

Sustained Activities and Retrieval in a Computational Model of the Perirhinal Cortex

Julien Vitay and Fred H. Hamker

Abstract

■ The perirhinal cortex is involved not only in object recognition and novelty detection but also in multimodal integration, reward association, and visual working memory. We propose a computational model that focuses on the role of the perirhinal cortex in working memory, particularly with respect to sustained activities and memory retrieval. This model describes how different partial informations are integrated into assemblies of neurons that represent the identity of an object. Through dopaminergic modulation, the resulting clusters can

retrieve the global information with recurrent interactions between neurons. Dopamine leads to sustained activities after stimulus disappearance that form the basis of the involvement of the perirhinal cortex in visual working memory processes. The information carried by a cluster can also be retrieved by a partial thalamic or prefrontal stimulation. Thus, we suggest that areas involved in planning and memory coordination encode a pointer to access the detailed information encoded in the associative cortex such as the perirhinal cortex. ■

INTRODUCTION

The perirhinal cortex (PRh), composed of cortical areas 35 and 36, is located in the ventromedial part of the temporal lobe. It receives its major inputs from areas TE and TEO of inferotemporal cortex (IT), as well as from the entorhinal cortex (ERh), parahippocampal cortex, insular cortex, and orbito-frontal cortex (Suzuki & Amaral, 1994). As part of the medial-temporal lobe system (with the hippocampus and the ERh), its primary role is considered to be object-recognition memory, as shown by impairments in delayed matching-to-sample (DMS) or delayed nonmatching-to-sample (DNMS) tasks following PRh cooling or removal (Buffalo, Reber, & Squire, 1998; Meunier, Bachevalier, Mishkin, & Murray, 1993; Zola-Morgan, Squire, Amaral, & Suzuki, 1989; Horel, Pytko-Joiner, Voytko, & Salsbury, 1987). It is thought to be particularly involved in the representation and learning of novel objects (Pihlajamäki et al., 2003; Wan, Aggleton, & Brown, 1999; Brown & Xiang, 1998), with a greater activation for these objects than for familiar ones. Ranganath and D'Esposito (2005) suggest that novel objects do not have a strong preexisting representation in the IT, and traces of long-term memory in the PRh could be used to manipulate these objects.

Despite the huge amount of evidence for a mnemonic role of the PRh, some recent findings suggest that it is also involved in high-level perception (for a controversy, see Buckley & Gaffan, 2006; Hampton, 2005), such

as object categorization and multimodal integration, by integrating different sources of information about the identity of an object (Taylor, Moss, Stamatakis, & Tyler, 2006). The PRh indeed receives connections from the insular cortex (somatosensory information) and the dorsal bank of the *superior temporal sulcus* (vision/audition coordination), therefore, being at a central place for integrating different modalities of an object. Interestingly, monkeys with lesions of the PRh are unable to select a visible object first sampled by touch (Goulet & Murray, 2001) or by a partial view of that object (Murray, Gaffan, & Mishkin, 1993).

Accordingly, the PRh is neither a purely mnemonic nor a perceptual area: It is a multimodal area which is presumably involved in the goal-directed guidance of perception. This link to the goals of the task at hand is reflected by the modulation of PRh activity by reward association (Mogami & Tanaka, 2006), which strongly depends on D2 dopamine (DA) receptors (Liu et al., 2004). Also, the PRh is involved in visual working memory, which is known to use integrated representations of objects rather than individual features (Lee & Chun, 2001; Luck & Vogel, 1997). Hölscher and Rolls (2002) showed that PRh cells are more active during a DMS task when their preferred stimulus is the sample (the object to be remembered) than when it is the match (the target), and that this property is actively reset between trials, supporting the evidence of a higher cognitive involvement. Some PRh cells also exhibit sustained activity between sample and match: Their proportion has been estimated to be 35% compared to 22% in the IT or 71% in the ERh (Naya, Yoshida, Takeda, Fujimichi,

Wilhelms-Universität Münster, Germany

& Miyashita, 2003; Nakamura & Kubota, 1995). However, contrary to the ERh, these sustained activities are not robust to the presentation of distractors between sample and match (Suzuki, Miller, & Desimone, 1997; Miller, Li, & Desimone, 1993). The exact mechanism and purpose of these sustained activities are still unknown. Are they only provoked by feedback connections from the prefrontal cortex, where sustained activities are robust to distractors (Miller, Erickson, & Desimone, 1996), or does the prefrontal cortex just control the maintenance or suppression of these sustained representations that are created with intrinsic mechanisms in the PRh?

This article presents a computational model of the PRh focused on the involvement of this cortical area in visual working memory processes by emphasizing the effect of DA modulation on perirhinal cell activation. Our aim is neither to model every aspect of PRh functioning nor to explore the biophysical properties of sustained activation. We rather propose a new interpretation at the functional level of these sustained activities in the framework of multimodal object identification or categorization. The model demonstrates how different aspects of an object or a category are linked into a neural assembly according to their co-occurrence through time and how this assembly can be reactivated for memory retrieval.

METHODS

Context

There are only few computational models of the PRh. One of the most famous is the *perceptual-mnemonic feature conjunction* (PMFC) model by Cowell, Bussey, and Saksida (2006) and Bussey and Saksida (2002). As its name indicates, it is primarily concerned with the interplay of perceptual and mnemonic processes in the PRh. The PRh is represented by a feature-conjunction layer that integrates individual features and learns to represent effectively objects in concurrent discrimination or configural learning tasks. Learning occurs either through a Rescorla–Wagner rule (Bussey & Saksida, 2002) or through self-association in Kohonen maps (Cowell et al., 2006). Despite its good predictions about the effects of PRh lesions on discrimination and configural learning tasks, it is a purely static model that cannot deal with sustained activities. The model by Bogacz, Brown, and Giraud-Carrier (2001) is much more detailed and dynamic (spiking neurons) but only deals with familiarity discrimination: Its Hopfield-like structure allows it to tell rapidly if an object has already been seen, but it does not allow to recollect its details. It is a purely mnemonic view of the PRh. The model we propose is original with regard to the functions it describes (autoassociative memory, sustained activation, memory retrieval) and its dynamical structure.

Architecture of the Model

To keep the model as simple as possible, we do not consider the precise timing of spikes but use mean-rate artificial neurons whose activity is ruled by a dynamical differential equation. This positive scalar activity represents the instantaneous firing rate, which is directly derived through a transfer function from the membrane potential, without using a spike-generation mechanism. As a consequence, the neurons used in this model exchange only this time-varying scalar activity through their connections, similar to dynamical neural fields (Taylor, 1999; Amari, 1977).

The neural network (Figure 1A) is composed of a population of excitatory pyramidal cells interconnected with a population of inhibitory interneurons. In order to reflect approximately the relative number of GABAergic interneurons in the cerebral cortex, the excitatory population is four times bigger than the inhibitory one (Beaulieu, 1993). Each inhibitory cell receives excitatory inputs from a subset of excitatory cells, with a Gaussian connectivity kernel centered on the corresponding neural location. Reciprocally, each excitatory cell receives connections from a subset of inhibitory cells with a broader Gaussian connectivity kernel. Additionally, inhibitory cells are reciprocally connected with each other in an all-to-all manner, with the connection strength decreasing with the distance between cells. Excitatory cells are also reciprocally connected in an all-to-all manner, but the strength of these connections is modifiable with experience.

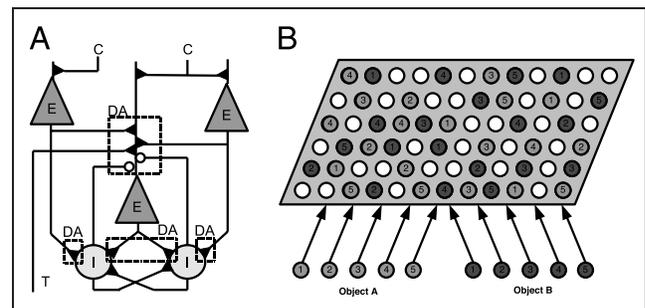


Figure 1. (A) Architecture of the model. It is composed of $n \times n$ excitatory cells (E) and $\frac{n}{2} \times \frac{n}{2}$ inhibitory cells (I). Excitatory and inhibitory cells are reciprocally connected through Gaussian connectivity kernels. Inhibitory cells are also reciprocally connected with each other with a strength decreasing with the distance. Excitatory cells are reciprocally connected with each other, but the strength of the connections is learned. Each excitatory cell receives a cortical input C from other areas. Additionally, some excitatory cells receive a thalamic input T. All connections, except the cortical ones, are modulated by DA (hatched squares). (B) Feedforward connectivity for excitatory cells. Two different objects have to be learned by the model: Object A (light gray) and Object B (dark gray, hatched) are each represented by five parts (numbered from 1 to 5), corresponding to different views or modalities. Each part is represented by a cortical input to four cells, what makes each object being represented by a cluster of 20 cells.

Each excitatory cell receives a cortical input that could originate in a visual area such as TE or in the multimodal parahippocampal cortex. Erickson, Jagadeesh, and Desimone (2000) showed that neighboring cells in the PRh tend to represent the same objects after visual experience. This finding could be explained by a self-organization of receptive fields, that is, the modification of feedforward connections. Our model does not include this feedforward learning but is rather designed to show how the gathering of these different informations can occur in the PRh. The cortical input to a cell will therefore be a time-varying scalar value, reflecting the weighted sum of the activity of its afferent cells, without any information about its origin. The basic idea of the model is that the perirhinal neurons representing a given object or category have receptive fields selective for a particular aspect of that object or category, either in visual space (different views of an object or different exemplars of a category sharing some visual features) or in multimodal space (some neurons are preferentially activated by the sound associated to this object, or its touch). In the following, we will not distinguish between the learning of different views or modalities of an object, or the learning of a category represented by different exemplars: The mechanism remains the same and we will use the term “object” for either a real object or a category. The increase in the strength of the lateral reciprocal connections between excitatory cells will provoke a clustering effect: The representation of an object will be distributed over several cells (forming what is called a *cluster* or an autoassociative pattern), which are individually selective for a particular aspect.

In our simulations, an object is represented by five parts corresponding each to a particular aspect. Each part provides a cortical input to four excitatory cells in the PRh (randomly chosen in the population), meaning that the representation of all aspects of an object forms a cluster of 20 neurons (Figure 1B). During learning, each object will be successively presented during a certain amount of time (250 msec here), but each of its parts will be randomly active with a probability of .6. The random activation of parts means that each presentation of an object will be incomplete in most cases. The goal of the learning in the lateral connections will be to correlate the different parts, even if they do not constantly appear together. Unless stated otherwise, all the simulations have been done with two different objects.

Dopamine Modulation

DA modulation is a very important feature of the model, responsible for most of its interesting properties. Unfortunately, little is known about its effects in the PRh. We will therefore assume that DA modulation in the PRh is similar to what occurs in the prefrontal cortex, given the fact that the PRh has a similar ratio of D1/D2 receptors, even if their density is higher (Hurd, Suzuki, & Sedvall,

2001). An exhaustive review about DA effects on prefrontal cells can be found in Seamans and Yang (2004). The picture that emerges from experimental observations is very heterogeneous. However, there is some accumulating evidence for the following properties:

- The effect of DA is strictly modulatory: It does not induce excitatory postsynaptic currents by itself (Yang & Seamans, 1996);
- DA modulates both pyramidal and fast-spiking inhibitory interneurons (Gorelova, Seamans, & Yang, 2002);
- DA modifies the cell’s excitability by modulating intrinsic ionic currents such as Na^+ and K^+ (Yang & Seamans, 1996);
- The effect of DA is dose-dependent: D1 receptor activation can have opposing functional effects depending on the level of stimulation, following an inverted U-shape (Goldman-Rakic, Muly, & Williams, 2000);
- The effect of DA is neurotransmitter receptor-dependent: *N*-methyl-D-aspartic acid [NMDA] (excitatory activity-dependent)- and gamma-aminobutyric acid [GABA] (inhibitory)-mediated currents are enhanced by DA, but alpha-amino-3-hydroxy-5-methyl-4-isoxazole propionic acid [AMPA] (excitatory)-mediated ones are decreased (Momiya, Sim, & Brown, 1996; Cepeda, Radisavljevic, Peacock, Levine, & Buchwald, 1992);
- The effect of DA is dendrite-dependent: DA reduces more strongly the excitatory postsynaptic potentials (EPSPs) generated in apical dendrites (long-distance cortical inputs) than in the basal ones (neighboring pyramidal cells), through a reduction of dendritic Ca^{2+} currents (Zahrt, Taylor, Mathew, & Arnsten, 1997; Yang & Seamans, 1996);
- The effect of DA is activity-dependent: The more the cell is active, the more DA modulates its inputs (Calabresi, Mercuri, Stanzione, Stefani, & Bernardi, 1987);
- DA levels are long-lasting in the target area (Huang & Kandel, 1995). The phasic DA bursts in the dopaminergic cells are therefore not relevant: We will only consider the tonic component of DA activity, not its phasic component.

Existing models of dopaminergic modulation of sustained activities in the prefrontal cortex do not all make the same hypothesis about the exact influence of DA. A detailed model by Durstewitz, Kelc, and Güntürkün (1999) supposes that DA enhances the persistent Na^+ ionic currents, reduces the slowly inactivating K^+ ionic currents, reduces the efficiency of apical inputs, reduces the amplitude of glutamate-induced EPSPs (including NMDA, even if they admit this is controversial) and increases the spontaneous activity of GABAergic cells as well as the amplitude of inhibitory postsynaptic potentials (IPSPs) in pyramidal cells. In their respective models, Brunel and Wang (2001), as well as Deco and Rolls (2003), suppose that DA only enhances NMDA-mediated currents in the basal dendrites in coordination

with a simultaneous increase of the amplitude of IPSPs. On the contrary, Dreher, Guigon, and Burnod (2002) consider that DA momentarily restricts excitatory inputs on apical dendrites. More recently, Chadderdon and Sporns (2006) considered that DA only modifies the gain of cells by increasing their firing threshold, without being more specific about synaptic currents.

The major link between most of these models is that they distinguish the effects of DA on apical dendrites and on basal dendrites of pyramidal cells: The influence of long-distance cortical inputs is reduced by DA, whereas the influence of neighboring pyramidal cells is increased. This last assumption is coherent with the fact that basal dendrites are primarily NMDA-mediated (Schiller, Major, Koester, & Schiller, 2000). The reduction of apical currents allows the network to be momentarily insensitive to external inputs, increasing the robustness of sustained activities when they appear. In the case of the PRh, as we know that sustained activities are not robust to the appearance of distractors (Miller, Gochin, & Gross, 1993), we neglected this effect. Accordingly, the major influences of DA we consider in our model are therefore the increase of the efficiency of lateral connections between excitatory cells (on an activity-dependent manner, as they are mainly mediated by NMDA receptors), the increase of the amplitude of IPSPs (by increasing the efficiency of the connections from inhibitory to excitatory cells), and the increase of the activity of the inhibitory cells through an increase in the efficiency of the connections from excitatory to inhibitory cells. These assumptions are summarized in Figure 1A. The modification of the excitability of cells through modulation of ionic currents has not been taken into account because the effects of this mechanism are thought to be similar to the selective modulation of synaptic currents. The differential effects of D1-like and D2-like receptors have not been considered because there exists no sufficient experimental evidence to draw a precise line between them.

Equations for Updating the Activity

The model consists of a single map of $n \times n$ excitatory units and $\frac{n}{2} \times \frac{n}{2}$ inhibitory units. We use $n = 20$ for the results in this article, but the properties of the model do not depend on this particular size: It has been tested from $n = 10$ to $n = 40$, showing that distributed computations and flexible learning can induce scalability. We used a mean-field approach, where the activity of each unit follows an ordinary differential equation, discretized with a time step of 1 msec. In the mean-field approach, a unit represents a population average of a certain number of single cells. Because the true underlying circuitry is not well known, we do not explicitly derive the mean-field solution but describe the dynamics at the macroscopic population level. Nevertheless, for the sake of

simplicity, we use the term “cell” for a unit. The mean activity $I_i(t)$ of an inhibitory cell at time t is ruled by Equation 1:

$$\tau_I \cdot \frac{dI_i(t)}{dt} + I_i(t) = \sum_{j \neq i} W_{ij}^{II} \cdot I_j(t) + (1 + K^{EI} \cdot \text{DA}) \times \sum_k W_{ik}^{EI} \cdot E_k(t) + \eta_i^I(t) \quad (1)$$

where $\tau_I = 10$ msec is the net time constant of the unit. W^{II} is the set of connections between inhibitory cells, decreasing with the distance between the cells and W^{EI} is the set of connections from the excitatory cells (activity denoted $E_k(t)$) to the inhibitory cell (formulas given in the Appendix). The DA level in the network (represented by the scalar value DA between 0 and 1) increases the gain of inputs from excitatory cells. K^{EI} is a fixed scaling parameter. Finally, $\eta^I(t)$ is a noise added to the cell that randomly fluctuates in the range $[-0.1, 0.1]$. The resulting activity is restricted to positive values.

The mean activity $E_i(t)$ of an excitatory cell at time t is ruled by Equation 2:

$$\tau_E \cdot \frac{dE_i(t)}{dt} + E_i(t) = f((1 + K^{EE} \cdot \sigma^{\text{lat}}(\text{DA}) \cdot \sigma^{\text{EE}}(E_i(t))) \cdot \sum_{j \neq i} W_{ij}^{\text{EE}} \cdot E_j(t) + (1 + K^{\text{IE}} \cdot \sigma^{\text{GABA}}(\text{DA}) \cdot E_i^2(t)) \cdot \sum_k W_{ik}^{\text{IE}} \cdot I_k(t) + W_i^{\text{C}} \cdot C_i(t) + (1 + K^{\text{T}} \cdot \sigma^{\text{T}}(\text{DA})) \cdot T_i(t) + \eta_i^{\text{E}}(t)) \quad (2)$$

where $\tau_E = 20$ msec is the net time constant of the unit. This value is chosen twice as large as in the inhibitory units to reflect the ratio of membrane time constants between pyramidal cells and inhibitory interneurons in the cortex (McCormick, Connors, Lighthall, & Prince, 1985). $f(x)$ is a transfer function, ensuring that the activity of the cell does not reach too high values. It is linear in the range $[0, 1]$ and then saturates slowly to a maximum value of 1.5 (formula given in the Appendix). There are five terms inside this transfer function. The first term denotes the influence of the lateral connections between excitatory cells W^{EE} . Its gain depends on DA through a sigmoidal term σ^{lat} and a fixed scaling parameter K^{EE} , but also on the activity of the cell itself through another sigmoidal function σ^{EE} . For these predominantly NMDA-mediated lateral connections, the influence of DA is therefore activity-dependent. These two sigmoids are independent to ensure that DA only modulates active cells and that effective transmission of activity through NMDA-mediated connections between excitatory cells only occurs in the presence of DA. The

second term represents the influence of the connections from the inhibitory cells with a negative strength W^{IE} . Their efficiency also increases with DA (sigmoidal function σ^{GABA} and fixed scaling parameter K^{IE}) and the activity of the cell. The feedforward inhibition produced by the increase of the efficiency of IPSPs by high levels of DA on pyramidal cells, as proposed by Goldman-Rakic et al. (2000), is realized through a square of the activity of the cell itself. The third term is the contribution of the cortical input $C_i(t)$ through a random weight W_i^{C} , without any dopaminergic modulation as they are considered to reach apical dendrites (see the Dopamine Modulation section). When the cell is stimulated, we set $C_i(t) = 1.0$. The fourth term is the contribution of a possible thalamic input $T_i(t)$, increased by DA through σ^{T} and the scaling parameter K^{T} . This term is clearly distinct from the cortical inputs: Although the PRh is dysgranular—with a very thin layer IV (Rempel-Clower & Barbas, 2000)—thalamocortical afferents from the dorsal and medial geniculate nuclei target layers I, III/IV, and VI (Furtak, Wei, Agster, & Burwell, 2007; Linke & Schwegler, 2000), therefore on both apical and basal dendrites of pyramidal cells, as well as on various interneurons. We therefore assume that the thalamic input has a driving force through apical dendrites, similar to the cortical input, and a dependence on DA through the basal dendrites. The last term $\eta^{\text{E}}(t)$ is a noise randomly fluctuating in $[-0.5, 0.5]$. The resulting activity is restricted to positive values (details about the sigmoidal functions and other parameters are given in the Appendix).

Although the general properties of DA modulation are largely supported by the discussed observations, the exact parameters and sigmoid functions have been determined through trial-and-error processes to enable sustained activities. Although the results we present here quantitatively depend on these choices, the global properties we intend to highlight admit some variations in the values of the parameters.

Learning Rule

The lateral reciprocal connections between excitatory cells W^{EE} are subject to learning. We considered a covariance rule combining input- and output-dependent long-term potentiation and an output-dependent only long-term depression:

$$\tau_{\text{W}} \cdot \frac{dW_{ij}^{\text{EE}}(t)}{dt} = (E_i(t) - \hat{E}_i(t))^+ \cdot \left((E_j(t) - \hat{E}_j(t))^+ - \alpha_i(t) \cdot W_{ij}^{\text{EE}}(t) \cdot (E_i(t) - \hat{E}_i(t))^+ \right) \quad (3)$$

where $E_i(t)$ is the presynaptic activity of cell i , $E_j(t)$ the postsynaptic activity of cell j . $()^+$ is the positive part

function. $\hat{E}_k(t)$ is a temporal sliding-mean of the activity $E_k(t)$ over a window of T msec defined by:

$$\hat{E}_k(t) = \frac{(T-1) \cdot \hat{E}_k(t-1) + E_k(t)}{T} \quad (4)$$

with $T = 5000$ msec in this model. This term ensures that learning occurs only when presynaptic or postsynaptic activities are significantly higher than their baseline value, ruling out learning of noise. However, the final weights determined by this rule alone are strongly dependent on the value of the parameter α_i , which is constant in classical covariance rules. If α_i is set too high, weights will never increase enough to produce postsynaptic activity, but if α_i is too low, the postsynaptic cell will have maximal activity for a too large set of stimuli. As we want our model to deal with different cluster sizes, we had to use a more flexible approach for the learning rule. We therefore focused on homeostatic learning, where the learning rule uses as a constraint that the activity of a cell should not exceed a certain value, in order to save energy (Turrigiano & Nelson, 2004; van Rossum & Turrigiano, 2001). Homeostatic learning is possible when the parameter α_i can vary with the experience of the cell, in our case, when the cell's activity exceeds a certain threshold. The following rule is studied in Hamker and Wiltshut (in preparation):

$$\tau_{\alpha} \cdot \frac{d\alpha_i(t)}{dt} + \alpha_i(t) = K_{\alpha} \cdot H_i(t) \quad (5)$$

$$\tau_{\text{H}} \cdot \frac{dH_i(t)}{dt} + H_i(t) = K_{\text{H}} \cdot ((E_i(t) - E_{\text{max}})^+)^2 \quad (6)$$

with $H_i(t)$ and $\alpha_i(t)$ restricted to positive values and $\alpha_i(0)$ equal to 10.

When $E_i(t)$ exceeds E_{max} (1.0 in our model), $H_i(t)$ becomes rapidly highly positive, leading to a slow increase of $\alpha_i(t)$. The inhibitory part of Equation 3 becomes preponderant and all the weights decrease. The reason why $H_i(t)$ is introduced is that $\alpha_i(t)$ must have a slow time constant so that learning is stable. This learning rule is similar to the classical BCM rule (Bienenstock, Cooper, & Munroe, 1982) but is more stable, as the inhibitory term in Equation 3 represents a constraint both on a short time scale—by its dependence on $E_i(t)$ and $W_{ij}^{\text{EE}}(t)$ —and on a long time scale with $\alpha_i(t)$. The effect of this learning rule is that weights will rapidly increase at the beginning of learning (the Hebbian part of Equation 3 is preponderant), but when the cells begin to overshoot, $\alpha_i(t)$ increases and forces the cell to find a compromise between increasing its afferent weights and activity overshooting. When learning is efficient, $\alpha_i(t)$ stabilizes to an optimal value that depends on the mean activity of the cell.

RESULTS

We will first show the consequence of learning the lateral connections between excitatory cells on the formation of clusters and the propagation of activity within the cluster. We then demonstrate the effect of DA modulation on sustained activities in the network and show that the model follows the classical inverted U-shaped curve. After introducing these basic properties, we then demonstrate the specific properties for memory recall such as the dependence of the propagation of activity between two clusters on the strength of their reciprocal connections, as well as the effect of thalamic stimulation on memory retrieval.

Learning and Propagation of Activity within a Cluster

During learning, a sequence of stimuli is shown to the network. The first object is presented for 250 msec, activating a random number of parts of the corresponding cluster. No stimulation is given to the network for the next 250 msec, followed by the second object for 250 msec, and further on. This sequence is repeated for 100 times. Please note that this is one particular learning protocol, but that other protocols ensuring that each object is sufficiently often presented also work. The DA level is set to a low value of 0.1 during learning, for reasons explained in the Discussion.

After learning, each cell has built connections with the cells representing other parts of an object. Figure 2A shows the 25 highest connection values for a randomly selected cell in the first cluster. One can observe that this cell has formed positive connections with the 19 other cells of the cluster. The weights within a cluster are not all equal, reflecting the probability of co-occurrence of the different parts during learning. Oppositely, the con-

nections with cells of another cluster have been reduced to negligible values.

After learning, how do we functionally retrieve the information about the correlation between different parts? Our hypothesis is that the activation of a sufficient number of parts should provoke activity in the remaining parts, at least under certain DA levels. Figure 2B shows the mean activity of the remaining parts dependent on the number of parts that receive cortical activation. When DA has too low (0.2) or high (0.8) levels, the remaining parts show only little activation, even if four out of five parts are stimulated. When DA has an intermediate level (0.4 or 0.6) and three or more parts are activated, the remaining parts show strong activity, as if they actually received cortical input. This shows that under intermediate DA levels, the network is able to retrieve all the parts of a cluster if a majority of them is stimulated. We also simulated clusters of bigger size (up to 20 parts of four cells, i.e., 80 cells) and observed that this minimum proportion of stimulated parts is slightly decreasing with the cluster size, but it is always superior to one third.

Sustained Activities and Intermediate Values of Dopamine

In the following experiments, we stimulate only three parts of a cluster (12 cells out of 20) and record two different neurons, one belonging to these three parts and called the “stimulated” cell, the other to one of the two remaining parts and called the “unstimulated” cell.

To determine the adequate range of DA levels, it is interesting to look at the sustained activities observable in the network. Figure 3A shows the time course of the activity of two cells during the successive presentation of the two objects. With a low DA level (0.1), only the stimulated cell shows significant activity (around 1.0) during

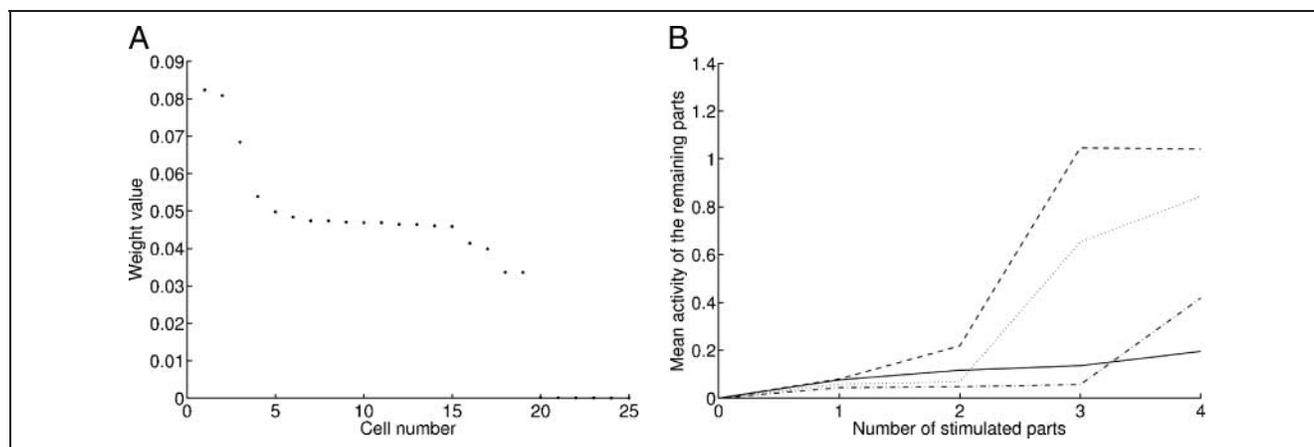


Figure 2. (A) Weight values for a given cell in the first cluster. Only the 25 highest values are represented in descending order. We observe that this cell has positive connections with the 19 cells that form the cluster and none with other cells. (B) Mean activity of unstimulated parts relative to the number of stimulated parts. We observe that for low (0.2) or high (0.8) DA levels, the remaining parts are only poorly activated. For intermediate levels (0.4 or 0.6), three stimulated parts are sufficient to provoke a high activity in the remaining two unstimulated parts.

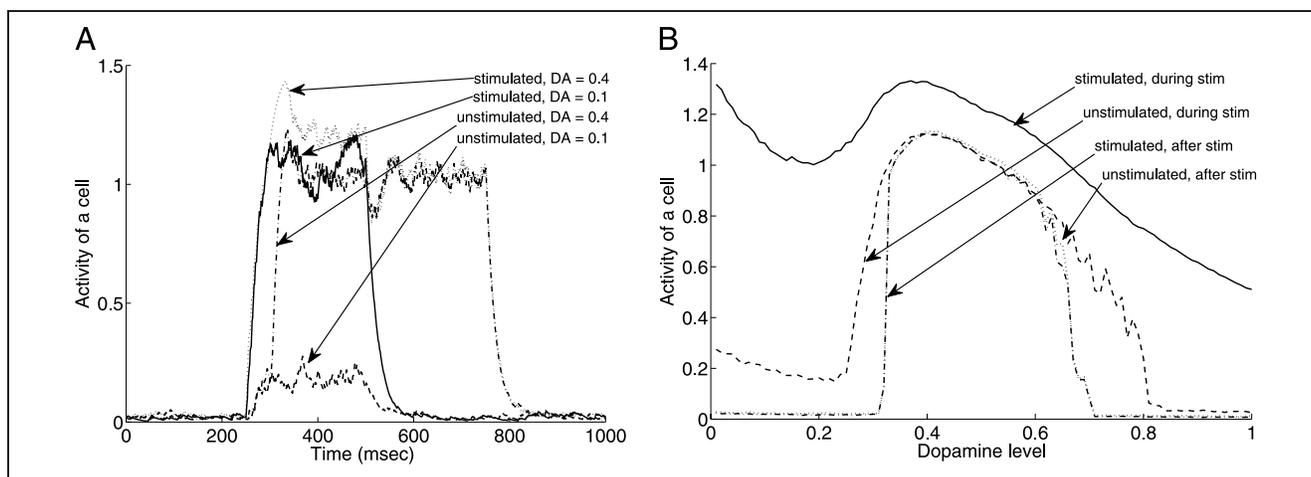


Figure 3. (A) Time course of the activity of two different cells in the same cluster. The first one (“stimulated cell”) belongs to one of the three parts that receive cortical input, the other one (“unstimulated cell”) receiving no cortical input. When the DA level is low (DA = 0.1), the stimulated cell responds strongly to the presentation of the object but not the unstimulated one. When the stimulation ends, the activity of these two cells returns to baseline. When the DA level is intermediate (DA = 0.4), the two cells respond equally strong to the presentation of the object. After disappearance, they show sustained activity until a new object is presented. (B) Effect of DA on two cells in the same cluster. The two upper curves represent the activity of the stimulated and unstimulated cells during stimulation, 200 msec after the corresponding object onset. With intermediate levels of DA, the activity of the unstimulated cell is high and only slightly inferior to the stimulated one (difference of 0.2). With large DA levels (>0.6), the activity of the two cells is drastically reduced because of the enhancement of inhibition by DA. The two lower curves (which seem identical) represent the activity of these two cells 100 msec after the end of the stimulation. We observe an inverted U-shape meaning that the level of DA necessary to observe sustained activities is between 0.3 and 0.7.

the presentation of the object. With an intermediate DA level (0.4), both cells become highly active (around 1.2 and 1.0, respectively) during the stimulation, with a little time lag due to the propagation of activity within the cluster. When the stimulation ends, their activity does not fall back to baseline but stays at a high level (1.0). This sustained activity is only due to the reciprocal interactions between excitatory cells and their modulation by DA.

When the second object is presented, its representation competes with the sustained activation. If the two representations are equally distributed on the map, which is the case here, some of their excitatory cells will be connected to the same inhibitory cells, leading to enhanced inhibition and disruption of the sustained activities. If the two representations are spatially segregated on the map (corresponding, for example, to two objects from very different categories, like a face and a tree), the two representations can exist in parallel. Data from Miller, Li, et al. (1993) about the robustness of sustained activities in the PRh does not deal with the distribution of competing stimuli on the surface of the cortex, allowing this property to be a prediction of the model. However, if the distracting stimulus has a low intensity [$C_i(t) < 0.4$] or is not represented by more than two parts, the sustained representation can resist its appearance, thanks to the increased activity of inhibitory cells.

Figure 3B shows the influence of the DA level on the activities of the two considered cells during and after stimulation. When the cluster is partly stimulated, DA globally enhances the activity of the stimulated cell when DA is inferior to 0.4 but then begins to depress it. For the unstimulated cell, one can observe a strong

enhancing effect when DA is around 0.25 due to the propagation of activity within the cluster. When DA exceeds 0.8, the activity of this cell falls abruptly to zero, showing that propagation of activity is not possible under high levels of DA, because of the enhancement of the reciprocal connections between inhibitory and excitatory cells. The two lower curves of Figure 3B show the sustained activity of the two cells 100 msec after the end of the stimulation. They have an inverted U-shape, which is typical for dopaminergic modulation of working memory in the prefrontal cortex (Goldman-Rakic et al., 2000). The graph shows that the values of DA in our model that allow to observe sustained activities range between 0.3 and 0.7. The amplitude of the sustained activities is relatively high (up to 80% of the activity during stimulation depending on the DA level) but is coherent with cellular recordings (Curtis & D’Esposito, 2003; Naya et al., 2003; Ohbayashi, Ohki, & Miyashita, 2003). Due to the balanced background inhibition, we can also change the parameters of the model to obtain lower sustained activities.

Propagation of Activity between Clusters

The propagation of activity within a cluster is an interesting property in the framework of multimodal object categorization and identification. However, contrary to the preceding experiments, where the two learned objects do not share any parts, learning in the real world does not ensure that parts of two different objects are not activated at the same time in the PRh, for example, because these objects share these parts. Consequently, the weights

between two clusters are not necessarily equal to zero. What happens to the propagation of activity if two clusters are reciprocally connected with small weight values?

Figure 4 shows the influence of these intercluster connections. After the two clusters have been learned, we artificially increase the strength of connection between the two groups of cells. As each cell does not receive the same amount of cortical input because of the random weights W_i^C , their lateral connections W_{ij}^{EE} are not equal. We therefore computed the mean value of these lateral connections for each cell of the second cluster (called the intracenter connection value) and set the connections from the first cluster to the corresponding cell in the second cluster proportional to this value (intercluster connection value).

We then stimulate three parts of the first cluster and record the mean activity of the second cluster. Under low or high DA levels, intercluster connections can be equal to the intracenter connections (meaning that they form one bigger cluster) without observing any propagation of activity to the second cluster. Under intermediate DA levels, the ratio between these connections must be below 40% to avoid that the activation of one cluster propagates without control to other weakly connected clusters. This result ensures a reasonable tradeoff between stability of object representation and propagation of activity.

Thalamic Stimulation

The preceding results show that our model is able to learn to correlate different parts of an object through lateral connections and to propagate activity between these parts under intermediate DA levels. It also exhibits

sustained activity after an object is presented, but which is easily disrupted by similar distractors. What can be the interest of such unrobust sustained activities in the more general framework of visual working memory? Our conviction is that this high-level representation of an object does not need to be actively maintained through time but only regenerated when needed. A cluster describes quite exhaustively the different aspects of an object: What needs to be remembered is more the location of the cluster in the PRh than the details of its representation. Propagation of activity within a cluster seems a useful mechanism in the sense that external activation of parts of a cluster can be sufficient under intermediate DA levels to retrieve the whole information carried by the cluster. This external activation can take its origins either from the prefrontal cortex or from the basal ganglia—through the dorsal nucleus of the thalamus—where sustained activities are robust.

Figure 5 shows the influence of partial thalamic stimulation of the cells of a cluster. For this experiment, the network learned simultaneously four clusters of different sizes: 12 cells (3 parts), 20 cells (5 parts), 28 cells (7 parts), and 36 cells (9 parts). A learning cycle (the successive presentation of the four partially stimulated objects) is therefore two times longer (2 sec) and learning is stopped after 200 cycles. For each cluster, we feed a certain percentage of cells with thalamic input ($T_i = 1.0$) and we record the mean activity of the remaining cells. Using an intermediate DA level (0.5), one can observe that, for the cluster of 12 cells, a thalamic stimulation of at least 35% of its cells is sufficient to propagate activity in the cluster. This proportion is even smaller with clusters of bigger sizes. This property allows the *retrieval* of the encoded information in the cluster without

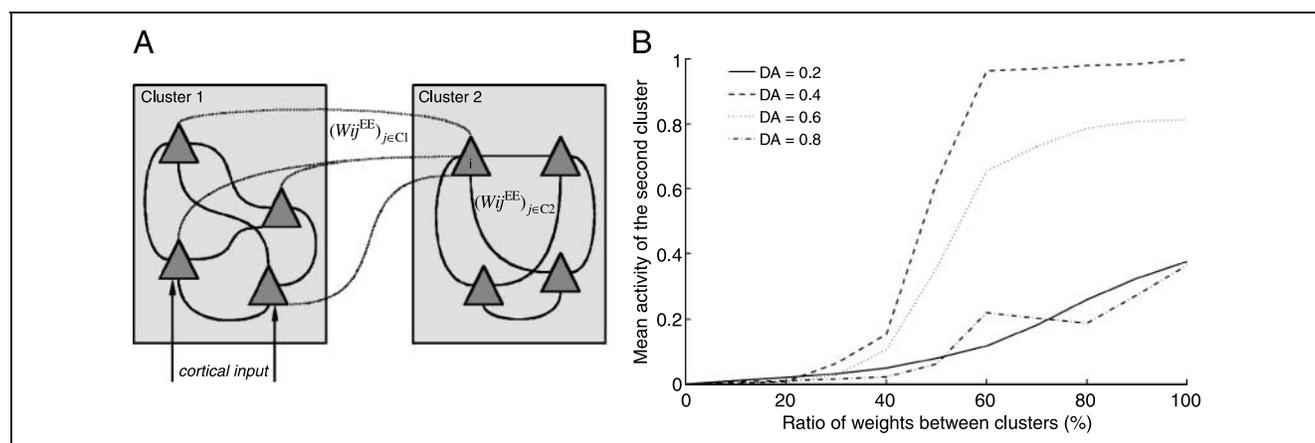


Figure 4. (A) Influence of the connections between different clusters on the propagation of activity. For simplicity, only four excitatory cells by cluster and just a few connections are shown on the figure. Two clusters C1 and C2 are learned. Each excitatory cell i of the cluster C2 receives connections $(W_{ij}^{EE})_{j \in C1}$ from excitatory cells of the cluster C1, but they are very low after learning. In this experiment, the weights of these intercluster connections are artificially set proportional to the mean value of the intracenter connections to the corresponding cell in the second cluster $W_i^{\text{mean}} = \frac{1}{n} \times \sum_{j \in C2} W_{ij}^{EE}$. (B) Results. A cluster is then stimulated and we plot the mean activity of the second cluster after 200 msec. When DA is low (0.2) or high (0.8), the second cluster becomes only poorly activated by the first cluster, even when the connections have equal strengths. When DA is intermediate, the intercluster weights must be below 40% of the intracenter weights to avoid the propagation of activity.

connections (which increases faster than the dopaminergic modulation term of excitatory connections). If the DA level is too high during learning, the afferent weights cannot increase enough because the homeostatic rule impairs learning when the activity of the cell exceeds a threshold. Thus, the lateral connections will not compensate the disappearance of the cortical input: There will be no sustained activity. However, they remain strong enough to propagate activity within the cluster. Therefore, this model cannot handle high constant levels of DA during the whole learning process (that would be, however, unrealistic), but only some increases to high levels for a finite period of time. These transient increases (which are not, however, phasic bursts) could momentarily signal the behavioral importance of certain objects and favorize their learning, but on the long-term, DA should show habituation to these objects.

The sustained activation in this model relies on the reciprocal interactions between excitatory cells. This concept has already been used in the previously cited computational models of working memory in the prefrontal cortex (Chadderdon & Sporns, 2006; Deco & Rolls, 2003; Dreher et al., 2002; Brunel & Wang, 2001; Durstewitz et al., 1999). The major differences with most of these models is that, in our model, these lateral connections are primarily relevant for memory recall and that they adapt to the experience of the system so that the attractors of the network can evolve through time. Another remarkable property is that the cells of a cluster do not need to receive input at the same time: A partial activation is enough to propagate activity and to create sustained activities in the whole cluster. It could be possible that the sustained activities in the PRh have no direct purpose but they occur as a side effect of the propagation of activity for memory retrieval.

What do the clusters of cells in the PRh exactly represent? We used the term “object” in a very broad sense, as a collection of parts that frequently appear together during learning. This could relate to spatial arrangements of parts of an object (the back, the seat, and the feet of a chair, for example) that do not all appear at the same time depending on the point of view to the object, but partly view-invariant cells are already present in the IT (Booth & Rolls, 1998). However, when the PRh is functional, learning to discriminate a set of visual objects under a certain viewpoint can be easily transferred to the same objects under another viewpoint, whereas this capacity is severely impaired without the PRh (Buckley & Gaffan, 1998). Another level of abstraction for the PRh is multimodal integration, that is, linking the visual representation of an object with its tactile information, its sound, or the associated action (grasping, pushing, sitting, etc.).

A cluster could also represent a subordinate-level category in the sense of Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976): Different objects sharing a sufficient number of sensory features (parts) would be rep-

resented by the same cluster. For example, a cluster could be generic for different espresso cups but not mugs, lacking the genericity of the “cup” basic-level category but providing a minimal sensory abstraction. This is coherent with the study by Tyler et al. (2004) that indicates that the PRh is only involved in fine-grained categorization. Such narrow categories could be used as “templates” to guide attention to the corresponding target through feedback connections to the ventral pathway (Hamker, 2005), as broader categories have been shown to be useless in visual search (Smith, Redford, Gent, & Washburn, 2005).

Our primary aim has been to extend the concept of visual working memory to association areas where the detailed visual properties of an object are stored. Most computational models of working memory make no such distinction and primarily deal with sustained activities in the prefrontal cortex. We propose that memory retrieval is achieved through a loop between the PRh, the basal ganglia, and the thalamus. The PRh receives thalamocortical connections from the dorsal and medial geniculate nuclei of the thalamus and, in turn, projects heavily to the caudate–putamen, a part of the main input structure of the basal ganglia, the striatum (Furtak et al., 2007). When a given object has to be retrieved, the basal ganglia can selectively disinhibit the thalamus and, therefore, favorize the thalamic stimulation of the cluster to be retrieved.

This pathway through the basal ganglia significantly compresses the information encoded in the cerebral cortex and cannot represent its rich and detailed representations: As Joel, Niv, and Ruppin (2002) pinpoints, the number of neurons projecting to the striatum is two orders of magnitude greater than the number of striatal neurons (Kincaid, Zheng, & Wilson, 1998). We propose that the basal ganglia acts as a pointer that allows the retrieval of the detailed representation when necessary through the disinhibition of the thalamus. Similarly, the prefrontal cortex is probably not encoding the content of memory, but rather a rule to retrieve this content. In a realistic DMS task, the basal ganglia and the prefrontal cortex have to learn which object has to be retrieved and which should be forgotten. This work is facilitated by the fact that the exact content of a cluster in the PRh does not need to be known by this external loop: Stimulating 35% of its cells (or even less for bigger clusters) is sufficient to retrieve its details.

APPENDIX: DETAILS OF THE MODEL

All equations described in the Methods section are numerized according to the finite difference method, with a time step of 1 msec. Their evaluation occurs asynchronously: Cells are randomly evaluated and their new activity is immediately used in the rest of the computations in order to emphasize the competition between neuronal representations (Rougier & Vitay, 2006).

The model is composed of 20×20 excitatory cells and 10×10 inhibitory cells. Excitatory and inhibitory cells are reciprocally connected through Gaussian connectivity kernels. We thus defined a distance between cells: Let the excitatory cell E_i have coordinates $(x_i, y_i) \in [0 \dots 20]^2$ on the map and the inhibitory cell I_j have coordinates $(x_j, y_j) \in [0 \dots 10]^2$. The distance $d_{EI}(i, j)$ between the two cells is therefore given by:

$$d_{EI}(i, j) = \sqrt{(x_i - 2 \times x_j)^2 + (y_i - 2 \times y_j)^2}$$

Similarly, the distance $d_{II}(i, j)$ between two inhibitory cells I_i with coordinates $(x_i, y_i) \in [0 \dots 10]^2$ and I_j with coordinates $(x_j, y_j) \in [0 \dots 10]^2$ is given by:

$$d_{II}(i, j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$$

We then define the Gaussian connectivity kernels by:

$$W^{IE}(i, j) = -0.12 \times \exp\left(-\left(\frac{d_{EI}(i, j)}{2.5}\right)^2\right)$$

$$W^{EI}(i, j) = 0.3 \times \exp\left(-\left(\frac{d_{EI}(i, j)}{2}\right)^2\right)$$

The connections between two inhibitory cells are given by:

$$W^{II}(i, j) = \begin{cases} 0.02 \times \exp\left(-\left(\frac{d_{II}(i, j)}{5}\right)^2\right) & \text{if } i \neq j \\ 0 & \text{else.} \end{cases}$$

The parameters of Equation 1 are the same for each inhibitory cell: $\tau_I = 10$ msec, $K^{EI} = 1.2$ and $\eta_i^I(t)$ is a random value uniformly distributed between -0.1 and 0.1 . The parameters of Equation 2 are: $\tau_E = 20$ msec, $K^{EE} = 3.0$, $K^{IE} = 3.0$, $K^T = 1.0$ and $\eta_i^E(t)$ is a random value uniformly distributed between -0.5 and 0.5 . Cortical weights W^C are randomly chosen in the range $[0.8, 1.2]$. The sigmoidal functions $\sigma^{\text{lat}}(x)$, $\sigma^{\text{EE}}(x)$, $\sigma^{\text{GABA}}(x)$, $\sigma^T(x)$ all have the same shape:

$$\sigma(x) = \frac{1}{1 + \exp(-l \cdot (x - c))} - \frac{1}{1 + \exp(l \cdot c)}$$

with l and c being: for $\sigma^{\text{lat}}(x)$, $c = 0.3$, $l = 20$; for $\sigma^{\text{EE}}(x)$, $c = 0.3$, $l = 20$; for $\sigma^{\text{GABA}}(x)$, $c = 0.5$, $l = 10$; for $\sigma^T(x)$,

$c = 0.5$, $l = 10$. The transfer function $f(x)$ is defined as follows:

$$f(x) = \begin{cases} 0 & \text{if } x < 0 \\ x & \text{if } 0 \leq x \leq 1 \\ \frac{0.5}{1 + \exp(-10.0 \cdot (x - 1))} + 0.75 & \text{if } x > 1 \end{cases}$$

The parameter of Equations 3, 5, and 6 are: $\tau_W = 50,000$ msec, $\tau_\alpha = 50,000$ msec, $K_\alpha = 100$, $\tau_H = 100$ msec, $K_H = 200$, $E_{\text{max}} = 1.0$.

Acknowledgments

This work was supported by the HA2630/4-1 grant of the German research foundation (Deutsche Forschungsgemeinschaft, DFG).

Reprint requests should be sent to Fred H. Hamker, Allgemeine Psychologie, Psychologisches Institut II, Westf. Wilhelms-Universität Münster, Fliehdnerstr. 21, Münster, 48149, Germany, or via e-mail: fhamker@uni-muenster.de.

REFERENCES

- Amari, S.-I. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Beaulieu, C. (1993). Numerical data on neocortical neurons in adult rat, with special reference to the GABA population. *Brain Research*, 609, 284–292.
- Bienenstock, E. L., Cooper, L. N., & Munroe, P. W. (1982). Theory of the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2, 32–48.
- Bogacz, R., Brown, M. W., & Giraud-Carrier, C. (2001). Model of familiarity discrimination in the perirhinal cortex. *Journal of Computational Neuroscience*, 10, 5–23.
- Booth, M. C., & Rolls, E. T. (1998). View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cerebral Cortex*, 8, 510–523.
- Brown, M. W., & Xiang, J. Z. (1998). Recognition memory: Neuronal substrates of the judgement of prior occurrence. *Progress in Neurobiology*, 55, 149–189.
- Brunel, N., & Wang, X. J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11, 63–85.
- Buckley, M. J., & Gaffan, D. (1998). Perirhinal cortex ablation impairs visual object identification. *Journal of Neuroscience*, 18, 2268–2275.
- Buckley, M. J., & Gaffan, D. (2006). Perirhinal cortical contributions to object perception. *Trends in Cognitive Sciences*, 10, 100–107.
- Buffalo, E. A., Reber, P. J., & Squire, L. R. (1998). The human perirhinal cortex and recognition memory. *Hippocampus*, 8, 330–339.
- Bussey, T. J., & Saksida, L. M. (2002). The organization of visual object representations: A connectionist model of

- effects of lesions in perirhinal cortex. *European Journal of Neuroscience*, *15*, 355–364.
- Calabresi, P., Mercuri, N., Stanzione, P., Stefani, A., & Bernardi, G. (1987). Intracellular studies on the dopamine-induced firing inhibition of neostriatal neurons in vitro: Evidence for d1 receptor involvement. *Neuroscience*, *20*, 757–771.
- Cepeda, C., Radisavljevic, Z., Peacock, W., Levine, M. S., & Buchwald, N. A. (1992). Differential modulation by dopamine of responses evoked by excitatory amino acids in human cortex. *Synapse*, *11*, 330–341.
- Chadderton, G. L., & Sporns, O. (2006). A large-scale neurocomputational model of task-oriented behavior selection and working memory in prefrontal cortex. *Journal of Cognitive Neuroscience*, *18*, 242–257.
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2006). Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *Journal of Neuroscience*, *26*, 12186–12197.
- Curtis, C., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*, 415–423.
- Deco, G., & Rolls, E. T. (2003). Attention and working memory: A dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience*, *18*, 2374–2390.
- Dreher, J.-C., Guigon, E., & Burnod, Y. (2002). A model of prefrontal cortex dopaminergic modulation during the delayed alternation task. *Journal of Cognitive Neuroscience*, *14*, 853–865.
- Durstewitz, D., Kelc, M., & Güntürkün, O. (1999). A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience*, *19*, 2807–2822.
- Erickson, C. A., Jagadeesh, B., & Desimone, R. (2000). Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. *Nature Neuroscience*, *3*, 1143–1148.
- Furtak, S. C., Wei, S.-M., Agster, K. L., & Burwell, R. D. (2007). Functional neuroanatomy of the parahippocampal region in the rat: The perirhinal and postrhinal cortices. *Hippocampus*, *17*, 709–722.
- Goldman-Rakic, P. S., Muly, E. C., & Williams, G. V. (2000). D1 receptors in prefrontal cells and circuits. *Brain Research Reviews*, *31*, 295–301.
- Gorelova, N., Seamans, J. K., & Yang, C. R. (2002). Mechanisms of dopamine activation of fast-spiking interneurons that exert inhibition in rat prefrontal cortex. *Journal of Neurophysiology*, *88*, 3150–3166.
- Goulet, S., & Murray, E. A. (2001). Neural substrates of crossmodal association memory in monkeys: The amygdala versus the anterior rhinal cortex. *Behavioral Neuroscience*, *115*, 271–284.
- Grace, A. A. (1991). Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: A hypothesis for the etiology of schizophrenia. *Neuroscience*, *41*, 1–24.
- Hamker, F. H. (2005). The reentry hypothesis: The putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement. *Cerebral Cortex*, *15*, 431–447.
- Hamker, F. H., & Wilschut, J. (in preparation). Homeostatic scaling and Hebbian learning in dynamic rate-coded neurons.
- Hampton, R. R. (2005). Monkey perirhinal cortex is critical for visual memory, but not for visual perception: Reexamination of the behavioural evidence from monkeys. *Quarterly Journal of Experimental Psychology, Series B*, *58*, 283–299.
- Hölscher, C., & Rolls, E. T. (2002). Perirhinal cortex neuronal activity is actively related to working memory in the macaque. *Neural Plasticity*, *9*, 41–51.
- Horel, J. A., Pytko-Joiner, D. E., Voytko, M. L., & Salsbury, K. (1987). The performance of visual tasks while segments of the inferotemporal cortex are suppressed by cold. *Behavioural Brain Research*, *23*, 29–42.
- Huang, Y. Y., & Kandel, E. R. (1995). D1/D5 receptor agonists induce a protein synthesis-dependent late potentiation in the ca1 region of the hippocampus. *Proceedings of the National Academy of Sciences, U.S.A.*, *92*, 2446–2450.
- Hurd, Y. L., Suzuki, M., & Sedvall, G. C. (2001). D1 and D2 dopamine receptor mRNA expression in whole hemisphere sections of the human brain. *Journal of Chemical Neuroanatomy*, *22*, 127–137.
- Joel, D., Niv, Y., & Ruppín, E. (2002). Actor-critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Networks*, *15*, 535–547.
- Kincaid, A. E., Zheng, T., & Wilson, C. J. (1998). Connectivity and convergence of single corticostriatal axons. *Journal of Neuroscience*, *18*, 4722–4731.
- Lee, D., & Chun, M. M. (2001). What are the units of visual short-term memory, objects or spatial locations? *Perception & Psychophysics*, *63*, 253–257.
- Lehky, S. R., & Tanaka, K. (2007). Enhancement of object representations in primate perirhinal cortex during a visual working-memory task. *Journal of Neurophysiology*, *97*, 1298–1310.
- Linke, R., & Schwegler, H. (2000). Convergent and complementary projections of the caudal paralaminar thalamic nuclei to rat temporal and insular cortex. *Cerebral Cortex*, *10*, 753–771.
- Liu, Z., Richmond, B. J., Murray, E. A., Saunders, R. C., Steenrod, S., Stubblefield, B. K., et al. (2004). DNA targeting of rhinal cortex D2 receptor protein reversibly blocks learning of cues that predict reward. *Proceedings of the National Academy of Sciences, U.S.A.*, *101*, 12336–12341.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*, 279–281.
- McCormick, D. A., Connors, B. W., Lighthall, J. W., & Prince, D. A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology*, *54*, 782–806.
- Meunier, M., Bachevalier, J., Mishkin, M., & Murray, E. A. (1993). Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *Journal of Neuroscience*, *13*, 5418–5432.
- Miller, E. K., Erickson, C., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, *16*, 5154–5167.
- Miller, E. K., Gochin, P. M., & Gross, C. G. (1993). Suppression of visual responses of neurons in inferior temporal cortex of the awake macaque monkey by addition of a second stimulus. *Brain Research*, *616*, 25–29.
- Miller, E. K., Li, L., & Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *Journal of Neuroscience*, *13*, 1460–1478.
- Mogami, T., & Tanaka, K. (2006). Reward association affects neuronal responses to visual stimuli in macaque TE and perirhinal cortices. *Journal of Neuroscience*, *26*, 6761–6770.
- Momiyama, T., Sim, J. A., & Brown, D. A. (1996). Dopamine D1-like receptor-mediated presynaptic inhibition of excitatory transmission onto rat magnocellular basal forebrain neurones. *Journal of Physiology*, *495*, 97–106.

- Murray, E. A., Gaffan, D., & Mishkin, M. (1993). Neural substrates of visual stimulus-stimulus association in rhesus monkeys. *Journal of Neuroscience*, *13*, 4549–4561.
- Nakamura, K., & Kubota, K. (1995). Mnemonic firing of neurons in the monkey temporal pole during a visual recognition memory task. *Journal of Neurophysiology*, *74*, 162–178.
- Naya, Y., Yoshida, M., Takeda, M., Fujimichi, R., & Miyashita, Y. (2003). Delay-period activities in two subdivisions of monkey inferotemporal cortex during pair association memory task. *European Journal of Neuroscience*, *18*, 2915–2918.
- Ohbayashi, M., Ohki, K., & Miyashita, Y. (2003). Conversion of working memory to motor sequence in the monkey premotor cortex. *Science*, *301*, 233–236.
- Pihlajamäki, M., Tanila, H., Hänninen, T., Könönen, M., Mikkonen, M., Jalkanen, V., et al. (2003). Encoding of novel picture pairs activates the perirhinal cortex: An fMRI study. *Hippocampus*, *13*, 67–80.
- Ranganath, C. (2006). Working memory for visual objects: Complementary roles of inferior temporal, medial temporal, and prefrontal cortex. *Neuroscience*, *139*, 277–289.
- Ranganath, C., & D'Esposito, M. (2005). Directing the mind's eye: Prefrontal, inferior and medial temporal mechanisms for visual working memory. *Current Opinion in Neurobiology*, *15*, 175–182.
- Rempel-Clower, N. L., & Barbas, H. (2000). The laminar pattern of connections between prefrontal and anterior temporal cortices in the rhesus monkey is related to cortical structure and function. *Cerebral Cortex*, *10*, 851–865.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Rougier, N. P., & Vitay, J. (2006). Emergence of attention within a neural population. *Neural Networks*, *19*, 573–581.
- Schiller, J., Major, G., Koester, H. J., & Schiller, Y. (2000). NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature*, *404*, 285–289.
- Seamans, J. K., & Yang, C. R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Progress in Neurobiology*, *74*, 1–58.
- Smith, J. D., Redford, J. S., Gent, L. C., & Washburn, D. A. (2005). Visual search and the collapse of categorization. *Journal of Experimental Psychology: General*, *134*, 443–460.
- Suzuki, W. A., & Amaral, D. G. (1994). Perirhinal and parahippocampal cortices of the macaque monkey: Cortical afferents. *Journal of Comparative Neurology*, *350*, 497–533.
- Suzuki, W. A., Miller, E. K., & Desimone, R. (1997). Object and place memory in the macaque entorhinal cortex. *Journal of Neurophysiology*, *78*, 1062–1081.
- Taylor, J. G. (1999). Neural bubble dynamics in two dimensions: Foundations. *Biological Cybernetics*, *80*, 393–409.
- Taylor, K. I., Moss, H. E., Stamatakis, E. A., & Tyler, L. K. (2006). Binding crossmodal object features in perirhinal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *103*, 8239–8244.
- Turrigiano, G. G., & Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, *5*, 97–107.
- Tyler, L. K., Stamatakis, E. A., Bright, P., Acres, K., Abdallah, S., Rodd, J. M., et al. (2004). Processing objects at different levels of specificity. *Journal of Cognitive Neuroscience*, *16*, 351–362.
- van Rossum, M. C. W., & Turrigiano, G. G. (2001). Correlation based learning from spike timing dependent plasticity. *Neurocomputing*, *38–40*, 409–415.
- Wan, H., Aggleton, J. P., & Brown, M. W. (1999). Different contributions of the hippocampus and perirhinal cortex to recognition memory. *Journal of Neuroscience*, *19*, 1142–1148.
- Yang, C. R., & Seamans, J. K. (1996). Dopamine D1 receptor actions in layers V–VI rat prefrontal cortex neurons in vitro: Modulation of dendritic-somatic signal integration. *Journal of Neuroscience*, *16*, 1922–1935.
- Zahrt, J., Taylor, J. R., Mathew, R. G., & Arnsten, A. F. (1997). Supranormal stimulation of d1 dopamine receptors in the rodent prefrontal cortex impairs spatial working memory performance. *Journal of Neuroscience*, *17*, 8528–8535.
- Zola-Morgan, S., Squire, L. R., Amaral, D. G., & Suzuki, W. A. (1989). Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. *Journal of Neuroscience*, *9*, 4355–4370.