Biol Cybern (2006)
DOI 10.1007/s00422-006-0054-z

ORIGINAL PAPER

Biological Cybernetics

Miruna Szabo · Gustavo Deco · Stefano Fusi
Paolo Del Giudice · Maurizio Mattia · Martin Stetter

Learning to attend: modeling the shaping of selectivity in infero-temporal cortex in a categorization task

Received: 22 November 2005 / Accepted: 22 December 2005 / Published online: ■■

Abstract Recent experiments on behaving monkeys have shown that learning a visual categorization task makes the neurons in infero-temporal cortex (ITC) more selective to the task-relevant features of the stimuli (Sigala and Logothetis in Nature 415:318–320, 2002). We hypothesize that such a selectivity modulation emerges from the interaction between ITC and other cortical area, presumably the prefrontal cortex (PFC), where the previously learned stimulus categories are encoded. We propose a biologically inspired model of excitatory and inhibitory spiking neurons with plastic synapses, modified according to a reward based Hebbian learning rule, to explain the experimental results and test the validity of our hypothesis. We assume that the ITC neurons, receiving feature selective inputs, form stronger connections with the category specific neurons to which they are consistently associated in rewarded trials. After learning, the top-down influence of PFC neurons enhances the selectivity of the ITC neurons encoding the behaviorally relevant features of the stimuli, as observed in the experiments. We conclude that the perceptual representation in visual areas like ITC can be strongly affected by the interaction with other areas which are devoted to higher cognitive functions.

Electronic Supplementary Material: Supplementary material is available in the online version of this article at http://dx.10.007/s00422-006-0054-z.org/.

M. Szabo · M. Stetter (✉)
aff Siemens AG, Corporate Technology,
Information and Communications, Otto-Hahn-Ring 6,
81739 Munich, Germany
E-mail: stetter@siemens.com
Tel.: +49-89-63655734
Fax: +49-89-63649767

M. Szabo
Technical University of Munich, Department of Computer Science,
85747 Garching, Germany

G. Deco
Institució Catalana de Recerca i Estudis Avançats (ICREA),
08010 Barcelona, Spain

G. Deco
Universitat Pompeu Fabra, Department of Technology Computational Neuroscience, Passeig de Circumval.lació 8, 08003 Barcelona, Spain

S. Fusi
ETH / University Zürich, Institute of Neuroinformatics,
8057 Zürich, Switzerland

P. D. Giudice · M. Mattia
Istituto Superiore di Sanitá, Department of Technologies and Health,
v.le Regina Elena 299, 00161 Roma, Italy

1 Introduction

Perceptual learning represents an important cognitive process that involves structural and functional modifications of the brain following sensorial experience, which leads to improvements in performance with training or practice (Goldstone 1998). Different studies show that neurons from higher stages of visual processing become tuned to some particular patterns of the input. These changes in the response properties of cortical neurons, that are supposed to be mediated by higher cognitive top-down inputs and attention, are associated with perceptual learning (Fine and Jacobs 2002).

Infero-temporal Cortex (ITC) and pre-frontal cortex (PFC) are two interconnected cortical areas thought to be involved in visual tasks, such as visual recognition, categorization and memory, although the contribution of each of these two areas for visual processing is not fully understood. In this context, recent studies have suggested that PFC is mainly associated with cognitive processing (such as categorization), while ITC is more associated with feature processing, (Freedman et al. 2003). In a recently performed neurophysiological experiment on behaving monkeys, Sigala and Logothetis have studied how the representation of visual stimuli in ITC is affected by their behavioral relevance, measuring the activity level of single infero-temporal cortical neurons during a visual categorization task (Sigala and Logothetis 2002). Their results show an enhancement in neuronal tuning for the values of the diagnostic features (Fig. 1b, top panel) as compared to the responses to the non-diagnostic features (Fig. 1b, lower panel). Further studies suggest that top-down signals from PFC could

partially determine ITC neuronal responses (Freedman et al. 2003; Tomita et al. 1999).

Taking into account all these findings on perceptual learning, higher visual processing and the tuning of ITC neurons after learning a visual categorization task (as reported by Sigala and Logothetis 2002), we hypothesize that the enhancement of selectivity to the diagnostic features in ITC (see Fig. 1b) might emerge, in the behavioral context, through a higher level cognitive feedback originating from category encoding neurons, possibly residing in the PFC.

In order to test our hypothesis and account for the experimental results, we propose a neurodynamical computational model in the framework of biased competition and cooperation. By biased competition we assume that multiple activated populations of neurons engage in competitive interactions that are biased by external interactions in favor of specific groups of neurons highlighting the attended or task-relevant stimuli or stimulus-features. Neurodynamical models developed within the conceptual framework of the *Biased Competition Hypothesis* (Moran and Desimone 1985; Chelazzi et al. 1993; Desimone and Duncan 1995; Chelazzi 1998; Reynolds and Desimone 1999) have been proved to successfully account for different aspects of visual attention (Rolls and Deco 2002; Corchs et al. 2003) and working memory context dependent tasks (Deco and Rolls 2003; Deco et al. 2004). Cooperation, on the other hand, promotes the co-activation of neuronal populations that represent stimuli or features associated with each other (Szabo et al. 2004; Almeida et al. 2004). For these models, competition is implemented through weak excitatory lateral connectivity and is mediated through the global inhibitory signal and cooperation is implemented through stronger than average excitatory lateral interactions.

The proposed biologically inspired two-layer model of excitatory and inhibitory spiking neurons simulates two small interconnected areas in the brain comprizing the visual responsive units from ITC and the category encoding units from PFC. Preliminary work (Szabo et al. 2005) explored network's behavior using the simple meanfield formulation and showed that the two-layer architecture can account for the referred tuning effect for a specific structure of the interlayer connectivity. For an already trained network, the connections between the ITC neurons and PFC category encoding neurons are structured in such a way that the bottom up input from the populations in ITC coding for relevant features activates the corresponding category neurons. Once active, these neurons, by top-down influence, bias the activity of the ITC neurons so that their representation becomes tuned for the behaviorally relevant features. In this work, the model is extended to include a learning scenario that reliably produces the correct associations. The network learns to attend: learning to categorize correctly enhances the representation of behaviorally relevant information in the sense that during learning, the model gradually 'attends' only to the relevant features which in turn improves the performance of categorizing correctly. The attentional signals that modulate the ITC representations are thus given by the top-down information coming from PFC.

The network is trained using a reward-based Hebbian learning algorithm which is compatible with a learning prescription that quantitatively reproduces the behavior of the recorded cortical activity of monkeys trained to learn visuo-motor associations in a continuously changing environment (Asaad et al. 1998; Fusi et al. 2005). We show that the adopted learning prescription robustly converges to a stable fixed point of the learning dynamics where the tuning effect is correctly reproduced.

2 Methods

2.1 The network model

We use a recurrent network of excitatory and inhibitory neurons that was introduced in earlier works (for example see Amit and Brunel 1997a; Brunel and Wang 2001), and that has been extended and successfully applied to explain several experimental paradigms (Szabo et al. 2004; Deco and Rolls 2003; Deco et al. 2004). We assume that a proper level of description at the neuronal level is captured by the spiking and synaptic currents dynamics of one-compartment *integrate-and-fire* neuron models. The details of the mathematical formulation are summarized in previous works (Brunel and Wang 2001; Szabo et al. 2004), and are provided in the supplementary material for the convenience of the reader. Following the results of Freedman et al. (2003) we structured the model as a two layer network: first layer modeling a small part of ITC where visual features are represented and the other modeling a small PFC region where according to their results, the categories should be represented. The ITC model layer, consisting of 1,000 spiking neurons, receives external information about the presented stimulus features and interacts with the PFC model layer, consisting of 650 spiking neurons, where the learned categories are encoded. The parameters for each layer are chosen the same as the ones opted for in previous works (Brunel and Wang 2001).

Both layers are organized into a discrete non-overlapping set of populations, as depicted in Fig. 2. Populations are defined as groups of excitatory or inhibitory neurons sharing the same inputs and connectivities. The specific populations have a specific function in the present task. In addition there are one 'Non-specific' population which groups all other excitatory neurons in the modeled brain area not involved in the present tasks and one 'Inhibitory' population grouping the local inhibitory neurons in the modeled brain area. The latter population regulates the overall activity and implements competition in the network by spreading a global inhibition signal.

In the Sigala and Logothetis experiment, the monkeys learned to categorize a set of images (schematic face and fish stimuli) into two categories, each associated with one lever (the left lever or the right lever) that the monkeys had to pull when the corresponding stimulus was presented. The stimuli were characterized by several features, each varying in a discrete small set of values as shown in Fig. 1a. Only

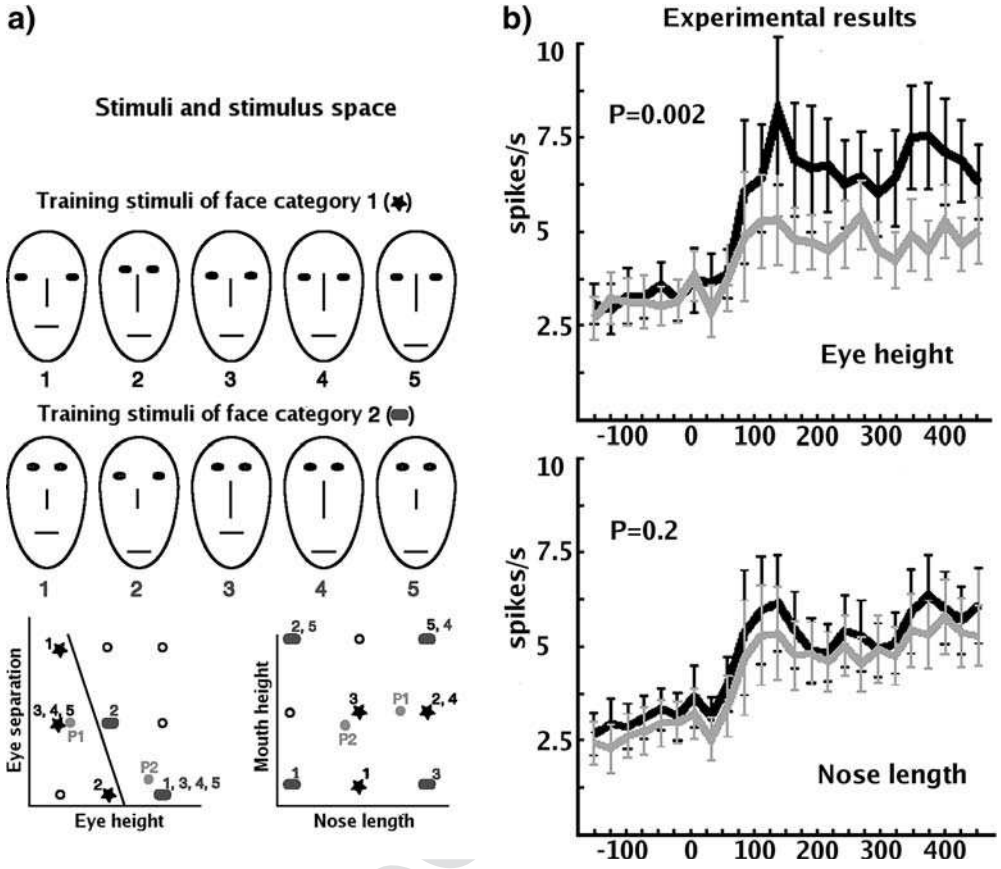

Fig. 1 **a** Stimuli and stimulus space for the visual categorization task (Sigala and Logothetis 2002). The stimuli have four varying features: 'Eye height', 'Eye separation', 'Nose length' and 'Mouth height', and can be linearly separated in two categories along two of the four dimensions: 'Eye height' and 'Eye separation'. **b** Experimental results adapted from Sigala and Logothetis (2002). Shown are the average spiking rates of all recorded ITC visually responsive neurons (a total of 96 units). For each neuron, the responses were sorted by the presented stimulus feature values and averaged over many trials. The resulting average activity levels reflect the feature values which excite a given neuron most and least, respectively. The population average activation trace was calculated by grouping these average activity levels according to their best (*black lines*) and worst (*gray lines*) responses to the levels of diagnostic feature 'Eye height' (*top panel*) and non-diagnostic feature 'Nose length' (*bottom panel*)

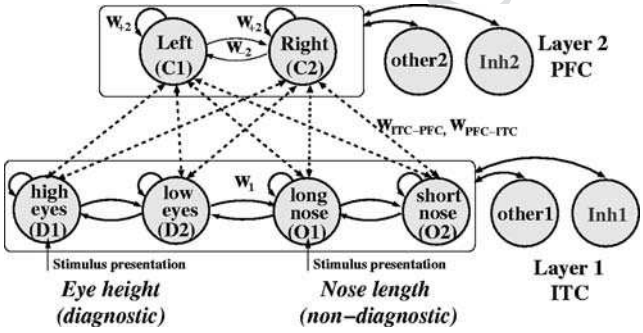

Fig. 2 Schematic representation of the two-layer model architecture. The weights between the two layer's specific populations (*dashed lines*) are modified during learning. For further details see text

two of the features, referred to as diagnostic, were informative for solving the categorization task. The two categories could be linearly separated along the two diagnostic features in the stimulus space, as depicted in Fig. 1a. The other two features, referred to as non-diagnostic, gave no information about the stimulus associated category and were irrelevant for the task. In our minimal model, we assume that the presented stimuli are characterized by only two features, 'Eye height' and 'Nose length', each with two discrete values, and that the two categories are determined exclusively only by one feature: the diagnostic feature 'Eye height'. Thus there are four specific populations in the ITC layer, denoted according to the specific input that they receive: one population receives input when the stimulus is characterized by the 'Eye height' feature being in the first state ('high eyes' population, or D1), one when the 'Eye height' feature is in the second state ('low eyes' population, or D2), one when the 'Nose length' feature is in the first state ('long nose' population, or O1) and one when the 'Nose length' feature is in the second state ('short nose' population, or O2). They encode the same type of stimulus and are differentiated only by their specific preferences to the stimuli feature values.

The neural activity of the PFC model layer is designed to reflect the category to which the presented stimulus corresponds. Thus, the specific populations in this layer encode

the two learned categories associated in the behavioral task with the motor actions of pulling one of the two levers: left lever ('Left' population, or C1) and right lever ('Right' population, or C2). Each category is defined through the set of stimuli associated with one required action. The stimuli with the diagnostic feature in the first state, 'high eyes', belong to category 1 and the ones with diagnostic feature in the second state, 'low eyes', belong to category 2, irrespective of the value of the non-diagnostic feature 'Nose length'.

Each individual population is affected by a spontaneous background excitatory input from outside the module along with the recurrent activity inside the module. The neurons in the four specific populations from the ITC layer additionally receive external inputs encoding stimulus specific information. They are assumed to originate from lower areas which process the visual scene such as to provide these signals. We assume that along the visual pathway, all the perceived stimulus features are processed and encoded in the same way, so that the 'bottom-up' signals coming to the ITC layer encoding the presented stimulus feature values on average have the same strength. When stimulating an ITC population, the rate of the background external input to the neurons of this population is increased by a fixed value of 150 Hz.

We modulate the conductance values for the synapses between pairs of neurons by connection weights, which can deviate from their default value 1. The structure and function of the network is achieved by differentially modulating these weights within and between populations of neurons. The labeling of the weights is defined in Fig. 2. Structurally, the network is fully connected within layers by excitatory and inhibitory synapses. Between the two layers, only the specific neurons are fully connected by excitatory synapses.

In our approach we assume, for simplicity, that the intralayer connections are already formed, e.g., by earlier self organization mechanisms, thus we fix their strengths as follows. Cooperation is implemented in the ITC model layer by setting all interconnecting and recurrent weights for the specific neurons equal to the default value w_1=1. This might arise from possible correlations between these neurons, which are all responsive to face stimuli. Second, we assume that the two categories associated to the two actions 'pull right lever' and 'pull left lever', are already encoded in the PFC model layer, in the sense that the monkey is already trained that pulling one or the other lever, but not both, might bring him some reward. The populations encoding for these two categories are likely to have anti-correlated activity in this behavioral context, resulting in weaker than average connections between them, denoting competition in this layer. We choose the extreme case w_{-2}=0 where there is no direct excitatory connection between the two category populations in the ITC. Within one category population, we set all connections to the default value w_{+2}=1. The weights from and to the non-specific neurons were computed such that the average of all excitatory connection weights to each specific neuron is 1 (see Brunel and Wang 2001), resulting in the value w_n=0.93 for both layers. All the connections from and to the inhibitory populations as well as the recurrent connections for the non-specific and inhibitory populations in both layers are set to the default value w=1.

The connections between the ITC and PFC are modeled as plastic synapses. Their absolute strengths, which we choose as free parameters, are learned using the reward-based Hebbian learning algorithm presented in the next section.

We will characterize network's different modes of operation corresponding to different parameter regimes by exploring the connecting weights between the two layers. Although explicit simulations of the network dynamics accurately capture the temporal dynamics and any order of the spike statistics, they are computationally expensive and not easy to use for systematic parameter explorations. Therefore we use mean-field models, which represent a well-established means for efficiently analyzing the approximate network behavior (Amit and Brunel 1997b; Stetter 2002; Del Giudice et al. 2003), at least for the stationary conditions (i.e., after the dynamical transients), in order to systematically explore parameter regimes of qualitatively different network response. For this, we use a recent derivation by Brunel and Wang (Brunel and Wang 2001), which is consistent with the type of conductance-based neuron networks we simulate (also provided in the supplementary material).

2.2 Learning mechanism

Our goal for the learning procedure is to enable the network capability of associating a set of stimuli, characterized by different combinations of the feature values, with a certain category, in a biologically plausible manner. The relevance to behavior of different features of the presented stimuli is determined through the consequences – receiving reward or not – of the selected action. The monkey learns to associate specific values of the diagnostic features with the corresponding category by evaluating the received reward. To model this learning strategy, we construct a reward-based Hebbian learning prescription that modifies, after each trial, the synaptic efficacies between the ITC and PFC layers according to the resulting network activities and reward signal using a simple regulatory mechanism. We show in the result section that this learning prescription robustly modifies the network free parameters to reach a configuration where the desired association is correctly performed.

On the network level, assuming that the monkey learns to categorize correctly using a trial-and-error strategy, the reward based model of learning, as in reinforcement learning models (Sutton and Barto 1998), is the biologically plausible choice from the well established learning strategies in the artificial neural networks field.

On the single synapse level, we consider a biologically inspired Hebbian learning scheme, as described in Fig. 3. Follwing reward, a synapse is potentiated if the presynaptic and post-synaptic neurons are simultaneously active; depressed if the presynaptic neuron is active but the post-synaptic neuron is inactive; and not modified otherwise. After non-reward, the synapse is depressed if both the presynaptic

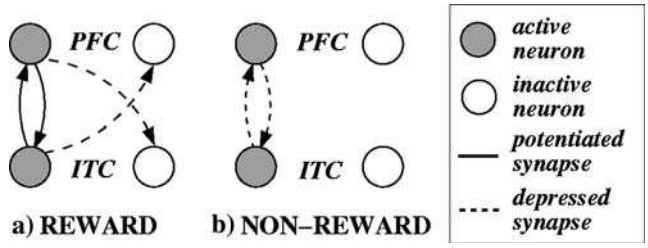

Fig. 3 Modified Hebbian learning scheme applied for reward (**a**) and non-reward (**b**) cases. The *filled blobs* represent 'active' neurons and the *empty blobs* represent 'inactive' neurons. The potentiated synapses are represented by *bold lines* and the depressed synapses by *dotted lines*. The connections that are not marked are the ones not modified

and post-synaptic neurons are simultaneously active and not modified otherwise.

We describe the learning dynamics through a "synaptic mean field" approximation, which captures, for computational convenience, the average synaptic dynamics between two given populations. For the hidden single-synapse dynamics, each synapse is binary (one 'potentiated' and one 'depressed' state) and undergoes stochastic transitions between the two available states depending on the pre- and post-synaptic neural activity (Fusi et al. 2000; Fusi 2002; Amit and Fusi 1994). At the mean-field level adopted here, all the synapses connecting neurons belonging to the same pair of pre- and post-synaptic populations are forced to have the same weight; the latter is updated by first computing the fraction of synapses that would get potentiated or depressed on the basis of the estimated firing rate distributions, as detailed in the next section; then computing the resulting average weight and, finally, assigning that value as the common new weight for all the synapses of the group, after additional normalization. In order for this to be a good description of what would result from the detailed synaptic dynamics, the non-trivial distribution of firing rates inside each neural population plays an important role; this is one reason why we deemed important to keep the detailed spiking dynamics of neurons in the face of the more abstract description of learning.

During learning, the network configuration changes with the modification of the interlayer synaptic weights from an initial, chosen configuration to a final, learned, configuration. The stimuli are presented to the network in a random order. We first reset all network internal variables and then simulate the spike dynamics for 500 ms of spontaneous activity followed by 800 ms under specific input encoding the presented stimulus. The simulation uses the present network configuration, given by the synaptic weight variables. For the period of time when the stimulus is presented to the network, the first 300 ms are regarded as transient time, and only the last 500 ms are used to acquire the time-averaged spiking rates of each simulated neuron. These are used to calculate the population firing distributions, the reward variable and finally the synaptic modifications as presented in more detail in the following subsection. The simulation runs until convergence to a stable configuration is reached.

For the typical average firing rates in our simulations, the 500 ms time window used to estimate single neuron rates implies non-negligible fluctuations in the estimated values. As a consequence, despite the full synaptic connectivity and the common value of the synaptic efficacies for each synaptic population, a wide distribution of estimated firing rates in each neuron population arises for the trial. This brings non-trivial consequences in the mean-field learning dynamics, in that superimposing tails of the rate distributions for different pairs of populations can induce unwanted potentiations or depressions, thereby pushing the learning trajectory to wrong directions. We decided to keep this feature to show the robustness of the model to the finite-size effects of various kinds that would affect the dynamics in a less constrained and more realistic setting.

2.3 Implementation of the learning algorithm

We first calculate the fraction of active neurons, n_i^a in each population i, by comparing the previously computed time-averaged spiking rate of each neuron inside this population with a chosen threshold: above 8 Hz for the ITC model layer and 14 Hz for the PFC model layer a neuron is considered to be active. When the population encoding the correct category has more than half neurons active and also there are more than twice neurons active in this population than in the other category population, the trial is assigned a reward, otherwise no reward is given. Next, for each ITC–PFC pair of specific populations, the fraction of synapses *to be* potentiated, N^{p}, and *to be* depressed, N^{d}, as the result of this stimulus presentation are evaluated.

Consider a presynaptic population with n_{pre} neurons and n^a_{pre} active neurons and a post-synaptic population with n_{post} neurons and $n^{\mathrm{a}}_{\mathrm{post}}$ active neurons. In case of reward, all synapses between pairs of active neurons are potentiated and all synapses from an active neuron to an inactive neuron are depressed (as described in Fig. 3a). Thus the fraction of synapses *to be* potentiated and depressed in the reward case are given by:

$$N^{\mathrm{p}}_{\mathrm{pre-post}} = \frac{n^{\mathrm{a}}_{\mathrm{pre}} \cdot n^{\mathrm{a}}_{\mathrm{post}}}{n_{\mathrm{pre}} \cdot n_{\mathrm{post}}}. \quad (1)$$

$$N^{\mathrm{d}}_{\mathrm{pre-post}} = \frac{n^{\mathrm{a}}_{\mathrm{pre}} \cdot (n_{\mathrm{post}} - n^{\mathrm{a}}_{\mathrm{post}})}{n_{\mathrm{pre}} \cdot n_{\mathrm{post}}}. \quad (2)$$

In case of non-reward, all synapses between pairs of active neurons are depressed and there are no synapses potentiated (as described in Fig. 3b), thus we have:

$$N^{\mathrm{p}}_{\mathrm{pre-post}} = 0. \quad (3)$$

$$N^{\mathrm{d}}_{\mathrm{pre-post}} = \frac{n^{\mathrm{a}}_{\mathrm{pre}} \cdot n^{\mathrm{a}}_{\mathrm{post}}}{n_{\mathrm{pre}} \cdot n_{\mathrm{post}}}. \quad (4)$$

For the adopted mean-field approximation, the corresponding learning scheme is designed to use information from the firing distributions of the populations, embodied in the N^{p} and N^{d} variables. Considering the variable C_{ij} as being the fraction of potentiated synapses from a specific

population i in one layer to a specific population j in the other layer, we update its value after each trial as follows:

$$C_{ij}(t+1) = C_{ij}(t) + (1 - C_{ij}(t))N_{ij}^{\mathrm{p}} q_{+} - C_{ij}(t)N_{ij}^{\mathrm{d}} q_{-}, \quad (5)$$

where i and j generally denote the pre- and post-synaptic population, respectively, for both feed-forward and feedback synaptic connections, thus $(i; j)$ or $(j; i) \in (\{\mathrm{D1}, \mathrm{D2}, O1, O2\}, \{\mathrm{C1}, \mathrm{C2}\})$; q_+ and q_- are the transition probabilities for potentiation and depression, respectively; $[1 - C_{ij}(t)]$ and $C_{ij}(t)$ are the fractions of depressed and potentiated synapses, respectively; and t is the trial number. The same Eq. (5) applies both in the reward and non-reward case but different learning rates can be used. Throughout this work, we used the values q_+^{reward}=q_-^{reward}=0.01 and $q_-^{\mathrm{non-reward}}$=0.05. We chose the learning rate in the non-reward case to be greater than the learning rate in the reward case. The difference is mostly motivated by previous experimental studies on the learning and forgetting rates of a monkey performing a visuo-motor task (Asaad et al. 1998; Fusi et al. 2005). In these studies, non rewarded trials led to a quick reset of the previously memorized associations, and learning the new ones required 20–30 trials. In order to reproduce this behavior the modifications in the case of no reward had to be significantly larger than in the case of reward.

The modified average synaptic weight between the ITC and PFC layers, can then be determined for each pair (ij) ITC–PFC of specific populations:

$$w_{ij} = w_{+} C_{ij} + w_{-}(1 - C_{ij}), \quad (6)$$

where w_+ and w_- are the values corresponding to the connection strength between two populations when all synapses are potentiated or depressed, respectively. Different values can be used for the feed-forward and feedback connections in the network.

As remarked above, the wide firing rates distributions can provoke unwanted drifts in the learning history of some synaptic populations. One of the most dangerous among such effects, is that the non-diagnostic weights, which should fluctuate around their initial value in the ideal case, start to increase and spoil the learning process. Such cases might occur, in particular, when coherent effects on non-diagnostic, weights accumulate, as it happens for unusual sequences of stimuli belonging to the same category. Several regulatory mechanisms might, in principle, help to keep under control, the effects of fluctuations in the synaptic dynamics (Miller 1994; Stetter et al. 1994, 1998). The solution we adopt here, is to keep the sum of the synaptic weights to each neuron constant.

We therefore apply a subtractive normalization of the total afferent synaptic connectivity – calculated over all presynaptic inputs coming to each given post-synaptic neuron (Miller 1994). We normalize the average synaptic weight for all connections between the pre-synaptic population i and post-synaptic population j as follows:

$$w_{ij}^{\mathrm{norm}}(t) = w_{ij}(t) - \frac{1}{N}\left(\sum_{k=1}^{N} w_{kj}(t) - \sum_{k=1}^{N} w_{kj}(t-1)\right), \quad (7)$$

where N is the number of pre-synaptic populations connected to the post-synaptic population j. We recompute the values for the C_{ij} variables based on the new w_{ij} values after normalization in order to keep valid the equality in Eq. (7). For the next stimulus presentation during the learning process, the synapses between each two ITC–PFC populations are set to the calculated average value w_{ij}.

In the next section, we report our choice of parameters and the simulation results for the network and learning dynamics.

2.4 Parameter choices

To ensure network's stability for all points in the learning process, we have to choose the connection weights between the two layers as being not too small so that there is information exchange between the two modeled areas and not too high so that the network does not evolve into an amplification regime where neurons lose their selectivity. Also the biological constraint of achieving realistic neuronal activities for the modeled neurons needs to be considered (see Szabo et al. 2005).

In the simulations presented here, and for the feed-forward synapses connecting two populations from ITC to PFC, we choose the values w_+^{ff}=0.8 and w_-^{ff}=0 for the strengths in the potentiated and depressed states, respectively. For the synapses in the feedback direction we choose the corresponding strengths w_+^{fb}=0.4 and w_-^{fb}=0. The feed-forward connections are chosen double in strength on average than the feedback connections. The above choice is inspired by the idea that between cortical areas, feed-forward projections have a strong driving role, while feedback projections have a weaker modulatory role.

We start learning from an unbiased initial network configuration, by choosing all connections between the two model layers equal to the average synaptic strength: w_{ij}^{ff}=$(w_+^{\mathrm{ff}} + w_-^{\mathrm{ff}})/2 = 0.4$ and $w_{ji}^{\mathrm{fb}} = (w_+^{fb} + w_-^{fb})/2 = 0.2$ with $(i; j) \in (\{\mathrm{D1}, \mathrm{D2}, \mathrm{O1}, \mathrm{O2}\}, \{\mathrm{C1}, \mathrm{C2}\})$. Thus the initial network configuration corresponds to half of the synapses being potentiated between each pair ITC–PFC of specific populations.

3 Results

We start by outlining the general description of the learning process, followed by a more detailed description of our results. In the untrained system, all weights between ITC and PFC have equal strengths. When a stimulus is presented to the ITC layer of the untrained network, one of the populations in

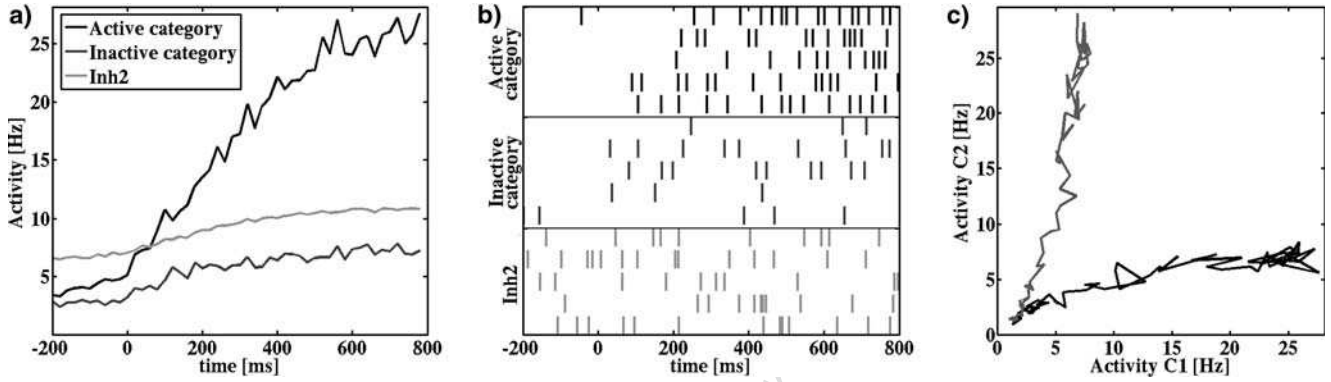

Fig. 4 Dynamics of the PFC model layer for the initial network configuration, when all the connecting weights between the two layers are set to be equal to $(w_+ + w_-)/2$. Driven by the fluctuations, one of the two category populations is active irrespective of the stimulus presented to the network. The first two graphs show average spiking rates over 50 trials for specific and inhibitory neurons in the model PFC layer. The responses of the specific populations were grouped based on their activity level: the higher responses were averaged into the active category response (*black lines*) and the lower responses were averaged into the inactive category response (*dark gray lines*). Population averages are depicted in (**a**) and spike raster plots of five neurons for each population in (**b**). The *light gray lines* represent the averaged activity over all 50 trials of the inhibitory neurons. The right most graph (**c**) plots the PFC layer dynamics in the phase space

the category layer will, by chance, receive a stronger input, mediated through network's fluctuations. These fluctuations, which are a finite-size effect, are another needed dynamic element of the model that requires the explicit description of neural dynamics at the spiking level.

Because there is strong competition between the category populations, mediated through the recurrent connectivity inside the category layer, the category being driven slightly stronger will win this competition and thus will be more strongly activated (Fig. 4). If this category happens to be the correct one, the network is rewarded. Synaptic populations that contributed most to this category's input, are hence potentiated, the ones that were driving the wrong category are weakened. As a consequence, the next presentation of that stimulus will be more likely to activate the correct category in the future. If the chosen category was wrong, the system is not rewarded, and the synaptic populations that were driving the wrong category are weakened. As a consequence, this category is less likely to be activated by that stimulus in the future. By repeated stimuli presentations, the ITC populations representing the diagnostic features will be consistently associated with the correct output population. In contrast, the ITC populations representing non-diagnostic features will be associated with each output population with the same probability. Consequently, the network will learn to perform better and better the categorization task.

We will report in the following sample learning histories of simulated networks that successfully develop both a forward ITC $\rightarrow$ PFC synaptic structure, able to support correct classification, and a backward PFC $\rightarrow$ ITC synaptic structure producing a task-dependent modulation of ITC response. To this end, we will show both the time course of the average synaptic efficacies for the populations of interest, and the manifestation of the plastic synaptic rearrangement in the ITC and PFC neural activities during the task, providing evidence of a qualitative agreement with the findings of Sigala and Logothetis.

Before illustrating how learning proceeds in the system, we start showing the properties of the initial network, which should exhibit the capability to decide stochastically with 50% probability in response to a stimulation. Figure 4 presents the activities of the specific and inhibitory neurons in PFC model area, averaged over 50 trials, for the initial network configuration. The strong competition between the populations encoding for the two categories and the stochastic fluctuations present in the network, ensure that even in the beginning of the learning process, when both categories are identically connected with the ITC layer, one of them randomly wins the competition. Hence, even the untrained network always reaches a clear decision.

Figure 5 presents average network activities (over 50 consecutive trials) in three moments of the learning process: at the beginning of learning, at an intermediate point (after 200 trials) and after the convergence of the synaptic parameters (after 1,500 trials). The plots in the first row were obtained by performing the same calculations as for the experimental data (Fig. 1b). For each specific neuron in the ITC model layer, the spiking rates for all 50 consecutive trials were grouped based on the presented stimulus values and were averaged. Each specific neuron has a different response level to the two values of each feature. The highest responses for the diagnostic feature of all specific neurons in ITC model area were averaged producing the 'best Diagnostic' response. The lowest responses for the diagnostic feature of all specific neurons in ITC model area were averaged to generate the 'worst Diagnostic' response. Similar calculations were done for the non-diagnostic feature.

These average activities over all ITC specific neurons are presented for three points in time in Fig. 5 top row. In the beginning of learning, there is no bias in the input to the PFC layer, the 'Left' (C1) and 'Right' (C2) populations are activated randomly with the same probability (Fig. 5a, bottom). Thus there is no difference between the tuning of the diagnostic and non-diagnostic features (Fig. 5a, top). As learning

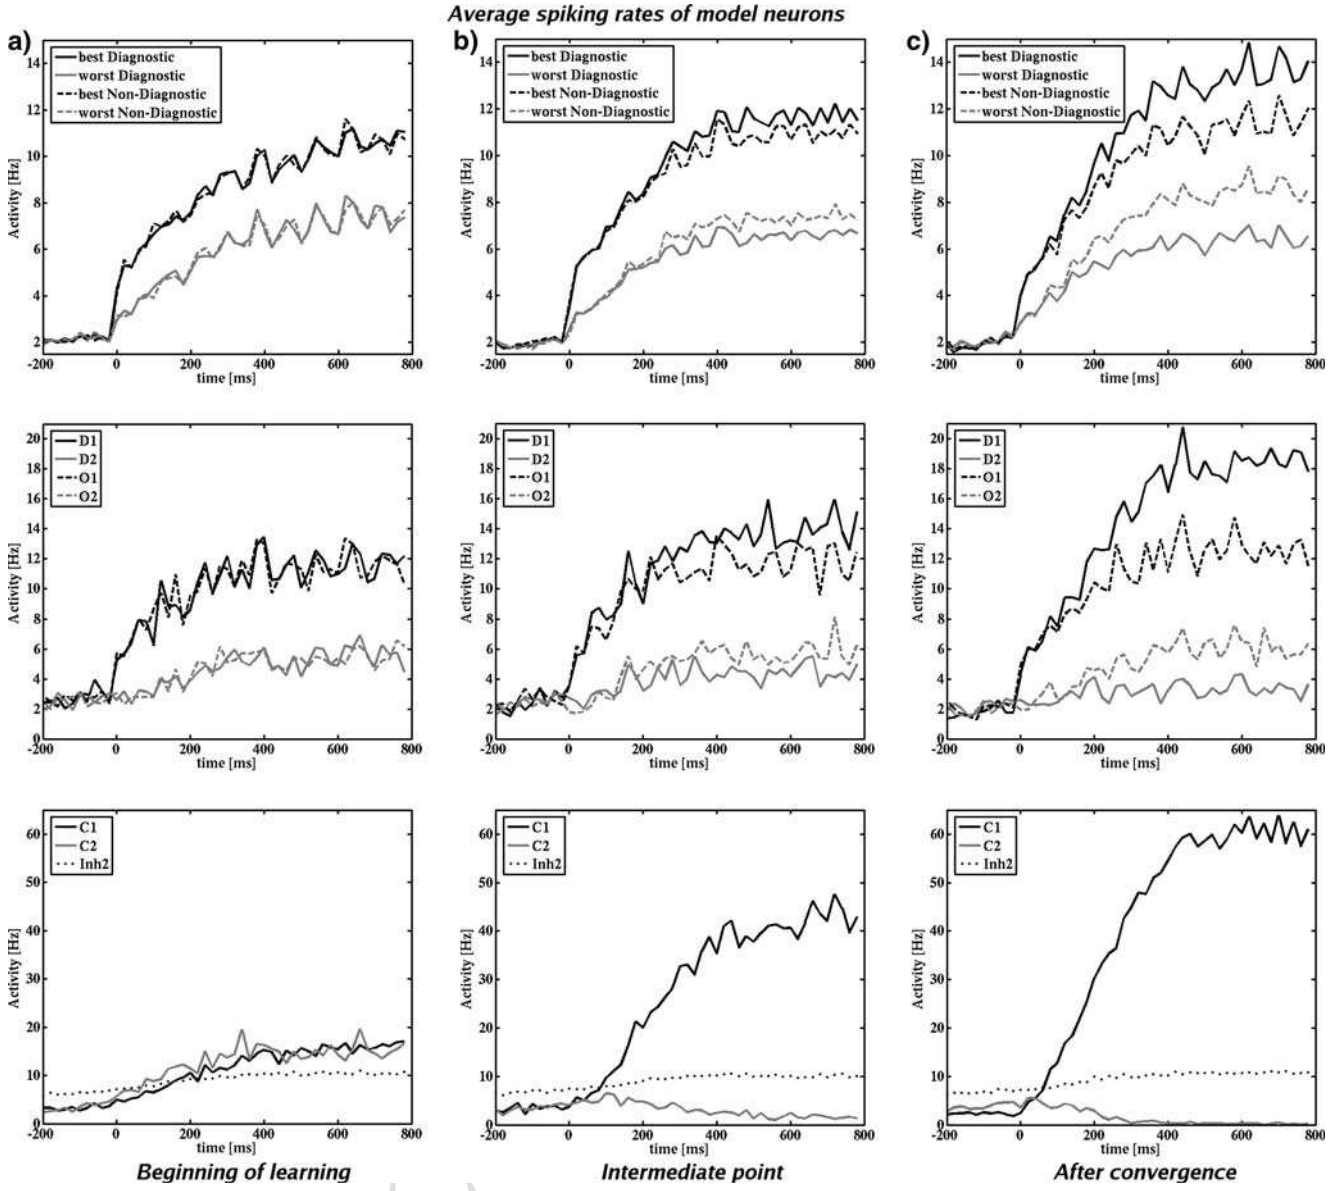

Fig. 5 Simulation results for a spiking network averaged over 50 successive trials for three points in the learning process: **a** in the beginning of learning; **b** an intermediate point during learning (after 200 steps); **c** after the weights converged to a stable configuration (1,500 steps). The *top row* shows the average spiking rates of stimulus responsive neurons, grouped according to their best and worst responses to the levels of diagnostic and non-diagnostic features. The *middle* and *bottom rows* show the average spiking rates of the specific populations in the ITC layer (D1, D2, O1, O2) and the PFC layer (C1, C2), respectively, for the trials where the presented stimulus was characterized by: high eyes and long nose (external input to the populations D1 and O1) among the 50 successive trials

progresses and the synaptic weights evolve, the network now correctly solves the categorization task (Fig. 5b, bottom). At the same time, we notice the beginning of the tuning process that will be enhanced in time (Fig. 5b, top). After convergence, the selectivity for the level of the diagnostic feature is enhanced, as compared to the non-diagnostic feature (Fig. 5c, top). The activities for the best and worst diagnostic feature values are more separated than those for the best and worst non-diagnostic feature values. This result is in good qualitative agreement with the experimental results, (Fig. 1b), that reflect the ITC activity after the monkeys had learned to categorize the stimuli.

The middle and bottom rows in Fig. 5 show the average spiking rates of the specific populations in the two layers for the selected trials among the 50 successive trials where the presented stimulus was characterized by 'low eyes' and 'long nose' (populations D1 and O1 stimulated). Since there is no structure in the model ITC layer, the enhancement of selectivity emerges due to the top-down input from the PFC layer, which encodes the previously learned stimulus

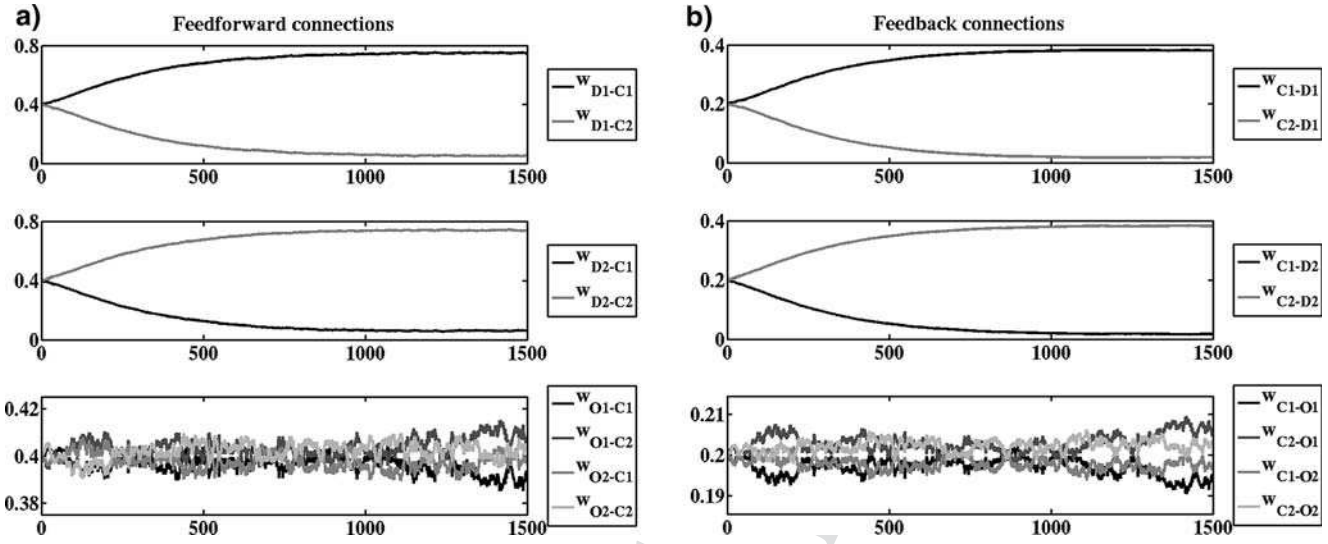

Fig. 6 Evolution of the average synaptic weight between the populations in the process of learning for **a** feed-forward and **b** feedback connections, starting from equal connectivity between the two layers

categories. The right-most column, Fig. 5c, corresponding to the point in the learning process where the weights converged to a stable configuration. From the time when the stimulus is presented to the network (time = 0 ms in Fig. 5), the selectivity of the category specific populations (Fig. 5c, bottom row) emerges through the competition biased by feed-forward inputs (ITC → PFC) from the specific populations of the ITC layer. Through the feedback modulatory inputs (PFC → ITC), this selectivity is transmitted afterwards to the feature-specific populations in ITC (Fig. 5c, middle). It can be seen that in the first 100 ms after the stimulus onset the D1 and O1 (stimulated) or D2 and O2 (non-stimulated) populations do not differ in activity. Hence there is no diagnostic tuning. Only after the correct category population acquires activity, does the diagnostic tuning build up.

The evolution of the synaptic weights between the two layers is presented in Fig. 6. For both feed-forward (Fig. 6a) and feedback connections (Figure 6b) the links between the diagnostic features and the visual object categories are selectively modified. Weights between a diagnostic feature population and the correct category population are increased, those with the wrong category population are weakened. The connections between non-diagnostic feature populations and category populations remain around the starting point, corresponding to a value for C around 0.5. This case corresponds to the network learning the task from scratch. This initial condition is referred to as 'unbiased' learning history.

Figure 7 shows learning trajectories for two other initial network configurations. In Fig. 7a, the network is previously tuned for the 'Nose length' which is non-diagnostic in our task protocol. In Fig. 7b, the network is previously tuned for both 'Eye height' and 'Nose length', thus both features were previously diagnostic, and also the 'Eye height' was differently associated to the two categories. Both cases correspond to a modification in the task protocol. The results show that the connection weights with the diagnostic feature are selectively modified in the direction of increasing the synaptic strength with the corresponding category population and decreasing the synaptic strength with the other category population. The connection weights transmitting signals from the non-diagnostic feature converge to the average synaptic strength between the two layers corresponding to the unbiased situation of equal connectivity with both category populations. Because here the network needs to react to a task switch, the corresponding initial conditions are called 'biased'.

For the three runs whose learning trajectories are presented in Figs. 6 and 7, we define and calculate quantities describing key feature of the network performace along the learning process. The results are presented in Fig. 8. We use the time-averaged activities over the last 500 ms of stimulation for each model neuron to calculate, in the same manner as described for the results in Fig. 5a, the best and worse responses for the diagnostic and non-diagnostic features of all specific neurons in the ITC model layer. The evolution during learning of these responses are presented in Fig. 8, top row. From these results, the tuning of the two features 'Eye height' (diagnostic) and 'Nose length' (non-diagnostic), is evidenced through a 'feature selectivity index' calculated as the difference between the best and worst activities for the corresponding feature divided by their sum. The time evolution of the tuning for both features is presented in Fig. 8, middle row.

The classification performance during learning, depicted in Fig. 8, bottom row, was estimated through a 'category selectivity index' calculated as the difference between the average activities of the population encoding the presented category and the population encoding for the other category, divided by their sum. It can be seen that for the run where the network was initially unbiased, the selectivity of the diagnostic feature starts building over time, whereas the selectivity for the non-diagnostic feature remains constant at a

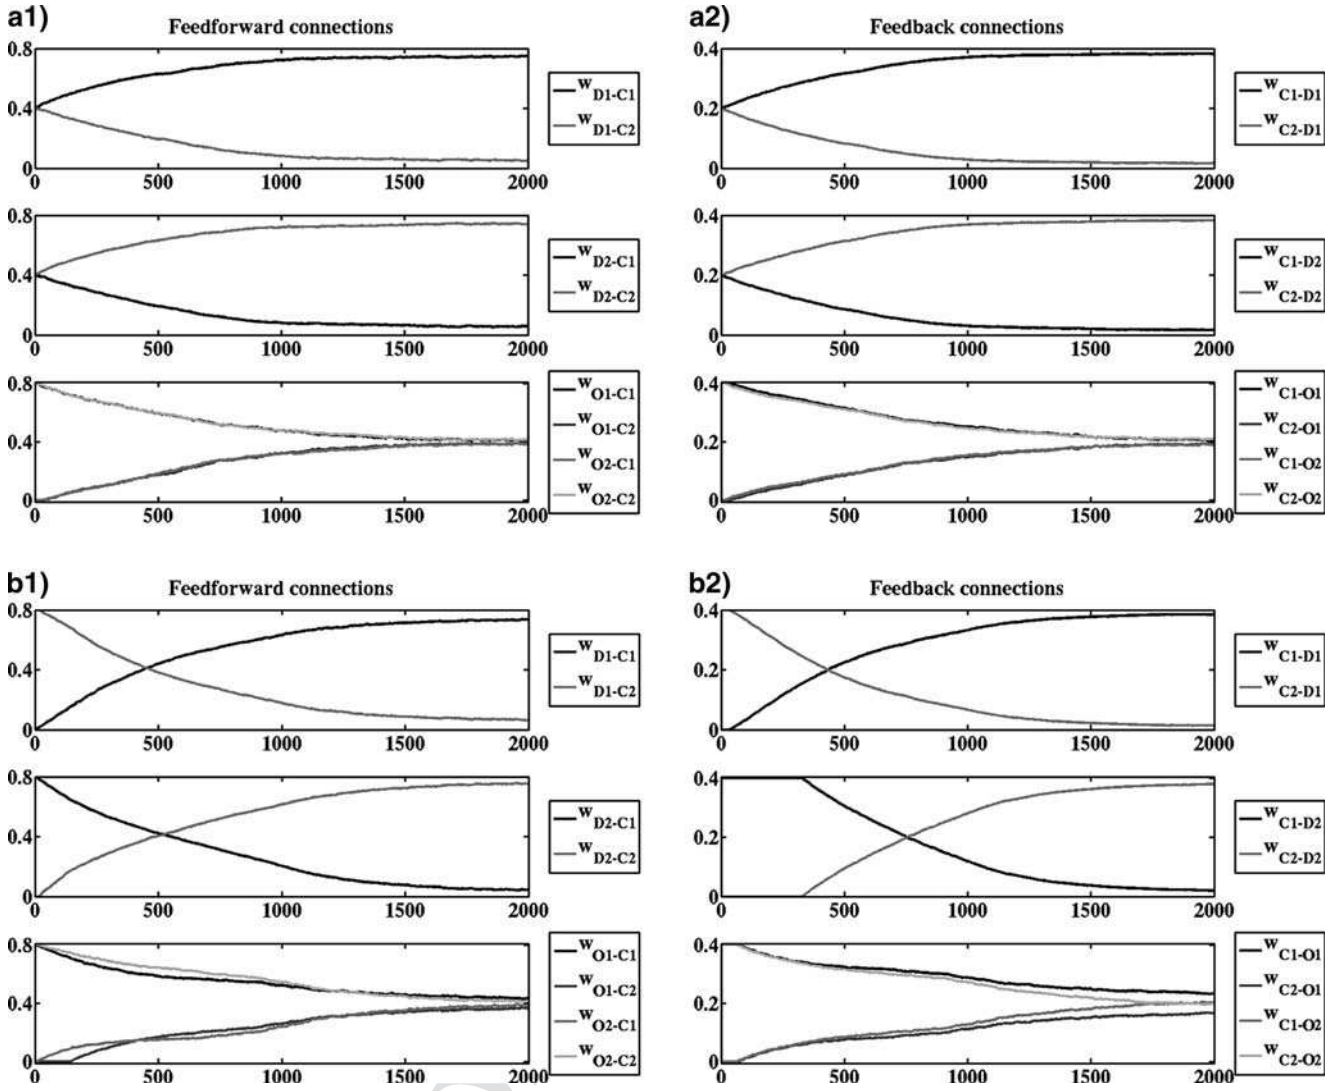

Fig. 7 Learning trajectories for the average synaptic weight between the populations in the two layers, starting from two other initial points corresponding to a switch in the behavioral task. In the first case, the diagnostic feature becomes non-diagnostic and the non-diagnostic feature becomes diagnostic (**a**). In the second case, both features were important for categorization before the task switch (**b**). The weights for both feed-forward (*a1, b1*) and feedback connections (*a2, b2*) are presented

low value (Fig. 8a, middle row). In the case of a switch in the behavioral task, where the network was previously tuned for the 'Nose length' feature, the selectivity for the non-diagnostic feature ('Nose length' in our task protocol) decreases while the selectivity for the diagnostic feature ('Eye height' in our task protocol) builds up (Fig. 8b, middle row). In the last run, the network was previously tuned for both features and also the 'Eye height' feature was differently associated to the two categories. As can be seen from Fig. 8c, middle row, the tuning of the diagnostic feature initially goes down, as for the present task protocol, the feature was previously erroneously associated to the two categories. After 500 stimulus presentations, the tuning of the diagnostic feature starts building up in accord to the chosen task protocol. Also the tuning of the non-diagnostic feature goes down, as it becomes irrelevant for behavior. From the bottom row in Fig. 8 we remark that the network performance in classification reaches a high value after 300 stimulus presentations, for the case when the network was initially unbiased, after 600 stimulus presentations, in the case of a modification in the behavioral task of only one variable (non-diagnostic previously tuned) and after 1,000 stimulus presentations, in the case of a modification in the behavioral task of two variables (non-diagnostic initially tuned and diagnostic erroneously tuned). In the latter case, the diagnostic features of the task to be learned at present actually were diagnostic before as well, but with the opposite mapping to the categories. This is reflected by a negative category selectivity index (the network is worse than guessing) in the initial phase of learning. It is remarkable that even this severe re-orientation towards a completely new task is robustly achieved by the learning network.

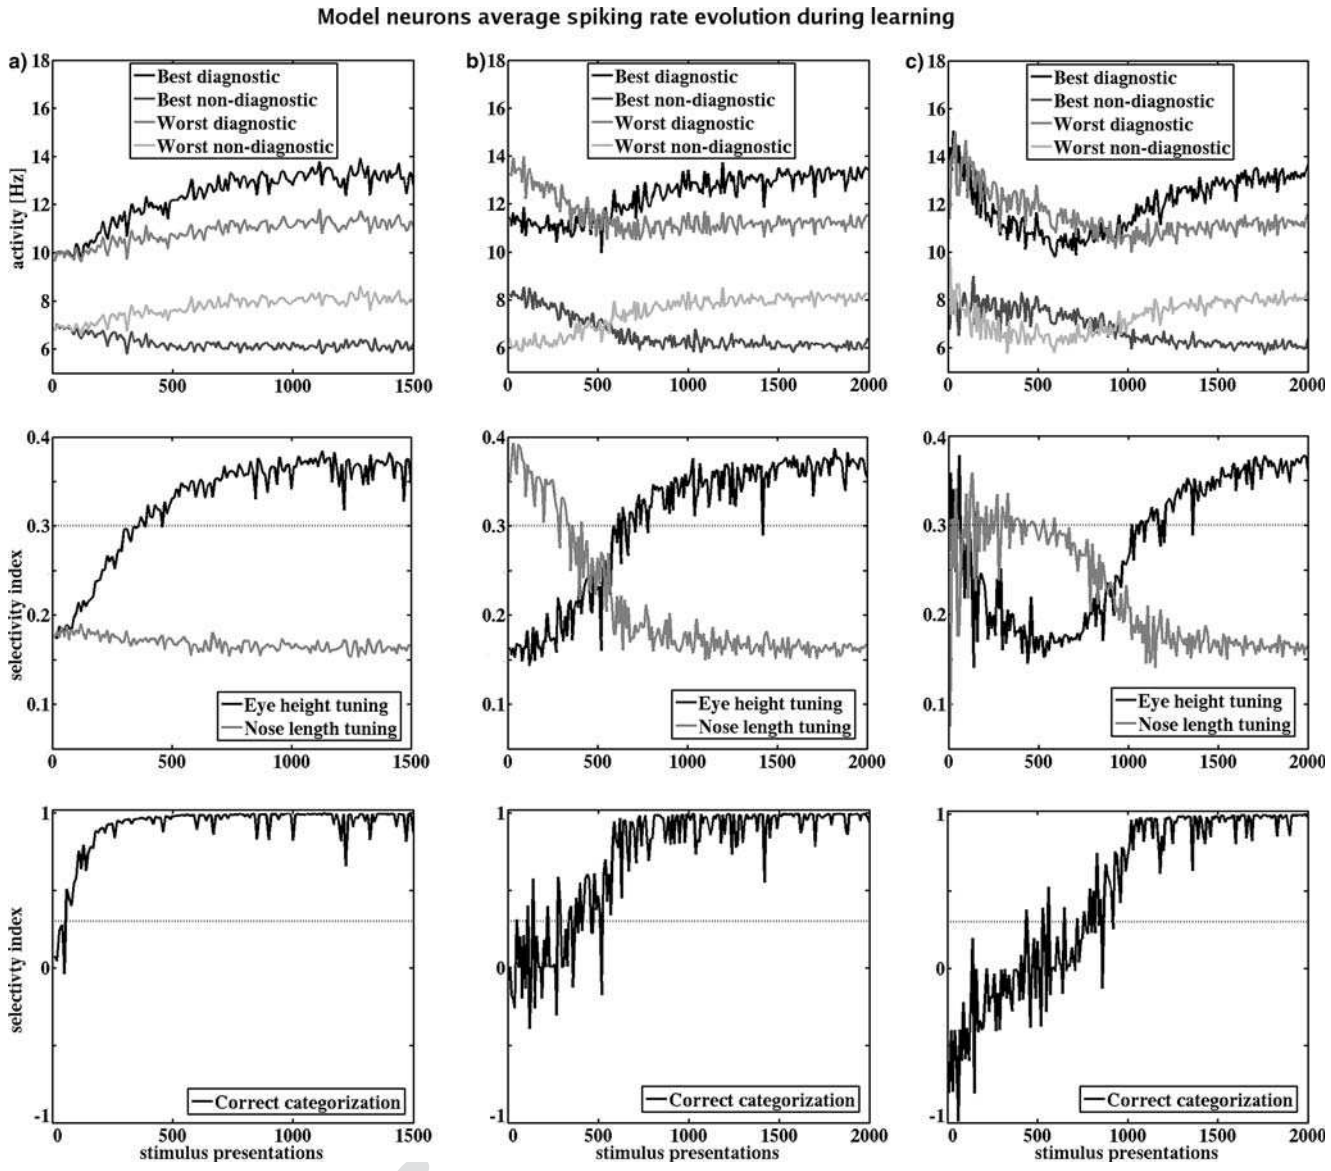

Fig. 8 Task performance evolution during learning for the three learning histories of Figs. 6 and 7: **a** learning from scratch (Fig. 6) and **b, c** learning after task switches (Fig. 7). The *top row* shows the time evolution of the best and worst responses for the diagnostic and non-diagnostic features of all specific neurons in the ITC model area. The *middle row* shows the time evolution of the two feature tuning values. The *bottom row* shows the classification performance during learning. For details see text

As can be seen in Figs. 6 and 7, some weights change in a similar manner. We conclude that we do not need all 16 free parameters (all feed-forward and feedback synaptic populations between the two layers' specific populations) to describe network's behavior and reduce the parameter space to the important dimensions only. We define four effective weights: w_d, w_i, w_{o1} and w_{o2}. w_d relates to the average connection weight between 'high eyes' and 'Left' and 'low eyes' and 'Right', which for our task protocol corresponds to the connections between the diagnostic feature values and the corresponding categories. w_i relates to the average connection weight between 'high eyes' and 'Right' and 'low eyes' and 'Left', corresponding to the connections between the diagnostic feature values and the non-corresponding categories. Similarly w_{o1} and w_{o2} are defined for the connections of the 'Nose length' feature with the two categories. Because the feed-forward weights are on average twice as big as the feedback ones, we weighted the feedback connection strengths with a factor of 2. This simplifies the representation by making a simple correspondence between the effective weights and the feed-forward or feedback ones.

$$w_d = \frac{1}{4}(w_{\mathrm{D1-C1}} + w_{\mathrm{D2-C2}} + 2 \cdot w_{\mathrm{C1-D1}} + 2 \cdot w_{\mathrm{C2-D2}}).$$

$$w_i = \frac{1}{4}(w_{\mathrm{D1-C2}} + w_{\mathrm{D2-C1}} + 2 \cdot w_{\mathrm{C2-D1}} + 2 \cdot w_{\mathrm{C1-D2}}).$$

$$w_{o1} = \frac{1}{4}(w_{O1-C1} + w_{O2-C2} + 2 \cdot w_{C1-O1} + 2 \cdot w_{C2-O2}).$$

$$w_{o2} = \frac{1}{4}(w_{O1-C2} + w_{O2-C1} + 2 \cdot w_{C2-O1} + 2 \cdot w_{C1-O2}).$$

The network shows robustness to the starting point in the parameter space. All three learning trajectories converge to the same final network configuration as shown in Fig. 9 black lines. The two axes reflect the tuning of the two features 'Eye height' and 'Nose length' expressed through the variables $w_d - w_i$, denoting 'Eye height' tuning, and $w_{o1} - w_{o2}$, denoting 'Nose length' tuning, calculated using the formulae above. The zero values correspond to equal connectivity of the feature values to the two categories, which is equivalent to no tuning for that feature. High values correspond to different connectivities between the two values of the feature with the two categories, which is equivalent to feature tuning. All traces converge to the area where there is no selectivity for the 'Nose length' feature and high selectivity for the 'Eye height' diagnostic feature.

To characterize different modes of operation of the network, we explore the effective parameters describing the excitatory weight setting between the model cortical layers. We use a mean-field approximation, fully consistent with the spiking neuron model used, that allows an exhaustive analysis of the regimes as a function of the parameter space. Each point in the parameter space, described by the chosen effective weights, was simulated for all four possible stimulus presentations. In order to determine the network's operational modes, the resulting mean firing rates of all specific populations were evaluated by calculating three suggestive parameters, as described below. From the activity of the four specific populations in the ITC layer, two parameters that measured the *'Eye height' Tuning* and *Diagnostic Tuning* were calculated.

'Eye height' Tuning measures, for a specified set of effective weights, the selectivity for the 'Eye height' feature through a selectivity index calculated as the difference between the mean firing rates for the best 'Eye height' feature value and worst 'Eye height' feature value, divided by their sum. Note that the best and worst values are calculated in the same way as for the results presented in Figs. 5 and 8, with the only difference that the population average activities are used instead of single neuron activities.

Diagnostic Tuning measures the difference between the selectivity index for the diagnostic feature (which in our case corresponds to 'Eye height' tuning) and the selectivity index for the non-diagnostic feature (which in our case corresponds to 'Nose length' tuning). The 'Eye height' and 'Nose length' tuning were also calculated for the single neuron activities, as presented in Fig. 8 middle-row.

From the activity of the two category populations in PFC, another parameter *correct categorization* was calculated as the difference between the mean firing rates of the population encoding the correct category of the presented stimulus and the population encoding the other category, divided by their sum. This parameter was also calculated for the single neuron activities, as presented in Fig. 8 bottom-row. It measures the level of association of the presented stimulus and corresponding category.

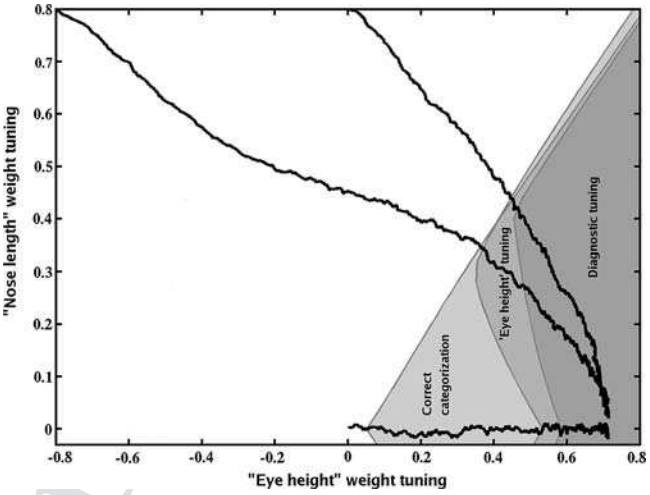

Fig. 9 Evolution of 'Eye height' and 'Nose length' tuning during learning, for three different initial network configurations. They all converge to the same final network configuration corresponding to high selectivity of the diagnostic feature 'Eye height' and low selectivity for the non-diagnostic feature 'Nose length' (bottom right corner of the graph). We characterize network's performance through an extensive exploration of network's effective parameters using a mean-field formulation. The network performed correctly the task in the *light gray area*, presented 'Eye height' tuning in the *medium gray area* and showed Diagnostic tuning in the *dark gray area*. For details see text

For each of these parameters, we chose a threshold that marked the limit where the requirements of having the respective selectivity or categorization are still satisfied, as shown in Fig. 8 middle and bottom rows by the horizontal dotted lines. The network was defined to show 'Eye height' Tuning when the Best 'Eye height' value response is twice or greater than the Worst 'Eye height' value response. We say that the network shows Diagnostic Tuning when the selectivity for the diagnostic feature is twice or greater than the selectivity for the non-diagnostic feature. Also a correct categorization corresponds to an activation of the correct category more than twice greater than the activation of the other category. In Fig. 9, we plotted the areas where these three performance criteria were satisfied. We notice that the learning trajectories converge to the area in the explored parameter space where all three conditions were fulfilled (the darkest gray area in the graph).

4 Discussion

Considering the challenging task of explaining the neural substrate of perceptual learning, we present in this work, a biologically inspired two layer network model of spiking neurons that accounts for the enhancement of ITC neurons' selectivity to stimulus features which are relevant for a learned visual categorization task (Sigala and Logothetis

2002). An alternative "attentional-gated reinforcement learning" paradigm was recently introduced for the Sigala and Logothetis perceptual learning task (Roelfsema and Van Ooyen 2005). We are not excluding the possibility, suggested by these authors, that the selective tuning of ITC neurons could arise from learning the feedforward connections coming to ITC from lower visual processing areas. For learning category tuning, some information must get back from a place in the brain where categories are represented to the site where learning occurs. In a feed-forward model, this can be done in the form of a generic attentional feedback that modulates the learning of the feedforward synaptic weights, as shown by Roelfsema and Van Ooyen (2005). Here we consider a recurrent network in which the activities of the category encoding neurons directly affect the ITC neuronal activities, and we proved that this top-down bias can explain the ITC tuning effect. An experimental scenario that could distinguish between the two network predictions is PFC cooling. After learning the categorization task, the influence of the top-down signals from the category encoding neurons could be measured: in case of a feedforward learning scenario, the effect would reside with the same strength, while in the case of a recurrent network, the effect is predicted to decrease or even vanish.

The experimental measurements in ITC (Sigala and Logothetis 2002) showed that, after training, the neuronal selectivity to the diagnostic features was enhanced as compared to the selectivity to the other, non-diagnostic, features (Fig. 1b). Hence the ITC activity not only encodes the presence and properties of visual stimuli but it is also tuned to their behavioral relevance. Different studies of perceptual learning and visual encoding suggest that the tuning of sensory neurons can be mediated by top-down information and that PFC can be associated with categorization processing (Freedman et al. 2003). In the present work, we hypothesize and test the assumption that the enhancement of selectivity to the behavioral relevant features in ITC might be determined by the higher level cognitive feedback from category encoding neurons residing in PFC and we demonstrate a learning scenario which produces such selective enhancement.

A key new feature of the proposed computational model is that the attentional biases needed to produce the competition inside the network are internally generated using the recurrent signals produced inside the network. As a stimulus is presented to the network, the sensory inputs (coming from lower visual processing areas) activate the neurons in the ITC model layer and are propagated through feed-forward connections to the PFC model layer. This bottom-up input from ITC biases the competition between the category encoding populations. The winning category expresses the monkey's decision (as in Wang 2002) and influences through feedback connections, the activity of the neurons in the ITC model layer such that, after a successful learning, they become selective for the behaviorally relevant features. Thus, in contrast to previous work (Szabo et al. 2004), the attentional biases needed to guide the competition in the PFC layer are produced autonomously in the model.

By construction, having identical inputs from the lower sensory areas encoding for the presented diagnostic and non-diagnostic feature values and no structure in the connectivity of ITC specific model neurons, a single layer model (only ITC) would induce identical tuning for the diagnostic and non-diagnostic features. The selective tuning of the ITC model area emerges then only through top-down modulatory signals from the PFC model area where the learned categories are encoded. A side effect of the identical inputs to the ITC model neurons, as can be seen from the results in Fig. 5, is that our model shows some selectivity also for the non-diagnostic feature as compared to the experimental results that show almost no selectivity for the different values of the non-diagnostic feature (Fig. 1b).

The network is trained using a biologically inspired reward-based Hebbian algorithm, that robustly ensures convergence for different initial network configurations (as shown by Figs. 7, 8). Reinforcement learning algorithms were previously shown to efficiently solve input–output mappings as in the trial-and-error interactions of the operant conditioning experiments (Williams 1992; Seung 2003). In extension to their models, we combine the concepts of reinforcement learning in connectionist and spiking networks with the biologically inspired concept of Hebbian learning. Our simple model is constructed with a small number of excitatory and inhibitory neurons for each cortical area. The finite size effect creating random fluctuations in the population firing rates, enables the spontaneous transition, in the beginning of learning when the two categories are equally connected to the ITC specific populations, of one category to win the competition. Increasing the network size reduces the probability of these spontaneous transitions.

Based on our simulation results, we find that the described effect could result from reward-based Hebbian learning that robustly modifies the connections between the feature encoding layer (ITC) and the category encoding layer (PFC) to a setting where the neurons activated by the level of a feature determinant for categorization are strongly connected to the associated category and weakly connected to the other category, and the neurons which receive input specific for a task-irrelevant feature, are connected to the category neurons with an average weight, not significantly changed during training. This structure of the interlayer connectivity seems to be a stable fixed-point of this learning dynamics and is able to reproduce the experimental data, by achieving a high selectivity of the ITC neurons for the diagnostic feature and low selectivity for the non-diagnostic feature.

The modeling approach taken in the present work could be used to generate some experimentally testable predictions, of which we name the following. Learning to provide the correct categorization, by modification of the ITC $\rightarrow$ PFC synapses, occurs before the selective modulation of the ITC responses. We infer this from the results presented in Fig. 5: for an intermediate point in the learning process (between time steps 100 and 300 in Fig. 8a), the network categorizes correctly but the responses in ITC are not yet tuned to the behaviorally relevant features. Also the other two cases of

contextual task change show that the ITC tuning occurs after correct categorization is achieved (Fig. 8b, c). The reward-based learning strategy modifies the weights in the recurrent network in a consistent way to achieve correct categorization. As the network starts to categorize correctly, a secondary effect starts to build up: the tuning of the ITC neuronal responses. Such a scenario would be consistent with the prediction that the tuning effect is an epiphenomenon of the primary synaptic process that allows to achieve the correct categorization. An observation here is that the plastic reorganization of the interlayer connections occurs in the same time for feedforward and feedback synapses, as can be seen in Figs. 6 and 7, this is due to the identical conditions adopted for the learning algorithm. The correct categorization occurs faster not because the synaptic modifications occur faster for the feedforward connections, but because the structural differences between the two layers. Implementing competition, the PFC model layer needs only a small bias to drive its activity. In the ITC model layer implementing cooperation, a more consistent bias is needed to modulate its activity.

Another interesting experimental scenario is suggested by the analysis presented in Figs. 7 and 8, where different initial conditions corresponding to switches in the behavioral task, are chosen: in one case exchanging the role of the diagnostic feature and in the other case switching from both features being diagnostic to only one being so. We can infer that the number of stimuli presentations needed until convergence to the final learned network configuration, i.e., the time needed by the network to learn the new association, increases with the number of modifications made in the task protocol. Figure 8 shows that for convergence in the case of unbiased starting configuration, around 500 trials are needed (Fig. 8a); in the case of a simple change where the non-diagnostic feature was previously associated to the two categories, around 900 trials are needed (Fig. 8b); and for a more complex change, where in addition to the change in the latter experiment, the diagnostic feature was also differently associated to the two categories, and around 1,300 trials are needed (Fig. 8c).

A natural extension of the above exercise is a task-switching scenario, in which the animal has to infer the correct rule from a changing context [analysis of the constraints on learning in a task-reversal scenario has recently been performed (Fusi et al. 2005)]. The experimental counterpart of such an extension would be most challenging and interesting for modeling.

Acknowledgements The authors wish to thank the European Union for support through EU-grant (IST-2001-38099).

References

Almeida R, Deco G, Stetter M (2004) Modular biased-competition and cooperation: a candidate mechanism for selective working memory. Eur J Neurosci 20:2789–2803

Amit DJ, Brunel N (1997a) Dynamics of a recurrent network of spiking neurons before and following learning. Netw Comput Neural Syst 8:373–404

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

Amit J, Fusi S (1994) Dynamic learning in neural networks with material synapses. Neural Comput 6:957

Asaad WF, Rainer G, Miller EK (1998) Neural activity in the primate prefrontal cortex during associative learning. Neuron 21:1399–1407

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

Chelazzi L (1998) Serial attention mechanisms in visual search: a critical look at the evidence. Psychol Res 62:195–219

Chelazzi L, Miller E, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex, Nature 363:345–347

Corchs S, Stetter M, Deco G (2003) System-level neuronal modeling of visual attentional mechanisms. Neuroimage 20:143–160

Deco G, Rolls ET (2003) Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. Eur J Neurosci 18(8):2374–2390

Deco G, Rolls ET, Horwitz B (2004) "What" and "Where" in visual working memory: a computational neurodynamical perspective for integrating fmri and single-neuron data. J Cogn Neurosci 16:683–701

Del Giudice P, Fusi S, Mattia M (2003) Modeling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. J Physiol Paris 97:659–681

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

Fine J, Jacobs R (2002) Comparing perceptual learning across tasks: a review. J Vis 2:190–203

Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2003) A comparison of primate prefrontal and inferior temporal cortices during visual categorization. J Neurosci 23:5235–5246

Fusi S (2002) Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. Biol Cybern 87:459–470

Fusi S, Annunziato M, Badoni D, Salamon A, Amit D (2000) Spike-driven synaptic plasticity: theory, simulation, VLSI-implementation. Neural Comput 12:2227–2258

Fusi S, Asaad WF, Miller EK, Wang X-J (2005) Learning and forgetting visuo-motor associations in changing environment (submitted)

Goldstone R (1998) Perceptual learning. Annu Rev Psychol 49:585–612

Miller K (1994) A model for the development of simple cell receptive fields and orientation columns through activity-dependent competition between on- and off-center inputs. J Neurosci 14:409–441

Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. Science 229:782–784

Reynolds J, Desimone R (1999) The role of neural mechanisms of attention in solving the binding problem. Neuron 24:19–29

Roelfsema P, Van Ooyen A (2005) Attention-gated reinforcement learning of internal representations for classification. Neural Comput 17(10):2176–2214

Rolls ET, Deco G (2002) Computational neuroscience of vision. Oxford University Press, Oxford

Seung H (2003) Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. Neuron 40(6):1063–1073

Sigala N, Logothetis N (2002) Visual categorization shapes feature selectivity in the primate temporal cortex. Nature 415:318–320

Stetter M (2002) Exploration of cortical function. Kluwer, Dordrecht

Stetter M, Müller M, Lang EW (1994) Neural network model for the coordinated formation of orientation preference and orientation selectivity maps. Phys Rev E 50:4167–4181

Stetter M, Lang EW, Obermayer K (1998) Unspecific long-term potentiation can evoke functional segregation in a model of area 17. NeuroReport 9:2697–2702

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. MIT Press, Cambridge

Szabo M, Almeida R, Deco G, Stetter M (2004) Cooperation and biased competition model can explain attentional filtering in the prefrontal cortex. Eur J Neurosci 19(7):1969–1977

Szabo M, Almeida R, Deco G, Stetter M (2005) A neuronal model for the shaping of feature selectivity in IT by visual categorization. Neurocomputing 65–66:195–201

Tomita H, Ohbayashi M, Nakahara K, Hasegawa I, Miyashita Y (1999) Top-down signal from prefrontal cortex in executive control of memory retrieval. Nature 401:699–703

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

Williams R (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine Learn 8:229–256