Neurodynamics of the Prefrontal Cortex during Conditional Visuomotor Associations

Marco Loh[1], Anitha Pasupathy[2], Earl K. Miller[3], and Gustavo Deco[1,4]

Abstract

■ The prefrontal cortex is believed to be important for cognitive control, working memory, and learning. It is known to play an important role in the learning and execution of conditional visuomotor associations, a cognitive task in which stimuli have to be associated with actions by trial-and-error learning. In our modeling study, we sought to integrate several hypotheses on the function of the prefrontal cortex using a computational model, and compare the results to experimental data. We constructed a module of prefrontal cortex neurons exposed to two different inputs, which we envision to originate from the inferotemporal cortex and the basal ganglia. We found that working memory properties do not describe the dominant dynamics in the prefrontal cortex, but the activation seems to be transient, probably progressing along a pathway from sensory to motor areas. During the presentation of the cue, the dynamics of the prefrontal cortex is bistable, yielding a distinct activation for correct and error trails. We find that a linear change in network parameters relates to the changes in neural activity in consecutive correct trials during learning, which is important evidence for the underlying learning mechanisms. ■

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

INTRODUCTION

Humans and animals can learn to associate stimuli with arbitrary responses, recall these associations, and adapt them when changes in the behavioral context occur. This type of learning is called conditional visuomotor learning (Passingham, 1993). Several brain structures take part in this task (for reviews, see Hadj-Bouziane, Meunier, & Boussaoud, 2003; Murray, Bussey, & Wise, 2000; Passingham, Toni, & Rushworth, 2000). These are the prefrontal cortex (PFC), the premotor cortex, the hippocampus, and the basal ganglia. In our modeling study, we focus on the PFC.

The PFC is important for the processing of conditional visuomotor mappings. Frontal lesions in humans cause deficits (Petrides, 1985, 1990, 1997). The ventrolateral PFC is crucial for the fast learning of novel associations (Bussey, Wise, & Murray, 2001; Murray et al., 2000; Wang, Zhang, & Li, 2000). The interaction with the inferotemporal cortex (IT) also contributes to the processing, as a disconnection using cross-lesion procedures of the IT and the PFC (Bussey, Wise, & Murray, 2002; Parker & Gaffan, 1998; Gaffan & Harrison, 1988), and transsection of the uncinate fascicle (Eacott & Gaffan, 1992) causes deficits in learning. However, lesions of the dorsolateral PFC cause no or mild impairments (Wang et al., 2000; Gaffan & Harrison, 1989; Petrides, 1982). This is consistent with human imaging studies which show an involvement of the ventrolateral PFC but not of the dorsolateral PFC (Toni, Ramnani, Josephs, Ashburner, & Passingham, 2001; Toni & Passingham, 1999), at least in simple versions of this task (Boettiger & D'Esposito, 2005). Electrophysiological recordings have identified neurons, which were selective to the stimulus identity, the response, and to combinations of both, that is, the associations (Asaad, Rainer, & Miller, 1998). Simultaneous recordings in the caudate nucleus and the PFC haven shown that the neural activity, which reflected the response, appeared earlier in the caudate nucleus than in the PFC (Pasupathy & Miller, 2005). In addition, the activity in the caudate nucleus changed more rapidly during learning compared to the slower changes in the PFC.

Based on the literature outlined above, we sought to integrate and discuss the following findings and hypotheses about the PFC activity in a computational model: First, the model neurons should be selective to combinations of stimuli and responses, namely, single associations in the experiment (Asaad et al., 1998). Second, stimulus information reaches the PFC from the IT. The connection of the two brain regions is crucial for learning visuomotor associations (Bussey et al., 2002; Parker & Gaffan, 1998; Gaffan & Harrison, 1988). Third, information about movement direction influences the PFC from the basal ganglia. The basal ganglia are connected to the PFC in a cortico-basal ganglionic loop (Graybiel, 1998; Houk & Wise, 1995) and have a dominant role in

[1]Universitat Pompeu Fabra, Barcelona, Spain, [2]University of Washington, [3]Massachusetts Institute of Technology, [4]Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

© 2008 Massachusetts Institute of Technology

the learning of new associations (Pasupathy & Miller, 2005; Bar-Gad, Morris, & Bergman, 2003; Houk & Wise, 1995). Fourth, we assume that the directional information is influencing the PFC throughout the delay period, as it has been hypothesized that the cortical–thalamic loop maintains this information (Houk & Wise, 1995). Fifth, one of the most prominent properties of the PFC is working memory related neural activity (Fuster & Alexander, 1971; Kubota & Niki, 1971). We model the PFC module using a recurrent network model of working memory to maintain the extracted association during the delay period.

Our goal is to integrate these aspects in a computational model and draw conclusions about the processes underlying arbitrary visuomotor mappings by comparing the results to experimental data.

METHODS

Experimental Background

We base our modeling approach on experiments featuring electrophysiological recordings in the PFC (Pasupathy & Miller, 2005; Asaad et al., 1998). In both experiments, monkeys were presented a cue (A or B) for 500 msec and, after a subsequent delay period of 1000 msec, they had to make a leftward or rightward saccadic eye movement (L or R). Selected combinations of stimuli and responses were rewarded and thereby reinforced. To enforce learning, the associations between stimuli and responses were reversed after the monkeys had reached high performance. For example, if the Stimuli A and B were associated with a leftward and rightward eye movement (L and R), respectively, then the association of A with a rightward (R) and B with a leftward (L) eye movement yielded reward after the reversal. Because the monkeys did not switch instantaneously, it was possible to investigate the neural activity during learning.

Neurodynamical Model

We use a standard recurrent network model as the basis for working memory in the PFC (Brunel & Wang, 2001). It consists of 1000 integrate-and-fire neurons which are all-to-all connected. The integrate-and-fire neurons are described by

$$C_{\mathrm{m}}\frac{\mathrm{d}V(t)}{\mathrm{d}t} = -g_{\mathrm{m}}(V(t)-V_{\mathrm{L}}) - I_{\mathrm{syn}}(t) \qquad (1)$$

where $V(t)$ is the membrane potential, C_{m} the membrane capacitance, g_{m} the leak conductance, and V_{L} the resting potential. Eight hundred of these neurons are excitatory and two hundred are inhibitory. The synaptic input I_{syn} is made up of four parts. An external excitatory input via AMPA-type synapses to the network model and recurrent input from the other neurons of the network. The latter one consists of AMPA, NMDA, and GABA currents. Thus, I_{syn} reads

$$I_{\mathrm{syn}}(t) = I_{\mathrm{AMPA,ext}}(t) + I_{\mathrm{AMPA,rec}}(t) + I_{\mathrm{NMDA,rec}}(t) + I_{\mathrm{GABA}}(t) \qquad (2)$$

The asynchronous external input $I_{\mathrm{AMPA,\ ext}}(t)$ can be viewed as originating from 800 external neurons firing at an average rate of $\nu_{\mathrm{ext}} = 3$ Hz per neuron, and thus, has a rate of 2.4 kHz in total. The recurrent input currents are summed over all neurons with weights depending on the connectivity. The synaptic dynamics is described by an exponential decay upon the arrival of a spike for AMPA and GABA currents and alpha-function including a rise time and extra voltage dependence for the NMDA current. The detailed mathematical description and parameters are provided in the Supplementary Material. The parameters of the integrate-and-fire neurons and the synaptic channels for AMPA, NMDA, and GABA are chosen according to biological data. We gain a dynamical system, which has biologically realistic properties and can have several distinct attractor states. These are envisioned to correspond to items stored in working memory.

We incorporate the five cornerstones of our modeling as mentioned in the Introduction as follows: First, the model neurons should show nonlinear object-and-direction selective activity. In the experimental paradigm of Pasupathy and Miller (2005) and Asaad et al. (1998), there are four associations between the two stimuli A and B, and two responses L and R: AL, BR, AR, and BL. Accordingly, our model contains four populations of 100 neurons each (Figure 1). The rest of the excitatory neurons comprise the nonspecific pool (NS). These neurons show spontaneous firing rates and introduce noise in the system, which aids in generating Poisson-like firing patterns throughout the system.

Second, stimulus information reaches the PFC from the IT. As we are interested in the dynamics of the PFC, we model this input by an external influence to the model. The external firing rate is increased from $\nu_{\mathrm{ext}} = 3$ Hz to $\nu_{\mathrm{stim}} = \nu_{\mathrm{ext}} + \lambda_{\mathrm{stim}}$ for the two pools associated with a stimulus, Stimulus A to AL and AR, Stimulus B to BL and BR. We use $\lambda_{\mathrm{stim}} = 0.1$ Hz throughout the simulations.

Third, the basal ganglia map stimuli onto actions and their output influences the activity of the PFC. We model the direction input which codes for response direction left (L) or right (R) by increasing the external input to (AL, BL) or (AR, BR), respectively ($\nu_{\mathrm{dir}} = \nu_{\mathrm{ext}} + \lambda_{\mathrm{dir}}$). In contrast to the stimulus input, we do not assume λ_{dir} to be constant during learning because the associations change during the blocks; that is, we increase the external input λ_{dir} from 0.0 to 0.1 Hz to simulate the changes that occur during learning.

Fourth, we assume that the input coding the response direction is influencing the PFC throughout the delay period, as it has been hypothesized that the cortical–thalamic loop maintains this information (Houk & Wise,

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

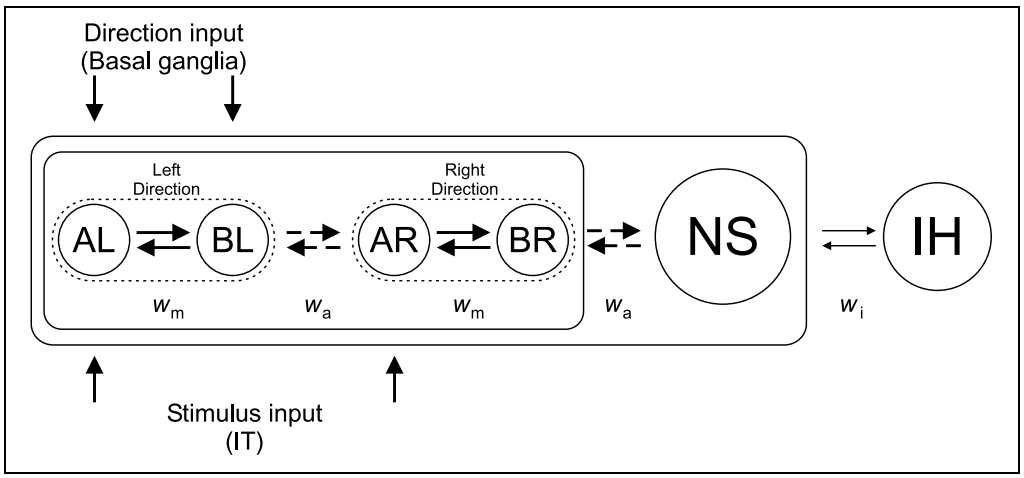

Figure 1. The cortical network model. The network is fully connected and the connection parameters between parts of the network are shown. The excitatory part of the network consists of selective pools (AL, BL, AR, BR) and a nonselective one (NS). The inhibitory neurons are clustered in the pool IH. The each neuron receives external input of Poisson spike trains firing at 2.4 kHz, which can be viewed as originating from 800 neurons firing at 3 Hz. This external input is selectively increased for the stimulus and response-direction input for the corresponding pools. The input for Stimulus A and the direct context is indicated as used in the simulations.

1995). This is implemented by applying the response-direction input from the cue period on throughout the delay period (0–1500 msec), whereas the stimulus input is applied only in the cue period (0–500 msec).

Fifth, we hypothesize that the PFC module holds the extracted association in working memory. Thus, we use a neural network model constructed to show working memory properties (Brunel & Wang, 2001) and target our analysis to parameters which show delay period activity. In our model, the stimulus–response associations are held in working memory (Asaad et al., 1998).

The connection weights between the neurons of the same population are called intrapool connection strength w_+ (Figure 1). In addition, we group the weights of the interpool connections into two parameters. The connection weights between neurons coding the same response direction, that is, between the pools AL–BL and AR–BR, are denoted w_m. The second parameter w_a applies to all other connections between excitatory pools including the nonspecific pool. The connection strength w_a is calculated using the other two connection parameters w_+ and w_m, so that the average connection strength input to a neuron equals to 1 ($fw_+ + fw_m + (1 - 2f)$ $w_a = 1$, where $f = 0.1$ is the fraction of the number of neurons in a selective pool in respect to all excitatory neurons). The excitatory connections regard both the AMPA and NMDA synaptic currents. All inhibitory GABA connections are 1 both to themselves and to all other neurons. The complete connection matrices are pictured in the Supplementary Material.

Analysis

To analyze the network, we use two kinds of techniques, which complement each other: spiking and mean-field simulations. Spiking simulations calculate the dynamics of every neuron in time and yield a detailed temporal evolution of the system including fluctuations. We start the spiking simulations with a precue period of 500 msec. In this precue period, neither the stimulus nor the response-direction input is applied to the PFC module. It represents the period before the stimulus onset in which the monkey does not have any information about the stimulus. Thereafter, we activate the stimulus input for 500 msec by increasing the external input ν_{ext} by 0.1 Hz (80 Hz in total). The response-direction input ν_{dir} is also present from the beginning of the cue period but lasts throughout the delay period until 1500 msec after stimulus onset with varying strength (0–0.1 Hz).

We use the mean-field formulation (Brunel & Wang, 2001) to assess the dependencies of the network parameters and identify the region of interest. The mean-field approach calculates the attractor states of the network. The attractor states represent the behavior of the network in terms of neural firing rates to which the system would converge to in the absence of fluctuations and external changes. We focus our mean-field analysis on the delay period of the system. During the delay period (500–1500 msec), the response-direction input is present, so we apply this input ν_{dir} in the mean-field simulation. The stimulus input has been active during the cue period and has increased the firing rates of the pools (AL, AR) or (BL, BR). Thus, we start the search for the attractor state at a higher initial frequency (40 Hz) for these pools. The other pools have an initial frequency of 3 Hz corresponding to the spontaneous firing.

We calculated the direction selectivity using a linear regression model. As dependent variable Y, we used the data of the pools AL and AR as these have the same stimulus selectivity. The data consisted of the firing rates

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

(averages of the spiking data in 20 msec bins for each trial) of 100 simulated trials for each parameter configuration and point in time. The independent variable X consisted of a vector denoting the response direction (e.g., 1 for L and −1 for R). In Figure 7, the data were additionally categorized in correct and error trials to allow a separate analysis of these trial types.

$$Y = Xb + e \tag{3}$$

We calculated the least-squares estimate of the linear regression coefficients,

$$b = \left(X^T X\right)^{-1} X^T Y \tag{4}$$

and used this estimate to calculate the proportion of the explainable variance by the sum of squares:

$$\mathrm{PEV}_{\mathrm{dir}} = \frac{(Xb)^T Xb}{Y^T Y} \tag{5}$$

The proportion of the explainable variance $\mathrm{PEV}_{\mathrm{dir}}$ represents the direction selectivity as also used in the experimental study (Pasupathy & Miller, 2005).

RESULTS

We used the mean-field technique to identify parameter regions, which fit our hypothesis. We looked for parameter configurations, which correspond to the properties of object-and-direction selective neurons. For example, when Stimulus A is associated with a leftward eye movement, the neurons in the pool AL should respond with a higher firing rate, but not the neurons of all other combinations (AR, BL, BR). We used a threshold of 10 Hz to determine high or low activation. Note that the network is symmetric and therefore it is sufficient to check just one combination of stimulus and response direction. We used Stimulus A and a leftward response direction for all simulations.

Figure 2 shows six slices of the three-dimensional parameter space (w_+, λ_{dir}, w_{m}) along the w_{m}-axis. The parameter w_+ connects neurons, which have the same selectivity and are thus highly correlated. The parameter w_{m} represents the connection weight between populations of neurons, which share the same directional selectivity. The populations which are connected by w_{a} are of opposite directional selectivity (see Figure 1). Within the gray region, the neurons show the property of object-and-direction selective neurons, that is, high activation (>10 Hz) in one pool and low activation (<10 Hz) in all other pools. The interdependency on the fourth parameter w_{a}, which is calculated using w_+ and w_{m}, shows that in all parameter configurations, in which the network neurons are object-and-direction selective, w_+ is larger than w_{a}, and w_{a} is larger than w_{m} ($w_+ > w_{\mathrm{a}} > w_{\mathrm{m}}$). This order of the connection strengths seems to be necessary to achieve the object-and-direction selective activity as demanded in our model description.

To evaluate the influence of the response-direction input, we choose a line in the parameter space along the direction input axis by fixing the two other parameters $w_+ = 1.8$ and $w_{\mathrm{m}} = 0.7$ (see dashed line in Figure 2). We

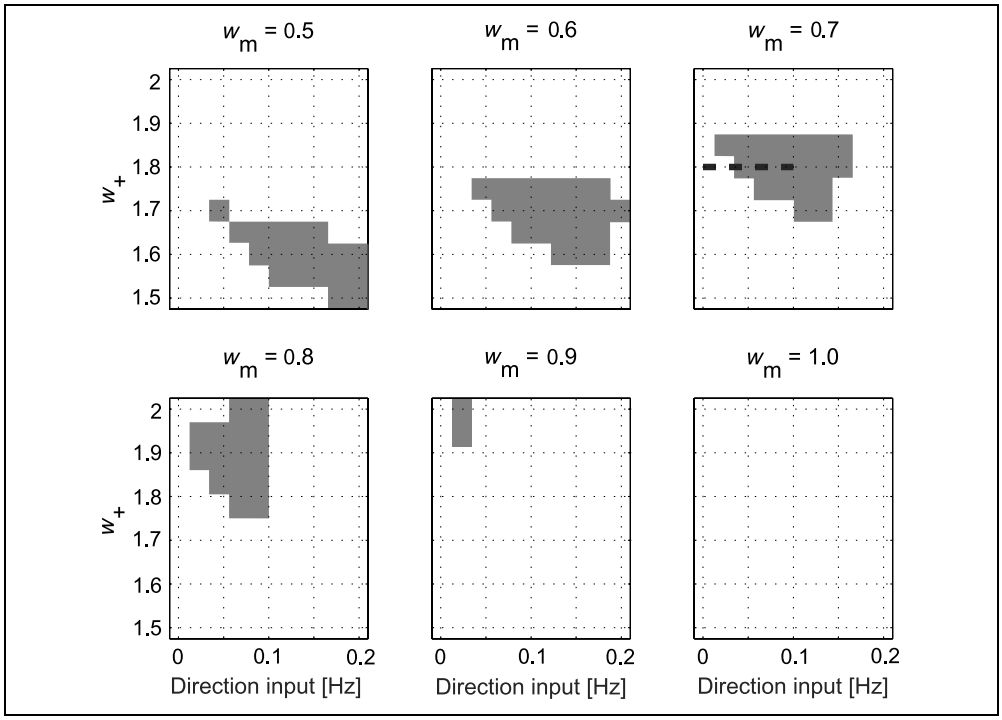

Figure 2. Mean-field analysis for object-and-direction selective patterns. The six panels show this property for slices of the parameter space with different values of w_{m} while varying w_+ and λ_{dir}. The mean-field analysis was performed for the parameter space w_+ ranging from 1.0 to 3.0, w_{m} from 0.5 to 1.0 and the response-direction input from 0.0 to 0.4 Hz. The parameter w_{a} was calculated as described above. We show just the part in which firing rates of the pools correspond to object-and-direction selective neurons. The dashed line corresponds to the parameters, which are discussed using spiking simulations ($w_+ = 1.8$, $w_{\mathrm{m}} = 0.7$, λ_{dir} from 0 to 0.1 Hz).

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

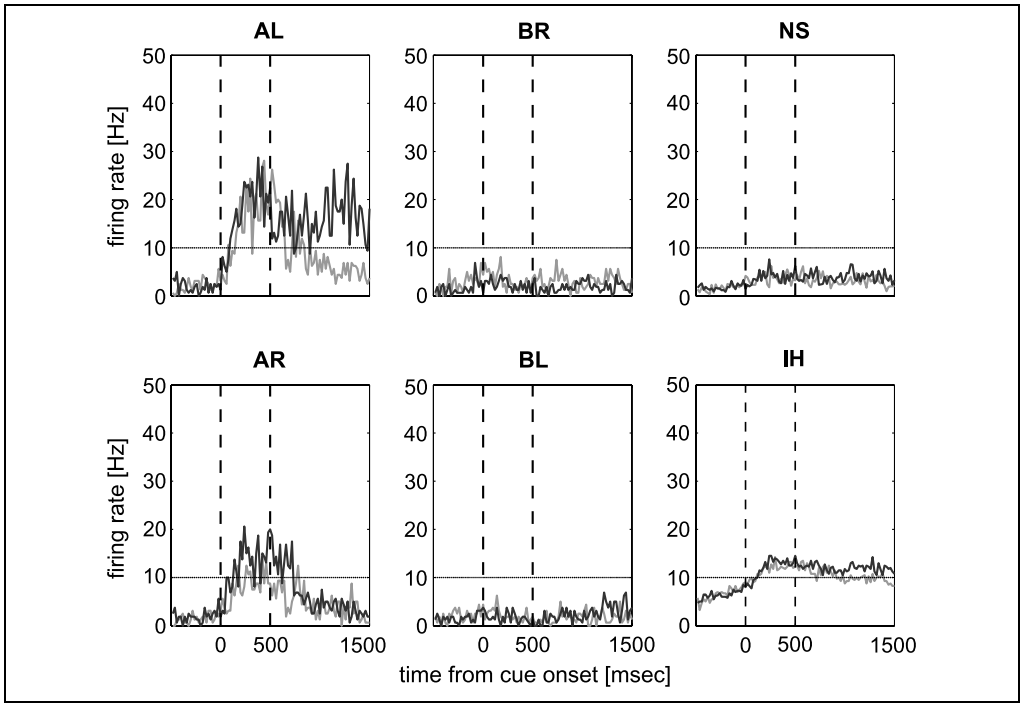

Figure 3. Two selected trials with same parameter configuration ($w_+ = 1.8$, $w_m = 0.7$, $\lambda_{dir} = 0.05$ Hz). The six graphs show the activity of all pools in time. The 10-Hz threshold is indicated and used to assess if a trial would identify an association correctly.

analyzed these parameters using multiple spiking trial simulations ($n = 100$ per parameter configuration).

Figure 3 shows an example of two simulated trials using the same parameter configuration ($w_+ = 1.8$, $w_m = 0.7$, and $\lambda_{dir} = 0.05$ Hz). Due to noise in the system, trials can end up in different attractor states. One trial shows a high activation in AL and the other one a low activation. Using a threshold criterion at 10 Hz to the end of the delay period, the two trials can be categorized as hit and misses, respectively.

First, we discuss the modeling data during the cue period. To facilitate the comparison to the experimental data, we also use a window of 300 msec, here between 100 and 400 msec after cue onset, which we call analogously to the experimental work ''peri-cue period.'' Pasupathy and Miller (2005) used a different offset, that is, the period between 300 and 600 msec after cue onset, due to the time delay until the visual information reaches the PFC. Figure 4A shows the average activity (and standard errors) during the peri-cue period as a function of the response-direction input of the pools AL and AR. These two pools represent the two conflicting directions for Stimulus A. The difference in firing rate increases with stronger response-direction input. Moreover, the activity for AL is steadily increasing, whereas the activity of AR stays at the same level. An analogous

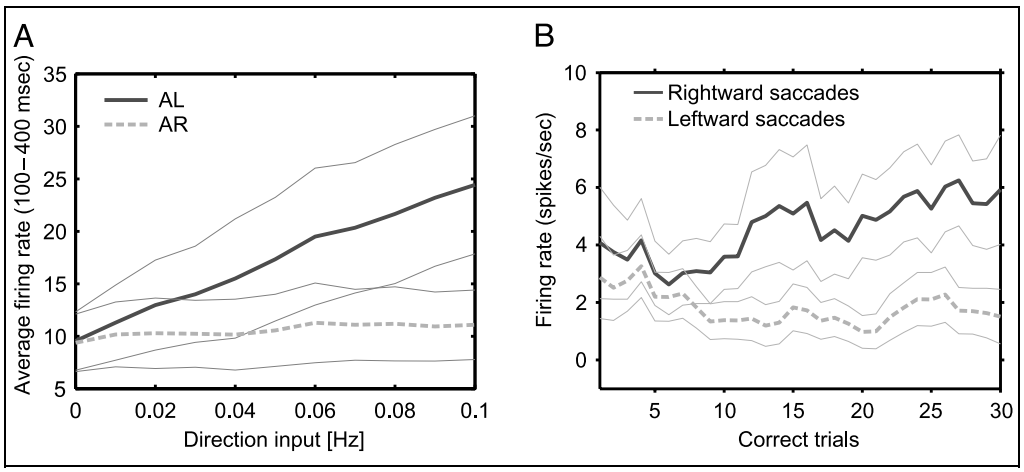

Figure 4. Average activity (and *SEM*) during peri-cue period for both model and experiment. (A) The average activity of the model calculated for the pools AL and AR between 100 and 400 msec as a function of direction input strength. The two pools show the same stimulus selectivity and code for the direction when Stimulus A is presented. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and $w_a = 0.9375$. (B) Average activity as a function of the correct trial number for the experimental data (adapted from Pasupathy & Miller, 2005).

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

observation is made in the experimental data during the learning phase in the experiment in consecutive correct trials (see Figure 4B).

To further quantify the activity in the peri-cue period, we calculated the average direction selectivity. Figure 5A shows that the dependency of the average direction selectivity on the response-direction input is approximately linear. This is consistent with the linear correlation between direction selectivity and consecutive correct trials as observed in the experimental data (Figure 5B). Note that in both Figure 4 and Figure 5, we already implicitly related the strength of the response-direction input to the correct trials in the experiment. The changes occurring during the cue period in consecutive correct trials in the experimental data correspond well to a linearly increasing response-direction input in the model.

The experimental data show that the monkeys' behavior correlates much better with the slower gradual changes in the neural activity in the PFC than with the almost bistable changes in the basal ganglia (see Figure 5B). Pasupathy and Miller (2005) concluded that the behavior is related more closely to the PFC activity. Motivated by this finding, we introduced a behavioral measure in our model: An association is correctly identified, if the mean of the last 500 msec of the simulation (1000–1500 msec after cue onset) of pool AL is above 10 Hz and of the other pools is below 10 Hz, analogously to the mean-field analysis. Figure 6 shows the evolution of the percent correctly identified associations with increasing response-direction input. With no direction input (0 Hz), the association cannot be identified, and all pools stay at low activation levels. An increasing direction input leads to a higher percentage of trials in which the AL pool is activated until the percentage reaches almost 100% at around 0.1 Hz. Note that when the system has not identified an association, the monkey could still guess. Thus, the performance could be at 50% (choice between L, R) when the identification level is at 0%.

We use this measure to compare the direction selectivity between correctly and incorrectly identified associations. Figure 7A shows the direction selectivity for a single direction bias (λ_{dir}), separated in correct and error trials. The progression of both cases is almost identical until 1000 msec after cue onset. After a rise due to the stimulus and response related input, the direction selectivity drops shortly around the end of the cue period (500 msec) before it increases again. The underlying mechanism of this phenomenon is based on the network dynamics. Stimulus and response-direction input are applied to excitatory populations. These drive the inhibitory population, which increase its activity with some time delay. Once the inhibitory effect increases, it reduces the direction selectivity of the network. After a transient low, the direction selectivity rises again, as the system settles in its stationary attractor, which has high direction selectivity. After 1000 msec from cue onset, the progression of the correct and error trials separates. This is because we categorize the trials in correct and error using the average firing rate between 1000 and 1500 msec. We used this period because the monkeys had to issue their response at the end of the delay period and, at this point in time, the information about the response direction must be available.

When comparing the time course of the direction selectivity of our model, which is based on hypotheses of the PFC, to the experimental data (Figure 7), we see that correct and error trials have a different time course. Whereas the direction selectivity in the model separates

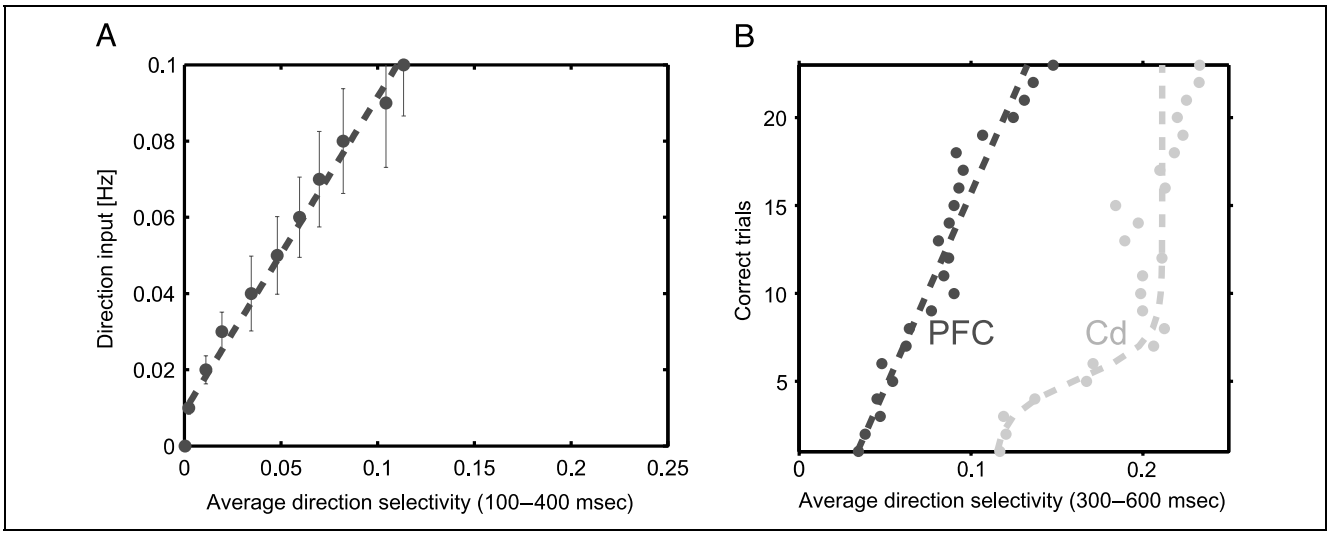

Figure 5. Average direction selectivity. (A) Average peri-cue period direction selectivity (100–400 msec) calculated from model data as a function of direction input. The dashed line shows a linear fit. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and $w_a = 0.9375$. (B) Average direction selectivity of experimental data during peri-cue epoch (adapted from Pasupathy & Miller, 2005). As in the model data, the progression of the PFC selectivity is almost linear.

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

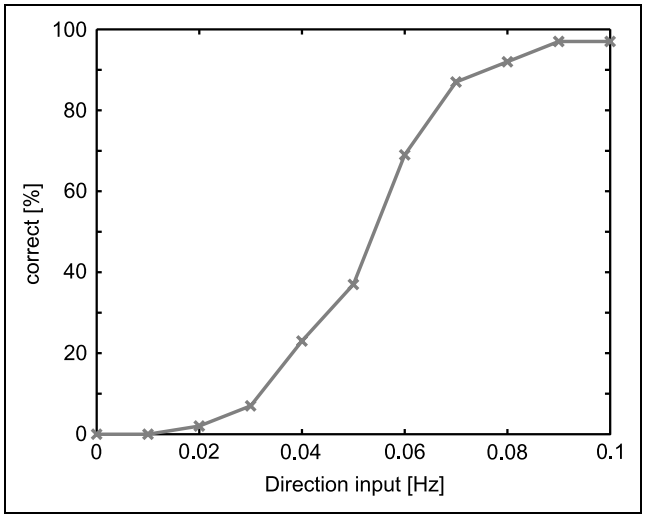

Figure 6. Analysis of multiple spiking simulations. The values of the parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and accordingly, $w_a = 0.9375$. The direction input strength is varied. The graph shows the percentage of correctly identified association. The mean of the last 500 msec of the delay period is calculated for every pool. An association is identified correctly if this mean of the pool connected to it (e.g., AL) has a value of above 10 Hz and the one of other three selective pools of below 10 Hz. For every data point, 100 spiking simulations were performed.

in the model in the middle of the delay period, this separation appears in the experimental data already in the cue period and disappears at the end of the delay period. After the good fit of the model data in the cue period, this result is surprising but also interesting. By comparing experimental data and hypotheses, we can discuss the hypotheses on which the computational model was constructed.

DISCUSSION

In this work we implemented a hypothesis-driven approach to modeling. Based on experimental evidences and hypotheses, we selected five cornerstones about the properties of a PFC module in arbitrary visuomotor mappings and integrated them in a computational model. Thereafter, we compared the results to experimental data. Note that a priori we did not seek to fit our model to experimental data. Rather, we wanted to test the validity of the hypotheses about the PFC by comparing the modeling results to experimental data. In the following, we will first picture the dynamics of our computational model and compare it to the possible neurodynamics underlying the experimental data.

We sketch the dynamical behavior of the model system in simplified cartoons of a hypothetical energy landscape (Figure 8). The energy landscape shows two valleys, which correspond to two attractors, one at low and one at high firing rates. By influence of fluctuations or external input, the system behavior can switch from one attractor state to another, analogously to a ball moving in the landscape which is at rest in the bottoms of the valleys and can be moved by external forces. Figure 8 shows the behavior over the three phases of the simulation: prestimulus, stimulus, and delay period. In the prestimulus period, the system is at rest in the low attractor state. Due to the stimulus and response-direction input, the system jumps up to a high attractor state both for correct and error trials. This can be seen by the increase in direction selectivity in the cue period (see Figure 7). When the direction selectivity increases, the difference in firing rate between two pools (in our simulations AL and AR) increases. Because we set up the network as a single attractor system, the system jumps to a high attractor state (in our simulations a high activity

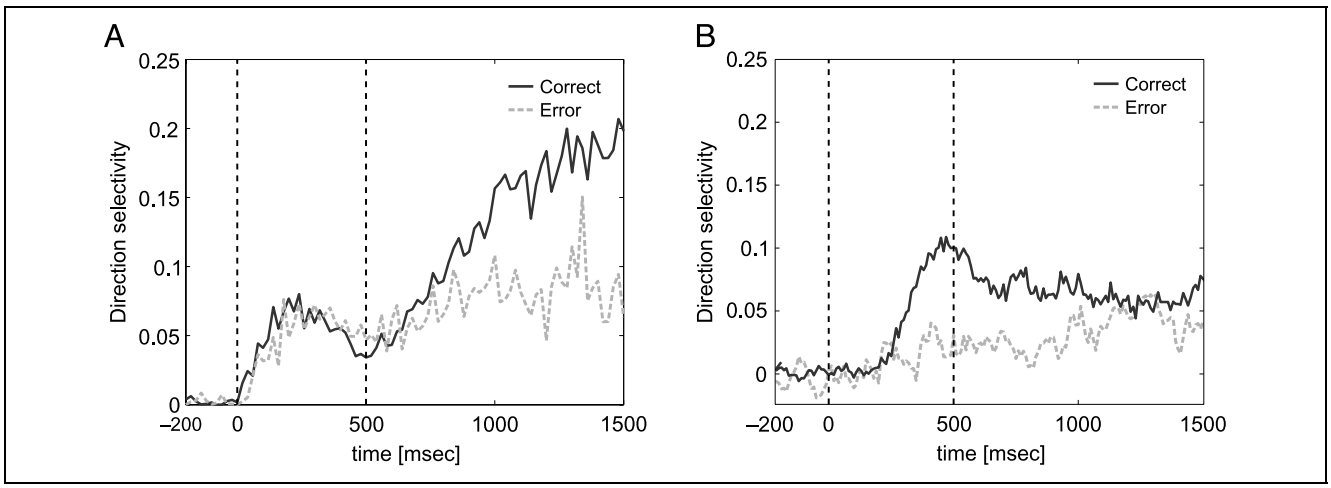

Figure 7. Direction selectivity of correct and error trials as a function of time. (A) Modeling data using a threshold criterion to categorize the trials: A trial is identified as correct if the average activity of the last 500 msec of the delay period (1000–1500 msec) is above 10 Hz for pool AL and below 10 Hz for all other pools. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, $w_a = 0.9375$, and $\lambda_{dir} = 0.06$ Hz, which yields a performance of 69%. (B) Experimental data of PFC direction selectivity (adapted from Pasupathy & Miller, 2005).

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

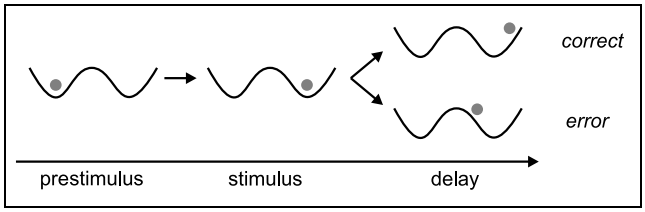

Figure 8. Sketch of the dynamical behavior of the model over time. We show in a hypothetical two attractor energy landscape for low and high activation the behavior of the model. In the prestimulus period, the system is at rest. The stimulus and response-direction input excite the system so that the firing rates increase by jumping to a high attractor state. The distinction between correct and error trials is caused by fluctuations in the high attractor state.

state of AL). After the cue period, the direction selectivity stays at high levels for both correct and error trials, indicating that, in both cases, the system stays in the high attractor state. The distinction between correct and error trials come about due to fluctuations among the fixed point of this attractor state. The system is stochastic and the firing rate can be lower or higher from trial to trial. The fluctuations in the system cause higher direction selectivity in some trials, and thus, the association to be extracted or not.

The analogous analysis of the experimental data yields a different view of the neurodynamics. The direction selectivity of correct and error trials already parts in the stimulus period (Figure 7B). This suggests that the system either jumps up to a high attractor state or stays in the low attractor state in the stimulus period, depending on a correct or an error trial (Figure 9). Thereafter, the direction selectivity of the correct trials drops and eventually reaches the same level as the error trials at the end of the delay period. Thus, the dynamical system drops back to the low activity state in the delay period for the correct trials.

Overall, there are differences in the neurodynamical properties. First, we will review the five assumptions which we stated in the Introduction and discuss their validity. Note that we achieved a particularly good fit to the correct trials in the cue period (see Figures 4 and 5), which suggests a relation between the strength of the response-related input and the learning in correct trials.

First, the model neurons should show nonlinear object-and-direction selective activity (Asaad et al., 1998). Using the mean-field technique, we specifically selected parameters which yielded an object-and-direction selective activity. We find that a specific ordering of the connection strength, namely, $w_+ > w_a > w_m$, is necessary to achieve this property. We analyzed a widespread parameter space, which contained all possible orderings of connection weights. The nonlinear object-and-direction selective activity is responsible for the directional selectivity observed in the model.

Second, stimulus information reaches the PFC from the IT. The IT is implicated in the processing of visual information about objects and anatomical studies suggest the presence of projections to the PFC (Pandya & Kuypers, 1969). Moreover, a disconnection of the PFC and the IT using cross-lesion procedures (Bussey et al., 2002; Parker & Gaffan, 1998; Gaffan & Harrison, 1988) or transaction of the uncinate fascicle (Eacott & Gaffan, 1992) causes deficits in learning arbitrary visuomotor mappings. Stimulus information itself is also present in the PFC during arbitrary visuomotor mappings (Asaad et al., 1998). Thus, the stimulus input to our module might originate either directly from the IT or from other PFC neurons or both. Because we do not model the originating brain regions explicitly, the exact origin of this input is not crucial for the simulation results. We assume this input to be constant in respect to learning. The perceived objects remained constant within the blocks, and thus, we also assumed their perception to be constant. Furthermore, imaging studies do not indicate a change in correlation between the IT and the PFC. The good match between our model and the experimental data during the cue period for the correct trials also supports this hypothesis.

Third, the basal ganglia work as a context detector and their output influences the activity of the PFC. The basal ganglia could generate a signal reflecting the response direction (Hadj-Bouziane et al., 2003). This would also be consistent with the idea of context extraction by the basal ganglia in terms of action planning (Houk & Wise, 1995). The basal ganglia are connected to the PFC in a cortico-basal ganglionic loop (Graybiel, 1998; Houk & Wise, 1995; Middleton & Strick, 1994). There are two hypothesis about the role of the basal ganglia in learning: One is that the basal ganglia have a dominant role in the learning of new associations (Pasupathy & Miller, 2005; Bar-Gad et al., 2003; Houk & Wise, 1995), whereas the other one suggests that they are responsible for consolidating associations, and thus, habits (Packard & Knowlton, 2002; Graybiel, 1998). We build upon the first hypothesis. We do not model the basal ganglia explicitly, but include the characteristics of the activity as identified by experi-

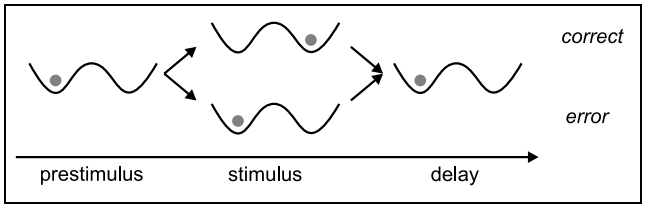

Figure 9. Sketch of the dynamical behavior of the experimental data over time based on Figure 7B. We show in a hypothetical two attractor energy landscape for low and high activation the behavior of the model. In the prestimulus period, the system is at rest. The stimulus and the response-direction input excite the system in a way that only correct trials jump to a high activity state, whereas error trials remain in the low activity state. Because the difference in direction selectivity disappears at the end of the delay period, we hypothesize that the correct trials jump out of the high attractor.

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

mental results in our modeling process. Thus, we focus on the input characteristics to the PFC and less on how the basal ganglia and the thalamus generate this input.

The response-direction input is applied at the same time as the stimulus input (onset at 0 msec). The simultaneous onset of direction and stimulus input is a first approximation, as the exact timing of the two potential inputs is not known. We suggest that stimulus information might reach both the PFC and the basal ganglia directly from the IT. Thereafter the basal ganglia extract the direction and transmit it via the thalamus back to the PFC. In addition, we assumed the onset of the response-direction input to be constant with learning. Pasupathy and Miller (2005) showed that the onset of the direction selectivity in the basal ganglia moved to an early position rapidly as learning progressed and reached a steady point in time after a few trials. The direction selectivity identified in the basal ganglia could be the origin of a response-direction input signal mediated by the thalamus. We varied the strength of that input. The underlying hypothesis is that slow synaptic plasticity in the synapses, which target PFC neurons from the thalamus, could enhance the connection strength, causing an increase in the response-direction input (Houk & Wise, 1995). This is consistent with observations of imaging studies: Both prefrontal areas and the basal ganglia were identified to take part in the arbitrary visuomotor associations (Toni, Ramnani, et al., 2001; Toni, Rushworth, & Passingham, 2001; Toni & Passingham, 1999) and, more importantly, the correlation between the two areas increased during learning (Boettiger & D'Esposito, 2005; Toni, Rowe, Stephan, & Passingham, 2002). The linear increase in the direction bias causes changes, which correspond well to the experimental observations during learning. Especially in the cue period, both the direction selectivity and average firing rate show good fits to experimental data.

Conceptually, the direction bias is different from the context bias implemented by earlier modeling approaches (Deco & Rolls, 2003, 2005; Loh & Deco, 2005). The context input would go to two pools, which represent the association active during an experimental trial, for example, AL and BR. This would mean that the two associations would not be learned independently. Learning the association AL would also favor BR. In the extreme case, the monkey could perform a one-trial reversal, that is, the monkey would switch context after the first error. The experiment by Thorpe, Rolls, and Maddison (1983) shows this evidence, which is modeled by Deco and Rolls (2005). Although such rule would greatly facilitate the execution of the task, there is no clear evidence in this conditional visuomotor experiment that the monkeys actually did make use of this relation.

Fourth, we assume that sustained cortical–thalamic loop feedback serves as a working memory of the directional information extracted by the basal ganglia. Houk and Wise (1995) explain the mechanism as follows: Striatal neurons in the basal ganglia receive input from the several cortical regions. By virtue of reinforcement signals originating from midbrain dopamine neurons, the striatal neurons learn to recognize complex activity patterns such as the association in a conditional visuomotor task. Via disinhibitory mechanism implemented by in the striatum–pallidum–thalamus connection, a detected direction might cause a positive feedback in the reciprocal cortical–thalamic loop. This activity could be self-sustained and thereby guide future actions such as an upcoming motor response to a stimulus. Our modeling data show that the stimulus and the response-direction input compute the correct association, so that the correct response can be issued. Furthermore, the association is maintained in the PFC module (see Figure 7). However, the experimental data show a significant decrease in direction selectivity after 1000 msec. These could be due to a variety of reasons, which are also related to the fifth hypothesis.

Fifth, we hypothesize that the PFC module holds the extracted association in working memory. There is much evidence that the PFC is involved in working memory related processes implemented by delayed neural activity (Fuster & Alexander, 1971; Kubota & Niki, 1971). Based on the fact that any neural system that contributes to visuomotor learning must retain information about the IS-response mapping (Buch, Brasted, & Wise, 2006), we reasoned that the PFC would be in the ideal position. The underlying neural model is set up for working memory properties (Brunel & Wang, 2001) and we chose the network parameter accordingly (see Methods section). However, the direction selectivity in the experimental data drops out in the middle of the delay period. This suggests that the information about the movement direction is not kept in the prefrontal area or that relatively few neurons may be needed to maintain this information during the delay period. Thus, either the working memory hypothesis of the PFC or the sustained activity in the cortico-thalamic loop could be a partly incorrect assumption.

The dynamical behavior could also be explained by other neurodynamical mechanisms. For example, the activity could propagate to other regions of the brain involved in the response. This idea is also supported by experimental evidences which show that the number of object-selective neurons decreases in the PFC from the beginning to the end of a trial, whereas the number of direction-selective neurons increases (Asaad et al., 1998). Propagation between layers has been observed in multilayer dynamical systems as studied by Loh and Deco (2005) and Deco and Rolls (2003). In addition, in these models, the activity of correct and error trials differ from the beginning of the cue period, which is consistent with the progression of the direction selectivity identified by Pasupathy and Miller (2005). Loh and Deco related the multistability in a multilayer neural network to response-space exploration in arbitrary visuomotor tasks. Thus,

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

the errors might not be caused by working memory failures in the PFC but by computational properties of the complete system supporting arbitrary visuomotor associations. In this sense, one could also view our model as a representation of a larger dynamical system compressed in one layer.

The model in this article was specifically constructed to target the activity in the PFC upon influences from connecting brain areas. Fusi, Asaad, Miller, and Wang (2007) presented a model based on the same experimental data (Pasupathy & Miller, 2005; Asaad et al., 1998), which features a learning rule and a reversal mechanism. We did not address this in our model. The two models make different claims about how the activity in the PFC arises. The model by Fusi et al. (2007) features one input to the PFC, which is modulated by learning. The learning rule is responsible for acquiring the associations and the reset after the reversals. We propose that the activity of the PFC arises due to the influence of several brain areas, the IT, and the basal ganglia. Thereby, we propose how the object-and-direction selective neurons might come about. If the input from the basal ganglia to the PFC indeed exists, then this could be tested directly experimentally by disconnecting the basal ganglia from the PFC. This experiment was proposed by Nixon, McDonald, Gough, Alexander, and Passingham (2004).

Our analysis shows that the dynamics of the PFC can arise due to several external influences, namely, the stimulus bias and the direction bias which we envision to originate from the IT area and the basal-ganglionic–thalamic–cortical loop, respectively. Neuropsychological and physiological evidences point in this direction. Furthermore, the dominant dynamics in the PFC seem to be transient because the comparison of model and experimental data does not suggest that a great population of neurons showing stable working memory properties exist. This does not imply that the PFC does not have working memory capabilities, but merely, in this task, it is not used to a great extent. The activity propagates along one or multiple processing pathways (Loh & Deco, 2005; Hadj-Bouziane et al., 2003) and thereby passes through the PFC. What is the role of the PFC in the processing pathway? Lesion studies suggest that the PFC is important for the fast learning (Bussey et al., 2001; Murray et al., 2000; Wang et al., 2000) of arbitrary visuomotor associations. Perhaps in conjunction with hippocampal regions, the PFC adds a fast learning component and thereby speeds up learning. The proposed plasticity in the cortico-thalamic loop (Houk & Wise, 1995) is a possible candidate. A second processing pathway from the IT and the basal ganglia to the premotor cortex might provide slower learning mechanisms as disconnection of the basal ganglia and the premotor cortex affects the execution of well-established associations but spares fast learning mechanisms (Nixon et al., 2004). An experiment disconnecting the basal ganglia and the PFC could clarify this issue.

In our model, we feature a hypothesis-driven approach. This means that, a priori, we did not seek to fit model data but wanted to test several existing hypotheses by integrating them in a computational model. Computational modeling is ideal for this approach because it allows us to both integrate several hypotheses in one system and to study the contributions of each hypothesis separately. Moreover, the language of computational models is explicit, and thus, forces one to make concrete assumptions. We believe that the presented way of modeling, namely, to test existing hypotheses instead of seeking model fitting to experimental data, might contribute significantly to the theoretical understanding of brain dynamics, as it emphasizes the idea of integrating neuroscientific evidences toward comprehensive models of brain functions.

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021

Acknowledgments

Marco Loh was supported by the Boehringer Ingelheim Fonds. We thank Ralph G. Andrzejak and Anders Ledberg for discussions and insights, which improved our work considerably.

Reprint requests should be sent to Marco Loh, Universitat Pompeu Fabra, Passeig de Circumval·lació, 8, 08003 Barcelona, Spain, or via e-mail: marco.loh@upf.edu.

REFERENCES

Asaad, W., Rainer, G., & Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron, 21,* 1399–1407.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Progress in Neurobiology, 71,* 439–473.

Boettiger, C. A., & D'Esposito, M. (2005). Frontal networks for learning and executing arbitrary stimulus–response associations. *Journal of Neuroscience, 25,* 2723–2732.

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.

Buch, E. R., Brasted, P. J., & Wise, S. P. (2006). Comparison of population activity in the dorsal premotor cortex and putamen during the learning of arbitrary visuomotor mappings. *Experimental Brain Research, 169,* 69–84.

Bussey, T. J., Wise, S. P., & Murray, E. A. (2001). The role of ventral and orbital prefrontal cortex in conditional visuomotor learning and strategy use in rhesus monkeys (Macaca mulatta). *Behavioral Neuroscience, 115,* 971–982.

Bussey, T. J., Wise, S. P., & Murray, E. A. (2002). Interaction of ventral and orbital prefrontal cortex with inferotemporal cortex in conditional visuomotor learning. *Behavioral Neuroscience, 116,* 703–715.

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.

Deco, G., & Rolls, E. T. (2005). Synaptic and spiking dynamics underlying reward reversal in the orbitofrontal cortex. *Cerebral Cortex, 15,* 15–30.

Eacott, M. J., & Gaffan, D. (1992). Inferotemporal–frontal disconnection: The uncinate fascicle and visual associative learning in monkeys. *European Journal of Neuroscience, 4,* 1320–1332.

Fusi, S., Asaad, W. F., Miller, E. K., & Wang, X. J. (2007). A neural circuit model of flexible sensorimotor mapping: Learning and forgetting on multiple timescales. *Neuron, 54,* 319–333.
Fuster, J., & Alexander, G. (1971). Neuron activity related to short-term memory. *Science, 173,* 652–654.
Gaffan, D., & Harrison, S. (1988). Inferotemporal–frontal disconnection and fornix transection in visuomotor conditional learning by monkeys. *Behavioural Brain Research, 31,* 149–163.
Gaffan, D., & Harrison, S. (1989). A comparison of the effects of fornix transection and sulcus principalis ablation upon spatial learning by monkeys. *Behavioural Brain Research, 31,* 207–220.
Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiology of Learning and Memory, 70,* 119–136.
Hadj-Bouziane, F., Meunier, M., & Boussaoud, D. (2003). Conditional visuo-motor learning in primates: A key role for the basal ganglia. *Journal of Physiology (Paris), 97,* 567–579.
Houk, J. C., & Wise, S. P. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: Their role in planning and controlling action. *Cerebral Cortex, 5,* 95–110.
Kubota, K., & Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. *Journal of Neurophysiology, 34,* 337–347.
Loh, M., & Deco, G. (2005). Cognitive flexibility and decision making in a model of conditional visuomotor associations. *European Journal of Neuroscience, 22,* 2927–2936.
Middleton, F. A., & Strick, P. L. (1994). Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science, 266,* 458–461.
Murray, E. A., Bussey, T. J., & Wise, S. P. (2000). Role of prefrontal cortex in a network for arbitrary visuomotor mapping. *Experimental Brain Research, 133,* 114–129.
Nixon, P. D., McDonald, K. R., Gough, P. M., Alexander, I. H., & Passingham, R. E. (2004). Cortico-basal ganglia pathways are essential for the recall of well-established visuomotor associations. *European Journal of Neuroscience, 20,* 3165–3178.
Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annual Review of Neuroscience, 25,* 563–593.
Pandya, D. N., & Kuypers, H. G. (1969). Cortico-cortical connections in the rhesus monkey. *Brain Research, 13,* 13–36.
Parker, A., & Gaffan, D. (1998). Memory after frontal/temporal disconnection in monkeys: Conditional and non-conditional tasks, unilateral and bilateral frontal lesions. *Neuropsychologia, 36,* 259–271.
Passingham, R. E. (1993). *The frontal lobes and voluntary action*. Oxford: Oxford University Press.
Passingham, R. E., Toni, I., & Rushworth, M. F. (2000). Specialisation within the prefrontal cortex: The ventral prefrontal cortex and associative learning. *Experimental Brain Research, 133,* 103–113.
Pasupathy, A., & Miller, E. K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature, 433,* 873–876.
Petrides, M. (1982). Motor conditional associative-learning after selective prefrontal lesions in the monkey. *Behavioural Brain Research, 5,* 407–413.
Petrides, M. (1985). Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. *Neuropsychologia, 23,* 601–614.
Petrides, M. (1990). Nonspatial conditional learning impaired in patients with unilateral frontal but not unilateral temporal lobe excisions. *Neuropsychologia, 28,* 137–149.
Petrides, M. (1997). Visuo-motor conditional associative learning after frontal and temporal lesions in the human brain. *Neuropsychologia, 35,* 989–997.
Thorpe, S. J., Rolls, E. T., & Maddison, S. (1983). The orbitofrontal cortex: Neuronal activity in the behaving monkey. *Experimental Brain Research, 49,* 93–115.
Toni, I., & Passingham, R. E. (1999). Prefrontal–basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: A PET study. *Experimental Brain Research, 127,* 19–32.
Toni, I., Ramnani, N., Josephs, O., Ashburner, J., & Passingham, R. E. (2001). Learning arbitrary visuomotor associations: Temporal dynamic of brain activity. *Neuroimage, 14,* 1048–1057.
Toni, I., Rowe, J., Stephan, K. E., & Passingham, R. E. (2002). Changes of cortico-striatal effective connectivity during visuomotor learning. *Cerebral Cortex, 12,* 1040–1047.
Toni, I., Rushworth, M. F., & Passingham, R. E. (2001). Neural correlates of visuomotor associations. Spatial rules compared with arbitrary rules. *Experimental Brain Research, 141,* 359–369.
Wang, M., Zhang, H., & Li, B. M. (2000). Deficit in conditional visuomotor learning by local infusion of bicuculline into the ventral prefrontal cortex in monkeys. *European Journal of Neuroscience, 12,* 3787–3796.

Downloaded from http://direct.mit.edu/jocn/article-pdf/20/3/421/1759404/jocn.2008.20031.pdf by guest on 07 May 2021