissenbach U, Philipp SE, Zhu MX, Bruns D, Cavalié A. 2009. TRPC5 is a Ca^{2+}-activated channel functionally coupled to Ca^{2+}-selective ion channels. J Biol Chem. 284:34423–34432.

Gutkin BS, Laing CR, Colby CL, Chow CC, Ermentrout GB. 2001. Turning on and off with excitation: the role of spike-timing asynchrony and synchrony in sustained neural activity. J Comput Neurosci. 11:121–134.

Haj-Dahmane S, Andrade R. 1998. Ionic mechanism of the slow afterdepolarization induced by muscarinic receptor activation in rat prefrontal cortex. J Neurophysiol. 80:1197–1210.

Hansel D, Mato G. 2013. Short-term plasticity explains irregular persistent activity in working memory tasks. J Neurosci. 33:133–149.

Hempel CM, Hartman KH, Wang XJ, Turrigiano GG, Nelson SB. 2000. Multiple forms of short-term plasticity at excitatory synapses in rat medial prefrontal cortex. J Neurophysiol. 83:3031–3041.

Itskov V, Hansel D, Tsodyks M. 2011. Short-term facilitation may stabilize parametric working memory trace. Front Comput Neurosci. 5:1–19.

Jahr CE, Stevens CF. 1990. Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. J Neurosci. 10:3178–3182.

Kalmbach BE, Chitwood RA, Dembrow NC, Johnston D. 2013. Dendritic generation of mGluR-mediated slow afterdepolarization in layer 5 neurons of prefrontal cortex. J Neurosci. 33:13518–13532.

Katona I, Sperlágh B, Sk A, Käfalvi A, Vizi ES, Mackie K, Freund TF. 1999. Presynaptically located CB1 cannabinoid receptors regulate GABA release from axon terminals of specific hippocampal interneurons. J Neurosci. 19:4544–4558.

Kulkarni M, Zhang K, Kirkwood A. 2011. Single-cell persistent activity in anterodorsal thalamus. Neurosci Lett. 498:179–184.

Laing CR, Chow CC. 2001. Stationary bumps in networks of spiking neurons. Neural Comput. 13:1473–1494.

Lim S, Goldman M. 2013. Balanced cortical microcircuitry for maintaining information in working memory. Nat Nerosci. 16:1306–1314.

Liu YH, Wang XJ. 2001. Spike-frequency adaptation of a generalized leaky integrate-and-fire model neuron. J Comput Neurosci. 10:25–45.

Llano I, Leresche N, Marty A. 1991. Calcium entry increases the sensitivity of cerebellar Purkinje cells to applied GABA and decreases inhibitory synaptic currents. Neuron. 6:565–574.

Machens CK, Romo R, Brody CD. 2010. Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. J Neurosci. 30:350–360.

Major G, Tank D. 2004. Persistent neural activity: prevalence and mechanisms. Curr Opin Neurobiol. 14:675–684.

Marder E, Goaillard JM. 2006. Variability, compensation and homeostasis in neuron and network function. Nat Rev Neurosci. 7:563–574.

Marsicano G, Lutz B. 1999. Expression of the cannabinoid receptor CB1 in distinct neuronal subpopulations in the adult mouse forebrain. Eur J Neurosci. 11:4213–4225.

Matveev V, Wang XJ. 2000. Differential short-term synaptic plasticity and transmission of complex spike trains: to depress or to facilitate? Cereb Cortex. 10:1143–1153.

Miller EK, Erickson CA, Desimone R. 1996. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J Neurosci. 16:5154–5167.

Miller P, Brody CD, Romo R, Wang XJ. 2003. A recurrent network model of somatosensory parametric working memory in the prefrontal cortex. Cereb Cortex. 13:1208–1218.

Mongillo G, Barak O, Tsodyks M. 2008. Synaptic theory of working memory. Science. 319:1543–1546.

Murray JD, Anticevic A, Gancsos M, Ichinose M, Corlett PR, Krystal JH, Wang XJ. 2014. Linking microcircuit dysfunction to cognitive impairment: effects of disinhibition associated with schizophrenia in a cortical working memory model. Cereb Cortex. 24:859–872.

Ohno-Shosaku T, Maejima T, Kano M. 2001. Endogenous cannabinoids mediate retrograde signals from depolarized postsynaptic neurons to presynaptic terminals. Neuron. 29:729–738.

Partridge LD, Valenzuela CF. 1999. Ca^{2+} store-dependent potentiation of Ca^{2+}-activated nonselective cation channels in rat hippocampal neurones in vitro. J Physiol. 521:617–627.

Pitler TA, Alger BE. 1992. Postsynaptic spike firing reduces synaptic $GABA_A$ responses in hippocampal pyramidal cells. J Neurosci. 12:4122–4132.

Renart A, Song P, Wang XJ. 2003. Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. Neuron. 38:473–485.

Rolls ET, Loh M, Deco G, Winterer G. 2008. Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. Nat Rev Neurosci. 9:696–709.

Romo R, Brody CD, Hernández A, Lemus L. 1999. Neuronal correlates of parametric working memory in the prefrontal cortex. Nature. 399:470–473.

Sidiropoulou K, Lu FM, Fowler MA, Xiao R, Phillips C, Ozkan ED, Zhu MX, White FJ, Cooper DC. 2009. Dopamine modulates an mGluR5-mediated depolarization underlying prefrontal persistent activity. Nat Neurosci. 12:190–199.

Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J. 2013. Dynamic coding for cognitive control in prefrontal cortex. Neuron. 78:364–375.

Strübing C, Krapivinsky G, Krapivinsky L, Clapham DE. 2001. TRPC1 and TRPC5 form a novel cation channel in mammalian brain. Neuron. 29:645–655.

Szatmary B, Izhikevich EM. 2010. Spike-timing theory of working memory. PLoS Comput Biol. 6:e1000879.

Tegnér J, Compte A, Wang XJ. 2002. The dynamical stability of reverberatory neural circuits. Biol Cybern. 87:471–481.

Trettel J, Levine ES. 2003. Endocannabinoids mediate rapid retrograde signaling at interneuron pyramidal neuron synapses of the neocortex. J Neurophysiol. 89:2334–2338.

Tsodyks M, Sejnowski T. 1995. Associative memory and hippocampal place cells. Int J Neural Syst. 6:81–86.

Tsodyks MV, Markram H. 1997. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc Natl Acad Sci USA. 94:719–723.

Tuckwell H. 1988. Introduction to theoretical neurobiology. Cambridge (UK): Cambridge University Press.

Wang M, Yang Y, Wang CJ, Gamo NJ, Jin LE, Mazer JA, Morrison JH, Wang XJ, Arnsten AFT. 2013. NMDA receptors subserve persistent neuronal firing during working memory in dorsolateral prefrontal cortex. Neuron. 77:736–749.

Wang XJ. 2008. Decision making in recurrent neuronal circuits. Neuron. 60:215–234.

Wang XJ. 2002. Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 36:955–968.

Wang XJ. 1999. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J Neurosci. 19:9587–9603.

Wang XJ. 2001. Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci. 24:455–463.

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.

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.

Wilson RI, Kunos G, Nicoll RA. 2001. Presynaptic specificity of endocannabinoid signaling in the hippocampus. Neuron. 31:453–462.

Wilson RI, Nicoll RA. 2001. Endogenous cannabinoids mediate retrograde signalling at hippocampal synapses. Nature. 410:588–592.

Yoshida M, Hasselmo ME. 2009. Persistent firing supported by an intrinsic cellular mechanism in a component of the head direction system. J Neurosci. 29:4945–4952.

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

Zucker RS. 1989. Short-term synaptic plasticity. Annu Rev Neurosci. 12:13–31.