To appear in Brain Research 2011

Modelling fast stimulus-response association learning along the occipito-parieto-frontal pathway following rule instructions.

Guido Bugmann

Centre for Robotics and Neural Systems.
University of Plymouth. Drake Circus. Plymouth PL4 8AA. United Kingdom.

Abstract

On the basis of instructions, humans are able to set up associations between sensory and motor areas of the brain separated by several neuronal relays, within a few seconds. This paper proposes a model of fast learning along the dorsal pathway, from primary visual areas to pre-motor cortex. A new synaptic learning rule is proposed where synaptic efficacies converge rapidly towards a specific value determined by the number of active inputs of a neuron, respecting a principle of resource limitation in terms of total iput synaptic efficacy available to a neuron. The efficacies are stable with regards to repeated arrival of spikes in a spike train. This rule reproduces the inverse relationship between initial and final synaptic efficacy observed in long-term potentiation (LTP) experiments. Simulations of learning experiments are conducted in a multilayer network of leaky integrate-and-fire (LIF) spiking neuron models. It is proposed that cortical feedback connections convey a top-down learning-enabling signal that guides bottom-up learning in "hidden" neurons that are not directly exposed to input or output activity. Simulations of repeated presentation of the same stimulus-response pair, show that, under conditions of fast learning with probabilistic synaptic transmission, the networks tends to recruit a new sub-network at each presentation to represent the association, rather than re-using a previously trained one. This increasing allocation of neural resources results in progressively shorter execution times, in line with experimentally observed reduction in response time with practice.

Keywords: Synaptic learning rule. Cortical feedback. Leaky integrate-and-fire (LIF) neuron model . Stimulus response association. Competitive learning. Synaptic resources.

1. Introduction

Humans can learn associations between visual stimuli and motor responses from a single verbal or visual instruction, for instance "when you see the letter A, press the red button". Humans show a low error rate from the first execution (Ruge and Wolfensteller, 2009, Cole et al., 2010). By what mechanisms can the brain implement such fast learning? In this paper I propose a solution in the form of a biologically-inspired network model that is able of fast learning. Its performance is compared to those of human subjects described in (Ruge and Wolfensteller, 2009).

Fast SR learning is traditionally presented as an associative learning process following the Hebbian principle "what fires together, wires together". For instance, if a neuron activated by the stimulus has weak connections with several neurons that can drive different responses, the activation of a specific response neuron at the same time as the presentation of the stimulus will lead to the strengthening of the connection linking the stimulus-specific neuron with the response-specific neuron. A connection with a strong enough "weight" will enable a future presentation of the stimulus to cause the production of the associated response. This principle is straight-forward to implement in simulated networks with direct synaptic connections between the s-layer[1] of neurons activated by the stimulus and the s-layer of neurons activating the response (see e.g. Fusi et al., 2007). Such two-layer models pose some problems with the separability between different stimuli, but pre- and post-synaptic activities are available to drive associative learning.

However, when multiple intermediate relays are involved, most neurons lack either pre-synaptic activity or postsynaptic activity or both. How to create drivers for associative learning in these so-called "hidden" neurons has been addressed in the work by Garagnani et al. (2009) through white noise current injection in hidden neurons. This then enables the creation of Hebbian cell assemblies which also include neurons in the input and output s-layers. This approach required several hundred presentations of all SR pairs to be learnt. It also assumed fully bi-directional connectivity between brain areas, whereas anatomical data tend to indicate hierarchical relations between most areas, with feed-forward connections to cortical layer 4 in one direction, and feedback projections to infra- and supra-granular cortical layers in the other direction. The model presented here differentiates these two types of connections and exploits them for guiding a fast learning process in hidden neurons.

To implement Hebbian learning a synaptic learning rule is required. A large number of experimental data show that synaptic efficacy changes are guided by the timing of pre- and post-synaptic spikes and by the potential of the postsynaptic membrane (see e.g Sjöström et al., 2001). Typically, for a synaptic efficacy to increase, pre-synaptic spikes must precede a post-synaptic spike by around 10ms, and the postsynaptic membrane must be depolarized by at least 2.3mv. This is reflected in Spike Time Dependent Plasticity (STDP) rules. However, early results have shown that a postsynaptic spike is not a pre-requisite for learning (Reiter and Stryker, 1988). Work by Artola, Bröcher and Singer (1990) showed how synaptic efficacy changes could be controlled by the potential of the postsynaptic membrane. Recent modeling work has shown that most experimental data, including spike-time-dependent results, can be explained by only considering the post-synaptic membrane potential, if one includes effects on the potential of the spike generation process (Mayr and Partzsch, 20010). Both "voltage-controlled" models by Artola et al (1990) and by Mayr and Partzsch (2010) use a threshold potential Θ_+ for the membrane, above which an input spike does lead to a Long-Term Potentiation (LTP) of the synapse, whereas Long-Term Depression (LTD) occurs if the potential is below that threshold[2]. The availability of voltage-controlled rules is key to the LTP model proposed here, as they allow thinking of learning as a process where activity is "pushed" through a network, provided that there is mechanism to pre-set the potential of target neurons above the LTP threshold. The synaptic efficacy to silent neurons can then

[1] An S-layer is a layer in an artificial neural network, often representing a cortical area, as opposed to one of the layers of the cortex.

[2] In the model of Artola et al. (1990), there is a second, lower, threshold Θ_- below which neither LTP nor LTD occurs, but this is not relevant to the LTP model proposed here.

increase under the "pressure" of firing inputs, until the weights are strong enough to actually cause firing. This eliminates the problem of lack of post-synaptic spiking activity in hidden neurons.

Existing learning rules have two problems that prevent their use for fast learning. Firstly, voltage-controlled learning rules are intrinsically unstable (Miller, 1996; Sjöström et al., 2008) because they have no natural upper limit for weight values. Practical implementations use ad-hoc saturation or normalization mechanisms (see e.g. Krichmar et al., 2005).

The second problem is linked to incremental learning. A standard approach in weight learning is to apply the inputs and let the learning rule increase the weights with small increments until some termination condition is met, e.g. the neuron starts firing. This is a very slow process (see eg. Standage and Trappenberg, 2007) that leads to neurons developing their pattern recognition capabilities at very different times. In the model presented here, neurons are trained to become selective pattern recognizers, where input weights need to take values $W(m) = W_0/m$, where W_0 is a constant and m is the number of inputs. When starting from a zero weight, in order to reach precisely, say $W_0/24$, but not $W_0/23$, increments in weights would need to be very small. However, if a weight of $W_0/2$or $W_0/3$ is needed, such small increments lead to impractically long learning times

One solution to this problem is to abandon the incremental approach to learning, and set the weights directly to their desired value W_0/m. The conceptual problem with this approach is that, in principle, individual synapses do not have access to the information of how many active inputs are present. Therefore, an indirect approach is proposed here inspired by models of the development of receptive fields in visual cortex. These are based on the assumption that the total available synaptic efficacy is a limited resource (for a review, see Miller, 1996; Tanaka and Miyashita, 2009). Experimental data also suggest that the total synaptic weight of a neuron is a limited resource (Royer and Pare, 2003). In this view, the learning process when a spike arrives at an individual synapse has to be treated as a *request* for a share of the total synaptic weight resources, and the more shares are requested, the smaller their values, which is what is required. This approach creates a natural upper limit to weights and normalizes the weights. A formal expression for such a rule is proposed here which produces neurons with pattern recognition capabilities and which is stable upon continuous arrival of spikes in a spike train.

The learning rule proposed here is, in a sense, complementary to existing rules. It aims at defining the saturation value of each weight, as a function of the history of input spikes and the number of active inputs. However, it does not model in details the dynamics of how such saturation values are reached. Biological measurements of saturation values can help test the rule.

In the next section 2, the anatomical basis of learning from instructions is described. Two stages of the process are identified: i) the conversion of verbal of visual instructions into a mental representation of the stimulus and of the response, and ii) actual learning of the SR rule through synaptic modifications in the dorsal route linking visual cortices and the pre-motor cortex. This paper deals only with modeling stage 2, focusing on the problem of fast encoding of a SR associations in a network with multiple intermediate cortical areas, modeled as "simulation layers" (s-layer) in an artificial neural network.

Section 3 describes in details the simulated spiking neuron model. This model incorporates features such as synaptic depression and probabilistic synaptic transmission. In section 4, the new learning rule and associated cortical feedback microcircuit model are described. An example of weight evolution after the onset of input spike trains is shown. It also describes an alternative approach to competitive learning, through a weight pruning scheme taking place after the end of the learning process proper. Section 5 gives details of the architecture of the simulated network, with reference to the modeled cortical network.

Section 6 describes the simulated learning procedure and two learning experiments inspired by Ruge and Wolfensteller (2009). The first experiment consists of learning four SR associations. Learning essentially develops independent sub-networks, linking neurons coding for a specific stimulus with neurons coding for the response. The second experiment consists of learning the same SR association repeatedly, to investigate effects of repeated practice. The results here depend strongly on i) the speed of propagation of the activation during learning compared to the speed of propagation of activity in an already trained sub-network and ii) the probabilistic nature of synaptic transmission. While slow learning with deterministic synapses leads to existing connections to be re-used, fast learning with a high synaptic failure probability leads to new sub-networks to be created at each presentation of the SR pair. This consumes neural resources, but increases the speed of response upon presentation of the stimulus alone. The latter result bears similitude with experimental results of Ruge and Wolfensteller (2009). This and other aspects of the model are discussed in section 7 and the conclusion follows in section 8.

2. Anatomical pathways and learning experiments

Anatomically, the inferior temporal lobe (IT), the prefrontal cortex (PFC) and the pre-motor (PM) cortex play a key role in the learning of conditional visuo-motor stimulus-response (SR) rules (Passingham and Toni, 2001). Lesions of either area impair the learning and execution of arbitrary SR rules (Bunge et al., 2005). These cortical areas define the PFC route for SR learning.

Another route often mentioned in the literature involves a sub-cortical structure, the basal ganglia (BG). The BG make a key contribution to conditional learning by trial-and-error, but probably less so to learning from instructions (Ruge and Wolfensteller, 2009). A useful review of the role of the BG in SR learning can be found in Ashby et al. (2010) who suggest that the long-term storage of SR associations consists of cortico-cortical representations.

A route rarely mentioned is via the cerebellum where lesions often have effects similar to PFC lesions (Middleton and Strick, 1998). The cerebellum has reciprocal connections with the PFC and PM, and could offer a bypass to the PFC-PM section of the PFC route.

Finally, the most direct route from the visual cortex to the PM is the dorsal visual stream via the posterior parietal cortex (PPC). It is generally assumed to convey only spatial information and is therefore an unlikely candidate for general SR mapping. However, there are reports on areas of the PPC (Toth and Assad, (2002) and motor areas (Zach et al., 2008) able to develop a selectivity for non-spatial attributes of stimuli.

There appear to be at least two stages in learning SR rules, the first being driven by the instruction and the second driven by actual or possibly mentally simulated practice. The first uses a network of PFC, PM and PPC areas (Ruge and Wolfensteller, 2009; Cole et al, 2010; Brass et al., 2009). The second seems to involve additional more rostral prefrontal areas (Cole et al., 2010), early visual areas (Ruge and Wolfensteller, 2009) or possibly the cerebellum (Brass et al., 2009). There may be a third stage of consolidation, which is assumed to take place during sleep. Consolidation has been shown to reduce interferences between acquired motor skills (Brashers-Krug et al., 1995), but consolidation of SR rules has not yet been documented.

Ruge and Wolfensteller (2009) describe an experiment where visual instructions are displayed on which hand should be used to respond to a given visual geometrical shape. They measured the variation of activity in brain areas immediately after instruction, during the 6-8 trials during which the execution time improved to a plateau level. The activity of some areas rapidly decreased to base-level, including PFC (probably BA46, BA: Brodman Area) and the lower bank of the intra parietal sulcus (IPS) of the PPC. The activity of some areas increased to a plateau level within 6-8 trials (PM, SMA, areas around the rolandic operculum and occipital visual areas). The activity of some areas decreased, but to a non-zero level (superior bank of the IPS, PFC (probably 9) and an area of the fusiform gyrus (probably BA37)). Interestingly, the plateau level of activity of the later is similar to that of the former. This suggests the areas with decreasing-then-plateau activity are involved in both the initial scaffolding for SR leaning and later maintenance of the SR mapping. All areas that have an increasing activity are part of the dorsal stream from primary visual areas, to MIP/PE in IPS and to PM, suggesting that the SR mapping takes place along the PPC route. Other experiments also describe increased activity in PM after learning (Wolfensteller et al., 2004; and references in Rizzolatti and Luppino, 2001). The role of areas in the Rolandic operculum remains unclear, possibly related to covert verbalization of the S-R rule. This paper will focus on modeling the second stage of fast learning through the PPC route.

Regarding the instruction-driven first stage, in another experiment, subjects were asked to realize a hand posture either through imitation or through verbal command (Makuushi et al., 2005). The experiment showed that IPS was activated in both cases. This suggests that, during instruction, movement specification is communicated to motor areas through existing mappings between PPC and PM (Pandya and Yeterian, 1984). Whether this occurs through direct input from the auditory cortex, or via the relay of the PFC is unclear. The PFC receives projections from all areas of all sensory streams, including visual input through IT, (Pandya and Yeterian, 1984) and is ideally placed to generate responses on the basis of any aspect of a stimulus. Another brain region, in the para-hippocampal gyrus receives multimodal inputs, but it appears to be mainly involved in the association between stimuli, not between stimuli and responses. It also provides input to the PFC (Pandya and Yeterian, 1984). Sensory information could be relayed from the PFC to the PM through a network of prefrontal relay areas (Passingham and Toni, 2001) or through driving projections back to the PPC and in particular to IPS (Medalla and Barbas, 2006), and then from PPC to PM. In this paper, I will assume that one effect of the instructions is to activate a representation of the desired response in PM. This process is not modeled in details here.

Regarding the stimulus component of the instructed task, when receiving task instructions, humans tend to mentally rehearse the task, which appears to be a limiting factor in the number of rules that can be specified in one instruction (Apel et al., personal communication). Rehearsal generates imagery of

the stimulus and the response which activates motor areas (e.g. Tettamanti, 2005) and sensory areas (e.g. Ganisa et al., 2004). Lower visual areas are activated during certain imagery tasks (Klein et al., 2000), but the detailed cortical mechanisms for this have not been identified.

Overall, SR learning can be treated as a supervised learning process where initial input activities and corresponding output activities are set by the instructions. Thereafter, each execution of the task will produce actual input and output activities that could also be used for further learning.

Several constraints are imposed on a model of stage 2 of learning. First, learning is a very fast process where one, or a very small number of rehearsal cycles are sufficient to establish the SR mapping within a few seconds (Nowak et al., 2007; Brass et al., 2009; Ruge and Wolfensteller, 2009). Secondly, learning generates a stimulus- and response-specific link between distant brain areas involving several neuronal relays. Thirdly, there is no reported "catastrophic interference" between newly learnt SR rules and existing ones.

These constraints of fast, interference-free, learning seriously challenges existing artificial network models with multiple hidden s-layers. For this reason, I have developed the model presented in this paper. Most aspects of the model are grounded in neurophysiology and neuro-anatomy. Some aspects rely on more speculative hypotheses introduced to explain fast learning. These are compatible with current knowledge, but would benefit from experimental verification. Two notable hypotheses concern the proposed synaptic learning rule where weights converge not to an arbitrary saturation strength, but to a value dependent on the number of active inputs (section 4.1), and the role of top-down feedback projections in the cortex as carrying driving signals for learning (section 4,2).

3. Spiking neuron model

3.1 Somatic potential

The neuron model used is a leaky integrate-and-fire (LIF) model. Such model comprise a soma, with capacity C, whose potential V(t) increases when electric charges are added by the synaptic current I(t) induced by input spikes. When the potential reaches the firing threshold V_{th}, the LIF produces an output spike and its potential is reset. The soma continuously leaks charges through its membrane represented by a leak resistance R. These two processes of charge and discharge are captured in the differential equation for the rate of change dV(t)/dt of the somatic potential:

$$\frac{dV(t)}{dt} = -\frac{V(t)}{RC} + \frac{I(t)}{C} \qquad (1)$$

Where the product RC is the time constant of the decay of the somatic potential through leakage and V(t)/RC is the rate at which the potential diminishes through leaks, in units of Volts/sec. The term I(t)/C is the rate at which the potential increases due to the input current, in units of Volts/sec. In this work I am using RC=50ms, V_{th}=15mV and C=1.

3.2 Synaptic current time course

The time course of the current $I_j(t,t_k)$ flowing into the soma due to the emission of one pre-synaptic spike at time t_k by source neuron j is modeled by an alpha function:

$$I_j(t,t_k) = I_{MAX,j}\alpha(t,t_k)$$

$$\alpha(t,t_k) = \frac{\exp(1)}{t_{MAX}}(t - t_k - \delta)\exp(-\frac{t - t_k - \delta}{t_{MAX}}) \qquad \text{if } t > t_k + \delta$$

$$\alpha(t,t_k) = 0 \qquad \text{if } t \le t_k + t_p \qquad (2)$$

The alpha function has a maximum amplitude of 1 at time t_{MAX} due to the term exp(1). Its integral over time is equal to e· t_{MAX}. There is an axonal propagation time δ between the pre-synaptic neuron and the synapse. The value of t_{MAX} reflects the dendritic distance between the synapse and the soma. In the simulations done here, only proximal synapses with t_{MAX} = 1ms are used. The propagation time δ is assumed to be 1 ms. The actual current amplitude if defined by $I_{MAX,j}$. In this paper $I_{MAX,j}$ is expressed as the product of the somatic capacitance C_i of the target neuron multiplied by the synaptic weight W_{ij}, where the weight represents the maximal *rate of potential increase* (in units of V/seconds) caused by the input current at the peak of the alpha function.

The total synaptic current in neuron i due to all its input M neurons is given by:

$$I_i(t) = C_i \sum_{j=1}^{M} W_{ij} \sum_{k=1}^{20} \alpha(t, t_{jk)}) \qquad (3)$$

The first sum in (3) runs over all M input neurons j. The second sum runs over all past input spikes produced by each neuron j. These are produced at time t_{jk}. In the simulation, only the 20 previous spikes are considered, as the alpha function becomes negligibly small for long delays.

3.3 Synaptic depression

The synaptic weight W_{ij} is proportional to the quantity of neurotransmitter released by a spike. Experimental evidence suggests that this quantity is temporarily reduced after an input spike and then recovers with a time constant τ_{REC}. This phenomena of short-term synaptic depression has been modeled by Tsodyks and Markram (1997). By assuming that the totality of available neurotransmitter is

released by each spike, the effective synaptic weight seen by a spike arriving at time $t_{j,k}$ depends only on the time of arrival $t_{j,k-1}$ of the previous spike[3]:

$$W_{ij}(t_{j,k}) = W_{MAX,ij}(1 - \exp(-\frac{t_{j,k} - t_{j,k-1}}{\tau_{REC}})) \qquad (4)$$

By setting τ_{REC} = RC, synaptic depression is balanced which confers optimal pattern recognition capabilities to a LIF neuron (Bugmann, 2002). In balanced conditions, temporal integration in the soma is cancelled by synaptic depression and the potential increases due to successive input spikes reach all the same peak potential, independently of the time separating the arriving spikes (Bugmann, 1992). This makes the postsynaptic effect of input spike trains independent on the frequency of these spike trains and allows a precise control of the selectivity of the target neuron (see an example in figure 1). In the simulations, τ_{REC} = RC = 50ms.

3.4 Probabilistic synaptic transmission

A spike generated by input neuron j has a probability p_{ij} of releasing neurotransmitter at its synapse with neuron i. Biological synapses can have a range of transmission probabilities from 100% to as low as zero (see e.g. Montgomery et al., 2001)). In the simulations, inputs spikes are processes with a probability p and ignored with a probability 1-p. All synapses have the same transmission probability.

One effect of probabilistic synapses is to reduce the effective input firing rate. The other is to create a different input spike train for each target neuron.

3.5 Spike production, partial membrane repolarisation and sustained firing.

The neuron model fires when its somatic potential V exceeds or equals the firing threshold V_{th}=15mV. It then enters into a refractory state where inputs keep being integrated, but spike production is disabled. A refractory period of 2ms is used here. At the same time, the somatic potential is reset to a value $V_r=\beta V_{th}$. The value of β allows to control the irregularity of the output spike train, with low values producing rather regular intervals and values close to 1 producing bursts of output spikes. I use the value β=0.91 which produces near-Poisson spike trains as physiologically observed (Bugmann et al., 1997). This ensures that spike trains remain of Poisson type throughout the network.

The neuron model can also emulate the effect of local recurrent excitatory feedback connections (Douglas and Martin, 1991), or the UP-state (Anderson et al., 2000) through a self-sustained firing mode. In this mode, once the neuron has produced its first output spike, it becomes a generator of Poisson spike trains with a preset rate. This process respects the refractory time and the rate of spike generation is set appropriately higher than the actual average rate of spiking. After each spike, the potential is reset

[3] Thus each input spike in a spike train sees a different weight, and the weight term needs to be moved into the second sum in (3).

to $V_r=\beta V_{th}$. The neuron still integrates inputs and will generate additional spikes if the potential increase is sufficient. The use of sustained firing has no effect on the selectivity of the neuron, but improves propagation speed and reliability in a multilayer network when the weights are small. Sustained firing is used in all network simulations presented here. The average firing rate generated is 100Hz.

3.6 Weights for pattern recognition

SR learning requires that a specific pattern of response activity is only generated by a specific pattern of stimulus activity. To that end, during learning, neurons must become selective pattern recognizers.

A neuron i has M input connections. A pattern is defined as a subset of m≤M active input neurons. A neuron realizes pattern recognition if it fires only when this subset is present. For that purpose, the synaptic weights from these m inputs need to be just strong enough to raise the somatic potential to the firing threshold when all m inputs are active. This can be achieved if each weight W_{ij} is set to the same value $W_{MAX,ij}(m)$, that has following dependence on the number m of active inputs in the pattern:

$$W_{MAX,ij}(m) = \frac{W_0}{m} \qquad (5)$$

This formula is correct when balanced synaptic depression is used, as the potential increase due to an input spike is a constant depending only on the weight (Bugmann, 1992). Such a neuron will only respond to the local pattern of activity for which it is trained, or any larger pattern that includes that local pattern.

In the simulations, the neuron starts responding for values of W_0 = 6.5 (found by trial and error[4]). With this minimum value, the neuron is very slow to respond, with average times-to-first output spike larger than 500ms (after the start of m input Poisson spike trains at 100Hz, which produce their first spike at random times after the start). The output firing rate is also low. For instance, figure 1 produced with W_0=7 shows 4 spikes produced in 200ms. The response speed can be improved by increasing W_0, but at the cost of a reduced selectivity S:.

$$S = 1 - \frac{(m - m')}{m} \qquad (6)$$

Where m' is a number of active inputs smaller than m that is able to cause firing. In the simulations, the responses speed was kept close to 10ms, as suggested for the visual system (Thorpe, 2002) , but at the cost of selectivities of 0.6-0.8.

[4] The theoretical $W_{OPTIMAL}$ is probably smaller, but would cause an impractically long average response times. The exact expression linking W_0 and the potential increase ΔV_j caused by an input spike can be calculated using (1) and using (3) and (5). As a first approximation (ignoring leaks), ΔV_j is equal to W_{ij} multiplied by the integral of the alpha function, i.e. $W_{ij} \cdot e^1 \cdot t_{MAX}$. The total potential increase is $m \cdot \Delta V_j = W_0 \cdot e^1 \cdot t_{MAX}$, which should be equal to V_{th}. We end up with = $W_{OPTIMAL} = V_{th}/(e^1 \cdot t_{MAX})$. In our case: 5.55….

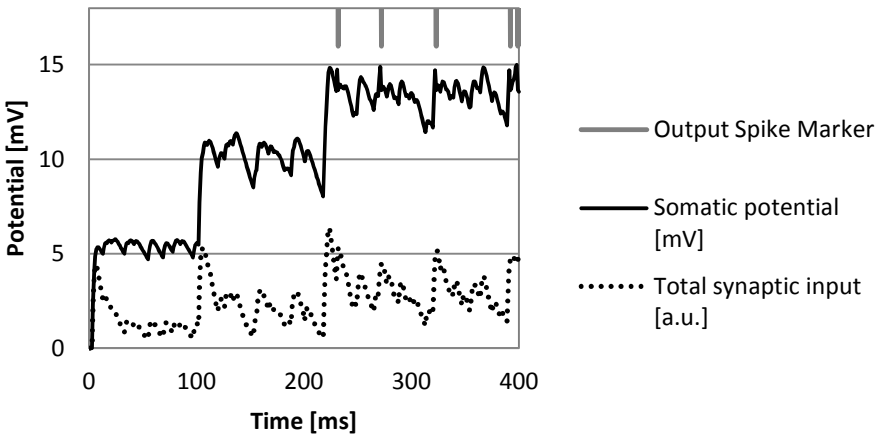

FIGURE 1. Example of the operation of the LIF neuron model used. The synaptic weights are set to $W_0/3$, with W_0=7. Three Poisson inputs (100Hz) start to become active at t=0, 100 and 200 respectively. Synaptic depression (4) causes the input current to peak when a synapse receives its first spike, then subsequent spikes have a smaller effect. The postsynaptic potential shows the characteristic saturation due to the use of balanced depression, with RC=τ_{REC}=50ms. The neuron fires when the potential reaches the firing threshold Vth=15 mV. It then enters a mode of sustained firing where the neuron produces spikes as a Poison process with refractory time.

4 Learning processes

4.1 Synaptic learning rule.

Learning is supposed to take place when an input spike arrives at a synapse which is enabled for learning, through a somatic potential above zero. The neuron has a fixed amount W_0 of total synaptic weight at its disposal. This could represent a fixed total number of receptors on the postsynaptic membrane. Initially this weight resource is located in a store W_{POOL}. Each input spike arriving at time t obtains for its synapse i a quantity ΔW_P of the weight available in the store and a quantity ΔW_{S+} from the other already trained synapses. The new weight $W_i(t+1)$ of the input synapse i then becomes:

$$W_i(t+1) = W_i(t) + \Delta W_p + \Delta W_{S+} \qquad (7)$$

The quantity gained from the store is a fixed fraction η of the content of the store.

$$\Delta W_p(t) = \eta \cdot W_{POOL}(t) \qquad (8)$$

The quantity available in the store is reduced by the same quantity:

$$W_{POOL}(t+1) = (1-\eta) \cdot W_{POOL}(t) \qquad (9)$$

Each of the other synapses j≠I releases a quantity ΔW_{js-}:

$$\Delta W_{jS-} = \frac{W_j(t)(W_j(t) - W_i(t))}{W_0 + W_j(t)} \qquad \text{if } W_j(t) > W_i(t)$$

$\Delta W_{js-} = 0$ otherwise (10)

And reduces its own value:

$$W_j(t+1) = W_j(t) - \Delta W_{js-} \qquad (11)$$

The quantity ΔW_{S+} is the sum of the contributions of the other synapses:

$$\Delta W_{S+} = \sum_{j \neq i} \Delta W_{jS-} \qquad (12)$$

Expression (10) defines the proposed learning rule. It ensures that all active input weights become equal to W_0/m, where m is the number of synapses that have received at least one spike as input. Formula (10) is designed so that all weights $W_0/(m-1)$ drop to W_0/m when the mth synapse is activated, and the quantity released by all active synapses is such that the new synapse receives a total weight of W_0/m. The formula is exact for η=1, and, for η<1, converges to the exact required weights (5) as soon as the weight store is empty.

The idea behind the rule is that spikes arriving at synapses make a "call" for weight resources. This is why the sign of (10) is constrained to be positive.

The learning rule (10) is derived as follows. Assuming that all existing weights W_j resulting from m activated synapses have the correct value W_0/m, then the activation of an additional synapse should result in the existing weights being reduced to the value $W_0/(m+1)$. Each existing weight thus looses a quantity ΔW_j^*:

$$\Delta W^*_{j-} = \frac{W_0}{m} - \frac{W_0}{m+1} = \frac{W_0}{m \cdot (m+1)} \qquad (13)$$

By replacing W_0/m by W_j, one finds:

$$\Delta W^*_{j-} = \frac{W_j^2}{W_0 + W_j} \qquad (14)$$

The newly activated synapse receives from the m previously activated synapses a total weight ΔW_{i+}:

$$\Delta W_{i+} = m \Delta W^*_{j-} = \frac{W_0}{m+1} \qquad (15)$$

When a spike arrives at a synapse that already has a non-zero weight W_i, it only needs to receive from another synapse a quantity proportional to the relative difference between their respective weights. To that effect a term is added "by hand" to the expression (14) for ΔW_{j-}:

$$\Delta W_{j-} = \frac{W_j^{\,2}}{W_0 + W_j} \frac{W_j - W_i}{W_j} = \frac{W_j \cdot (W_j - W_i)}{W_0 + W_j} \qquad (16)$$

This results in the expression for the learning rule (10). It respects the principle that synaptic modification only uses information available locally at the synapse. In particular, no explicit knowledge is needed of the number m of other activated synapses. The quantity of synaptic weight released by a synapse in (14) is only a function of the synaptic weight of the releasing synapse. The value of the weight W_i of the synapse that just received a spike can be inferred from the size of the generated post-synaptic potential which is available to the whole target neuron, and the value of W_0 is a global property of the target neuron.

The term $W_j(t)$-$W_i(t)$ in (10) ensures the stability of the weights upon arrival of subsequent spikes in a spike train. Therefore, learning can be continuously active and weights will only change when a yet untrained synapse starts receiving spikes from its input. Figure 2 shows an example of learning of synapses of a neuron connected to five active inputs.

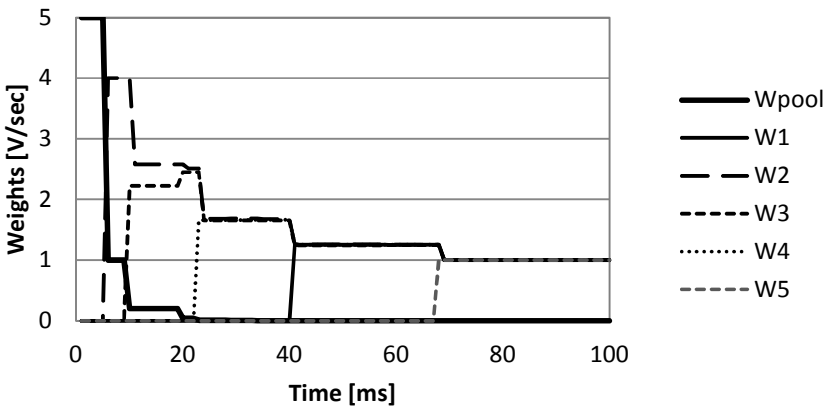

FIGURE 2: Weight learning example with 5 active inputs starting to fire at time 0. All weights are initially zero and are set by the learning procedure described above as soon as the first spike of an input neuron arrives. I used W_0=5 and the final weights are all equal to W_0/5. The thick line is the quantity of weight W_{POOL} in the store. Each step change in weights corresponds to an input spike arriving on a synapse with a weight that has not reached its final value W_0/m where m is the current number of activated synapses. Inputs are Poisson spike trains firing at 50Hz in this example. η = 0.8.

The fact that the final weight is independent of the initial weight leads to an inverse relation between the relative weight change LTP and the initial weight as observed experimentally (Montgomery et al, 2001; Sjöström et al., 2001). Using the notations of (Sjöström et al., 2001):

$$LTP = \frac{W_{FINAL}}{W_{INITIAL}} = \frac{W_0}{m}\frac{1}{W_{INITIAL}} \qquad (17)$$

This expression may allow the estimation of physiological values of W_0 from data where a known number m of input synapse are stimulated. As noted above, the dynamics of how synaptic resources are released and transported between synapses is not modeled here, but it is likely that the new synaptic weight becomes effective some time after the input spike has arrived. Such a delay has a practical influence on the time taken by a trained neuron to start responding to its inputs (figure 3). Learning experiments in section 6 explore this further at the network level.

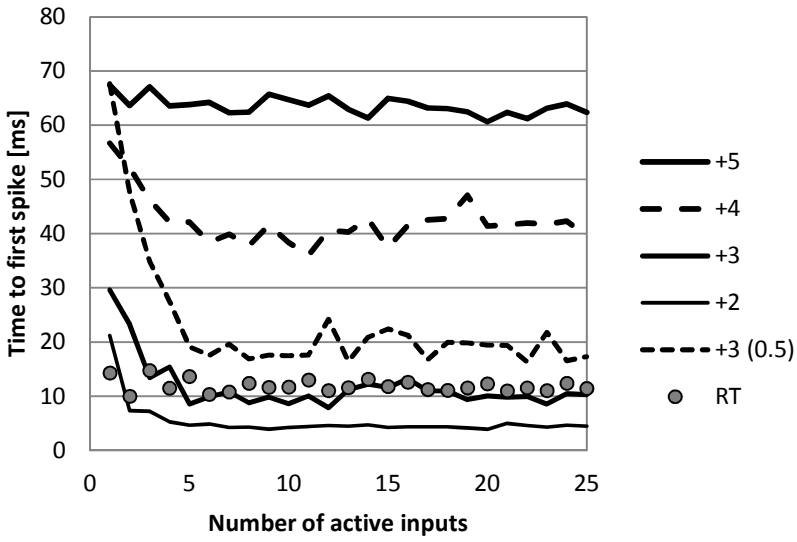

FIGURE 3: Average times to first spike of a leaky integrate-and-fire neuron (LIF) learning to respond to the indicated number of active inputs. Rule (10) was applied with W_0=8. The curves differ by the delay after spike production when the new weights are becoming effective. The number "+2" means that the weights change at the peak of the alpha function. "+5" means 3ms after the peak. All curves are obtained with η =1, except for one curve with η =0.5 (indicated as "+3 0.5"). RT is the response time of a neuron with weights preset to W_o/m. Input spike trains are Poisson spike trains, where the firing probability is set to 100Hz at time 0. The somatic potential and synaptic weights are initially zero.

In figure 3, rule (10) was simulated with a single leaky integrate-and-fire (LIF) neuron with various delays of the implementation of the rule. The results show that implementations at the peak of the alpha function (curve +2 in figure 3) lead to activation times faster than the response time of the trained neuron with weights preset to W_o/m (curve RT in figure 3). The reason is that the first input spikes see weights larger than W_o/m. Implementations 3ms after the peak can lead to significantly long

delays (see e.g. +5 curve in figure 3). The use of smaller values of η further increases the learning time. Interestingly, small numbers of inputs have a slight disadvantage in learning speed, but this is only in average.

4.2 Micro-circuit for synaptic learning

For the push mechanism (section 3.2) to operate, there need to be a mechanism for raising the post-synaptic membrane potential above the threshold potential for LTP prior to the onset of the inputs. In my model I assume that this is achieved by recurrent collaterals of layer 5 or 6 neurons that propagate backward through the network the information about which output is to be activated. Feedback from a neuron in a higher area targets all neurons in the area below that can provide feed-forward information to the neuron in the higher area (figure 4).

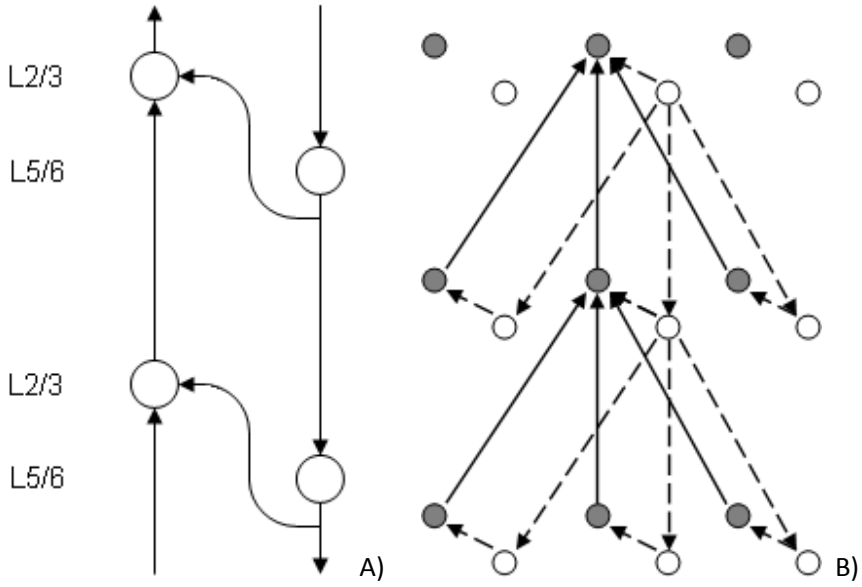

FIGURE 4: Microcircuit for learning. A) Principle of Top-down excitation from higher areas used to increase the membrane potential above the LTP learning threshold. Two parallel streams of information are assumed, flowing in opposite direction. Information from higher areas enable learning in a cascade of neurons able to provide feed-forward information to the higher area which originated the enabling signal. B) Illustration of the divergent top-down connectivity (dashed lines) to the sources of bottom-up inputs (full lines). Through the proposed mechanism, a higher area action neuron enables its input neurons, then the inputs of the inputs, and so on, until cascade of enabling signals reaches the lowest area that receives sensory inputs.

Anatomically, in visual areas, a feedback stream of projections links neurons in layer 5 from a higher area to layer 5 neurons in a lower area. This is paralleled by a feed-forward stream linking layer 4/3/2 neurons in successive areas (Martinez-Conde, 1999, and references therein). It has also been suggested that neurons in layer 6 may constitute such a parallel stream of feedback information (Shipp, 2007). For the purpose of my model, the feedback stream is assumed to convey information about target activity in

the response area, e.g PM, down to visual areas where the stimulus is represented. This feedback flow of information has a diverging pattern (Henry et al., 1991), defining a "cone of attraction" in the form of silent neurons with increased membrane potential that enables them to learn the incoming feed-forward information. The importance of this cone lies in that it guides feed-forward learning towards the desired target neurons in higher areas.

For consistency, any active neuron should be able project and enabling signal back to its input neurons, and to the inputs of its inputs, etc. Through these feedback signals, all inputs neurons of an active neuron will start learning their inputs. This is likely to re-ignite the initial learning process. I have not attempted to observe this in simulation, where only the neurons in the last s-layer are given the ability to start a cascade of enabled neurons. However, in future work, this is a process to examine.

4.3 Competitive learning through pruning.

Using the learning rule (10) all neurons in a s-layer that see a part of the stimulus will learn to represent that part (Fig 5A). Computationally, this is not an effective use of neural resources. Only a subset of neurons with disjoint RFs is really needed to represent the stimulus. The brain has solved a similar problem during development, where neurons in sensory cortices, for instance, have divided the input space in some optimal sense, e.g. information theoretic (Bell and Sejnowski, 1997) or neuronal resources preserving (Tanaka and Miyashita, 2009). There are numerous models of competitive learning, but there is little experimental information on the biological mechanism of competitive learning and current models are mainly computationally motivated (Miller, 1996).

Several mechanism can be considered, such as competition through lateral inhibition, through feedback inhibition, and a new one explored here, using pruning. They are quite equivalent from the end result point of view, resulting in few but essential connections being developed.

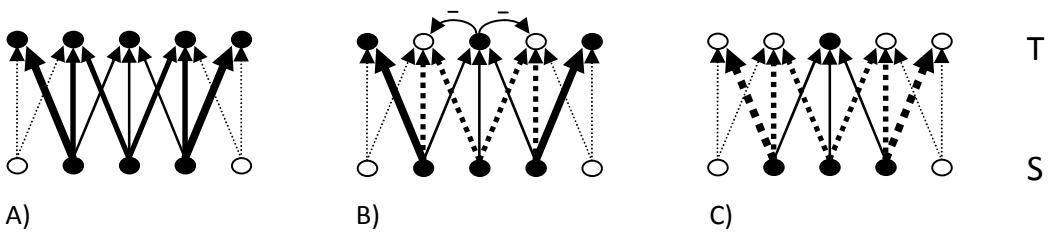

FIGURE 5: Approaches to competitive learning. A) No competitive learning. All neurons in the target s-layer T that can see an active source in the input s-layer S are connected to their inputs. The weight values are represented by the thickness of the connection line, reflecting rule (5). B) Lateral inhibition prevents neighboring neurons in s-layer T from developing their own connections. It is assumed here that the central neuron has been first to fire. C) Effect of pruning the connections in A). Each neuron in s-layer S keeps only the one output connection with the lowest weight. This automatically preserves in s-layer T the neuron with the most representation power, i.e. represents the largest number of active neurons in s-layer S.

A mechanism of competitive dynamics often referred to involves inhibitory lateral influence. The first neuron that starts responding to a patch of the stimulus, e.g. because it has the most inputs, or the most active ones, somehow prevents others in the same area from representing the same patch, e.g.

through lateral inhibition followed by LTD triggered by the lower membrane potential (figure 5B). However, there is no direct experimental data showing that this mechanism can operate on a short-time scale. Horizontal fibers show slow signal propagation and the fastest route for lateral inhibition would actually be through feedback from higher areas (Angelucci and Bullier, 2003, Gonchar and Burkhalter, 2003). There is evidence that feedback can have an inhibitory role in the Lateral Geniculate Nucleus (LGN) (Sillito and Jones, 2002) but it is less clear in other areas. Another problem is that the winning neuron is not necessarily the one with the most representational power (representing the most active input neurons) as neurons start firing at random time. This can prevent some elements of the input pattern from being represented.

A wider question is whether short-time scale competitive learning is actually required. There are good reasons for it to achieve a minimal use of neuronal resources, but the brain may not be that concerned about this at learning time, to deal with resource issues at consolidation time. In the PFC, for instance, more than a half of all randomly selected cells turn out to discriminate stimulus categories relevant for a current stimulus-target association task (Kuzunoki et al., 2010). This is an enormous number of neurons used to represent a few stimuli. Assuming that redundancy is acceptable, I will explore here a grow-then-prune approach to competition: In a first pass, all neurons that can see parts of the stimulus have their input weights trained using (10). In a second pass, neurons prune their output weights to minimize the representation cost.

Pruning operates as follows (fig 5C): Each neuron keeps only the output connection with the weakest non-zero weight. For instance, if a neuron has output connections with weights ½, 1/5 and 1/16. It will only keep its connection with the target that has 16 inputs. In this way, targets with many inputs stay preferentially connected, and, at the same time, the source neuron minimizes its weight cost. This pruning rule also ensures that every active input stays represented. If a target neuron looses all its input weights, it is also made to lose its output weights and is made available for learning new SR pairs. Neurons that survive the process are flagged as “consolidated” and their other input connections cannot take part in further learning. This ensures that learnt SR associations are stable and not interfered with by new learning. During pruning as describe above, each neuron preserves only one output weight, which is flagged as “consolidated”, but its other output connections stay available for learning further SR associations. Thus, over time, one input neuron can contribute to several SR associations. and have several consolidated output connections with different weights.

Several aspects of this part of this model and procedure are debatable. The concept of a two-phase competitive learning may seem artificial, but competition through slow lateral connections would also, in practice, remove connections *after* they have formed. The concept of pruning synapses with the strongest weights sounds counterintuitive, but it curiously achieves the same result as more conventional competitive dynamics. The concept of “flagged” frozen weights and neurons that do not take part in further learning seems a convenient ad-hoc feature introduced in this model to prevent interference, but there is evidence, in the field of motor skill learning, that skills become less prone to interference after a night sleep (Brashers-Krug et al., 1995). So, some form of freezing out must take place. Nevertheless, the lack of proper biological grounding of the competition mechanism implemented here is uncomfortable and more biological data on this process will be welcome.

In my simulation, the two passes are done successively to release resources immediately before a new SR pair is presented. In a simulation with a very large number of neurons, it may be possible to conduct pruning only after a number of SR pairs have been learnt, e.g. during a simulated “sleep” time.

5 Systems architecture

5.1 Anatomical basis

Anatomically, the shortest route from retina to motor output includes LGN, V1, PO(v6), PPC areas (dorsal bank of IST/V6a, 5, 7m), PM, and M1. Neurons in these areas have progressively larger receptive fields, increasing with eccentricity (Galletti et al., 1999). This behavior is captured in the hierarchical multi-layer model of Wallis and Rolls (1997), where RF sizes roughly increase by a factor 2.5 in each s-layer. The network used here follows that convergent multilayer model (figure 6), although cortical magnification and eccentricity effects are not included. The number of s-layers is determined by the degree of convergence in each s-layer, and, due to the discrete nature of my input grid, cannot match exactly the number of cortical relays. This is not a fundamental issue.

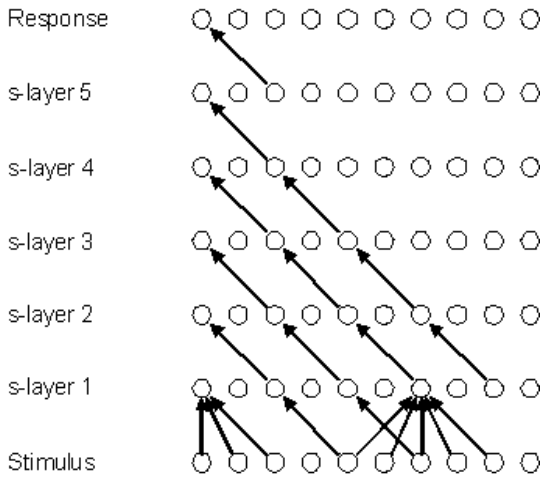

FIGURE 6: Network architecture. Each s-layer is a 10x10 square grid of neurons. Each neuron in an s-layer above Stimulus receives inputs from 5 x 5 neurons. The network is seen from the side.

Each s-layer in the model is constituted by the described learning microcircuit (section 4.2), except for the input and output.

The input s-layer, named “Stimulus”, represents the LGN which reflects the firing of retinal ganglion cells . No learning takes place between the retina and the LGN. So there is no need for enabling feedback to the LGN.

The primary motor cortex (M1) is the last stage of motor output specification. M1 is known to combine somato-sensory information and information on the desired action provided by PM to

generate actual motor commands (Shipp, 2005). The mapping between PM and M1 can be considered to reflect previously learnt basic motor skills. I therefore consider that the activity in the PM is the final specification of the response in the SR task. Anatomically, there is no definitive description of the cortical layers of termination of the feed-forward and feedback projections from PPC to PM. According to a rule based on the density of neurons of the involved areas (Medalla and Barbas, 2006), PPC drives PM in feed-forward mode and PM sends feedback projections to PPC. Given that Layer 5 neurons in PM project to M1, I will assume that layer 3 receives inputs and produces also feedbacks to PPC. Thus PM is modeled by a single L2/3 neuron in the Response s-layer. The learning signal is an activation of PM which can at the same time trigger actions, if deemed behaviorally appropriate by other prefrontal areas (Narayanan and Laubach, 2006), and generate feedback to drive learning in upstream areas.

The source of the initial learning signal in PM cannot be specified in my model as experimental data are missing on the precise mechanism of action initiation through verbal commands or instructions. I simply assume that the PM generated feedback signals corresponding to the desired pattern of response activation.

Data by (Ruge and Wolfensteller, 2009) show progressive performance and brain activity changes during repeated execution of a SR rule. It is possible that that execution of the action in itself generates learning signals so that the learning process is repeated during execution. The proposed architecture supports this, provided that the stimulus activity persists in lower visual areas to benefit from the teaching signal produced by the execution of the action.

As all the output spike trains have the same frequency, there is a question as to whether this quasi-binary output representation is suitable. The exact nature of the action representation in PM is not fully understood, however is it generally accepted that PM (and SMA) represent high-level task specification rather than execution details which may be more of an analog nature. In certain experiments, neurons seem to code for simple movements (Caminiti et al., 1991), sparsely for some neurons, and in a distributed way for others (Pesaran et al., 2006). This is not fundamentally incompatible with an ON/OFF interpretation of their output. In other experiments, PM neurons learn to code for an effector-independent abstract specification of the action (Wolfensteller et al., 2004), which can be thought of as binary information. Regarding the number of neurons activated in PM to encode an action, there is a patchy projection patterns from PM to M1 (Dancause et al., 2006) where the output of a PM neuron targets isolated groups of M1 neurons located at different points of the motor map. Therefore, a single PM neuron could activate a diverse mixture of M1 neurons and a corresponding complex muscle activation pattern. For generality however, I will assume that several neurons in PM need to be activated to specify a response. For generality also, I will assume that different stimuli could activate the same response, e.g. in a categorization task. This poses a consistency problem for the learning rule: If all synaptic resources have been used by a neuron to learn its input pattern, how could it learn a second one? One solution is to assume that there is redundancy in PM and that different neurons can code for the same action. In this view, the learning signal would activate multiple possible targets in PM, but learning would actually only recruit one of them. Learning during practice would then recruit more and more neurons in PM for representing the same action, up to using up the pool of redundant neurons, possibly explaining the increasing then saturating activity in PM with practice (Ruge and Wolfensteller, 2009). The redundancy hypothesis predicts a reduction in the capacity of association of new stimuli with the same response if a previous association with the same response has been repeatedly practiced.

The feedback from higher s-layers that enables learning in lower s-layers is sent by each learning microcircuit to the microcircuits in previous s-layers from which it receives feed-forward input. A cascade of feedbacks guides learning in the feed-forward mode towards targets in the PM. However, if a circuit in a higher level has already been recruited, there is no point for it to call for more input. Computationally, it therefore makes sense to disable recruited neurons from sending a learning signal to their inputs. This avoids attracting the feed-forward push into a possible dead end. It is unclear if biology offers a mechanism for implementing this. It is also not clear that this is actually needed. Given the large number of neurons in each area, it is unlikely that such a dead-end will ever occur. However, for my simulation with a small number of neurons, it is useful to remove recruited neurons from the feedback chain.

5.2 Network design

The simulated network has seven s-layers of square arrays of 10x10 LIF neurons models (neurons). Each s-layer projects to its higher target s-layer. The feedback activity is implemented by program rather than by connections. The first s-layer is an input s-layer called "Stimulus", corresponding to an early visual area, and the output s-layer is called "Response" corresponding to the PM. Intermediate s-layers are named s-layer 1 to s-layer 5 (upper halves of figure 7).

Each neuron receives inputs from a square 5x5 receptive field in the previous s-layer. Neurons on the edges and in corners see less input neurons, respectively 5x3 and 3x3. The number of s-layers used is determined by the size of the receptive field of each neuron. To ensure that any neuron in s-layer 5 can, through a cascade of relays in intermediate s-layers, receive inputs from the whole of the Stimulus s-layer, following number L of s-layers is required before the Response s-layer.

$$L = 1 + \frac{\sqrt{N} - 1}{0.5\sqrt{N_{RF}} - 0.5} \qquad (18)$$

Where N_{RF} is the number of neuron in the receptive field of each neuron and N is the number of neurons in each s-layer. The number given by (18) is to be rounded up to nearest integer, e.g. for N=100 and N_{RF} = 25, I find L= 6.

All weights are positive, except for some weights from Stimulus to s-layer 1 that are used to simulate the opponent mechanism underlying V1 tuning (Ferster and Miller, 2000). This improves the capacity of the SR learning model by reducing interferences between similar stimuli.

S-layers 1 to 5perform pattern recognition using a trained convergent tree structure. In my simulation I do not use multiple equivalent outputs. To save neurons, I allow one output in the Response s-layer to be activated by multiple stimuli in s-layers 5. Computationally, each Response neuron operates as an OR gate, able to respond if any of the trained input is active. As Response s-layer neurons share the same limited receptive field as all the other neurons, an extended response pattern will require several neurons in s-layer 5 to be trained to represent the same stimulus, to activate each a different subset of Response neurons.

6 Simulated learning experiments

6.1 Input coding

Stimuli and responses are stored as pairs of 10x10 binary bitmaps. For training, the values in the stimulus bitmap are used to set the spontaneous Poisson firing probability of neurons in the 10x10 Stimulus s-layer. All bright pixels set their corresponding neurons to a rate of 100Hz. Neurons corresponding to dark pixels stay silent. The Poisson process starts at time zero, when the simulation starts. The response bitmaps are used as described in 6.2.1.

6.2 Learning procedure

6.2.1 Activity pull

At the start of training, all target neurons in the Response s-layer are activated. Thereafter, one neuron in s-layer 5 is selected as the head of the pull cone. This neuron is selected on the basis on the number of neurons in the Response s-layer that it is able to activate. From s-layer 5, learning-enabling feedback is propagated down the network by a program routine looking up all the inputs of an enabled neuron, working backward from s-layer 5. Each enabled neuron has its “enabled” flag set to 1, representing the higher background potential of neurons receiving feedback from their targets.

Another option would be to use all activated neurons in Response as source of enabling signals, and increment the value of the “enabled” flag by one for each feedback received. Thereby, neurons in lower s-layers with a high value stored in their “enable” flag are more likely to contribute to a tree that is able to activate several neurons in Response. This “multi-headed pull” option gives an additional dimension to the learning signal, but is left to future work. The consequence of using one target neuron at a time in s-layer 5 is that several runs may be needed to train one SR association in the case of extended response patterns.

6.2.2 Representational push

Once the enabled cone is defined, learning starts in all s-layers, driven by the initial activity in Stimulus, using rule (10). Progressively, neurons become active in s-layer 1 and weights can start to increase between s-layer 1 and s-layer 2, and so on, until the desired s-layer 5 neuron is activated. This is a stochastic process, and some neurons in s-layer 1 may not have started firing yet, while other already make connections with s-layer 2. Progressively, learning progresses to higher s-layers, with the restriction that only enabled neurons are able to learn. A series of snapshots of the intermediate states of the network activity is given in figure 7. Once the single target identified in s-layer 5 starts firing, the connections with all accessible active neurons in Response are established (corresponding to the divergent pattern of connectivity seen e.g. in figures 7 and 8). This is the termination condition of learning in the simulation. After that, the pruning process takes place, as described in section 4.3. If not all neurons in the Response s-layer are connected, a new “head of cone” is selected in s-layer 5 and a second round of learning takes place. This process is taking place within one trial where several trees can be developed to service different groups of neurons in the response pattern. This ends when all

neurons in the Response have been connected to the stimulus by a tree of relays in intermediate s-layers.

A form of the opponent mechanism of Ferster and Miller (2000) is implemented here. For that purpose, all weights between inactive neurons in the Stimulus and consolidated neurons in s-layer 1 are set to -1. This prevents patterns that contain more active neurons than the trained one from activating the neuron. E.g. a neuron trained for “C” should not respond to “O”. It is a useful feature but is not key for the presentation of the concepts in this paper.

Note that, if different stimuli share common features, some consolidated neurons of a previously trained tree will be activated by the stimulus and re-used when constructing a new tree. If stimuli had many shared patterns, the pool of recruited neurons in s-layer 1 would eventually stop to grow.

(caption for the figure on the following page)

FIGURE 7. Snapshots of the network during the learning process. The figure shows the state of the network at times t = 10 ms, 20 ms, 30 ms,50 ms and after pruning has taken place. The stimulus is the A-shape and the response is a simplified representation of the right hand, symbolizing a right-hand response to the stimulus. The upper row of grids for each time represents the activity of neurons in the 7 s-layers of the network. Large white squares represent a spike and small white squares represent an active neuron, not spiking at the time the snapshot is taken. The lower row of grids has no active neurons. Its purpose is to display which neurons in the corresponding upper grid are in the enabled cone (large white squares) and which ones are recruited (large grey/red squares). The last grid on the right shows the target of learning in the Response s-layer. It is extinguished once the corresponding neurons in the Response s-layer are activated by their inputs in s-layer 5. Learning proceeds at a brisk pace and activity rapidly propagates through the s-layers. One can see in the lower row of grids that the number of neurons recruited rapidly increases. After pruning, only a few consolidated neurons remain. The blue lines superposed to the upper row of s-layers show the consolidated tree of connections. A few orphan neurons are not included in the tree, including the top pixel of the A. These are part of a secondary tree branch that has built up to s-layer 4, but has not managed to be connected to the neuron in s-layer 5 before the Response was activated

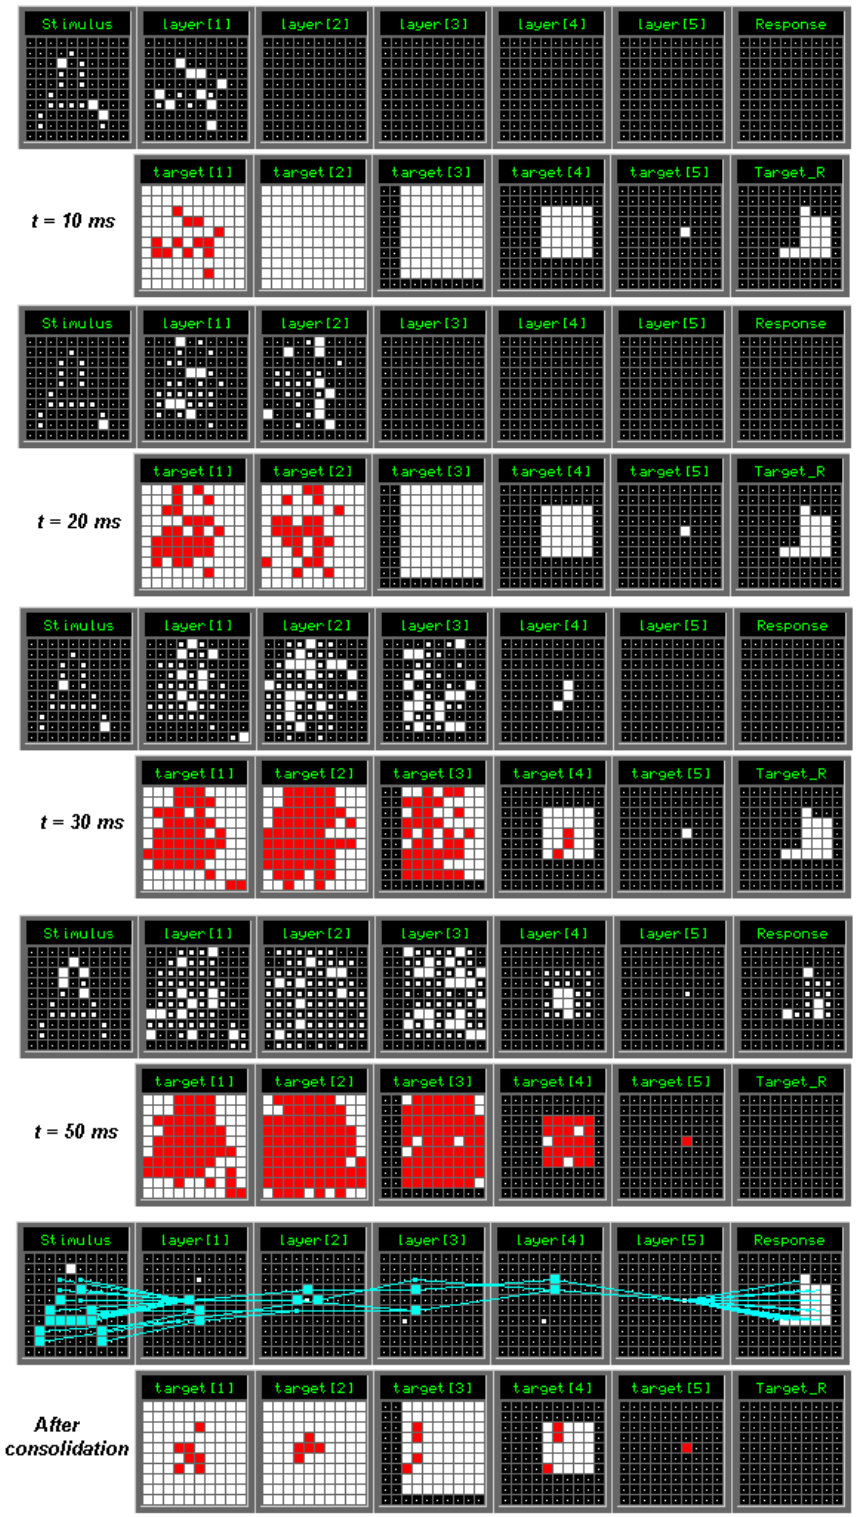
Stimulus
layer[1]
layer[2]
layer[3]
layer[4]
layer[5]
Response
t = 10 ms
target[1]
target[2]
target[3]
target[4]
target[5]
Target_R
t = 20 ms
t = 30 ms
t = 50 ms
After consolidation

6.3 Experiment 1: Learning four SR rules.

The first experiment was done to test the capability of the network to learn several different SR pairs. I trained the network to associate the stimuli A and B with a right-hand response, represented by a right hand icon in the Response s-layer (figure 8) and the stimuli C and D with a left hand response. The network was able to learn this without error, with the help of the opponent feature in s-layer 1 (section 6.2.2). Figure 8 shows the main connection trees developed during training.

During the initial phase of learning, almost all neurons that see some parts of the stimulus develop input weights that enable them to respond to the part of the stimulus in their receptive field. Pruning plays an important role in reducing the number of trained connections and neurons to those shown in figure 8. The competitive mechanism is here probably more important than the learning rule in shaping the trained network.

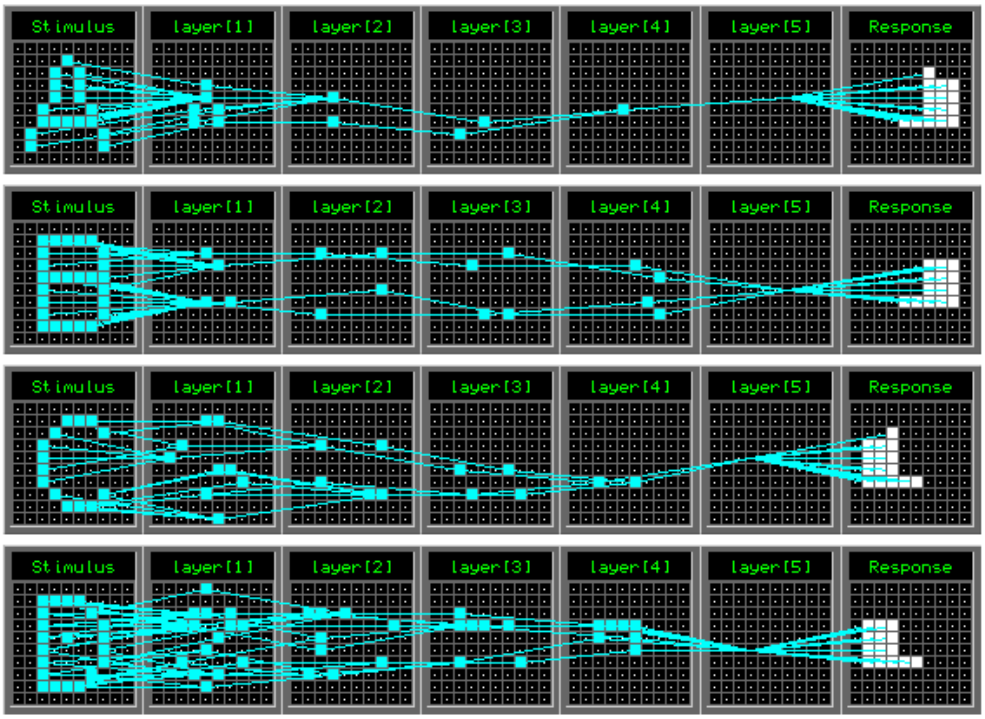

FIGURE 8. Main connection trees developed during training of four stimulus-response associations. A->right hand response, B->right hand response, C-> left hand response, D-> left hand response. Two other trees that link B and D to the “finger tip” of the left and right-hand responses are not shown. Parameters: Synaptic transmission probablility=1.0, eta=0.8, weights updates 2 ms after spike arrival, i.e. at the peak of the alpha function. W_0=8.

6.4 Experiment 2: Simulating the effect of practice.

Experiment 2 is inspired by (Ruge and Wolfensteller, 2009) and is designed to explore the effect of practice on performance and on level of MRI activity along the dorsal route. It is assumed that, during practice, the network is exposed to the same stimuli and responses as during initial instruction and is therefore repeatedly learning the practiced SR association. In the simulations, the network was presented 15 successive times with the SR pair A -> right hand response. After each pair has been trained, i.e. the response has been activated, pruning took place. The total number of neurons recruited in s-layer 1 was recorded after each presentation. This is taken to be a good measure of how many additional nodes and connections have recruited to represent the SR association.

After each single SR training, the response time (RT) of the network is measured by applying the stimulus and waiting until all the expected neurons in the response pattern have been activated.

The results show that redundant connections and neurons are recruited by repeated training (figure 9A and 9B), despite the creation of a functional tree already after the first exposure to the SR pair. Figure 10 shows an example of tree resulting from 15 training runs in the conditions of figure 9A.

The formation of redundant connections is enhanced by the use of stochastic synapses or by early updating of synaptic weights. In the first case, new connections often managed to activate the response before the activity following existing connections. In the second case, early weight updating is known to lead to a faster propagation of the activity during learning than during normal activation (see figure 3). With deterministic synapses, redundant connections can only be produced through early updating (Fig 9B). The response times are then shorter due to the larger effective input frequency, however, redundant connections have no effect on the response time. The reasons for this need to be investigated further, but an initial hypothesis is that the additional connections formed are in some sense not functional, i.e. do not form a full tree linking input to output.

The response time diminishes as the number of redundant connections increases (figure 9A) only when synapses are probalilistic. The RT reduction is especially clear when the number of redundant connections increases rapidly with practice. These curves bear similitude with experimental data by Ruge and Wolfensteller (2009). In these instruction experiments, the subjects exhibited a practice-related increase of the MRI signal in areas of the dorsal stream, similarly to the increase in the number of redundant connections observed here. At the same time, the response time of the subjects diminished by about the same time difference as the RT measured here.

In the simulations, these fast increases in the number of redundant connections were obtained with synapses with a 50% failure probability and with a fast recruitment of synaptic resources, during the peak of the alpha function, i.e. 1ms after the spike has arrived. Such a fast recruitment is probably not realistic, given that resources transfer between synapses is assumed in the model. However, beyond the details of the simulation, the results suggests that biological SR learning could be the result of a very disorganized rush of activation through the network.

The saturation number of neurons recruited in s-layer 1 is determined by the size of the stimulus. Essentially, every neuron in s-layer 1 that can see a part of the stimulus is eventually recruited.

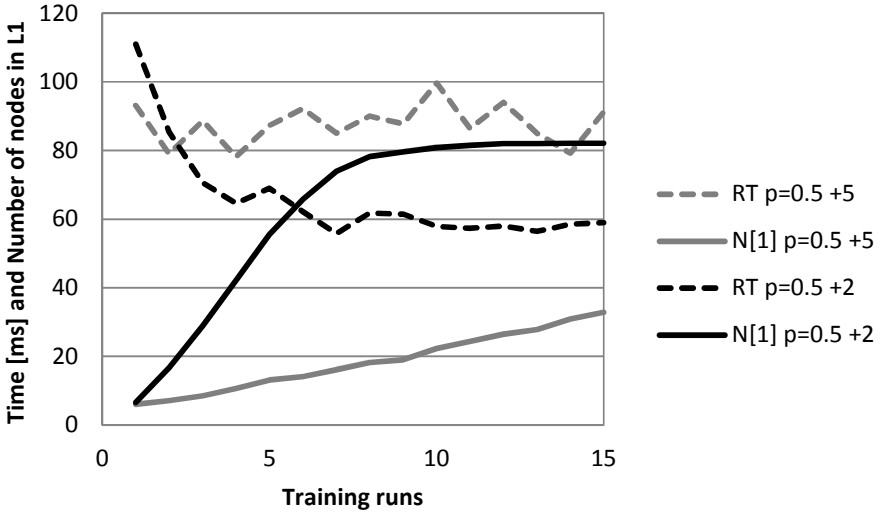

A)

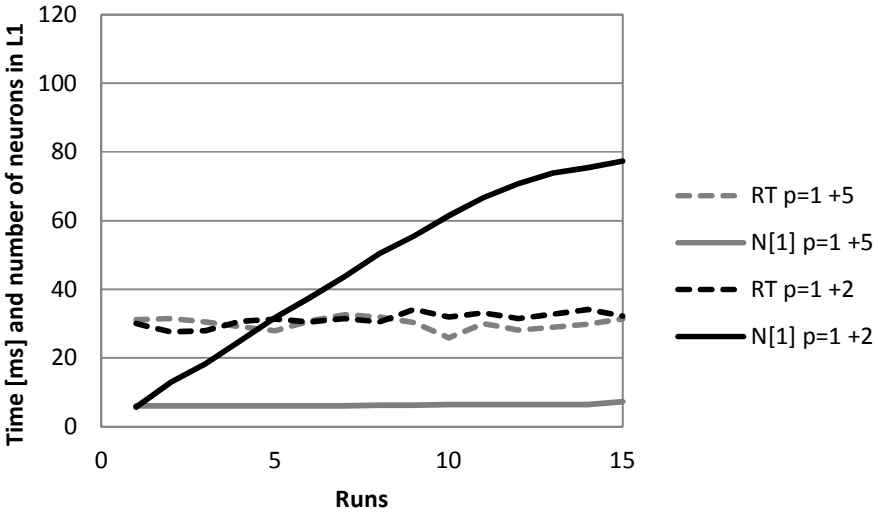

B)

FIGURE 9. Results of repeated learning of the same SR pair, e.g A->Right-hand, for two values of the synaptic transmission probability and of the weight updating delay. A) Synapses are probabilistic, with transmission probability p=0.5. Weights are either updated at the peak of the alpha function (+2) or in the tail of the alpha function, 3ms after its peak. B) Synapses are deterministic (p=1) with two updating times as in A). The continuous curves marked with N represent the total number of neurons recruited in s-layer 1. The dotted lines represent the response time of the network to presentation of the stimulus after the indicated number of runs.

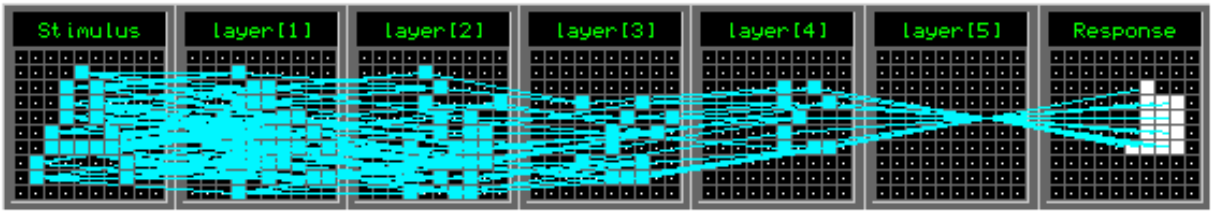

Figure 10: Connection tree developed after 15 training runs with the same A->right hand response SR pair. The network has recruited 4 nodes in layer 5, of which 3 show similarly dense trees. This figure can be compared with the top row of figure 8, showing a large number of redundant connections have been developed. The only difference in network parameters is the synaptic transmission probability p = 0.5. Simulations of 15 training runs with p=1 produce a single tree, as in the top figure 8.

7. Discusssion

Key features of the model: It was important that (Mayr and Partzsch, 2010) pointed to the fact that biological learning does not require postsynaptic spiking, This allowed to conceived of a method for "pushing" the stimulus activity through the network towards the response s-layer. To achieve fast learning with a learning speed that is independent on the number of active inputs, I had to break with the standard view of biological learning as a "blind" and slow progressive increase of weights. With the new rule, weights are all increasing at different paces to reach their specific target values at the same time. With a large value of η, learning can be almost instantaneous and still reach the right value. A hypothesis and prediction of the model is that trained synaptic efficacies take final values in an inverse relation to the number of synapses activated during training.

It should be noted that the proposed resource-limited learning rule naturally develops neurons as pattern recognizers. It eliminates the inherent instability of Hebbian learning rules offers a natural weight normalization mechanism. Whether the total quantity of synaptic efficiency is truly a constant is debatable, as recent work on homeostatic activity control suggest a possible adaption of that quantity (see e.g. Turrigiano, 2008).

A key hypothesis in the model is that interlayer feedbacks in the cortex are providing learning information. I call this the "pull" that guides the feed-forward representation "push". Whether this makes sense and how it can be tested is an open question. Learning effects described in my model should affect higher areas, beyond primary sensory areas.

Another important point is the use of a simulation of spiking neurons producing Poisson spike trains. Experimentation with the network has pointed to the fact that fast learning is an asynchronous stochastic process driven by the random arrival of input spikes.

Simulations have shown that competition is an important process. The push model of learning is quite indiscriminate in that it recruits a large number of neurons. The final shape of the network is strongly dependent on the competition mechanism. The pruning model of competition implemented here is functional but has no clear biological support. This can however be said of most models of competition. Much work is presently ongoing on homeostatic regulation of synaptic strengths and there

are tentative suggestions on the possible role of similar mechanisms for competition (see e.g. Watt and Desai, 2010) , but this is very much work in progress. Further work is needed to i) determine the precise form of biological competition, including the time scale at which it operates and ii) how details of competition mechanisms affect the final shape of trained networks.

The "problem" of multiple constructed trees was revealed by the simulation and could actually constitute an explanation for biological observations (Ruge and Wolfensteller, 2009). Biological systems would in principle benefit from fast SR learning (see also argument below), but this appears to be at the cost of redundant representations being built.

Timing questions: Training requires i) the pull signal to be present during the learning time and ii) the push signal to last as long as it takes for the information to find its way through the network. The requirement for persistence of the input in the first layer of a multilayer artificial neural system is due to the fact that decrements in activity propagate much faster than increments (Bugmann, 1992). As soon as the stimulus stops, the activity in all downstream layers stops too. If I assume that the duration of the input signal is solely determined by retinal persistence, then all training has to be completed in 40-80ms. This imposes the use of high learning rates with their side effect of redundant representations.

While the duration of input and pull signals may not be an issue in imagery mode, it is critical if learning is to take place during practice. First, in a feed-forward phase, the stimulus activates the response. According to my simulations, this takes around 80 ms. Once the response is active, it is able to send a learning enabling signal back in less than 20 ms. Such a short time consists essentially of the propagation time (approx 2 ms per synapse). Neurons in the feedback stream are not supposed to integrate inputs, but should be able to fire as soon as one input is active. So, 100 ms after stimulus onset, training can start. This can typically take 50-80ms. Overall, stimulus activity needs to stay available for up to 180ms. This cannot be achieved through retinal persistence alone and requires either a physical presence of the stimulus, of some internal mechanism for supporting sustained firing. The latter is not totally implausible given data by (Rolls and Tovee, 1994) showing neural firing outlasting a stimulus by close to 300ms. Even without such a mechanism, the delay to action initiation after PM stimulation is in the range 60 — 150 ms or longer (Weinrich and Wise, 1982), and the stimulus would normally stay in the field of view for a sufficient period of time. However, the role of persistence could be tested through masking experiments. Masking or early saccades could interfere with the improvement in performance during practice after instruction.

Open questions: One open question concerns the actual source of learning signal in the PM. How is the PM activated by the instructions? Several routes where mentioned and discussed in this paper, but none has definite experimental support.

Another question is about the proposed synaptic learning rule. It has the benefit of speed and precision, but needs experimental confirmation (see section 4.1).

As noted in section 4.3, more data are need on biological mechanisms of competition. Data on possible mechanisms for shielding trained neurons from recruitment for new tasks would also be welcome.

Finally, the concept of saturation of available connections during practice could be verified through behavioral experiments, especially regarding multiple stimuli associated with the same response (see section 5.1 and 6.4).

8. Conclusion

A model of synaptic learning has been proposed and demonstrated. The rule defines saturation values for synaptic weights with an inverse relation to the number of activated synapses, on the basis of the hypotheses that the total synaptic weight of a neuron is a limited resource. This rule produces naturally neurons with pattern recognition capabilities, in a binary neural coding scheme, where key information lies in the fact that a neuron is firing or not.

It is proposed that SR learning proceeds through a pull-push mechanism. Its use for the simulations of learning stimulus-response associations in a multilayer network of spiking neurons shows that competition is an important component of learning, as it ultimately determines the shape of the trained network. The proposed implementation, using synaptic pruning, is probably not the most biologically plausible, but its effect is probably similar to that of other implementations.

This work offers a stochastic view of SR learning, where activity races to trigger the desired response, generating multiple redundant paths in the process. The availability of these multiple alternative routes could explain the reduction in response time with practice. The results suggest that probabilistic nature of synaptic transmission plays an important role in practice-related performance improvement.

The main purpose of this work was to elaborate a computational model able of fast interference-free learning. I have taken inspirations from a range of biological data and have proposed a new computational description of some known structures. Models normally represent our understanding of a system. This work has revealed so many questions that it can, at best, be seen as an intermediate representation that hopefully will inform future work.

Acknowledgments:

I gratefully acknowledge very helpful comments by Roman Borisyuk and anonymous reviewers.

References

Anderson J, Lampl I, Reichova I, Carandini M, and Ferster D. (2000) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. Nature Neurosci, 3:617-621.

Angelucci A, Bullier J (2003) Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? J Physiol. (Paris) 2003, 97:141-154.

Artola, A., Bröcher, S., and Singer, W. (1990). Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. Nature 347, 69–72.

Ashby F.G., Turner B.O. and Horvitz J.C. (2010) Cortical and basal ganglia contributions to habit learning and automaticity. Trends in Cognitive Sciences Vol.14 No.5, 208-215.

Bell A. and Sejnowski T.J. (1997) The “Independent Components” of Natural Scenes are Edge Filters. Vision Res. 37(23): 3327–3338

Brashers-Krug T., Shadmehr R. and Todorov E. (1995) Catastrophic interference in human motor learning. In: Advances in Neural Information Processing Systems, vol. 7, G. Tesauro, D. S. Touretzky, T. K. Leen (eds), MIT Press, pp. 19-26.

Bugmann G. (1992) The neuronal computation time. in Aleksander I. and Taylor J.G. (eds) "Artificial neural networks II", Elsevier, 861-864.

Bugmann G. (2002) Synaptic depression increases the selectivity of a neuron to its preferred pattern *and binarizes the neural code. Biosystems: Special Issue on Neural Coding, 67, 17-26.*

Bugmann G., Christodoulou C. and Taylor J.G. (1997) Role of Temporal Integration and Fluctuation detection in the highly irregular firing of a Leaky Integrator Neuron Model with Partial Reset. Neural Computation, 9, pp. 985-1000.

Bunge SA, Wallis JD, Parker A, Brass M, Crone EA, Hoshi E, Sakai K (2005) Neural circuitry underlying rule use in humans and non-human primates. Journal of Neuroscience 25, 10347-10350.

Brass M., Wenke D., Spengler S., Waszak F. (2009) Neural Correlates of Overcoming Interference from Instructed and Implemented Stimulus–Response Associations. The Journal of Neuroscience, 29(6):1766-1772.

Caminiti R. Johson P.B., Galli C., Ferraina S. and Y. Burnod (1991) Making arm movements within different parts of space: the premotor and motor representation of a coordinate system for reaching to visual targets. J. Neurosc. 11(5) 1182-1197.

Cole M.W., Bagic A., Kass R. and Sccneider W. (2010) Prefrontal dynamics underlying rapid instructed task learning reverse with practice. J. Neurosci. 30(42): 14245-14254.

Dancause N., Barbay S., Frost S.B., Plautz E.J., Popescu M., Dixon P.M., Stowe A.M., Friel K.M. and Nudo R.J. (2006) Topographically Divergent and Convergent Connectivity between Premotor and Primary Motor Cortex. Cerebral Cortex 16:1057-1068.

Douglas R. and Martin K. (1991) A functional microcircuit for cat visual cortex. J. Physiol. 440:735-769.

Ferster, D., and Miller K.D. (2000) Neural mechanisms of orientation selectivity in the visual cortex, Ann. Rev. Neurosci., 23:441-471.

Fusi S., Asaad W.F., Miller E.K. and Wang X-J. (2007) A neural circuit model of flexible sensorimotor mapping: Learning and forgetting on multiple timescales. Neuron 54:319-333.

Galletti C., Fattori P., Gamberini M. and Kutz D.F. (1999) The cortical visual area V6: brain location and visual topography. European Journal of Neuroscience, 11: 3922-3936

Ganisa G., Thompson W.L., Kosslyn S. M. (2004) Brain areas underlying visual mental imagery and visual perception:an fMRI study. Cognitive Brain Research 20: 226–241.

Garagnani M., Wennekers T. and Pulvermuller F. (2009) Recruitment and consolidation of cell assemblies for words by way of hebbian learning and competition in a multilayer neural network. Journal of Cognitive Computation 1, 160-176.

Gonchar Y. and Burkhalter A. (2003) Distinct GABAergic Targets of Feedforward and Feedback Connections Between Lower and Higher Areas of Rat Visual Cortex. The Journal of Neuroscience, 23(34):10904 –10912

Henry G.H., Salin P. A., J. Bullier J. (1991) Projections from Areas 18 and 19 to Cat Striate Cortex: Divergence and Laminar Specificity. European Journal of Neuroscience, 3(2): 186–200.

Klein, I., Paradis, A.-L., Poline, J.-B., Kosslyn, S. M. and LeBihan, D. (2000) Transient activity in human calcarine cortex during visual imagery. Journal of Cognitive Neuroscience, 12, 15S-23S.

Krichmar, J. L., Seth, A. K., Nitz, D. A., Fleischer, J. G., and Edelman, G. M. (2005). Spatial navigation and causal analysis in a brain-based device modeling cortical-hippocampal interac-tions. Neuroinformatics 3, 197-221

Kusunoki M., Sigala N., Nilli H., Gaffan D. and Duncan J. (2010). Target Detection by Opponent Coding in Monkey Prefrontal Cortex, Journal of Cognitive Neuroscience, 22(4):751-60.

Makuushi A., Kaminaga T. and Sugishita M. (2005) Brain activation during ideomotor praxis: imitation and movement executed by verbal command. J. Neurol Neurosurg Psychiatry 76:25-33.

Martinez-Conde S., Cudeiro J., Grieve K.L., Rodriguez R., Rivadulla C., and Acuña C. (1999). Effects of feedback projections from area 18 layers 2/3 to area 17 layers 2/3 in the cat visual cortex. J. Neurophysiol. 82: 2667–2675.

Mayr CG and Partzsch J (2010) Rate and pulse based plasticity governed by local synaptic state variables. Front. Syn. Neurosci. 2:33.

Medalla M. and H. Barbas H. (2006) Diversity of laminar connections linking periarcuate and lateral intraparietal areas depends on cortical structure. European Journal of Neuroscience, 23: 161–179.

Middleton F.A. and Strick P.L. (1998) Cerebellar output: motor and cognitive channels.Trends in Cognitive Science, 2(9): 348-354.

Miller KD. (1996) Synaptic economics: competition and cooperation in synaptic plasticity. Neuron Vol. 17 No.3, pp.371-374.

Narayanan NS, Laubach M. (2006) Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. Neuron 52:921-931

Montgomery J.M., Pavlidis P. and Madison D.V. (2001) Pair Recordings Reveal All-Silent Synaptic Connections and the Postsynaptic Expression of Long-Term Potentiation. Neuron Vol. 29, 691–701

Nowak D.A., Koupan C., Hermsdoerfer J. (2007) Formation and decay of sensorimotor and associative memory in object lifting. EUr. J. Appl. Physiol. 100:719-726.

Royer S, Paré D (2003) Conservation of total synaptic weights through balanced synaptic depression and potentiation. Nature 422: 518-522.

Ruge H. and Wolfensteller U. (2009) Rapid formation of pragmatic rule representation in the human-brain during instruction-based learning. Cerebral Cortex, 20: 1656-1667.

Pandya, D. N. and Yeterian, E. H. (1984) Architecture and connections of cortical association areas. In A. Peters and E.G. Jones (Eds.), Cerebral Cortex, Association and Auditory Cortices, Vol. 4, Plenum Press, New York, 1984, pp. 3-61.

Passingham R.E. and Toni I. (2001) Contrasting dorsal and ventral visual systems: guidance of movement versus decision making. NeuroImage 14, S125-S131.

Pesaran B., Nelson M.J. and Andersen R.A. (2006) Dorsal Premotor Neurons Encode the Relative Position of the Hand, Eye, and Goal during Reach Planning. Neuron 51, 125–134

Reiter H.O.and Stryker M.P. (1988) Neural plasticity without postsynaptic action potentials: less-active inputs become dominant when kitten visual cortical cells are pharmacologically inhibited. PNAS, 85(10):3623-3627

Rizzolatti G. and Luppino G. (2001) The Cortical Motor System. Neuron, Vol. 31, 889–901

Rolls,E.T. and Tovee,M.J. (1994) Processing speed in the cerebral cortex and the neurophysiology of visual masking. Proceedings of the Royal Society, B, 257: 9-15.

Shipp S. (2005) The importance of being agranular: a comparative account of visual and motor cortex. Phil. Trans. R. Soc. B 360: 797-814

Shipp S. (2007) Structure and function of the cerebral cortex. Current Biology 17(12):R443-R449.

Sillito A.M. and Jones H.E. (2002) Corticothalamic interactions in the transfer of visual Information. Phil. Trans. R. Soc. Lond. B 357(1428): 1739-1752.

Sjöström P.J., Turrigiano G.G. and Nelson S.B. (2001) Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. Neuron, 32: 1149–1164.

Sjöström P.J., Rancz E.A., Roth A. and Häusser M. (2008) Dendritic Excitability and Synaptic Plasticity. Physiol Rev 88:769-840.

Standage D. and Trappenberg T (2007) The trouble with weight-dependent STDP. Proceedings of the International Joint Conference on Neural Networks 2007, 1359 - 1364

Tanaka S. and Miyashita M. (2009) Constraint on the Number of Synaptic Inputs to a Visual Cortical Neuron Controls Receptive Field Formation. Neural Computation 21, 2554–2580

Tettamanti M., Buccino G., Saccuman M.C., Gallese V., Danna M., Scifo P., Fazio F., Rizzolatti G., Cappa S.F., Perani D. (2005) Listening to action related sentences activates fronto-parietal motor circuits. J. Cogn Neurosci, 17:273-281.

Thomson A.M. and Lamy C. (2007) Functional maps of neocortical local circuitry. Frontiers in Neuroscience. 1(1):19-42.

Thorpe S. (2002) Ultra-Rapid Scene Categorization with a Wave of Spikes. H.H. Bülthoff et al. (Eds.): BMCV 2002, LNCS 2525, pp. 1–15. Springer-Verlag Berlin Heidelberg

Toth LJ, Assad JA (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. Nature 415:165–168.

Tsodyks M.V. , Markram H. (1997) The Neural Code between Neocortical Pyramidal Neurons Depends on Neurotransmitter Release Probability. PNAS, 94, 719-723.

Turrigiano G.G. (2008) The self-tuning neuron: synaptic scaling of excitatory synapses. Cell. 135(3):422-35.

Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. Progress in Neurobiology, 51, 167–194.

Wang Y, Markram H, Goodman PH, Berger TK, Ma J, Goldman-Rakic PS. (2006) Heterogeneity in the pyramidal network of the medial prefrontal cortex. Nature Neuroscience. 9(4):534-42.

Watt, A. J., and N.S. Desai. (2010) Homeostatic plasticity and STDP: keeping a neuron's cool in a fluctuating world. Frontiers in Synaptic Neuroscience 2:5.

Weinrich M. and Wise S.P. (1982) The premotor cortex of the monkey. The Journal of Neuroscience 2(9): 1329-1345

Wolfensteller U., Schubotz R.I. and von Cramon D.Y. (2004) "What" Becoming "Where": Functional Magnetic Resonance Imaging Evidence for Pragmatic Relevance Driving Premotor Cortex. The Journal of Neuroscience 24(46):10431–10439.