**Persistent Storage Capability Impairs Decision Making in a Biophysical Network Model**

Dominic Standage[a,b] and Martin Paré[a,b,c]

Canadian Institutes of Health Research Group in Sensory-Motor Integration[a], Department of Physiology[b], Department of Psychology[c], Queen's University, Canada

**Neural Networks (2011) 24:1062-1073**

**Abstract**

Two long-standing questions in neuroscience concern the mechanisms underlying our abilities to make decisions and to store goal-relevant information in memory for seconds at a time. Recent experimental and theoretical advances suggest that NMDA receptors at intrinsic cortical synapses play an important role in both these functions. The long NMDA time constant is suggested to support persistent mnemonic activity by maintaining excitatory drive after removal of a stimulus and to enable the slow integration of afferent information in the service of decisions. These findings have lead to the hypothesis that the local-circuit mechanisms underlying decisions must also furnish persistent storage of information. We use a local circuit cortical model of spiking neurons to test this hypothesis, controlling intrinsic drive by scaling NMDA conductance strength. Our simulations provide further evidence that persistent storage and decision making are supported by common mechanisms, but under biophysically realistic parameters, our model demonstrates that the processing requirements of persistent storage and decision making may be incompatible at the local circuit level. Parameters supporting persistent storage lead to strong dynamics that are at odds with slow integration, whereas weaker dynamics furnish the speed-accuracy trade-off common to psychometric data and decision theory.

**Introduction**

The length and variability of the time needed to discriminate visual stimuli and the susceptibility of this behaviour to errors indicate that decisions intervene between sensory and motor processing (see Schall, 2001 for review). The eye movement system has been invaluable as a model of decision making and experiments on non-human primates show that decisions can be decoded from neural activity in several cortical regions, including the lateral intraparietal area (LIP) of posterior parietal cortex (PPC) (Roitman and Shadlen, 2002; Thomas and Paré, 2007) and the frontal eye fields (FEF) (Schall and Hanes, 1993) and dorsolateral region of pre-frontal cortex (PFC) (Hasegawa and Matsumoto, 2000).

Abstract mathematical models have long provided phenomenological explanations of decision-making. Sequential sampling models assume that decision-making involves an integration process, where evidence is integrated until a threshold is reached (see Smith and Ratcliff, 2004). Because neural processing is noisy and evidence may be incomplete or ambiguous, integration is slower than the sampling rate, so decisions are based on an average of the evidence and not on momentary fluctuations in processing (see Bogacz, 2007). Several models have addressed the neural mechanisms underlying such a process (Usher and McClelland, 2001; Wang, 2002; Wong and Wang, 2006). The underlying premise of these models is that a discrete population of pyramidal neurons is selective for each decision option, and that competition between these populations is provided by a common pool of inhibitory

interneurons. Activity in each stimulus-selective population therefore comes at the expense of the other(s), providing a natural means of selection that scales with the number of decision options. Under constraints with biophysical correlates, mutual inhibition instantiates a calculation of the difference between the evidence favouring each option in two-choice tasks, a process known to optimise speed and accuracy with respect to one another with independent sequential samples (see Bogacz et al., 2006). Consistent with cortical processing (Douglas and Martin, 2004; 2007), intrinsic (recurrent) activity is crucial to these models, where the time constant of integration depends on a balance between the passive leakage of information and the amplification of information by recurrent activity (Usher and McClelland, 2001).

Biophysically-based models predict that NMDA receptors (NMDAR) at intrinsic synapses onto pyramidal neurons provide an important mechanism underlying the integration of evidence, where their long time constant enables the slow buildup of evidence (Wang, 2002; Wong and Wang, 2006). It is widely believed that intrinsic synapses also provide a mechanism for persistent mnemonic activity following the extinction of a stimulus, though this mechanism is just one of a number of mechanisms hypothesized to support persistent mnemonic activity (see the Discussion). Such activity is extensively correlated with working memory, the active retention of information for use in cognitive tasks (Goldman-Rakic, 1995; Wang, 2001). In this regard, NMDARs are hypothesized to provide an excitatory plateau (Fransén and Lansner, 1995; Lisman et al., 1998) while limiting network oscillations (Wang, 1999; Durstewitz et al., 2000), a hypothesis supported by observations that injection of NMDA blockers in PFC impairs working memory (Dudkin et al., 1997; Aura and Riekkinen, 1999). Because persistent mnemonic activity has been recorded in cortices correlated with perceptual decisions, including PFC (Fuster, 1973; Funahashi et al., 1989), FEF (Bruce and Goldberg, 1985) and PPC (Gnadt and Andersen, 1988), it has been proposed that intrinsic excitation strong enough to support persistent mnemonic activity is a property of decision circuits (Wang, 2002; Wong and Wang, 2006; Wang, 2008), similar in principle to suggestions that persistent storage (PS) capability may be required for coordinate transformations in PPC (Salinas and Sejnowski, 2001).

To address the hypothesis that decision making relies on local circuit PS capability (Wang, 2002; Wong and Wang, 2006; Wang, 2008), we model a decision-correlated circuit in LIP with a spiking implementation (Compte et al., 2000; Gutkin et al., 2001; Ma et al., 2006; Ardid et al., 2007; Furman and Wang, 2008) of a local circuit model widely used in population and firing rate simulations of cortical circuits (Wilson and Cowan, 1973; Pouget et al., 2000; Douglas and Martin, 2007), including visuospatial maps in PFC (Camperi and Wang, 1998), PPC (Standage et al., 2005) and frontoparietal cortex (Cisek, 2006). A spiking implementation provides synaptic resolution, enabling the manipulation of intrinsic NMDARs. Unlike earlier models with discrete stimulus-selective neural populations (Usher and McClelland, 2001; Wang, 2002), the model assumes a columnar organization where the strength of intercolumnar pyramidal interactions decreases with axial distance (see White, 1989; Abeles, 1991; Goldman-Rakic, 1995). Combined with unstructured or more broadly tuned synapses onto inhibitory interneurons, this synaptic profile creates centre-surround activity in which pyramidal neurons support each other locally via intrinsic projections and inhibit each other distally via interneurons. This family of networks is often used to model persistent mnemonic activity in visuospatial working memory tasks (Camperi and Wang, 1998; Trappenberg and Standage, 2005), where intrinsic excitation must be sufficiently strong for a selective population to drive itself over a memory interval, necessitating strong inhibition to limit the spread of excitation. These constraints lead to strong intrinsic dynamics that naturally cater to choice selection in decision making tasks, but are potentially at odds with the slow, simultaneous buildup of activity seen in decision-related cortices in multiple-choice tasks (see Schall, 2001).

In simulated visuospatial tasks, we control the network's intrinsic drive by scaling NMDA conductance at intrinsic synapses onto pyramidal neurons. In a simulated visuospatial working memory task, we determine values of this parameter that support (and do not support) PS. In a simulated two-choice visual search task, we measure the decision making abilities of the network for a range of values of this parameter. Our decision making task has no memory component, so there is no *a priori* requirement of local circuit PS capability for task completion. Under parameters consistent with biophysical data, we find that parameters supporting PS lead to intrinsic dynamics too strong for slow integration of evidence, amplifying momentary fluctuations and leading to hasty, inaccurate decisions. The model is a much more accurate decision maker under parameters that do not support PS. Indeed, the best decision-making network is far from the PS regime (Figure 2). In this case, the network enacts a speed-accuracy trade-off with increasing task difficulty (eg. Palmer et al., 2005), simulated reaction times and their distributions are consistent with those of psychophysical experiments, and simulated neural data are consistent with neural recordings in LIP during visual search tasks. This finding is different from that of earlier studies, but the mechanisms underlying it are much the same. Intrinsic processing fosters a balance between leakage and amplification of accumulated evidence (Usher and McClelland, 2001) that corresponds to a given NMDAR conductance strength (Wang, 2008). Under our parameters, the balance that best supports the task is outside the PS regime.

Our results fit with a distributed framework in which no single microcircuit is responsible for all aspects of a decision task, but where different functions (e.g. integration of evidence and choice selection) are mediated by different circuits (Beck et al., 2008). Our results further speak to the functions of decision-related cortical regions in distributed circuitry. For example, the prediction that decision circuits in LIP are characterized by weak intrinsic dynamics is consistent with reports that LIP represents the relative importance of items in the visual field (Serences and Yantis, 2006; Goldberg and Bisley, 2006), a function for which categorical dynamics are not well-suited. The difference between our findings and those of earlier studies is explained by consideration of the network's time constant of integration, optimization of which is parameter-dependent. We thus do not claim that PS capability cannot be a property of local circuits mediating decision processes, but under biophysically realistic parameters, we demonstrate the potential incompatibility of persistent mnemonic activity and decision making in the same local circuit at the same time.

**Materials and Methods**

We simulated a decision circuit in LIP with a fully connected recurrent network of leaky integrate-and-fire neurons (Tuckwell, 1988) with 1000 pyramidal neurons and 250 interneurons, depicted in Figure 1A. Intrinsic (recurrent) activity from pyramidal cells was mediated by AMPA and NMDA conductances and from interneurons by GABA conductances (Figure 1B). The strength or *weight* of pyramidal-to-pyramidal synapses was thus scaled by a decreasing function of spatial location (see White, 1989; Abeles, 1991; Goldman-Rakic, 1995), added to a baseline weight (Compte et al., 2000; Tegnér et al., 2002; Ardid and Wang, 2007) (Figure 1C). Combined with unstructured synapses between pyramidal cells and inhibitory interneurons, this synaptic profile creates a centre-surround network where pyramidal neurons support each other locally via intrinsic projections and inhibit each other distally via interneurons. In simulated visuospatial working memory and visual discrimination tasks, stimuli were simulated by Poisson spike trains where spike rates were drawn from a normal distribution and the mean corresponded to the centre of a Gaussian receptive field, depicted in Figure 1A for the discrimination task. Spike response adaptation among upstream, visually responsive neurons was mimicked by a decaying function of input rate (Trappenberg et al., 2001; Wong et al., 2007) with a 40ms (Thomas and Paré, 2007) response delay (Figure 1D). These selective inputs were superimposed on non-selective Poisson input spikes that lead to

background spike rates of 1 to 2 Hz among pyramidal neurons and 7 to 8 Hz among interneurons (Destexhe and Paré, 1999), depending on the strength of intrinsic NMDARs.

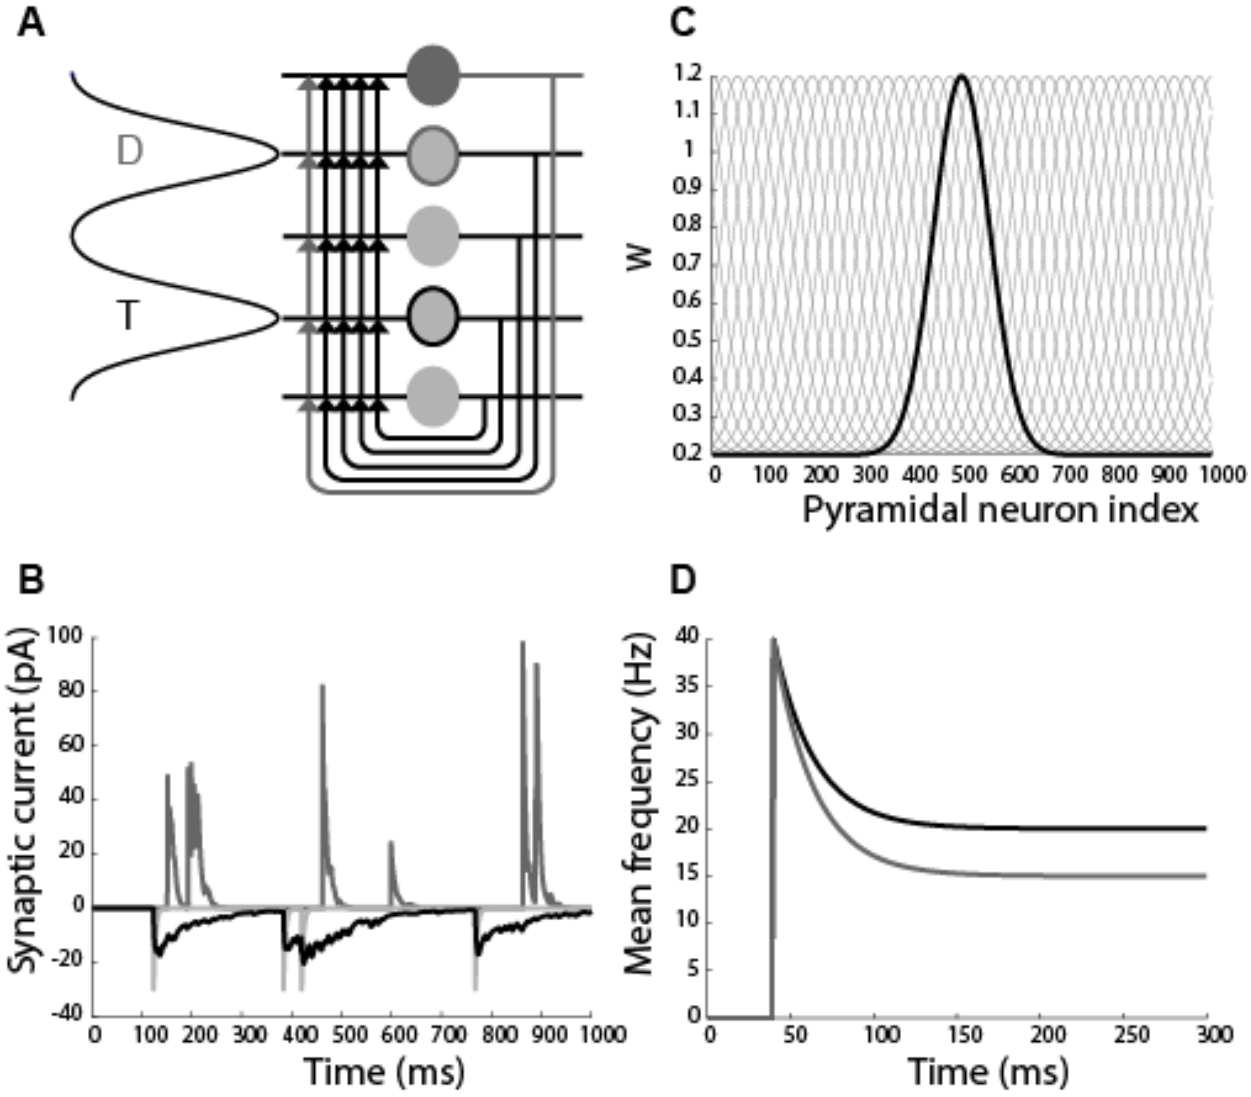


**Figure 1.** (**A**) Fully recurrent network of pyramidal neurons (light grey with black projections) and interneurons (dark grey neuron and projections). The 4/1 ratio of pyramidal neurons to interneurons preserves population sizes in the spiking model (1000/250). Target (T) and distractor (D) neurons are given Gaussian RFs and are outlined in black and dark grey respectively. (**B**) Intrinsic AMPA (light grey) and NMDA (black) currents onto a pyramidal neuron are consistent with *in vitro* recordings from cortical slices (e.g. Wang et al., 2008). Currents are shown at the peak of the connectivity structure in *C* and are thus mediated by the strongest intrinsic synapses in the model. GABA currents onto the same neuron are shown in dark grey and are also consistent with *in vitro* cortical data (e.g. Galarreta and Hestrin 1998). (**C**) Structured pyramidal-to-pyramidal weights *W*. Black curve shows weights to all pyramidal neurons from neuron 500. Grey curves show periodic shift-invariance of *W*. (**D**) Mean input frequency at the centre of the target (black) and distractor (grey) RFs during the stimulus interval. Curves are shown for the easiest level of target-distractor similarity ($\gamma_{NMDA} = 0.75$). Only the first 300ms is shown.

Each model neuron is described by

$C_m \, dV/dt = -g_L (V - E_L) - I$,

where $C_m$ is the membrane capacitance of the neuron, $g_L$ is the leakage conductance, $V$ is the membrane potential, $E_L$ is the equilibrium potential, and $I$ is the total input current. When $V$ reaches a firing threshold $\vartheta_V$, it is reset to $V_{res}$, after which it is unresponsive to its input for an absolute refractory period $\tau_{ref}$. For pyramidal neurons, $C_m$ = 0.5nF, $g_L$ = 25nS, $E_L$ = -70mV, $\vartheta_V$ = -50mV, $V_{res}$ = -60mV and $\tau_{ref}$ = 2ms. For interneurons, $C_m$ = 0.2nF, $g_L$ = 20nS, $E_L$ = -70mV, $\vartheta_V$ = -50mV, $V_{res}$ = -60mV and $\tau_{ref}$ = 1ms (Wang 1999; Compte et al., 2000).

The total input current $I$ is given by

$I = I_{AMPA}^{ext} + I_{AMPA} + I_{NMDA} + I_{GABA}$.

For each neuron, $I_{AMPA}^{ext}$ is AMPAR-mediated afferent current, and $I_{AMPA}$, $I_{NMDA}$ and $I_{GABA}$ are the summed AMPAR, NMDAR, and GABAR currents from recurrent synapses. These currents are each defined by

$$I_{syn} = G \cdot g(V - V_{syn}) \cdot \eta \cdot W \cdot \kappa, \quad (1)$$

where subscript *syn* refers to $\{AMPA^{ext}, AMPA, NMDA, GABA\}$. For each synapse contributing to the current, $G$ is the conductance strength, $g$ is the receptor activation, $V_{syn}$ is the synaptic reversal potential, $\eta$ captures the voltage-dependence of NMDARs (set to 1 for AMPARs and GABARs and described below for NMDARs) and $\kappa$ is a scale factor (set to 1 for all intrinsic synapses and described below for extrinsic activity). The constant $W$ distinguishes between structured and unstructured network connections (see below). For AMPA and NMDA synapses, $V_{syn}$ = 0mV. For GABA synapses, $V_{syn}$ = -70mV. For AMPARs and GABARs, the receptor activation $g$ follows a step-and-decay formula $dg/dt = -g/\tau_g + \delta(t - t_f)$, where $\delta$ is the Dirac delta function and $t_f$ is the time of firing of a pre-synaptic neuron. The decay constant $\tau_g$ is given values $\tau_{AMPA}$ = 4ms (Gonzalez-Burgos et al., 2007) and $\tau_{GABA}$ = 10ms (Salin and Prince, 1996) respectively for these synapses. For NMDA currents, $g$ has a slower rise and decay and is described by

$dg_{NMDA}/dt = -g_{NMDA}/\tau_{NMDA} + \alpha_{NMDA} \cdot x_{NMDA}(1 - g_{NMDA})$,

where $\tau_{NMDA}$ = 100ms (Wang et al., 2008), $\alpha_{NMDA}$ = 0.5kHz controls receptor saturation, and $x_{NMDA}$ is defined by

$dx_{NMDA}/dt = -x_{NMDA}/\tau_x + \delta(t - t_f)$.

Decay constant $\tau_x$ is set to 2ms. The voltage-dependence of NMDARs is captured by $\eta = 1/(1 + Mg \cdot \exp(-0.062 \cdot V)/3.57)$, where $Mg$ = 1mM describes the extracellular Magnesium concentration (Jahr and Stevens, 1990). Conductance strengths $G$ at recurrent synapses are $G_{AMPA,p}$ = $G_{AMPA,i}$ = 0.5nS, $G_{NMDA,p}$ = 4.1nS, $G_{NMDA,i}$ = 2.5nS, $G_{GABA,p}$ =6nS, and $G_{GABA,i}$ = 5.75nS, where subscripts $p$ and $i$ refer to synapses onto pyramidal neurons and interneurons respectively. The strength of extrinsic AMPA conductances onto pyramidal neurons and interneurons are $G_{AMPA,p}^{ext}$ = 2.75nS and $G_{AMPA,i}^{ext}$ = 2nS respectively. Example currents mediated by these synapses are shown in Figure 1B.

***Structured Pyramidal-to-Pyramidal Synapses***

All pyramidal-to-interneuron, interneuron-to-interneuron and interneuron-to-pyramidal and extrinsic synapses are unstructured, so in Equation 1, $W$ = 1 for all these cases. For pyramidal-to-pyramidal synapses, $W$ is a Gaussian function of the distance between neurons in a ring. The weight $W_{i,j}$ between any two pyramidal neurons $i$ and $j$ is thus given by

$W_{i,j} = \lambda + exp(-d^2/2\sigma^2)$,

where $\lambda$ = 0.2 provides an unstructured baseline weight between all pyramidal neurons (Compte et al., 2000; Tegnér et al., 2002; Ardid et al., 2007), $d = min(|i - j|dx, 2\pi - |i - j|dx)$ defines distance in the ring, $dx=2\pi/N_E$ is a scale factor and $\sigma$ = 0.35 (20°). Pyramidal-to-pyramidal weights are depicted in Figure 1C.

***Background Activity***

Extrinsic currents $I_{AMPA,\{p,i\}}^{ext}$ mediate non-selective background activity, simulating background firing from other brain regions and generating background activity in the network. We simulate the convergent activity of 1000 pre-synaptic neurons firing independent, homogeneous Poisson spike trains at 1Hz each with a single homogeneous Poisson spike train at 100Hz, where the conductance strength $G_{AMPA,\{p,i\}}^{ext}$ is scaled by $\kappa$ = 10, trading spatial and temporal summation (Prescott and deKoninck, 2003).

***Visual Discrimination Task***

Our discrimination task corresponds to a visual search task with two stimuli, one of which is designated the target by virtue of a feature contrast difference. Following a 300ms equilibration period (background activity only), the spike rates of the target and distractor stimuli were drawn from a normal distribution with mean $\mu$ corresponding to the centre of a Gaussian RF defined by $exp(-d^2/2\sigma_{ext}^2)$. Constant $d$ is given above for structured weights $W$ and $\sigma_{ext}$ = 0.52 (30°). The target and distractor RF centres were separated from each other by 180° in the ring network. Spike response adaptation in upstream visually-responsive neurons is modelled by a step-and-decay function (Trappenberg et al., 2001; Wong et al., 2007)

$d\mu/dt = (\mu - \mu_{init}/\mu_{div} \cdot \gamma_\mu)/\tau_\mu + \mu_{init}\, \delta(t - \alpha_\mu)$,

where $\mu_{div}$ = 2 determines the asymptotic input spike rate, $\tau_\mu$ = 25ms determines the rate of (upstream) spike response adaptation, and $\alpha_\mu = \Delta t$ ms is the onset of selective input following the 40ms visual response delay, prior to which $\mu=0$. Constant $\gamma_\mu$ is set to 1 for the target and to $\gamma_{ext}$ for the distractor, where $0.75 \leq \gamma_{ext} \leq 0.99$ determines target-distractor similarity. As with non-selective inputs, $G_{AMPA,p}^{ext}$ is scaled by $\kappa$ = 10. With an initial frequency $\mu_{init}$ = 400Hz, our selective inputs thus approximate a population of 100 upstream neurons firing at an initial rate of 40Hz, attenuating to rates of 20 Hz and 20 $\cdot \gamma_{ext}$ Hz at the target and distractor neurons respectively (Paré and Wurtz, 1997; Chafee and Goldman-Rakic, 1998), depicted in Figure 1D. Simulations were run with timestep $\Delta t$ = 0.1ms and the standard forward implementation of Euler's method, and were verified with the standard fourth order Runge-Kutta method.

**Results**

***Establishing Persistent Storage Parameters***

We determined the PS capability of the network as a function of the strength of intrinsic NMDARs, multiplying conductance strength at these synapses by a factor $\gamma_{NMDA}$. Figure 2 shows results from a simulated visuospatial working memory task in which the network was driven for 500ms by selective input (stimulus interval) followed by 5000ms without selective input (memory interval). Raster plots are shown on the top row and spike density functions (SDF, see Figure 2 caption) in the middle row. Spike rates are consistent with LIP data (e.g. Thomas and Paré 2007). As shown on the left of the figure (top two rows), for $\gamma_{NMDA}$ = 1, the network supports persistent neural activity during the memory interval at a much lower rate than during the stimulus interval (Paré and Wurtz, 1997). The network is thus inside the

PS regime, but close to the onset of PS dynamics, previously suggested to be optimal for decision circuits (Wang, 2002). We refer to this configuration as PS baseline. The network on the right does not support PS ($\gamma_{NMDA}$ = 0.8), as evidenced by the cessation of stimulus-selective activation during the memory interval.

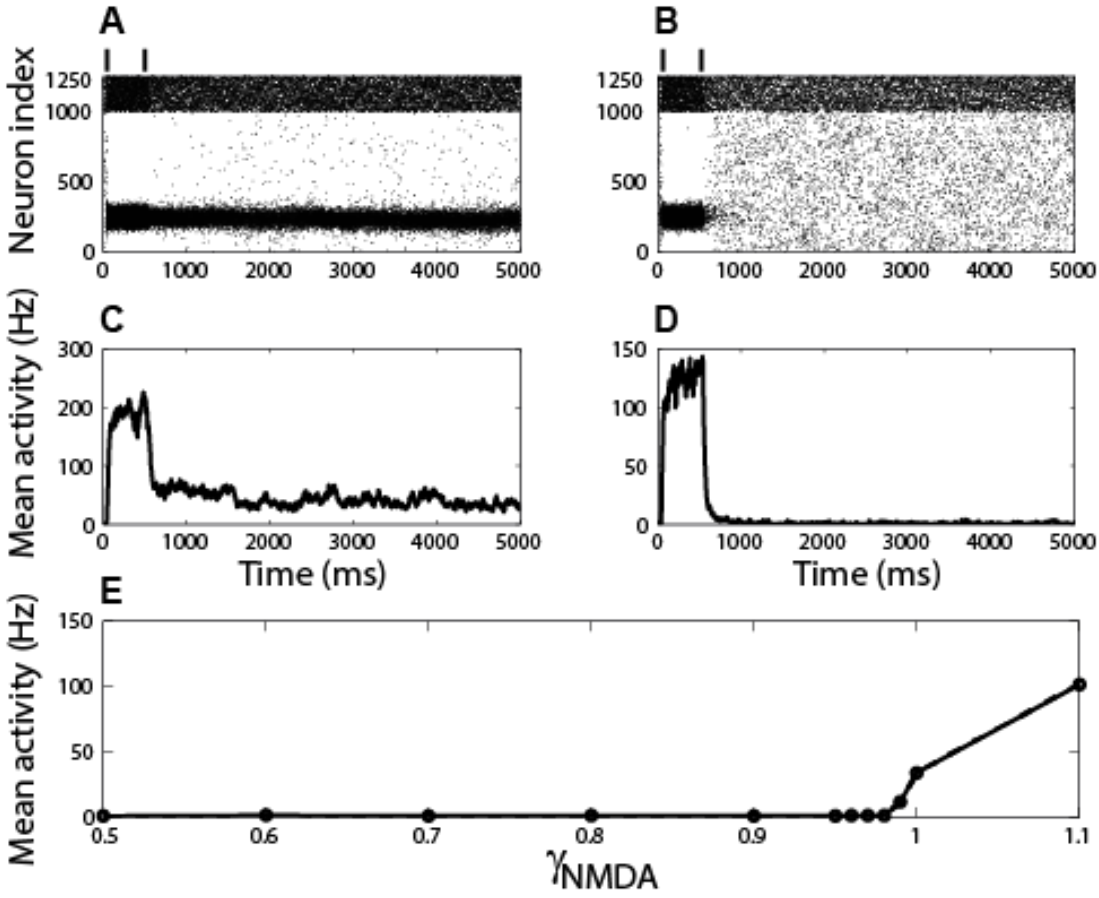


**Figure 2.** Establishing persistent storage (PS) parameters in the model. (**A-D**) Rasters and spike density functions (SDF, rounded to the nearest millisecond) for a simulated PS task in which the stimulus interval (first 500ms) was followed by a memory interval (5000ms, no selective input) under parameters that support (A,C, $\gamma_{NMDA}$ = 1) and do not support (B,D, $\gamma_{NMDA}$ = 0.8) PS. SDFs were built by convolving spike trains with a rise-and-decay function $(1 - exp(-t/\tau_r) \cdot exp(-t/\tau_d) / (\tau_d^2/(\tau_r + \tau_d))$ where $t$ is the time following stimulus onset, and $\tau_r$ = 1ms and $\tau_d$ = 20ms are the time constants of rise and decay respectively (Thompson et al., 1996). In rasters, pyramidal cells are indexed from 1-1000. Interneurons are indexed from 1001-1250. The beginning (40ms) and end (540ms) of the stimulus interval are indicated by the vertical bars at the top of the figure. Selective input was provided at the target only and SDFs are averaged over the target-selective population (41 neurons, see Discrimination Time). (**E**) Mean population SDF (solid curve) over the final 1000ms of the memory interval for $\gamma_{NMDA}$ = {1.1, 1.0, 0.99, 0.98, 0.97, 0.96, 0.9, 0.8, 0.7, 0.6, 0.5}. For $\gamma_{NMDA}$ = 1 (PS baseline) the network supports PS. For $\gamma_{NMDA} \leq$ 0.98, the network shows no mnemonic activity. $\gamma_{NMDA}$ = 0.99 is effectively the PS border. The dashed curve shows the population spike count rate over the same interval. The two curves are barely distinguishable.

The bottom row of the figure plots the respective firing rates for $0.5 \leq \gamma_{NMDA} \leq 1.1$ over the last 1000ms of the memory interval. For $\gamma_{NMDA} \geq 1$, the network supports PS, where higher rates are due to higher values of $\gamma_{NMDA}$. For $\gamma_{NMDA}$=1.1 (referred to as strong PS below), the network supports strong PS states, but beyond this parameter value (increments of 0.1), the network develops persistent, selective states from background noise. Consequently, we do not consider $\gamma_{NMDA} > 1.1$. Note that the bottom row does not show the rate of decline of activation following removal of the stimulus; among the non-PS networks, stimulus-selective activation drops off more slowly with higher values of $\gamma_{NMDA}$. For $\gamma_{NMDA} = 0.8$ (above right), the drop-off in activation is very sharp, indicating that the network is far from the PS regime. Unless otherwise specified, we refer to this configuration as 'the non-PS network' below.

***Simulated Visual Discrimination Task***

We simulated a two-choice visual discrimination task by centering extrinsic activity at two network locations for 1000ms, one designated the target and the other the distractor. Initial 'visual' responses were equal (Thomas and Paré, 2007) and the rate of target and distractor inputs diverged to separate steady states (Trappenberg et al., 2001; Wong et al., 2007), determining their similarity (and thus task difficulty). Average input spike rates at the target and distractor RF centres (henceforth the target neuron and distractor neuron respectively) are shown in Figure 1D. We ran 100 trials of the two-choice task for a range of target-distractor similarities $\gamma_{ext}$ and NMDA scale factors $\gamma_{NMDA}$, recording the spike times of each neuron in the network. Figure 3 (top row) shows mean activation of the target and distractor neurons on correct trials (see Discrimination Time) for the two values of $\gamma_{NMDA}$ shown for the visuospatial working memory task in Figure 2. For the non-PS network, target and distractor activation is slower to diverge and less distinguishable as input similarity increases, consistent with reaction times and recordings from PPC during visual discrimination tasks (Roitman and Shadlen, 2002). Regardless of task difficulty, activation follows an invariant profile in the PS baseline network.

***Discrimination Accuracy***

We used signal detection theory (Green and Swets, 1966) to determine how well an ideal observer could discriminate between the target and the distractor from spiking activity in the model, estimating the separation of the distributions of target and distractor activity at successive 1ms intervals. To this end, we calculated receiver operating characteristic curves (ROC) from the mean activation of the target and distractor neurons over the 100 trials. The area under the ROC (AUROC) quantifies the separation of their distributions (see Thompson et al., 1996). We quantified the probability of neuronal discrimination by a least squares fit of the AUROCs to a Weibull function

$$w(t) = \gamma - (\gamma - \delta) \cdot \exp(-(t/\alpha)^{\beta}) \qquad (2)$$

where $t$ is the time after stimulus onset, $\alpha$ is the time at which the function reaches 64% of its maximum, $\beta$ is the slope, and $\gamma$ and $\delta$ are the upper and lower limits of the function respectively. Discrimination magnitude (accuracy) was the upper limit of the function.

Figure 3 (bottom row) shows $w$ for all levels of target-distractor similarity for the PS baseline network and for the best performing non-PS network. Our ideal observer analysis shows that both networks quickly discriminate the target from the distractor on the easiest task, but as the task is made harder, only the non-PS network discriminates more slowly, reaching a high level of accuracy for all but the hardest task. The non-PS network thus trades speed for accuracy with increasing task difficulty, consistent with neural and pyschometric data from subjects performing visual discrimination tasks (eg. Roitman and Shadlen, 2002; Palmer et al., 2005). In contrast, the PS baseline network cannot enact this

trade-off, evidenced by the early asymptotes of these curves at low values. The strong PS network performs even worse (not shown).

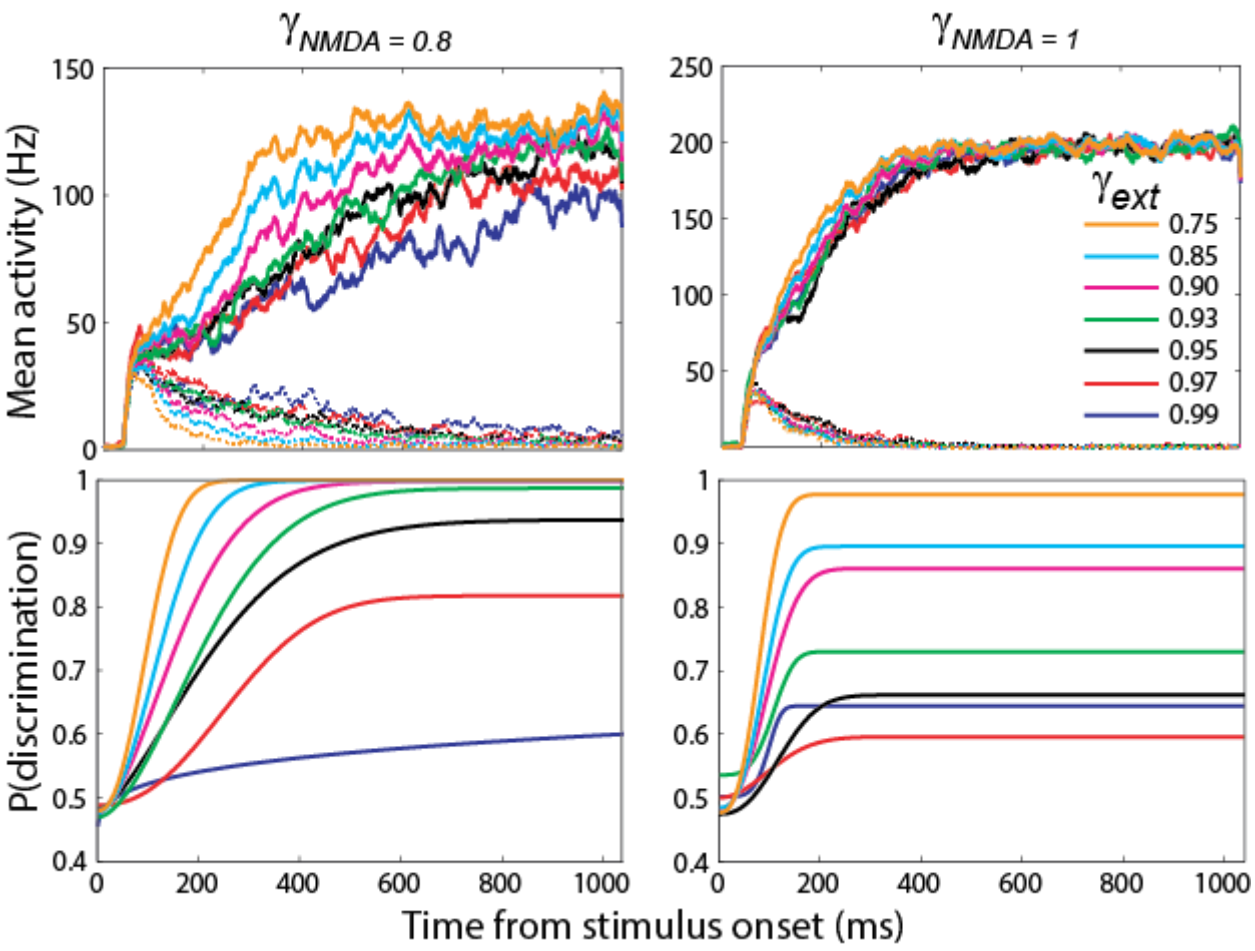


**Figure 3.** (**Top**) Mean spike density (100 trials) at the target and distractor neurons for each level of target-distractor similarity $\gamma_{ext}$ (see legend, top right). Correct trials only. In the non-PS network ($\gamma_{NMDA}$ = 0.8, left) target and distractor activation is slower to diverge with increasing task difficulty. Mean spike densities for the PS baseline network ($\gamma_{NMDA}$ = 1, right) are similar for all task difficulties. (**Bottom**) Quantifying the probability of target discrimination. Weibull fits to AUROCs (see text) calculated at the target and distractor neurons. With increasing task difficulty, curves from the non-PS network asymptote with increasing latency, trading speed for accuracy (left). All curves from the PS baseline network asymptote with similar latency, with corresponding loss of accuracy (right).

***Discrimination Time***

We quantified the timecourse of neuronal discrimination by calculating AUROCs for a population of target and distractor-selective neurons during each trial. The population $p$=41 included the target neuron, the distractor neuron, and an additional 20 neurons either side of these RF centres. The fitting procedure on each trial was thus equivalent to that described above for determining accuracy across all trials, but averaged over $p$ neurons instead of 100 trials. Additionally, because either the target or distractor population could dominate the network on any given trial (correct and error trials respectively), the AUROCs could be correspondingly fit with increasing or decreasing Weibull functions

*w*. On error trials (decreasing function), $\alpha$ in Equation 2 refers to the time at which *w* reached 64% of 1-min(*w*), and $\gamma$ and $\delta$ are the lower and upper limits respectively. The time at which *w* reached 0.75 was considered the discrimination time (Thompson et al., 1996) (0.25 on error trials). We averaged discrimination times (DT) over all trials to determine the speed of decision-making under each combination of task difficulty and strength of recurrent NMDARs, depicted in Figure 4. Trials on which *w* reached neither 0.75 nor 0.25 were discarded. For difficult tasks with weak recurrent NMDARs ($\gamma_{ext} \geq 0.97$ & $\gamma_{NMDA} < 0.7$) there were substantial numbers of discarded trials with a 1000ms stimulus interval (see Figure 6 caption), but results under these parameter values are informative about the discrimination abilities of the network because an ideal observer analysis does not require categorical choice. Quantifying DT on a trial-by-trial basis allows direct comparison with psychometric data, but if there is little to discriminate between target and distractor stimuli, it is reasonable not to choose at all, as in non-forced choice tasks. This important decision-theoretic issue is considered in the Discussion.

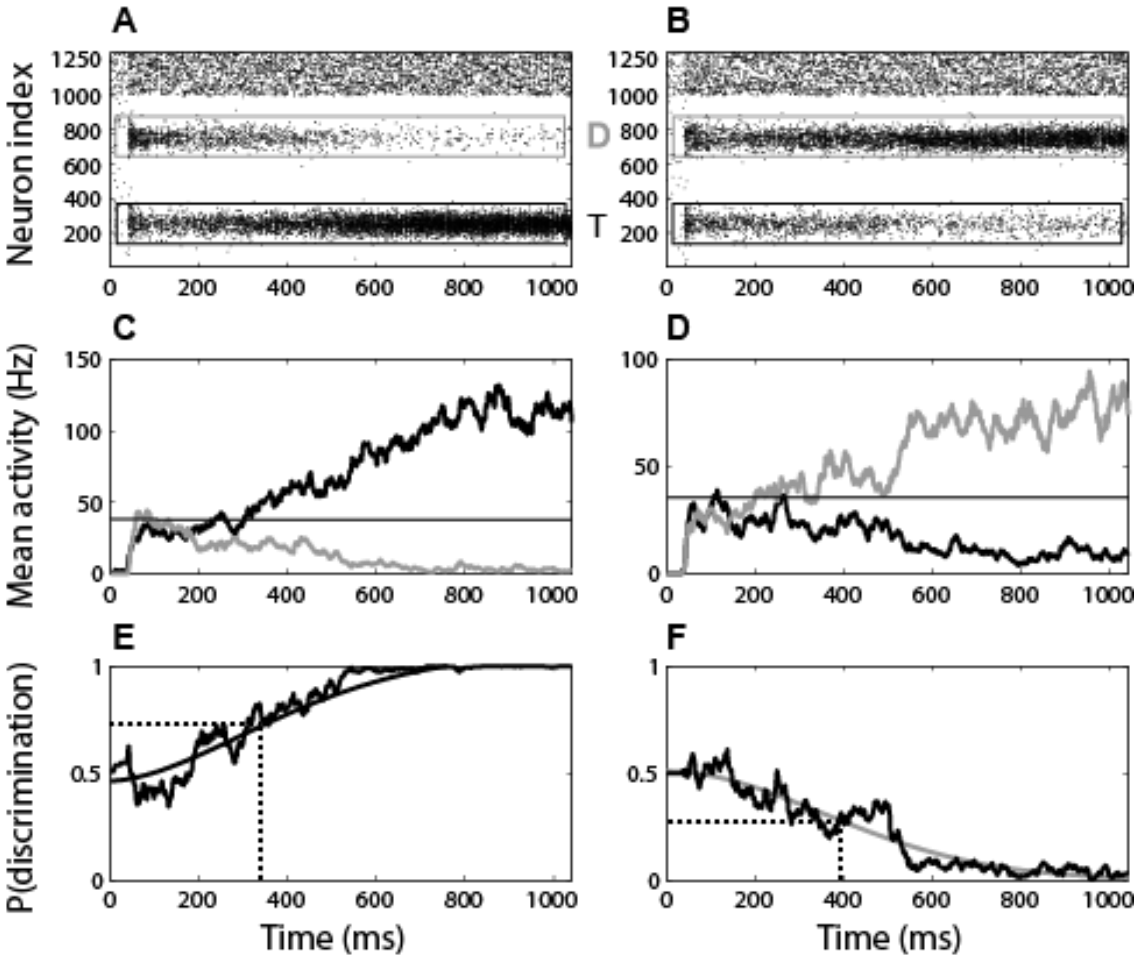


**Figure 4.** Determining discrimination time in the two-choice task. A correct trial (A,C,E) and error trial (B,D,F) where $\gamma_{NMDA} = 0.8$ and $\gamma_{ext} = 0.97$. (**A,B**) In raster plots, pyramidal neurons are indexed 1 - 1000. Interneurons are indexed 1001 - 1250. Target (T) and distractor (D) neurons are numbered 250 and 750 respectively. Stimulus interval begins at 40ms. (**C,D**) Mean spike densities for target (black) and distractor (grey) populations (41 neurons each). Horizontal black bars depict threshold method for checking discrimination time (see text). (**E,F**) AUROC and Weibull fits corresponding to spike densities above. Dotted black bars indicate discrimination time, when the Weibull function reaches 0.75 on correct trials and 0.25 on error trials.

We compared the above method for determining the timecourse of discrimination with an alternative method, similar to that of Wang (2002), comparing the mean activity of the target and distractor populations to a threshold frequency. Under this method, DT was considered to be the time at which the target activity (distractor activity for error trials) exceeded this threshold, with the additional constraints that 1) one activation function remain above threshold for 100ms and 2) the other concurrently remain below threshold for 100ms. These additional constraints served to distinguish the initial 'visual' response from the subsequent, decision-related activity. Results under the two methods were similar for thresholds between approximately 20 and 60Hz. We do not systematically investigate the similarities between the two methods here, though we note that their use captures a subtle distinction in the context of distributed processing. ROC analysis implements an ideal observer of the network, akin to a downstream circuit making decisions based on the network's activity (*e.g.* the superior colliculus reading out LIP activity). A neural threshold enables the network to make its own decisions, *i.e.* without an observer of its activity. However, in a model with competitive interactions between selective populations, the reaching of the threshold by one population entails a difference between its activation and that of the other, the same criterion used by ROC analysis to discriminate between the two populations. It is thus not surprising that the two methods yield similar results. See Standage et al. (2011) for a dynamic systems perspective on ROC with a network from the same family as this one.

As expected due to the longer latency of maximum discrimination as a function of task difficulty (Figure 3, bottom left), mean DTs for the best-performing network increased with target-distractor similarity, rising from 109ms to 283ms as $\gamma_{ext}$ was increased from 0.75 to 0.99. These DTs not only indicate a speed-accuracy trade-off on a trial-by-trial basis, but are consistent with reaction times in visual search tasks (Thomas and Paré, 2007). In contrast, mean DTs for the PS baseline network showed a very slight increase from 97ms to 107ms as task difficulty was increased.

Figure 5 shows cumulative distributions of DTs for the non-PS network and the PS baseline network. Not only does mean DT increase with task difficulty in the non-PS network (62% from the easiest to hardest task), but so does the standard deviation (84%), consistent with greater variability in reaction times with increasing task difficulty in monkeys performing a visual search task (Cohen et al., 2007). Furthermore, DT distributions show an increasingly long tail with increasing task difficulty ($0.75 \leq \gamma_{ext} < 0.99$), consistent with typical human reaction times in decision making tasks (see Smith and Ratcliff, 2004) and mathematical models known to be optimal for two choice tasks (Ditterich, 2006; Bogacz et al., 2006). Not only are mean DTs approximately constant across task difficulty in the PS baseline network (increasing by 9% from the easiest to hardest task), but their standard deviations *decrease* by 18%.

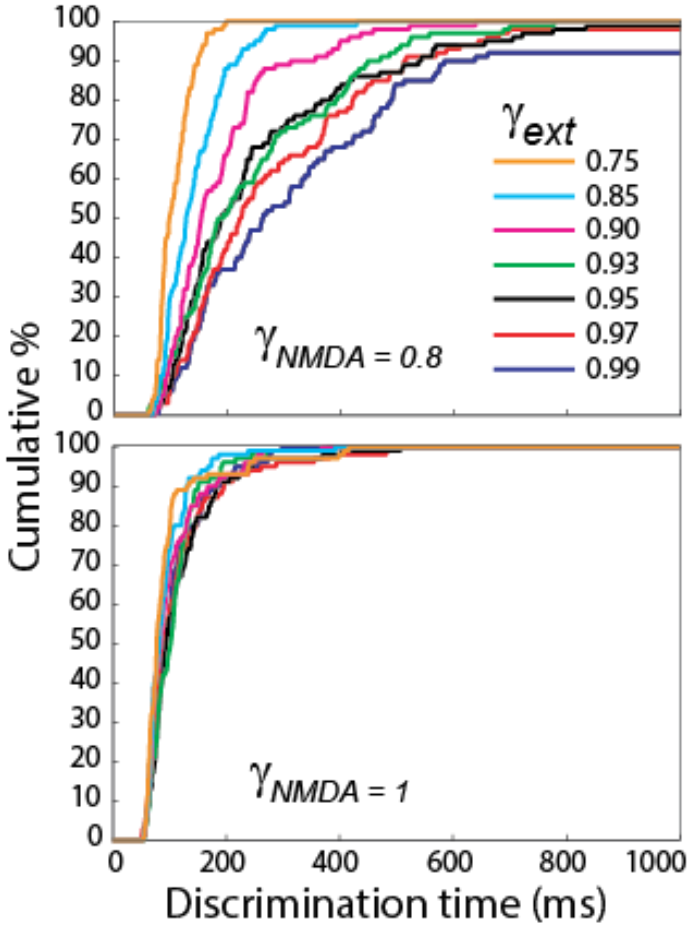


**Figure 5.** Cumulative distribution of discrimination times for the best-performing non-PS network ($\gamma_{NMDA}$ = 0.8, top) and the PS baseline network ($\gamma_{NMDA}$ = 1, bottom) for all levels of target-distractor similarity $\gamma_{ext}$. Discrimination times are broader with longer tails with increasing task difficulty for the non-PS network, but are approximately constant for the PS network.

***Target Discrimination with PS and non-PS Networks***

Figure 6 shows mean discrimination accuracy and time as functions of target-distractor similarity for a range of values of $\gamma_{NMDA}$. The top figure shows results for accuracy. The curves cluster for the easiest and most difficult tasks, though values of $\gamma_{NMDA}$ leading to the strongest and weakest recurrent dynamics furnish the least accurate networks. For task difficulties between these extremes, accuracy steadily improves as the strength of recurrent NMDARs is reduced from $\gamma_{NMDA}$ = 1.1 to $\gamma_{NMDA}$ = 0.8. This effect bottoms out at $\gamma_{NMDA}$ = 0.7 and network accuracy decreases for $\gamma_{NMDA}$ = 0.6, comparable to (though slightly better than) the PS baseline network. At $\gamma_{NMDA}$ = 0.5, the network is less accurate than baseline, performing comparably to the strong PS network.

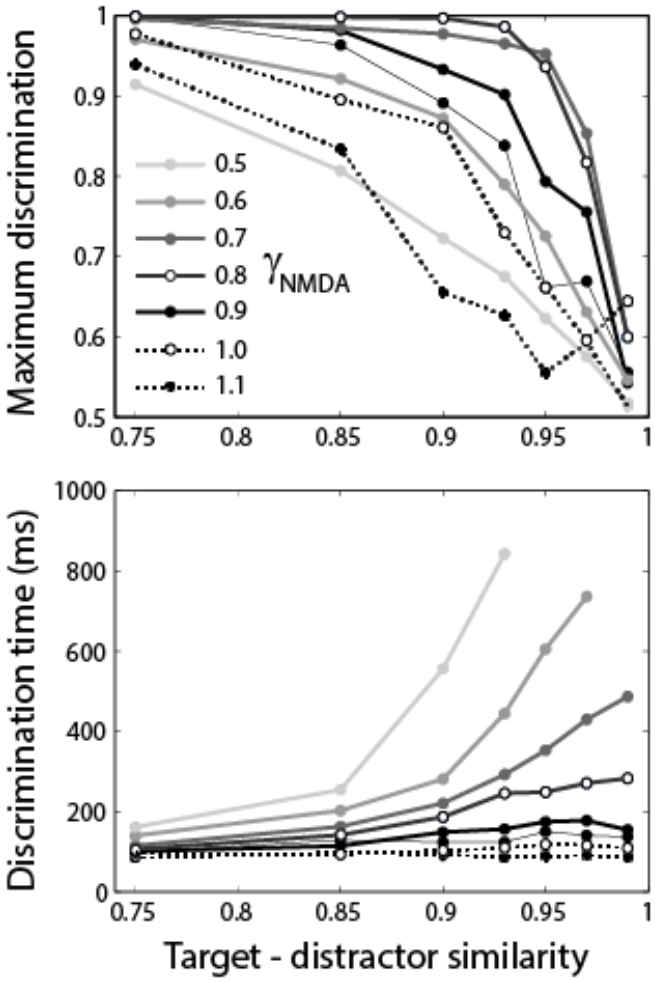


**Figure 6.** Accuracy (top) and speed (bottom) of target discrimination in the model. Dotted curves correspond to PS networks, where open and closed circles show PS baseline ($\gamma_{NMDA}$ = 1) and strong PS ($\gamma_{NMDA}$ = 1.1) respectively. Solid curves correspond to non-PS networks, where lighter curves correspond to lower values of $\gamma_{NMDA}$ ($0.5 \leq \gamma_{NMDA} \leq 0.95$, see colour legend). Results from the best performing network ($\gamma_{NMDA}$ = 0.8) are plotted with open circles. The thin black curves (one in each figure) correspond to $\gamma_{NMDA}$ = 0.95. The bottom figure shows that the weakest networks do not make decisions on the hardest tasks with a stimulus interval of 1s and discrimination criterion of 0.75 (see Results, Discrimination time). For $\gamma_{NMDA}$ = 0.5, the network makes decisions on 35% of trials for $\gamma_{ext}$ = 0.9, 8% for $\gamma_{ext}$ = 0.93 and 0% for $\gamma_{ext} \geq 0.95$. For $\gamma_{NMDA}$ = 0.6, the network makes decisions on 85% of trials for $\gamma_{ext}$ = 0.93, 44% for $\gamma_{ext}$ = 0.95, 7% for $\gamma_{ext}$ = 0.97 and 0% for $\gamma_{ext}$ = 0.99. For $\gamma_{NMDA}$ = 0.7, the network makes decisions on 75% of trials for $\gamma_{ext}$ = 0.97 and 44% for $\gamma_{ext}$ = 0.99. All other combinations of $\gamma_{NMDA}$ and $\gamma_{ext}$ yield decisions on more than 90% of trials.

The bottom figure shows results for DT. The curves cluster for the easiest level of target-distractor-similarity, much like the accuracy curves above. In general, as the task is made harder, all networks except the strong PS network show an increase in DT, though this effect is slight for PS baseline (an increase of ~10% from the easiest to the hardest task). This increase in DT with task difficulty is more pronounced as $\gamma_{NMDA}$ is reduced, though for $\gamma_{NMDA} \leq 0.6$, the network makes too few decisions to be

useful as a categorical decision-maker. Indeed, these two parameter values yield no decisions at all for the hardest task (see Figure 6 caption). Such a non-committal description of ambiguous evidence could be valuable to a downstream integrator reading the output of more than one circuit (see the Discussion). Note that longer DTs coincide with lower accuracy for $\gamma_{NMDA} \leq 0.6$ because the network is dominated by leakage of information (see Discussion and Supplementary Material).

We checked that decision-making followed a performance gradient near the onset of PS dynamics by running simulations with $\gamma_{NMDA}$ = {0.99, 0.97, 0.95, 0.93}. These parameter values lead to results for accuracy and DT that transitioned smoothly between PS baseline and $\gamma_{NMDA}$ = 0.9, shown by the thin solid curves in Figure 6 for $\gamma_{NMDA}$ = 0.95. Notably, for $\gamma_{NMDA}$ = 0.99, network performance was nearly indistinguishable from PS baseline (not shown).

***Robustness of the Results***

Our study is based on earlier work in which the long time constant of NMDARs at intrinsic synapses onto pyramidal neurons was hypothesized to support persistent mnemonic activity and the integration of evidence in decision processing (Wang, 2002; Wong and Wang, 2006; Wang, 2008). We followed these authors in using an NMDAR time constant of $\tau_{NMDA}$ = 100ms, but measurements of $\tau_{NMDA}$ differ according to experimental methods, receptor subtype and function (see Cull-Candy et al., 2001; Cull-Candy and Leszkiewicz, 2004). The value of $\tau_{NMDA}$ used here and in earlier work is consistent with the upper end of measurements for intrinsic cortical processing, so we also ran simulations with $\tau_{NMDA}$ = 50ms, consistent with measurements at the lower end of the data (Kumar and Huguenard, 2003; Wang et al., 2008). Under this parameter value, the PS network corresponded to $\gamma_{NMDA}$ = 1.2, while $\gamma_{NMDA}$ = 0.8 was far outside the PS regime. The shorter NMDAR time constant did not qualitatively effect our results (Figure 7, solid curves).

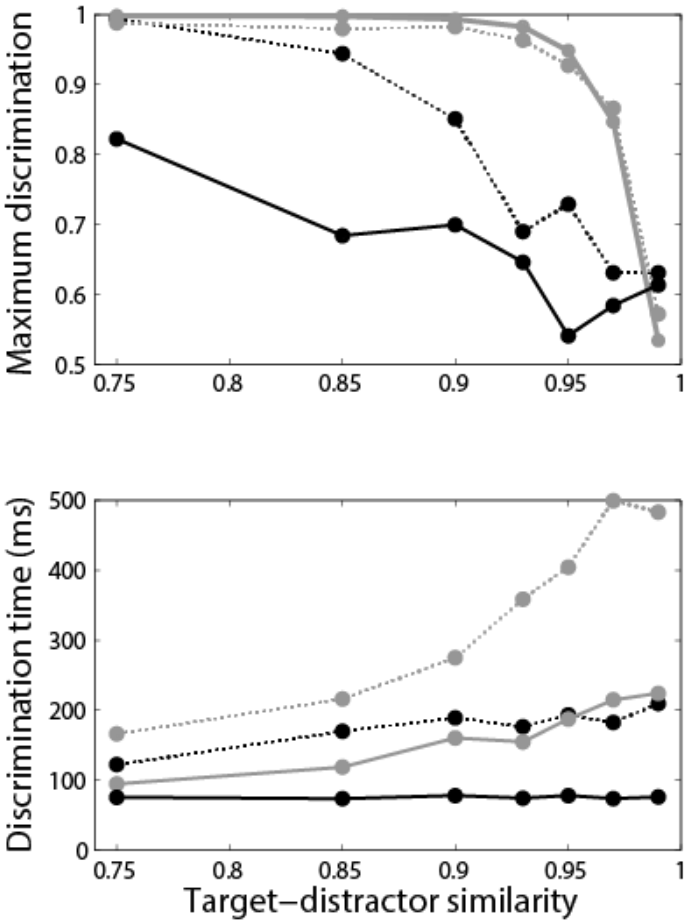


**Figure 7.** Accuracy (top) and speed (bottom) of target discrimination under alternative parameters. Persistent storage (PS) capability impaired decision making in the network when simulations were run with a shorter NMDA time constant ($\tau_{NMDA}$ = 50ms, solid curves) and when the signal-to-noise ratio of selective inputs and the strength of synaptic conductances were decreased (dotted curves, see Robustness of the Results). Black and grey curves correspond to PS and non-PS networks respectively.

To determine if the results depend on the details of selective input, we ran 100 trials with the PS baseline network and the non-PS network, where the initial input rate at the target RF centre was reduced from 400Hz to 200Hz - simulating 100 upstream visually responsive neurons at 20Hz instead of 40Hz. We also ran 100 trials with a step input (no decay) at 100 Hz. These alternate input parameters did not qualitatively affect the results (not shown). As long as the inputs were strong enough to furnish a transition from extrinsic to intrinsic processing (Wilson and Cowan, 1973; Douglas and Martin, 2004, 2007), results were qualitatively invariant.

We also examined the signal-to-noise ratio (SNR) of the selective input. This ratio is determined by the relative rates of background and selective input spike trains and the spatial and temporal profiles of selective input. At the centre of the target RF, the selective input SNR was 4 in the simulations above, decaying to 2 with a time constant of 25ms (see Methods). This ratio is within range of neuronal responses in LIP (Paré and Wurtz, 2001), but is two orders of magnitude higher than in earlier work in which PS capability was advantageous in a network model of a decision circuit (Wang, 2002). This

difference remains when the SNR is integrated over the RFs (Gaussian vs. square) and temporal input profiles (step-and-decay vs. step only) used in each study, normalized for network size and trial length. Synaptic conductance strength also differed markedly (one order of magnitude greater here), so we reduced both these differences to further test the robustness of our findings. All synaptic conductance strengths were divided by 2, the background rate was multiplied by 2 and selective rates were multiplied by 3/4. We ran 100 trials across all task difficulties for PS baseline and the non-PS network and we confirmed that these new configurations remained inside and outside the PS regime respectively. These simulations further demonstrated that the specific values of our original parameters are not crucial to our findings. Accuracy was very similar to Figure 6 for both configurations. DTs were longer than in Figure 6, but were well within range of DTs in visual search tasks (Figure 7, dotted curves). These simulations do not exhaust the parameter variations that may affect our results, but they address major differences between our model and earlier modelling work where results differed from ours (Wang, 2002; Wong and Wang, 2006). Further parameter dependencies are described in the Discussion.

**Discussion**

We have interpolated between strong and weak dynamics in a spiking network model of a decision circuit in LIP, investigating the hypothesized dependence of decision making on local circuit PS capability (Wang, 2002; Wong and Wang, 2006; Wang, 2008). Building on earlier studies demonstrating the potential importance of intrinsic NMDARs to persistent mnemonic activity (Fransén and Lansner, 1995; Lisman et al., 1998; Wang, 1999; Durstewitz et al., 2000; Compte et al., 2000; Tegnér et al., 2002) and decision making (Wang, 2002; Wong and Wang, 2006), we systematically controlled both these functions by manipulating the strength of NMDA conductance at intrinsic synapses onto pyramidal neurons. We controlled task difficulty by manipulating target-distractor similarity in a simulated visual search task, finding that parameters that support PS entail poor decision making. Non-PS parameters (far from the PS regime, see Figure 2) enable more effective integration of decision options for a broad range of task difficulties, yielding signature characteristics of reaction time distributions and reproducing the speed-accuracy trade-off predicted by decision theory (Ratcliff and Smith, 2004; Smith and Ratcliff, 2004) and shown by psychometric data from visual discrimination tasks (eg. Roitman and Shadlen, 2002; Palmer et al., 2005; Churchland et al., 2008).

While this finding is different from those of earlier studies, the principle mechanisms at play are not. The long NMDA time constant at intrinsic synapses allows the slow integration of decision options, instantiating the accumulators of classic sequential sampling models (Ratcliff and Smith, 2004; Smith and Ratcliff, 2004). A common pool of inhibitory interneurons among stimulus-selective pyramidal neurons effectively creates a diffusion process (Usher and McClelland, 2001; Bogacz, 2007), where evidence for one option accumulates at the expense of the other. If intrinsic NMDARs are too strong, however, recurrent dynamics dominate the network's input, amplifying noise and eliminating the advantage of slow evidential buildup. See Wong and Wang (2006) for an analysis of these dynamics in a simplified model. In the model used here, recurrent dynamics strong enough to furnish persistent mnemonic activity are too strong for slow integration under parameters determined by experimental data (see Materials and Methods). In this respect, the difference between earlier findings and ours is the onset of this ‘too strong’ regime.

There are a number of differences between our simulations and those of earlier studies (Wang, 2002; Wong and Wang, 2006) that may contribute to the difference in network dynamics. A difference of an order of magnitude in intrinsic synaptic conductance strength and a difference of two orders of magnitude in selective input SNR are described above. Additionally, different network architectures were used to simulate tasks where reaction times occur on different timescales and spike rates differ

markedly in LIP (visual search vs. random dot motion tasks). Ultimately, the decision making ability of a recurrent network depends on its time constant of integration, reflecting a balance between the leakage and amplification of accumulated input (Usher and McClelland, 2001). This balance has been expressed in terms of the strength and time constant of intrinsic NMDARs (Wang, 2008). As such, there is an optimal NMDAR conductance strength above and below which the network is dominated by amplification and leakage respectively (see the Supplementary Material). The more the conductance exceeds the optimum, the more quickly activation is amplified, leading to earlier choice selection (thus preventing further integration). The more the conductance falls short of the optimum, the more quickly activation reaches a level at which it leaks as fast as it accumulates (also preventing further integration). The optimum in our model is furnished by $\gamma_{NMDA} \approx 0.8$ (Figure 6), but this value depends on other model parameters, such as the relative strength of excitatory and inhibitory synaptic conductance, the spatial extent of intrinsic interactions governing tuning curves, and the level of background noise in the network. We thus do not claim that our results are general. Indeed, outside the PS regime, the model cannot make decisions if the stimuli offset before discrimination of the target or distractor. We have, however, clearly shown that local circuit PS capability is not a general requirement of biophysically-based models of decision circuits. Decision-theoretic analyses have identified biophysical parameters under which models with discrete stimulus-selective populations are equivalent to drift diffusion models in two-choice tasks (Bogacz et al., 2006). An extension of these analyses to the spatial continuum of interactions in centre surround models (here and e.g. Beck et al., 2008; Furman and Wang, 2008) is an important next step.

***Processing Requirements of a Salience Map***

The incompatibility between the processing requirements of decision-making and persistent storage in our PPC model is consistent with reports that LIP represents the relative importance (salience, priority) of items in the visual field (Serences and Yantis, 2006; Goldberg and Bisley, 2006). Our model belongs to a family of cortical models (Wilson and Cowan, 1973) in which PS capability entails winner-take-all dynamics in the limit of infinite time (Amari, 1977). The model can therefore support a single active region following stimulus offset, but the time over which the dynamics converge is parameter-dependent and potentially within the time constraints of many cognitive tasks (Trappenberg and Standage, 2005). The winner-take-all constraint does not necessarily apply during a stimulus interval, but nonetheless, the strong recurrent dynamics of PS circuits are potentially ill-suited to multiple stimulus-driven representations; strong inputs are required to dominate the network's recurrent dynamics, negating the role of intrinsic circuitry. As such, the dynamics of PS networks may be better suited to the representation of one item at a time. Under non-PS parameters, mutual inhibition between regions of the network facilitates competition between stimulus-selective populations, but the recurrent dynamics are weaker than for PS parameters, more easily permitting multiple items to be simultaneously represented, scaled in proportion to the strength of their inputs. In effect, weaker dynamics allow the `push-pull' of a diffusion process without imposing categorical choice and allow slow transitions between the respective populations dominating the network over time. Both these features would support the representation of salience in LIP. This prediction is consistent with the results of Standage et al., (2005), who found that under non-PS parameters, a population rate model of PPC could explain divergent experimental results on the distribution of visuospatial attention if driven by persistent mnemonic inputs (putatively from PFC). Gradual transitions in the non-PS network can be seen in Figure 4.

While this class of model is commonly used to simulate visuospatial working memory tasks (e.g. Camperi and Wang, 1998; Compte et al., 2000) due to its support for the spatially periodic configuration of items

in many such tasks (*e.g.* Funahashi et al, 1989; Chafee and Goldman-Rakic, 1998) and for its persistent storage regime, fast winner-take-all dynamics during a memory interval conflict with well-established multi-item capacity constraints of working memory (Miller, 1956; Luck and Vogel, 1997; Cowan, 2001). Our network (in the PS regime) can therefore represent each item of the periodic array, but it can only support one such item over a memory interval. The model can be adapted to support multi-item working memory by structuring the connectivity between pyramidal neurons and inhibitory interneurons (Macoveanu et al., 2006; Edin et al., 2009), effectively partitioning the network into functionally separate modules. Mechanisms proposed for the stabilization of persistent mnemonic activity have a similar effect (Trappenberg, 2003), but they do so at the expense of the competitive dynamics required of decision circuitry (Trappenberg and Standage, 2005). We do not expect a given local circuit to account for the full capacity of working memory, but rather, we envision interactions between such circuits. For example, if one circuit were to support the representation of the items in a visual array, persistent storage could be provided by coupled circuits, perhaps one per item. If so, interference between such 'caching' networks might account for working memory capacity constraints. Further work is required to address such possibilities.

***Other Mechanisms Hypothesized to Support Persistent Mnemonic Activity***

Earlier studies proposing that persistent storage capability is a requirement of local circuit decision-making focused on intrinsic (recurrent) network processing (Wang, 2002; Wong and Wang, 2006; Wang, 2008). Our model addresses this hypothesis and therefore continues in this vein, but recurrent synaptic activity is just one mechanism proposed to underlie persistent storage. Other mechanisms include intracellular calcium dynamics (Fransén, 2005; Fransén et al., 2006; Winograd et al., 2008); feedforward network oscillations (Lisman and Idiart, 1995); and inter-cortical and cortico-subcortical interactions (see Wang, 2001; Constantinidis and Wang, 2004). While our model is capable of encoding a continuum of feature values such as spatial location, persistent activity encoding these values over a delay period is limited to a single frequency for a given value of NMDAR conductance strength (Figure 2). Without some form of modulation, the model is therefore unable to simulate graded persistent activity. Such activity has been recorded in cortex during working memory tasks (Romo et al., 1999; Barak et al., 2010) and has been reproduced in a network model that makes decisions in a two-interval discrimination task (Machens et al, 2005). Similar line attractor dynamics have also been used to model the memory of eye position in the brain stem (Seung, 1996, 1998). It is worth noting that stable persistent activity is generally regarded as a minority case among forms of delay period activity (Durstewitz and Seamans, 2006). More commonly, up and down-ramping activity is seen during these intervals, hypothesized to support prospective and retrospective coding respectively (see Brody et al, 2003) as well as the encoding of temporal intervals (see Durstewitz and Seamans, 2006). Recently, ramping activity has been proposed to modulate local circuit dynamics during decisions, driving a transition from outside to inside the PS regime, where the rate of transition governs the speed-accuracy trade-off (Standage et al., 2011).

***Local and Distributed Processing***

We have presented our findings in the context of sequential sampling models, where in the visuospatial domain, a single decision circuit integrates evidence for spatial locations until one representation exceeds a threshold level of activity, leading to an eye movement. This role of eye movement decision-maker has been attributed to several cortical regions, including PPC (Roitman and Shadlen, 2002) and FEF (Hanes and Schall, 1996), as well as to the mid-brain superior colliculus (Paré and Hanes, 2003). We have further emphasized the decision making abilities of our model under parameters that either do or do not support PS, but it may not be necessary to draw categorical distinctions between PS and non-PS networks in the domain of decision making, nor between the physiological properties of cortical regions

correlated with eye-movement decisions (Chafee and Goldman-Rakic, 1998). By manipulating the strength of intrinsic NMDARs, we show a gradient of network dynamics and consequent speed-accuracy trade-offs (Figure 6). The PS border sits on this gradient, but may be epiphenomenal in this context. In a framework of distributed decision making, a range of computational properties are conferred along the gradient, including PS. At one extreme, the strongest dynamics yield PS states robust to distractor stimuli, whereas somewhat weaker dynamics (eg. PS baseline) yield PS states more readily subject to interference (Compte et al., 2000). Depending on task demands, networks with both these properties would provide useful input to other brain regions directly involved in eye movement production, such as FEF and SC. At the other extreme are non-committal noise filters with dynamics too weak to make categorical decisions in difficult tasks. Such slow, conservative `advice' would also be useful downstream, potentially balancing any errors from more decisive circuits. In our model, the behaviour of sequential sampling models is furnished by parameters between these extremes ($0.7 \leq \gamma_{NMDA} \leq 0.9$) where the speed-accuracy trade-off resembles experimental data (eg. Palmer et al., 2005). Clearly, networks in this range would be useful to downstream decision-related structures.

These results are consistent with a widely held view of eye-movement decisions where no single circuit is responsible for integrating sensory evidence or making decisions (Schall, 2001). We posit that networks in cortical regions such as PFC, FEF and PPC, frequently correlated with integration of decision options, may play different roles in integrating these options by virtue of the strength of their dynamics. For example, the winner-take-all dynamics of PS networks would appear ideal for caching decisions and providing categorical bias to circuits involved in sensory processing, consistent with the biased competition theory of attention (Desimone and Duncan, 1995). As discussed above, weaker dynamics may play a complementary role, more readily allowing multiple items to be simultaneously considered by downstream decision-related structures, consistent with the concept of a salience map (Treisman and Gelade, 1980; Koch and Ullman, 1985). There is, of course, the possibility that the dynamics of decision circuits are modulated by task demands. For example, the same circuit could mediate PS and a salience map by modulation of intrinsic NMDARs, consistent with reports of increased dopaminergic activity during working memory tasks and dopamine enhancement of NMDA conductance in PFC (see Durstewitz et al., 2000).

Finally, earlier proposals that decision circuits should support PS (Salinas and Sejnowski, 2001; Wang, 2002; Wong and Wang, 2006; Wang, 2008) may have been influenced by the memory component of delayed response tasks; in single-circuit models in which persistent mnemonic activity and decision making are supported by the same mechanism, PS capability is required *a priori* to simulate these tasks. Similarly, PS capability guarantees categorical choice in single-circuit simulations of delayed response tasks, even in the absence of information to guide that choice (Wang, 2002). Given the general acknowledgement that persistent mnemonic activity may be supported by inter-circuit mechanisms in addition to intra-circuit ones (Chafee and Goldman-Rakic, 1998, 2000), this possibility highlights the need for coupled circuit models to guide further experiments. For instance, if LIP is not a PS network *per se*, then persistent activity in this cortical area must either be driven by activity somewhere else, such as PFC (see Constantinidis and Wang, 2004), or distinct parietal networks including LIP may be differentiated by their dynamics, as demonstrated by our simulations. Fronto-parietal connectivity and the contributions of PFC and PPC circuits to working memory processing are receiving considerable attention (Chafee and Goldman-Rakic, 2000; Babiloni et al., 2004; Curtis et al., 2004; Edin and Klingberg, 2007; McNab and Klingberg, 2008). Similar attention is required in the domain of decision making.

If decision making relies on NMDARs for integration of evidence, it should be possible to interfere with decision making by pharmacological manipulation of NMDARs, as shown previously for persistent mnemonic activity and working memory more generally (Dudkin and Kruchinin, 1997; Aura and

Riekkinen, 1999). Recent experimental works supports this hypothesis, where the performance of monkeys in a visual discrimination task improved with low doses of the NMDA antagonist Ketamine, before deteriorating at higher doses (Shen et al., 2010). These findings are consistent with our simulations of PPC, where for a range of values of intrinsic NMDAR strength, network performance improved before deteriorating at lower values (Figure 6). Further experiments are required to constrain pharmacological impairment of NMDA to cortical regions such PFC and PPC during decision making tasks.

**Acknowledgements**

This work was supported by the Canadian Institutes of Health Research. We thank Kevin Johnston and Thomas Trappenberg for helpful comments.

**References**

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge, UK: Cambridge University Press.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern*., 27, 77–87.

Ardid, S., Wang, X.-J., & Compte, A. (2007). An integrated microcircuit model of attentional processing in the neocortex. *J Neurosci*., 27, 8486–8495.

Aura, J. & Riekkinen, P. J. (1999). Blockade of NMDA receptors located at the dorsomedial prefrontal cortex impairs spatial working memory in rats. *Neuroreport*, 10, 243–248.

Babiloni, C., Babiloni, F., Carducci, F., Cincotti, F., Vecchio, F., Cola, B., Rossi, S., Miniussi C., & Rossini, P. M. (2004). Functional frontoparietal connectivity during short-term memory as revealed by high-resolution EEG coherence analysis. *Behav Neurosci*., 118, 687–697.

Barak, O., Tsodkys, M., & Romo, R. (2010).Neural population coding of parametric working memory. *J Neurosci*., 30, 9424–9430.

Beck, J. M., Ma, W. J., Kiani, R., Hanks T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., & Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, 60, 1142-1152.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics or optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev*., 113, 700–765.

Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn Sci*., 11, 118–125.

Bruce, C. J. & Goldberg, M. E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol*., 53, 603–635.

Camperi, M. & Wang, X.-J. (1998). A model of visuospatial working memory in prefrontal cortex: Recurrent network and cellular bistability. *J Comput Neurosci*., 5, 383–405.

Chafee, M. V. & Goldman-Rakic, P. S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol*., 79, 2918–2940.

Chafee, M. V. & Goldman-Rakic, P. S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *J Neurophysiol*., 83, 1550–1566.

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nat Neurosci*., 11, 693–702.

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: A computational model. *J Neurosci*., 26, 9761–9770.

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex*, 10, 910–923.

Constantinidis, C. & Wang, X.-J. (2004). A neural circuit basis for spatial working memory. *The Neuroscientist*, 10, 553–565.

Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci.,* 24, 87–185.

Cull-Candy, S. G., Brickley, S., & Farrant, M. (2001). NMDA receptor subunits: diversity, development and disease. *Curr Opin Neurobiol*., 11, 327–335.

Cull-Candy, S. G. & Leszkiewicz, D. N. (2004). Role of distinct NMDA receptor subtypes at central synapses. *Sci. STKE*., 2004, re16.

Curtis, C. E., Rao, V. Y., & D'Esposito, M. (2004). Maintenance of spatial and motor codes during oculomotor delayed response tasks. *J Neurosci*., 24, 3944–3952.

Desimone, R. & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu Rev Neurosci*., 18, 193-222.

Destexhe, A. & Paré, D. (1999). Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol*., 81, 1531-1547.

Douglas, R. J. & Martin, K. A. C. (2004). Neuronal circuits of the neocortex. *Annu Rev Neurosci*., 27, 419–451.

Douglas, R. J. & Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology*, 17, R496–R500.

Dudkin, K. N., Kruchinin, V. K., & Chueva, I. V. (1997). Effect of NMDA on the activity of cortical glutaminergic structures in delayed visual differentiation in monkeys. *Neurosci Behav Physiol*., 27, 153–158.

Durstewitz, D., Seamans, J. K., Sejnowski T. J. (2000). Neurocomputational models of working memory. *Nat Neurosci*., 3, 1184–1191.

Durstewitz, D. & Seamans, J. K. (2006). Beyond bistability: biophysics and temporal dynamics of working memory. *Neuroscience*, 139, 119–133.

Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegnér, J., & Compte, A. (2009). Mechanism for top-down control of working memory capacity. *Proc Natl Acad Sci USA*., 106, 6802-6807.

Edin, F., Klingberg, T., Stodberg, T., & Tegnér, J. (2007). Fronto-parietal connection asymmetry regulates working memory distractibility. *Journal of Integrative Neuroscience*, 6, 567-596.

Fransén, E. & Lansner, A. (1995). Low spiking rates in a population of mutually exciting pyramidal cells. *Network: Computation in Neural Systems*, 6, 271–288.

Fransén, E. (2005), Functional role of entorhinal cortex in working memory processing. *Neural Networks*, 18, 1141–1149.

Fransén, E., Babak T., Egorov, A. V., Hasselmo, M. E., & Alonso, A. A. (2006). Mechanism of graded persistent cellular activity of entorhinal cortex layer 5 neurons. *Neuron*, 49, 735–746.

Funahashi, S., Bruce, C. J., Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol*., 61, 331–349.

Furman, M. & Wang, X.-J. (2008). Similarity Effect and Optimal Control of Multiple-Choice Decision Making. *Neuron*, 60, 1153–1168.

Fuster, J. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol*., 36, 61–78.

Galarreta, M. & Hestrin, S. (1998). Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. *Nat Neurosci*., 1, 587–594.

Gnadt, J. W. & Andersen, R. A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res*., 70, 216–220.

Goldberg, M. E., Bisley, J. W., Powell, K. D., & Gottlieb, J. (2006). Saccades, salience and attention: the role of the lateral intraparietal area in visual behavior. *Prog Brain Res*., 155, 157–175.

Goldman-Rakic, P. (1995). Cellular basis of working memory. *Neuron*, 14, 477–485.

Gonzalez-Burgos, G., Kroener, S., Zaitsev, A. V., Povysheva N. V., Krimer, L. S., Barrionuevo, G., & Lewis, D. (2008). Functional maturation of excitatory synapses in Layer 3 pyramidal neurons during postnatal development of the primate prefrontal cortex. *Cereb Cortex*., 18, 626-637.

Green, D. M. & Swets, J.A. (1966). Signal detection theory and psychophysics (New York: Wiley).

Gutkin, B. S., Laing, C. R., Colby, C. L., Chow, C. C., & Ermentrout, G. B. (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.

Hanes, D. P. & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, 274, 427–430.

Hasegawa, R. P., Matsumoto, M., & Mikami, A. (2000). Search target selection in monkey prefrontal cortex. *J Neurophysiol*., 84, 1692–1696.

Jahr, C. E. & Stevens, C. F. (1990). Voltage dependence of nmda-activated macroscopic conductances predicted by single-channel kinetics. *J Neurosci*., 10, 3178–3182.

Koch, C. & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, 4, 219–227.

Kumar, S. S. & Huguenard, J. R. (2003). Pathway-specific differences in subunit composition of synaptic NMDA receptors on pyramidal neurons in neocortex. *J Neurosci.*, 23, 10074–10083.

Lisman, J. E. & M. A. P. Idiart (1995). Storage of 7 +- 2 short-term memories in oscillatory subcycles. *Science*, 267, 1512-1515.

Lisman, J., Fellous, J., & Wang X.-J. (1998). A role for NMDA-receptor channels in working memory. *Nat Neurosci*., 1, 273-5.

Luck, S. J. & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-281.

Machens, C. K., Romo, R., & Brody, C. D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science*, 307, 1121–1124.

Macoveanu, J., Klingberg, T., & Tegnér, J. (2006). A biophysical model of multiple-item working memory: a computational and neuroimaging study. *Neuroscience*, 141, 1611–1618.

Miller, G. A. (1956).The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol Rev*., 63, 81–97.

McNab, F. & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nat Neurosci*., 11, 103–107.

Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5, 376–404.

Paré, M. & Hanes, D. P. (2003). Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *J Neurosci*., 23, 6480–6489.

Paré, M. & Wurtz, R. H. (1997). Monkey posterior parietal cortex neurons antidromically activated from superior colliculus. *J Neurophysiol*., 78, 3493–3497.

Paré, M. & Wurtz, R. H. (2001). Progression in neuronal processing for saccadic eye movements from parietal cortex area LIP to Superior Colliculus. *J Neurophysiol*., 85, 2545–2562.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nat Rev Neurosci*., 1, 125–132.

Prescott, S. A. & De Koninck, Y. (2003). Gain control of firing rate by shunting inhibition: Roles of synaptic noise and dendritic saturation. *Proc Natl Acad Sci USA*., 100, 2076–2081.

Ratcliff, R. & Smith P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev*., 111, 333–367.

Roitman, J. D. & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci*., 22, 9475–9489.

Romo, R., Brody, C. D., Hernandez, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, 399, 470–473.

Salin, P. A. & Prince, D. A. (1996). Spontaneous $GABA_A$ receptor-mediated inhibitory currents in adult rat somatosensory cortex. *J Neurophysiol*., 75, 1573-1588.

Salinas, E. & Sejnowski, T. J. (2001). Gain modulation in the central nervous system: Where behavior, neurophysiology, and computation meet. *The Neuroscientist*, 7, 430–440.

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nat Rev Neurosci*., 2, 33–42.

Schall, J. D. & Hanes, D. P. (1993). Neural basis of saccade target selection in frontal eye field during visual search. *Nature*, 366, 467–469.

Serences, J. T. & Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends Cogn Sci*., 10, 38–45.

Seung, H. S. (1996). How the brain keeps the eyes still. *Proc Natl Acad Sci USA*., 93, 13339–13344.

Seung, H. S. (1998). Continuous attractors and oculomotor control. *Neural Networks*, 11, 1253–1258.

Shen, K., Kalwarowsky, S., Clarence, W., Brunamonti, E., & Paré, M. (2010). Beneficial effects of the NMDA antagonist Ketamine on decision processes in visual search. *J Neurosci*., 30, 9947–9953.

Smith, P. L. & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci*., 27, 161–168.

Standage, D. I., Trappenberg, T. P., & Klein, R. M. (2005). Modelling divided visual attention with a winner-take-all network. *Neural Networks*, 18, 620–627.

Standage, D., You, H., Wang, D., & Dorris, M.C. (2011). Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Front Comput Neurosci.*, 5:7, *in press*.

Tegnér, J., Compte, A., & Wang, X.-J. (2002). The dynamical stability of reverberatory neural circuits. *Biol Cybern*., 87, 471–481.

Thomas, N. W. D. & Paré, M. (2007). Temporal processing of saccade targets in parietal cortex area lip during visual search. *J Neurophysiol*., 97, 942-947.

Thompson, K., Hanes, D., Bichot, N., & Schall, J. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field. *J Neurophysiol*., 76, 440–455.

Trappenberg, T. P. (2003). Why is our capacity of working memory so large? *Neural Information Processing - Letters and Reviews*, 1, 97-101.

Trappenberg, T. P., Dorris, M. C., Munoz, D. P., & Klein, R. M. (2001). A model of saccade initiation based on the competitive integration of exogenous and endogenous signals in the superior colliculus. *J Cognitive Neurosci*., 13, 256-271.

Trappenberg, T. P. & Standage, D. I. (2005). Multi-packet regions in stabilized continuous attractor networks. *Neurocomputing*, 65–66, 617–622.

Treisman, A. M. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychol*., 12, 97–136.

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

Usher M. & McClelland J. L. On the time course of perceptual choice: The leaky competing accumulator model. *Psychol Rev*. 108: 550–592, 2001.

Wang X.-J. (1999). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J Neurosci*., 19, 9587–9603.

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci*., 24, 455–463.

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36, 955–968.

Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, 60, 215–234.

Wang, H., Stradtman III, G. G., Wang, X.-J, Gao, W.-J. (2008). A specialized NMDA receptor function in layer 5 recurrent microcircuitry of the adult rat prefrontal cortex. *Proc Natl Acad Sci USA*., 105, 16791–16796.

White, E. L. (1989). *Cortical circuits: Synaptic organization of the cerebral cortex; structure, function and theory*. Boston: Birkhauser.

Wilson, H. R. & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13, 55–80.

Winograd, M., Destexhe, A., & Sanchez-Vives, M. V. (2008). Hyperpolarization-activated graded persistent activity in the prefrontal cortex. *Proc Natl Acad Sci USA*., 105, 7298–7203.

Wong, K.-F., Huk, A. C, Shadlen, M., & Wang, X.-J. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Frontiers in Computational Neuroscience*, 1, 1–11.

Wong, K.-F. & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci*., 26, 1314–1328.

**Supplementary Material**

***The Network Time Constant of Integration***

The decision making ability of the network is governed by its time constant of integration. As presented by Wang (2002), the difference $\Delta e$ between target and distractor input can be described by

$$d\Delta e/dt = -\Delta e/\tau_e + \mu_t - \mu_d + \eta_e \quad (3)$$

where $\mu_t$ and $\mu_d$ are the average firing rates of the target and distractor inputs, $\eta_e$ is noise, and $\tau_e$ determines the time over which the evidence can be integrated. Neural and behavioural data indicate that $\tau_e$ must be hundreds of milliseconds at the very least, much longer than the time constants of contributing biophysical processes, *e.g.* the time constants of membranes and synapses. Such long integration times can be furnished by intrinsic (recurrent) activation, so Equation 3 becomes

$$d\Delta e/dt = (-\Delta e + w_{syn}\,\Delta e)/\,\tau_{syn} + \mu_t - \mu_d + \eta_e$$

where $\tau_{syn}$ is the synaptic time constant of the dominant receptor and $w_{syn}$ refers to the strength of synapses mediating currents via this receptor (NMDAR in this case).

Under Wang's (2002, 2008) formulation, an estimate of the network's integration time constant is provided by

$$\tau_{eff} = \tau_{syn}/(|\varepsilon - w_{syn}|)\,.$$

Here, $\varepsilon$ is a parameter-dependent conductance influencing the attractor dynamics of the network. Integration times thus increase as $w_{syn}$ approaches $\varepsilon$, decreasing above and below this threshold. In our model, $w_{syn} = \gamma_{NMDA}\, G_{NMDA,p}$ and $\varepsilon \approx 0.8\, G_{NMDA,p}$.

Accuracy (Figure6 top) is thus explained as follows. As $\gamma_{NMDA}$ is increased from 0.5 to ~0.8, $w_{syn}$ approaches $\varepsilon$, lengthening the effective integration time constant $\tau_{eff}$ and consequently leading to higher accuracy. As $\gamma_{NMDA}$ is increased beyond ~0.8, we cross the threshold $\varepsilon$ and $\tau_{eff}$ is shortened, leading to lower accuracy.

The speed of decisions (Figure 6 bottom) is similarly explained for $w_{syn} > \varepsilon$. Smaller $\tau_{eff}$ yields faster (less accurate) decisions. However, for $w_{syn} < \varepsilon$, target discrimination takes longer with smaller $\tau_{eff}$. This occurs because $\tau_{eff}$ reflects a balance between the loss of accumulated evidence by leakage and the amplification of accumulated evidence by recurrent excitation (Usher and McClelland, 2001). When $w_{syn} > \varepsilon$, amplification dominates, serving as a mechanism for selection and thus preventing further integration (Wang 2002). When $w_{syn} < \varepsilon$, the network is dominated by leakage, only making decisions on difficult tasks (high target-distractor similarity) when sufficient levels of noise favour one of the inputs. The longer the trial progresses, the more likely this is to happen. As such, the network makes fewer decisions with increasing task difficulty (Figure 6 caption).

**Supplementary References**

Usher, M. & McClelland, J. L. On the time course of perceptual choice: The leaky competing accumulator model. *Psychol Rev*., 108: 550–592, 2001.

Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36: 955–968, 2002.

Wang, X.-J. Decision Making in Recurrent Neuronal Circuits. *Neuron*, 60:215–234, 2008.