Binding objects to cognition: A brain-like systems approach to the cognitive control of visual perception

Julien Vitay and Fred H. Hamker

***Abstract*— We describe a brain-like approach for the cognitive control of visual perception and present the first simulation results. We first describe how working memory is a key process in visual perception, allowing perception to be linked to action through time. We list the different challenges of this approach and propose a computational model of perirhinal cortex that allows memory retrieval in mnemonic tasks. We then discuss a functional model of working memory allowing to take advantage of memory retrieval by linking robust working memory to selective updating and reinforcement learning in prefrontal cortex and basal ganglia.**

I. INTRODUCTION

A reactive system determines its actions in response to external events or stimuli. Much research has been devoted to the development of learning rules that enable an agent to determine an action in order to achieve a maximum of reward, dependent on the sensor information. Despite much progress in this field, the flexibility to change behavior depending on particular context information, to selectively emphasize task-relevant sensory information and to ignore task-irrelevant stimuli is limited. We aim at developing cognitive systems that actively manipulate their working memory to flexibly adapt their strategy, particularly with respect to visual perception. Our research direction is twofold: we aim to understand human cognition and provide models for artificial cognitive agents. Part of this research goal is coordinated with the EU Consortium Eyeshots "Heterogeneous 3-D Perception Across Visual Fragments" which intends to investigate how an active exploration allows to acquire a knowledge of the surrounding environment. From an autonomous agent's point of view, such an active exploration requires that visual perception is as well an active process: the agent has to select which part of the incoming information has to be processed in priority to assure its survival. Visual attention is an example of the various cognitive processes involved in perception. Although it has a bottom-up component induced by the physical properties of the visual objects, visual attention has also a cognitive top-down component which guides the way visual information is processed, depending on the context or tasks requirements [1], [2]. This goal-directed behaviour needs internal representations of the objects that are expected in order to favorize their perception. We here address the fundamental questions of how these internal representations are learned, maintained and recalled when needed.

Working memory (WM) is a process that allows such representations to be momentarily maintained throughout the visual search phase (for example), directing attention to objects having similar properties. In delayed match-to-sample (DMS) tasks, where the presentation of a visual cue indicates which target has to be chosen after a certain delay, the internal representation of the cue or the associated target (if visually different from the cue) has to be maintained during the delay period to direct attention through top-down connections to the correct target. WM has often been related to states of persistent activity. Several computational models of WM in prefrontal cortex have been proposed. They typically consist of a pool of excitatory coupled pyramidal neurons and several inhibitory interneurons [3], [4], [5], [6]. Common effects are that the network can remain in an active stage even in the absence of an external input. Most of these models consider that dopamine (DA) is a key modulator for sustained activities during WM processes, but it is also critical for many more aspects of cognitive vision [7]. Dopamine is one of multiple neurotransmitters in the nervous system, such as noradrenaline, serotonin and acetylcholine. Dopamine has been linked to almost all aspects of behaviour, cognition, addiction and particularly to reward. The assumed modulatory influence of DA, such as an increase of the NDMA and GABA and a decrease of AMPA conductances, varies in these models. Neurons show a bi-phasic property (switching between low and high active states). An increase in DA, assumed to have a global, spatially non-selective effect, changes the properties of the switch from spontaneous activity to a state of persistent activity, such that memorization is facilitated. Limitations of these models lie in the functional context, since the timing of the signals for reseting the memory is not specified [8]. Thus, systems level aspects of working memory have to be considered as well.

However, DA has different effects depending on the receiving cell. Its effect on cortical cells is very slow and long-lasting [9], putting emphasis on the tonic firing mode of DA cells, whereas striatal cells respond to the phasic bursts of DA they receive. These bursts have been related to reinforcement learning since the discovery that the phasic bursts of DA cells occur at the time of the delivery of a reward before learning a DMS task, but progressively occur at the time of the presentation of the reward-predicting cue [10]. Moreover, when reward is omitted, these cells cease to fire. This firing pattern has been quite accurately modelled as the prediction error signal of the temporal-differences (TD) algorithm [11], but the neural mechanism behind it is still unclear. How does this DA firing pattern influence the learning and maintenance of WM?

Psychologisches Institut II, Westfälische Wilhelms Universität Münster, Fliednerstr. 21 D-48149 Münster, Germany. `{fhamker,jvita_01}@uni-muenster.de`

Another issue related to the link between sustained activities and WM is the content of the memorized pattern. In the case of visual search, does this pattern represent the visual features that have to searched for, or a more abstract code allowing only the manipulation of the object? The lack of robustness to distractors of sustained activities in visual areas (like in inferotemporal or perirhinal cortices) suggests that they are not the primary mechanism for memorizing the visual content. We propose to shift emphasis on memory retrieval as part of the visual WM process. In this article, we first describe how perirhinal cortex (PRh) can be part of the memory retrieval process in the context of delayed match-to-sample tasks and we propose a computational model. We will then sketch how this model can fit in a broader framework of WM manipulation by the reciprocal interactions between prefrontal cortex and basal ganglia.

II. PERIRHINAL CORTEX INVOLVED IN MEMORY RETRIEVAL

PRh (areas 35 and 36) is located in the medial temporal lobe and has strong reciprocal connections with the inferotemporal cortex (IT). It is involved in object recognition memory, novelty detection, object categorization, multimodal integration [12]. Its cells have a broader selectivity than IT cells and tend to gather in clusters representing the same object, perhaps with different modalities or viewpoints. They also show sustained activities during DMS tasks, but they are not robust to the presentation of distractors [13]. To explore their functional role, we propose that these sustained activities are only the consequence of the clustering of these cells during learning. They do not support by themselves the process of WM, but rather facilitate memory recall.

We designed a computational model of PRh emphasizing this hypothesis [14]. It is composed of a set of 400 excitatory cells reciprocally connected with another set of 100 inhibitory cells (Fig. 1). We 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. Each excitatory cell receives a cortical input C coming from other cortical areas like IT and representing different aspects of objects. In our simulation, one aspect of an object directly stimulates a few excitatory cells whereas one object is represented by multiple aspects that can occur randomly at each presentation of the object during the learning phase. Additionnaly, the excitatory cells are reciprocally connected in a all-to-all manner, but with a strength that can adapt to the experience of the network. We use for this a learning rule that is a mixture of the classical covariance rule [15] with a homeostatic regulation of synaptic strengths at the cell level [16].

The goal of the learning phase is that the multiple sets of excitatory cells representing the different aspects of the same object create strong connections between them to form a cluster. After learning, the modulation by DA of the different synaptic currents can switch the network in a different mode.

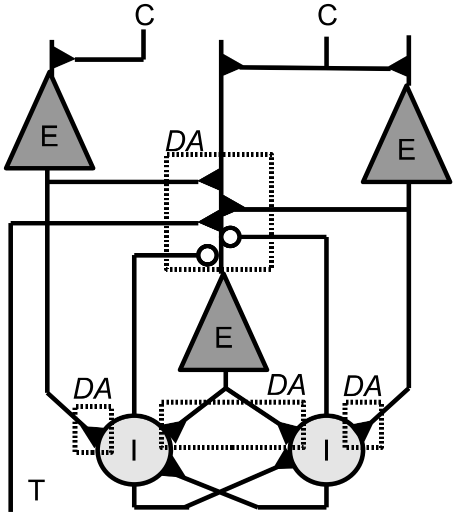

Fig. 1. Each excitory cell E receives connections from neighbouring excitatory cells, neighbouring inhibitory cells I, cortical input C from other areas and a thalamic input T. Inhibitory cells receive connections from neighbouring excitatory cells and other inhibitory cells. Every connection except the cortical ones (C) and inhibitory-to-inhibitory ones are modulated by DA.

Consistent with other models of dopaminergic modulation of WM in prefrontal cortex [3], [4], [5], [6], we considered that DA enhances reciprocal connections between excitatory cells (on an activity-dependent manner, as they are mostly mediated by NMDA receptors), reciprocal connections between excitatory and inhibitory cells as well as thalamic inputs. We reproduced classical results for this type of neural network, such as the inverted-U shape of sustained activities depending on the dopamine level [17] (DA must be in a certain range to observe sustained activities, see Fig. 2), propagation of activity within a cluster for intermediate levels of DA (when only half of the sets of cells receive a cortical input during the presentation of an object, the remaining cells exhibit strong activity due to the increase in the strength of the reciprocal connections with the stimulated cells) and the lack of robustness of sustained activities. Moreover, we showed that a thalamic stimulation of a little percentage of the cells of a cluster (lower than 35%) is sufficient to propagate this activity to the whole cluster.

The implications of this model is that PRh is a multimodal area gathering different multimodal informations on individual objects through the formation of clusters. Each individual cell of the cluster has a high selectivity on an aspect of an object through its cortical inputs, but the propagation of activity inside the cluster under intermediate levels of DA makes each cell also responsive for other aspects of the object. Thalamic stimulation of a small number of cells of a cluster can easily retrieve the content of the memory stored in the reciprocal connections within the cluster. In return, this retrieved memory can be used to bias visual processing in the visual stream, through feedback connections to IT. The core WM of the object can be located somewhere else, but through

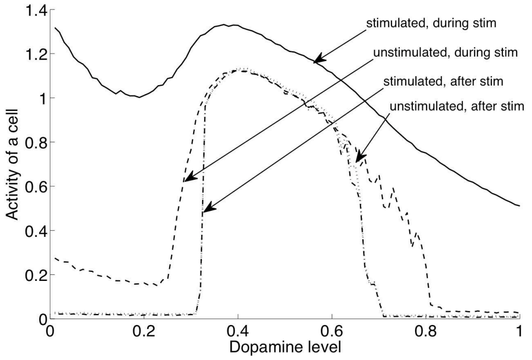

Fig. 2. Effect of the DA level on sustained activities. Two cells have been recorded: one receiving a cortical input, the other being in the same cluster but without stimulation. When the first cell receives stimulation, it exhibits a strong activity whatever the DA level. The unstimulated one shows significant activity only for intermediate DA levels (between 0.3 and 0.7, the maximum level of DA being 1.0). One hundred milliseconds after the end of stimulation, we record again the activity of the two cells. We can then observe that for, the intermediate range of DA level, the two cells exhibit sustained activities.

this thalamic stimulation the sensory features associated to the memorized object can be retrieved when needed (”GO signal”), not necessarily during the delay period.

III. PREFRONTAL CORTEX AND BASAL GANGLIA FOR MEMORY MANIPULATION

A. Context

Since the sensory details of an object can be retrieved through partial thalamic stimulation of a cluster in PRh, the robust sustained activities representing this object during the delay period do not need to be so detailed. They could rather be more focused on the manipulation of the objects with possible associations. Prefrontal cortex (PFC) and basal ganglia (BG) are good candidates to implement this robust WM. Present findings indicate that prefrontal cortex rather provides executive control than represents the top of a hierarchy of the representation of the visual world. For example, with lesions of the PFC, subjects tend to only react towards visual cues with behaviors that are associated with the cue itself[18]. This suggests that subjects represent the cue outside PFC, but with lesions of the PFC, they have difficulties in the coordination of new responses.

DA is directly involved in prefrontal sustained activities, through modulation of different ionic and synaptic currents [19]. However, its effects are very slow and long lasting, meaning that reward-predictive phasic DA bursts are not potential candidates to gate the updating of WM in PFC. DA rather provides a global tone for prefrontal processing and sustained activities. On the contrary, the striatum (the major input structure of BG) is directly modulated by DA phasic bursts. BG are known to play an important role in WM processes, as revealed by imaging studies [20], [21]. Striatal medium spiny neurons have been shown to exhibit a region of bistability (up/down states) under elevated dopamine levels [22]. This intrinsic mechanism of sustained activities is therefore under direct control of DA, whose effects have very quick latencies in the striatum, allowing phasic DA discharges to control the updating of working memory at this level. This information can then be sent back to PFC through direct connections between the striatum and the mediodorsal part of the thalamus [23] and/or indirectly through GPi (*globus pallidus pars interna*, one output structure of BG).

Anatomical investigations suggest that the basal ganglia is also part of cortical-basal ganglionic loops [24], [25] in which the cortex provides excitatory links to the striatum followed by disinhibition of the thalamus and back to cortex. It has been proposed that the basal ganglia participate in two circuits with skeletomotor and oculomotor areas of cortex: the primary motor cortex (M1) and the frontal eye field (FEF) and in three additional circuits which involve non-motor areas in the frontal lobe: dorsolateral prefrontal cortex, lateral orbitofrontal cortex and anterior cingulate cortex. The outputs to PFC are topographically organized and they are largely segregated from those to the motor areas [23]. The learned competition between the direct pathway (directly from striatum to GPi) and the indirect pathway (with a relay to GPe (*globus pallidus pars externa*)) can then selectively disinhibit a recurrent loop between a cortical area and the corresponding thalamic nuclei. In the case of WM, this disinhibition can either stabilize existing sustained representations or selectively gate the entry of new reprensentations into WM.

Several computational models have already tried to address the interplay between the prefrontal cortex and the basal ganglia in WM tasks. The PBWM model of Frank and O’Reilly [26] is designed to learn different cognitive tasks involving WM on the same substrate, thanks to a robust reward-prediction error algorithm. However, time is not explicitely considered, making no distinction between phasic and tonic modes of DA firing. The FROST model of Ashby et al. [27] uses BG to maintain information in prefrontal cortex, but does not deal with updating. Finally, the model of Gruber et al. [28] is the first that distinguishes between the tonic long-lasting effect of DA on prefrontal cells (increasing the cortical gain) and the phasic effect of DA on striatal cells (intrinsic bistability). Prefrontal cortex and BG in this model cooperate to provide a robust and dynamic WM depending on dopamine modulation. However, the model would need to be integrated into a realistic task, with an efficient learning algorithm for DA activation.

B. Proposed tasks

Our main hypothesis is that the content of WM is not located in PFC or in BG, but can be retrieved through external activation of multimodal areas like PRh. It can be tested by creating a computational model able to learn different tasks involving WM manipulation. Following [29], we use an intermix of delayed match-to-sample (DMS) and delayed pair-association (DPA) as depicted on Fig. 3. In the DMS task, the model is presented with a visual object (A or B), creating a distributed representation in PRh. After a certain delay, a cue representing the DMS task is shown, and the representation of A or B in PRh is reset due to

Fig. 3. Delayed match-to-sample (DMS) and delayed pair-association (DPA) tasks. When a visual object (A or B) is followed by the DMS cue, the subject has to choose the same object to obtain reward. When it is followed by the DPA cue, the subject has to chose the opposite object to obtain reward.

competition. The goal of the model in the DMS task is to retrieve the content of the previously shown item in order to obtain reward. In the DPA task, the system must on the contrary retrieve the content of the associated item to obtain reward.

In psychological experiments, the subjects would have to choose the correct target between A and B, but here, the retrieval of the content of the target in PRh is considered as sufficient to induce top-down attention on the visual features processed in the ventral pathway and consequently choose the correct target. The interest of this mixture of tasks is that the WM system does not need to systematically remember the visual features of the first presented object, but rather its identity in order to retrieve the correct information.

C. Proposed model

In Fig. 4 we propose a functional model that allows to learn the formerly described tasks. It relies on the ability of our PRh model to retrieve the multimodal representation of an object through thalamic stimulation. PRh is reciprocally connected to a thalamic nuclei, so that coordinated self-organization in these two structures leads to segregated recurrent loops for each cluster. Thalamus is under tonic inhibition of the output structure of BG (GPi) to control the thalamic stimulation of PRh. GPi is itself inhibited by the striatum (STR) and of other BG nuclei (GPe or STN (subthalamic nuclei, not included here)), meaning that activation of these structures can selectively disinhibit one or more thalamocortical loops in PRh and retrieve the content of a cluster.

PRh therefore projects to the striatum that will learn to associate PRh representations to the disinhibition of the corresponding thalamocortical loops thanks to the reward-prediction system represented by phasic DA releases to the striatum. These DA bursts can favorize either LTP or LTD in corticostriatal connections depending on the polarization state of the striatal cells [30]. They can also switch the striatal cells from their intrinsic bistable states, particularly into the ”up-state” where the cell exhibits robust sustained activity. This part of the model is able to learn the DMS task alone. In order to concurrently learn the DPA task, we have to introduce a mixing of sensory information and task-related context at the striatal level to disinhibit the correct thalamocortical loop. We consider then a prefrontal area (probably located in orbitofrontal cortex) that will receive information about the identity of the objects cues, and exhibit compound representations of objects in the context of the current task. This PFC area projects on the same striatal cells that PRh. Self-organization of striatal cells through lateral competition creates separate representations of objects according to the task requirements, allowing different thalamocortical loops to be disinhibited depending on the task. The interplay between direct and indirect pathways in the BG circuitry favorizes this selection.

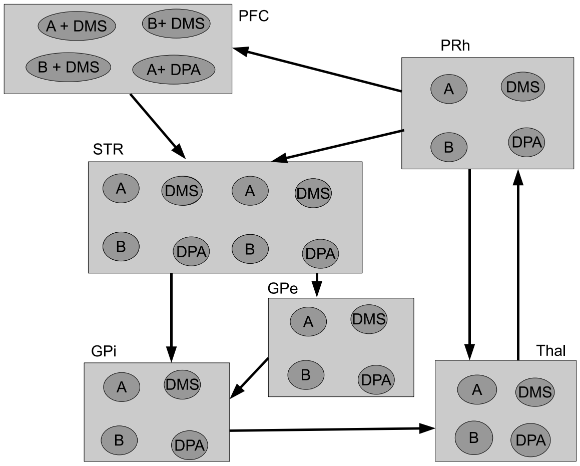

Fig. 4. Schematic architecture of the WM model. PRh provides the striatum (STR) and PFC with a distributed representation of the identity of the visual object (A, B, DMS or DPA). Self-organization in PFC and STR allows to create a representation of both, the identity of the object and the current task. Disinhibition of thalamus (Thal) can then retrieve the content of the memorized or associated object.

The key problem this model addresses is the timecourse of DA release and its differential effects on striatal and cortical cells. As stated earlier, cortical cells react very slowly to DA release, whereas DA can switch very quickly a striatal cell into an ”up-state”. At the beginning of the learning phase, DA fires for salient objects or for unexpected rewards, whereas later it only responds when an object predicts reward. In our experiments, objects and cues are not distinguishable as they are equally represented in PRh. However, if the system tries to retrieve the content of the cues (DMS or DPA), it will never receive reward, and this tendency will be reduced. The only reward-predicting stimuli will then be the objects A and B, that switch the striatal cells into the up-state, leading to a robust working memory of all possible combinations of this object. When the cue is presented, it will create a burst of activity in PFC which will stimulate the striatal cells that represent the correct association between an object and the task, and consequently disinhibit the correct thalamocortical loop.

D. Results

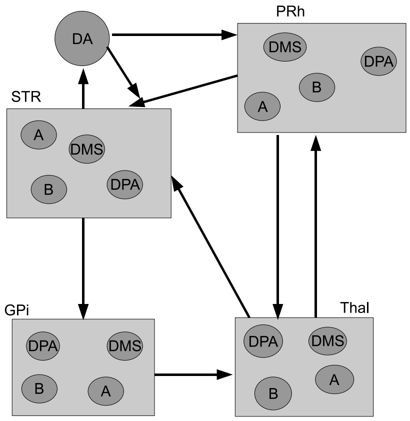

Fig. 5. Model for memory retrieval property. PRh provides STR with a distributed representation of the identity of an object (A, B, DMS or DPA). Self-organization in STR allows to create a different and reduced representation of this object. The disinhibition of thalamus by GPi can then retrieve the content of the memorized or associated object. DA is phasically activated by new patterns in STR and modulates the learning of the corticostriatal connections.

The preceding functional model is not fully implemented yet, but we have extended our PRh model [14] to show how BG can disinhibit thalamus to retrieve object memory. This preliminary version of the WM model is depicted on Fig. 5. Each area is represented by a set of mean-firing rate units whose activity evolves through time following a differential equation. PRh is stimulated by four different objects representing the A, B, DMS and DPA cues. It projects to the striatum in an all-to-all manner. The striatum inhibits GPi with a gaussian connectivity kernel, as well as the inhibitory connections from GPi to thalamus. PRh and the thalamus are reciprocally connected in an all-to-all manner. The thalamus also projects to the striatum with a gaussian connectivity kernel, for reasons explained later. The reciprocal connections between PRh and thalamus are learned according to the same learning rule we use for the lateral connections in PRh.

Each time an object is presented, it creates some new activity in striatum, which phasically activates the dopamine cell for approximatively 200 ms. DA then modulates the homeostatic learning of the corticostriatal connections, allowing LTP only at the time of a DA burst. Competition in the striatum and in thalamus ensures self-organization in these structures, so that only a few cells represent the same object. The topological projection from thalamus to striatum ensures that the two representations of the same object are similar but not identical, since thalamus has four times less cells than striatum in this model. Learning then converges to three different representations of the same object in PRh, striatum and thalamus, with a decreasing number of active cells.

During learning, GPi is always inhibited by striatal activity to ensure that thalamic cells can become active and learn their reciprocal connections with PRh. When the model has to learn DMS and DPA tasks, this inhibition will not be systematic anymore. Also, striatal cells are kept in the ”down-state” during learning, i.e. they only respond to cortical stimulation but do not show sustained activities. Fig. 6 highlights the properties of this model showing the simulation of a sequence of events. The object A is first shown for 200 ms. The corresponding cluster exhibits sustained activation after the disappearance of the object. We then artificially set the striatal cells into their ”up-state”, by self-stimulating them. Another object then appears (DMS, for example) which erases the sustained activation in PRh. When it disappears again, the striatal cells still fire and the thalamic cells corresponding to object A are disinhibited. This activates the thalamocortical loop and the cluster in PRh corresponding to object A shows again high activity, without any cortical input. Thus, a sustained activation of striatal cells allows memory retrieval in PRh.

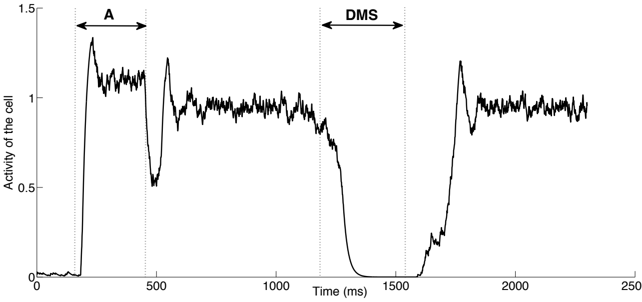

Fig. 6. Timecourse of the activity of a perirhinal cell. At approx. 200 ms after the start of the trial, the preferred object of this cell is presented for 200 ms. When the stimulation ends, the cluster exhibits sustained activation until a new object is presented at approx. 1200 ms. This new object disappears at 1600 ms and the cell becomes active again due to memory retrieval.

IV. CONCLUSIONS AND FUTURE WORKS

A. Conclusions

We have presented a general framework of cognitive control of visual perception, by emphasizing the role of working memory in perception. It relies on the distinction between the content of working memory (what is actually kept in mind) and the content of the memorized object (its visual features). We suggest that memory retrieval is a key process in everyday WM tasks which are scientifically explored by delayed match-to-sample or delayed pair-association tasks. Memory retrieval is essential in exploiting internally-generated goals that guide behaviour without any environmental anchoring. We also suggest that top-down attention to targets in these tasks are mainly managed by the subordinate categories that are described in high-level visual areas like PRh.

We propose a functional model of working memory in visual tasks that relies on the known anatomy and functionality of diverse cortical and subcortical structures. Dopamine is thought to play a great role in the modulation of cortical activity as well as in the learning of associations at the striatal level. The understanding of the distinction between tonic and phasic modes of DA firing should allow to differentiate reward-related processing in cortical and subcortical structures and highlight the relation between limbic information and cognition.

Autonomous agents should take benefit of this internal biasing of perception-action loops. The coupling of visual selective attention [2] with reinforcement-based maintenance of relevant information in WM will allow to generate goal-directed behaviours grounded on actual perception and past experience.

B. Future Works

Only a part of the proposed functional model has been implemented in a brain-like computational model. The problems still unsolved are the intrinsic bistability of striatal cells, the learning of compound representations in PFC and the development of an efficient learning algorithm for the DA signal. The classical TD algorithm [31] will not fit because it relies on strictly constant delays between cue onset and reward delivery. Realistic tasks can not make such a strong hypothesis. A distinction between primary rewards and learned values as in [32] seems more plausible, by using the sustained activation in striatum to avoid dependency on time intervals.

V. ACKNOWLEDGMENTS

This work has been supported by the HA2630/4-1 grant "A neurocomputational systems approach to modeling the cognitive guidance of attention and object/category recognition" of the german research foundation (Deutsche Forschungsgemeinschaft, DFG) and in part by the EU grant "Eyeshots".

REFERENCES

[1] H. Tomita, M. Ohbayashi, K. Nakahara, I. Hasegawa, and Y. Miyashita, "Top-down signal from prefrontal cortex in executive control of memory retrieval," *Nature*, vol. 401, pp. 699–703, 1999.

[2] F. H. Hamker, "The reentry hypothesis: the putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement." *Cereb Cortex*, vol. 15, no. 4, pp. 431–447, Apr 2005.

[3] N. Brunel and X. J. Wang, "Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition." *J Comput Neurosci*, vol. 11, no. 1, pp. 63–85, 2001.

[4] D. Durstewitz, M. Kelc, and O. Güntürkün, "A neurocomputational theory of the dopaminergic modulation of working memory functions." *J Neurosci*, vol. 19, no. 7, pp. 2807–2822, Apr 1999.

[5] J.-C. Dreher, E. Guigon, and Y. Burnod, "A model of prefrontal cortex dopaminergic modulation during the delayed alternation task." *J Cogn Neurosci*, vol. 14, no. 6, pp. 853–865, Aug 2002.

[6] S. Tanaka, "Dopamine controls fundamental cognitive operations of multi-target satial working memory," *Neural Netw*, vol. 15, pp. 573–582, 2002.

[7] J. Vitay and F. H. Hamker, "On the role of dopamine in cognitive vision," in *Attention in Cognitive Systems*, L. Paletta and E. Rome, Eds. Springer LNAI 4840, 2007, pp. 352–366.

[8] Y. E. Cohen and R. A. Andersen, "A common reference frame for movement plans in the posterior parietal cortex," *Nat Rev Neurosci*, vol. 3, pp. 553–562, 2002.

[9] Y. Y. Huang and E. R. Kandel, "D1/D5 receptor agonists induce a protein synthesis-dependent late potentiation in the ca1 region of the hippocampus." *Proc Natl Acad Sci U S A*, vol. 92, no. 7, pp. 2446–2450, Mar 1995.

[10] W. Schultz, "Predictive reward signal of dopamine neurons," *J Neurophysiol*, vol. 80, pp. 1–27, 1998.

[11] R. E. Suri and W. Schultz, "Temporal difference model reproduces anticipatory neural activity." *Neural Comput*, vol. 13, no. 4, pp. 841–862, Apr 2001.

[12] M. J. Buckley and D. Gaffan, "Perirhinal cortical contributions to object perception." *Tr Cogn Sci*, vol. 10, no. 3, pp. 100–107, Mar 2006.

[13] E. K. Miller, P. M. Gochin, and C. G. Gross, "Suppression of visual responses of neurons in inferior temporal cortex of the awake macaque monkey by addition of a second stimulus," *Brain Res*, vol. 616, pp. 25–29, 1993.

[14] J. Vitay and F. H. Hamker, "Sustained activities and retrieval in a computational model of perirhinal cortex," *J Cog Neurosci*, *in press*.

[15] E. Oja, "A simplified neuron model as a principal component analyze," *J Math Biol*, vol. 15, pp. 267–273, 1982.

[16] M. C. W. van Rossum and G. G. Turrigiano, "Correlation based learning from spike timing dependent plasticity," *Neurocomputing*, vol. 38-40, pp. 409–415, 2001.

[17] P. S. Goldman-Rakic, E. C. Muly, and G. V. Williams, "D1 receptors in prefrontal cells and circuits." *Brain Res Rev*, vol. 31, no. 2-3, pp. 295–301, Mar 2000.

[18] T. Shallice and P. W. Burgess, "Deficits in strategy application following frontal lobe damage in man," *Brain*, vol. 114, pp. 727–741, 1991.

[19] J. K. Seamans and C. R. Yang, "The principal features and mechanisms of dopamine modulation in the prefrontal cortex." *Prog Neurobiol*, vol. 74, no. 1, pp. 1–58, Sep 2004.

[20] B. R. Postle and M. D'Esposito, "Dissociation of human caudate nucleus activity in spatial and nonspatial working memory: an event-related fmri study." *Brain Res Cogn Brain Res*, vol. 8, no. 2, pp. 107–115, Jul 1999.

[21] S. J. G. Lewis, A. Dove, T. W. Robbins, R. A. Barker, and A. M. Owen, "Striatal contributions to working memory: a functional magnetic resonance imaging study in humans." *Eur J Neurosci*, vol. 19, no. 3, pp. 755–760, Feb 2004.

[22] C. J. Wilson and Y. Kawaguchi, "The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons." *J Neurosci*, vol. 16, no. 7, pp. 2397–2410, Apr 1996.

[23] F. A. Middleton and P. L. Strick, "Basal-ganglia 'projections' to the prefrontal cortex of the primate." *Cereb Cortex*, vol. 12, no. 9, pp. 926–935, Sep 2002.

[24] G. E. Alexander, M. R. DeLong, and P. L. Strick, "Parallel organization of functionally segregated circuits linking the basal ganglia and cortex," *Ann Rev Neurosci*, vol. 9, pp. 357–381, 1986.

[25] D. G. Beiser and J. C. Houk, "Model of cortical-basal ganglionic processing: encoding the serial order of sensory events," *J Neurophysiol*, vol. 79, pp. 3168–3188, 1998.

[26] M. J. Frank, B. Loughry, and R. C. O'Reilly, "Interactions between frontal cortex and basal ganglia in working memory: a computational model." *Cogn Affect Behav Neurosci*, vol. 1, no. 2, pp. 137–160, Jun 2001.

[27] F. G. Ashby, S. W. Ell, V. V. Valentin, and M. B. Casale, "Frost: a distributed neurocomputational model of working memory maintenance." *J Cogn Neurosci*, vol. 17, no. 11, pp. 1728–1743, Nov 2005.

[28] A. J. Gruber, P. Dayan, B. S. Gutkin, and S. A. Solla, "Dopamine modulation in the basal ganglia locks the gate to working memory." *J Comput Neurosci*, vol. 20, no. 2, pp. 153–166, Apr 2006.

[29] T. Gisiger and M. Kerszberg, "A model for integrating elementary neural functions into delayed-response behavior." *PLoS Comput Biol*, vol. 2, no. 4, p. e25, Apr 2006.

[30] P. Calabresi, B. Picconi, A. Tozzi, and M. Di Filippo, "Dopamine-mediated regulation of corticostriatal synaptic plasticity," *Trends in Neurosciences*, vol. 30, no. 5, pp. 211–219, 2007.

[31] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998, a Bradford Book.

[32] R. C. O'Reilly and M. J. Frank, "Making working memory work: A computational model of learning in the frontal cortex and basal ganglia," *Neur Comput*, vol. 18, pp. 283–328, 2006.