

A spike based learning rule for generation of invariant representations

Konrad Körding*, Peter König

Institute of Neuroinformatics, ETH/University Zürich, Winterthurerstr. 190, 8057 Zürich, Switzerland

Received 27 June 2000; accepted 19 October 2000

Abstract – For biological realism, models of learning in neuronal networks often assume that synaptic plasticity solely depends on locally available signals, in particular on only the activity of the pre- and post-synaptic cells. As a consequence, synapses influence the plasticity of other synapses exclusively via the post-synaptic activity. Inspired by recent research on the properties of apical dendrites it has been suggested, that a second integration site in the apical dendrite may mediate specific global information. Here we explore this issue considering the example of learning invariant responses by examining a network of spiking neurones with two sites of synaptic integration. We demonstrate that results obtained in networks of units with continuous outputs transfer to the more realistic neuronal model. This allows a number of more specific experimental predictions, and is a necessary step to unified description of learning rules exploiting timing of action potentials. © 2000 Elsevier Science Ltd. Published by Éditions scientifiques et médicales Elsevier SAS

1. Introduction

Interest in neural networks originates to a large degree in their ability to learn from examples. Present research concentrates on synaptic plasticity as a mechanism supporting adaptive behaviour [32]. However, it is not clear which variables at the synapse, in the neurones or in the network co-determine synaptic plasticity. Due to the lack of evidence of the contrary, it is often assumed that solely local variables determine synaptic plasticity. As a consequence in connectionist systems synaptic plasticity usually depends only on information available locally at the synapse and in particular on the averages of pre- and post-synaptic activities. However, recent experimental results suggest a re-examination of this issue.

The most abundant type of nerve cells in the cerebral cortex, pyramidal neurones, display a remarkable asymmetry. Their basal and the apical dendrites have distinct morphologies and functional properties. Research on the properties of apical dendrites suggests that it acts, in addition to the soma, as a second site of integration [6, 27, 28]. Indeed, the electrotonic length of the apical dendrite for passive propagation of post-synaptic potentials towards the soma is of the order of 0.25 mm leading to a strong attenuation of non-regenerative potentials (H. Lüscher, pers. comm.). In the reverse direction the electrotonic length is of

the order of 0.62 mm allowing a limited contribution of somatic potentials for depolarization of the distal dendrite (H. Lüscher, pers. comm.). Thus, each site integrates input from a subset of synapses defined anatomically via their position on the soma or the dendritic tree and is able to generate regenerative potentials [43]. The sites of integration exchange information in well characterized ways. Signals from the soma are transmitted to the apical dendrite via actively back-propagating dendritic action potentials [1, 9, 45] or, to a limited degree, passive current flow. Signals from the apical dendrite to the soma are sent via actively propagating slow regenerative calcium spikes, which have been observed in vitro [43] and in vivo [19, 21]. These calcium spikes are initiated in the apical dendrites, and cause a strong and prolonged depolarization, typically leading to bursts of action potentials [28, 46]. Experimental studies support the view that excitation to the apical dendrite is strongly attenuated on its way to the soma unless calcium spikes are induced [28, 29, 43, 45]. Importantly, it has been shown that the threshold for the generation of a dendritic calcium spike is lowered if the excitation of the apical dendrite is paired with an action potential and increased by local inhibitory activity [28]. In conclusion, a subset of synapses on the apical dendrite is able to induce rare discrete events of strong prolonged depolarization combined with bursts.

Investigating the influence of post-synaptic depolarization on plasticity of synapses the synaptic weight change was measured as the function of post-synaptic activity [2, 13]. Those experiments

* Correspondence and reprints.

E-mail addresses: koerding@ini.phys.ethz.ch (K. Körding), peterk@ini.phys.ethz.ch (P. König).

showed that at low firing rates/depolarization LTD occurs whereas at higher activities a switch to LTP takes place. Thus, without a calcium spike and, as a consequence, with a limited depolarization, the cell is in a regime where synapses are depressed. In contrast, when a calcium spike is triggered, the high depolarization leads to a potentiation of active synapses. Thus integrating the above two lines of research naturally leads to a plausible biological implementation for the hypothesized mechanism, where calcium spikes could correspond to ‘learning events’. Indeed experiments by Pike et al. [38] give further evidence that post-synaptic bursting is essential for the induction of LTP.

The generation of calcium spikes is highly sensitive to inhibitory activity. Even the activity of one single inhibitory neurone can significantly raise the threshold of calcium spike generation [28, 29]. As axons of inhibitory interneurons branch homogeneously within the range of a hypercolumn this implements an effective soft winner take all mechanisms on the level of calcium spikes. It seems reasonable to assume that the number of neurones generating calcium bursts upon presentation of a stimulus is limited.

To complete the picture we have to consider, which afferent is targeting the apical and basal dendritic tree respectively? The anatomy of a cortical column is complicated; nevertheless, some regular patterns can be discerned. The apical dendrites of the considered layer 5 pyramidal cells receive long range cortico-cortical projections [10, 50]. Furthermore, top-down projections from areas higher in the hierarchy of the sensory system usually terminate in layer 1, where many apical tufts can be observed [41]. This supports the idea that top-down connections from higher to lower areas preferentially terminate on the apical dendrites. The basal dendrites of the considered cells receive direct subcortical afferents (e.g. the koniocellular pathway in visual cortex) in addition to projections from layer 4 spiny stellate cells. These are the main recipients of afferents from sensory thalamus or from areas lower in the cortical hierarchy. Therefore, we use the approximation that the bottom-up input targets mainly the basal dendritic tree.

In previous work we proposed that these ingredients – layer 5 pyramidal neurones possess two independent sites of integration with separate threshold processes; the somatic site of integration receives bottom-up information on stimulus properties and dominates the level of cellular activity,

i.e. the spike rate; the integration site in the apical dendrite receives top-down information and gates synaptic plasticity of synapses in the whole neurone – amount to a learning rule with interesting properties and analysed an implementation in a network of units with continuous outputs [27]. From the experimental results described above, it is obvious that a network composed of spiking units is the much more relevant system to investigate, which is the subject of this article. Furthermore, we discuss the investigated learning rule as a possible implementation of algorithmic learning rules defined on a global level.

2. Methods

The simulated network mimics a cartoon visual system roughly as outlined in a previous work [27]. It consists of two processing streams, which can be interpreted either as representing non-overlapping patches of the visual field or as one stream with time unfolded into space (*figure 1*). This allows to exploit object continuity in space or time as described below. Each stream contains an input layer and a hierarchy of two more layers. Each layer contains excitatory neurones with two separate compartments (somatic and apical dendritic as described below) and local inhibition. Bottom-up projections originate in the input layer and terminate at the somatic compartment of the excitatory neurones of layer 1 within each stream. The activity of these neurones is then relayed to the somatic compartment of excitatory neurones of layer 2. Layer 2 neurones send top-down projections to layer 1 terminating at the apical dendritic compartment as well as to the apical dendritic compartment of neurones in layer 2 of the other stream. Furthermore, each layer contains local inhibitory mechanisms (*figure 2*).

A Poisson statistic with a rate set directly by the topographically mapped stimulus describes activities of input units. Excitatory cells are implemented as leaky integrate and fire neurones along with a mechanism for burst generation. They are described by two variables corresponding to the two sites of integration: A is referred to as the somatic potential, D represents the potential at the apical dendrite and O the output of the neurone. When the somatic potential of a neurone reaches threshold (θ_s) a spike is emitted and the somatic potential is reset to its resting value (θ_r):

$$\text{If } A_i(t) > \theta_s \text{ then } O_i(t) = 1$$

and

$$A_i(t) = \theta_v, \text{ or else } O_i(t) = 0$$

The subthreshold dynamics of A is described in layer 1 by:

$$\Delta A_i^1 = C_1 O^0 * \mathbf{W} / \underline{A}_i^1,$$

with $\underline{A}_i^1(t) = \underline{A}_i^1(t-1)\alpha + A_i^1(t) * (1-\alpha)$.

$C_1 = 50$ and $\alpha = 0.999$, resulting in a running average with a time-scale of 1000 iterations. The activity of layer 0 neurones (O^0) is written as a vector containing either 1s for active neurones and 0s for inactive neurones. The scalar product ($*$) implies simple summation of all inputs weighted by their respective synaptic efficacy (\mathbf{W}).

In the second layer the subthreshold dynamics follows:

$$\Delta A_i^2 = C_2 \max(O_{\text{psp}}^1 * \mathbf{W}_b) / \underline{A}_i^2,$$

with $\underline{A}_i^2(t) = \underline{A}_i^2(t-1)\alpha + A_i^2(t) * (1-\alpha)$.

$C_2 = 200$ and $\alpha = 0.999$. In order to mimic the finite decay of post-synaptic potential, a running average with an exponential decay of ten iterations of the output of layer 1 is taken as the effective input. The synaptic input, however, is not summed as above, but the maximum determined. The choice of max as an integration function has

been discussed in depth by Riesenhuber and Poggio [39].

The local inhibition has a multiplicative effect

$$\Delta A_i = -\beta O_i A_i$$

with $\beta = 0.001$.

The dendritic potential is determined by:

$$\Delta D_i = O_{\text{pre}} * \mathbf{W}_a + C_3 A_i - C_4 D_i.$$

Two factors influence the potential D at the apical dendrite. Activity of the respective pre-synaptic neurones (O_{pre}^2) increases D proportional to the corresponding synaptic weights (\mathbf{W}_a). Following the experimental finding that the threshold for generation of a dendritic calcium spike is lowered when the excitation of the apical dendrite is paired with a somatic action potential [28, 29], the activity A (the somatic variable) also increases D in the case of layer 1 ($C_3 = 1$, $C_3 = 0$ otherwise). D decays with a time constant of 100 time steps ($C_4 = 0.01$). Based on the experimental result that input by a single interneurone may abolish calcium spikes [28], we assume a strong local competition within each module. Only in the neurone with the highest dendritic potential calcium spikes are triggered.

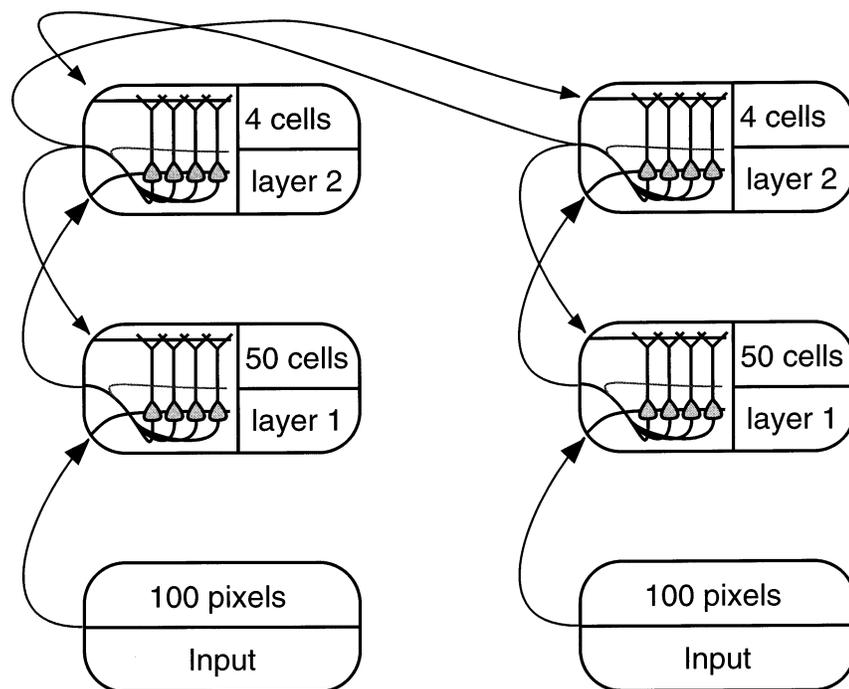


Figure 1. The network consists of two streams with three layers each: input layer, layer 1 and layer 2, each depicted as a box. In the left part of a box, a schematic wiring diagram with three connections to other boxes is shown. Units receive input to the apical and basal dendritic tree (upper and lower connection, respectively). The output via the axon is shown in between (middle connection). Local inhibition is indicated by the lightly shaded line and influences the somatic compartment. The size of each layer is indicated in the right part of the box. The two streams communicated by reciprocal connections onto the apical dendritic tree on the highest level.

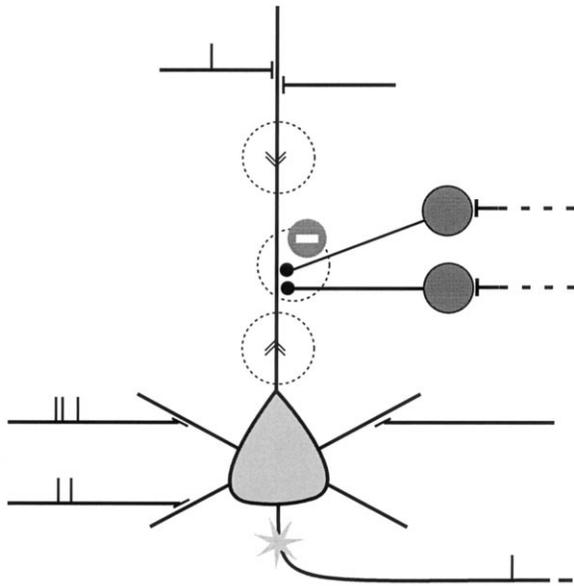


Figure 2. The diagram shows a principal with events influencing synaptic plasticity. An action potential triggered at the axon hillock (star) and propagates anterogradely along the axon (down to the right) and retrogradely invades the dendritic tree (double arrow pointing upwards). Synapses of inhibitory interneurons are located on the proximal dendrite (filled circles) and affect the backpropagating action potential. The backpropagating action potential may reach the distal apical dendrite in interact with local afferent input (a pre-synaptic action potential is indicated by a short vertical line riding on an axon) in the generation of calcium spikes. The calcium spike travels towards the soma (double arrow pointing downward) where it influences synaptic plasticity of all active synapses. See text for details.

Following the experimental results on properties of LTP and LTD summarized above changes of synaptic efficacy in our model depend on pre-synaptic activity and post-synaptic calcium spike activity. In those neurones, where input to the apical dendritic tree triggered calcium spikes all active synapses of apical as well as basal dendritic tree are updated:

$$\Delta W_{k1} = \eta \underline{A}_k^{\text{pre}} - W_{k1} + \varphi (T_c / N_c - 0.5),$$

with $\eta = 0.0001$, \underline{A} is a running average with exponential decay of ten iterations. As in many other neuronal network simulations we use normalization of the weight vectors, the second term, to avoid diverging weights [34]. Actually, there is some experimental support for such a biological mechanism. Blocking neural activity in tissue culture leads to increased miniature excitatory post-synaptic synaptic currents, whereas blocking

GABA mediated inhibition leads to higher firing rates and to smaller synaptic currents [48, 49]. Alternatively a BCM-like learning rule to normalize activity [7] could be combined with the calcium spike effect. Analysing such effects could help understand how calcium spikes and ‘normal’ firing could interact leading to efficient learning. The cells habituate with respect to calcium spikes, the third term, and change all their weights by an amount proportional to the number of iterations since the last calcium spike (T_c , with $\varphi = 0.000002$, and N_c as the total number of units in the layer). This implies that units, which have not learned for a long time increase their input.

Initially all weights are chosen randomly in the interval $(0 \dots 1)$. The system is simulated for 20 000 iterations.

Stimuli used were elongated Gaussians resembling ‘light bars’ as used in physiological experiments. Their luminance had a Gaussian profile orthogonal to the long axis with a length constant of 1 and along the main axis with a length constant of 4. The latter was introduced to reduce boundary effects. They were presented at different orientations and positions. Most importantly is the relation of stimuli presented to the two streams. As in natural images the orientation of contours in small patches is correlated over time, however, the precise position of the contours is not correlated (Weinhäuser, Körding, König, unpubl. data) the orientation of stimuli presented to both streams is chosen to be perfectly correlated and the position of stimuli uncorrelated. For quantitative evaluation position is parameterized on a one-dimensional axis orthogonal to the major axis of the stimulus pattern.

3. Results

As a first step we investigated the dynamics in the neuronal network. Presenting an oriented stimulus leads to activity of appropriately tuned neurones. Due to the two sites of integration, the temporal structure of activity varies. In *figure 3*, spike traces of two neurones in layer 1 are shown, which receive top-down input from layer 2 (upper trace) and which do not (lower trace). Although overall activity is not that different, the grouping of action potentials to bursts in the activity trace of the former neurones is obvious. In view of the learning rule employed, this is a decisive difference. In neurone # 1, the bursts of action potentials induced by the top-down input gate plasticity

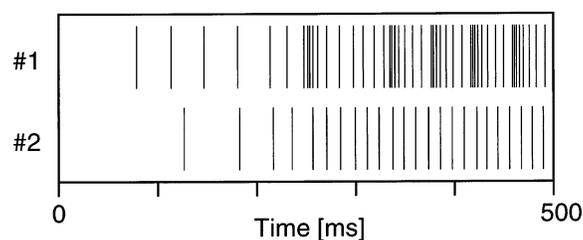


Figure 3. Examples of spike trains showing regular firing and bursting activity.

of all synapses and allow the neurone to learn the presented stimulus. This leads to a difference in the size of representation of a stimulus (many neurones → coarse coding) and the number of neurones adapting to a stimulus (few neurones → soft winner take all).

The network is trained with oriented stimuli presented at different positions. Importantly, stimuli to the two streams have correlated orientation, but uncorrelated positions. After training we investigate response properties of neurones in layer 1. Spike trains for with different orientation and position are shown in *figure 4*. Presenting stimuli at different orientations (*figure 4A*, upper panel) and positions (*figure 4A*, lower panel) induces different levels of activity in the neurone considered. Please note, that the induction of burst occurs at a delay of a few hundred milliseconds. Although the precise value of this delay is dependent on the

parameters of the simulation, a slight sharpening in the tuning of the neurones can be observed. For a further quantification the receptive fields of the neurones are plotted in visual space (*figure 4B*, left) and stimulus space (*figure 4B*, right). In visual space receptive fields are elongated with different orientations and located at different positions. As a consequence, plotted in stimulus space receptive fields are structured along both the orientation and position dimension. This property resembles ‘simple cells’ in primary visual cortex [22]. The union of all receptive fields covers evenly the whole stimulus space (*figure 4C*).

Neurones in the second layer show a qualitatively different behaviour after training. Upon stimulation with oriented bars at different positions their responses are orientation tuned (*figure 5A*, upper panel). Comparable to layer 1 tuning of mean firing rate is somewhat broader than tuning of burst responses. An important difference can be seen when presenting stimuli at different positions. The neurone responds to stimuli presented at any position and the size of the response drops only at the extremes, where the stimuli are located near one corner and few of the units in the input layer are activated. This property is a result of learned convergent connections from layer 1 to layer 2. The linear receptive field, as predicted from the feed-forward synaptic connections is shown in *figure 5B*. The little squares depict receptive fields in stimulus space of layer 1 neurones connected to

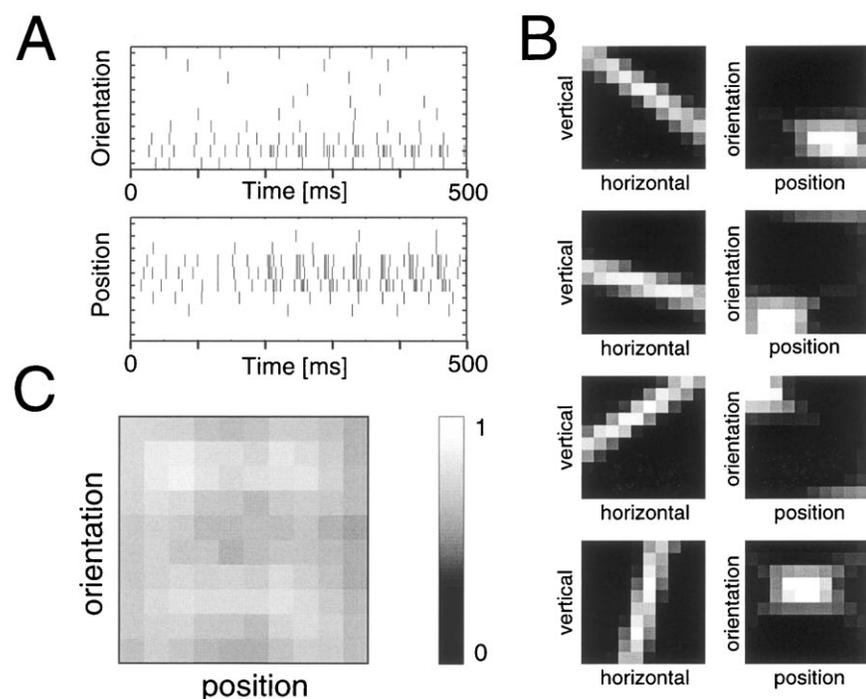


Figure 4. A. Trains of action potentials of a selected neurone in layer 1 for stimuli of different orientations and position. The eight orientations shown cover the whole range of 180 degree available. Position of the stimulus is parameterized on a one-dimensional axis orthogonal to the stimulus. B. Receptive fields of four layer 1 neurones shown in different rows. The left column depicts the receptive field in visual space as results from the synaptic connectivity. In the right column, strength of response to a stimulus is shown in stimulus space. Here all possible combinations of orientations and positions are shown in a two-dimensional diagram. C. Coverage of stimulus space by the union of all neurones in layer 1 is shown grey-level coded.

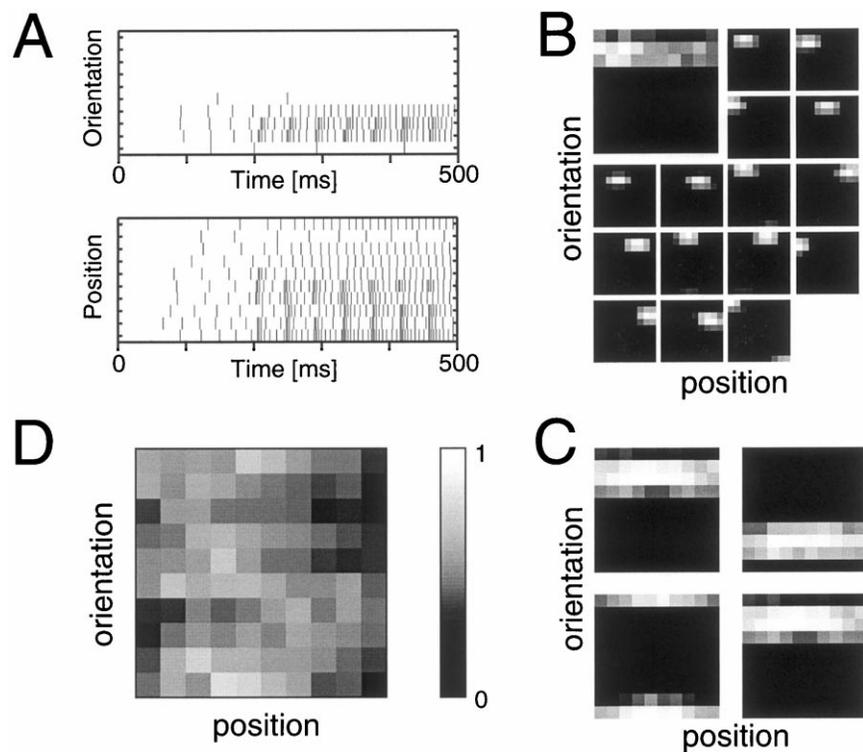


Figure 5. **A.** Trains of action potentials of a selected neurone in layer 2 for stimuli of different orientations and position. The eight orientations shown cover the whole range of 180 degree available. Position of the stimulus is parameterized on a one-dimensional axis orthogonal to the stimulus. **B.** Receptive field of the layer 2 neurone as predicted by the bottom-up afferent shown in stimulus space (big square). Receptive fields of layer 1 neurones strongly connected to the layer 1 neurone in stimulus space are shown in the little squares. Please note that receptive fields of these layer 1 neurones converging on the same layer 2 neurone have similar preferred orientation, but vary in preferred position. The resulting receptive field of the layer 2 neurones is therefore orientation selective and translation invariant. **C.** Effective receptive field of layer 2 neurones as determined from the cellular activity and thus taking the tangential connectivity between neurones in different streams into account. Please note that the response to stimuli at different positions is even more homogeneous as predicted from the feed-forward connectivity as shown in **B.** **D.** Coverage of stimulus space by the union of all layer 2 neurones is shown grey-level coded.

the same layer 2 neurone in sequence of decreasing synaptic efficacies. The receptive field of the latter is shown in stimulus space as well in the large square. As described above, receptive fields of layer 1 neurones are tuned in the orientation and position dimensions. As neurones of layer 1 tuned to similar orientations and dissimilar positions converge onto a neurone in layer 2, the receptive fields of the latter is still orientation tuned but little modulated in the position dimension. The effective receptive fields, when considering tangential interactions between the two streams are actually even smoother. Thus, it has position invariant responses and resembles in its properties complex neurones in primary visual cortex.

The union of receptive fields of layer 2 neurones also covers the complete stimulus space with a slight drop off at extreme positions. Thus, the learning rule leads to a homogeneous representation of the stimulus space combined with translation invariant responses.

4. Discussion

In this work, we demonstrate that properties of a previously proposed learning rule for learning invariant representations in a network of neurones with continuous input/output functions hold up in a more realistic simulation incorporating spiking neurones.

Here we studied the learning on invariant representations using the example of translation invariance as it occurs in the step from simple cells to complex cells in primary visual cortex [22]. In the biological system, the emergence of invariant responses is observed in many instances. In parallel to the increasing sophistication of receptive field properties in higher visual areas, some aspects of visual stimuli are less and less important. Neurones in area V4 show some degree of colour constancy, which is equivalent to invariance with respect to the spectral composition of illumination. Furthermore, in inferotemporal cortex translation,

scaling and rotation invariant responses can be observed [8, 35, 40]. An even more extreme combination of specificity and invariance can be found in premotor cortex. Neurones may respond with high specificity to a stimulus, irrespective of it being heard, seen or felt [17]. Thus, a highly specific response is combined with invariance with respect to modality.

Computing systems obtain invariant properties by several means [3]. First, appropriate preprocessing can supply a neuronal network with invariant input data. Another option is to build the invariance into the structure of the system so that invariances gradually increase. Finally, a system may learn invariances from the presented stimuli following principles of supervised or unsupervised learning. The former need labelled training data, and actually for large networks, a lot of these. With such data at hand, it is possible to learn invariant recognition training the network with a variant of the backpropagation algorithm (e.g. [20]). To alleviate the problem of getting enough training data, applying the desired invariance operators to individual examples can enlarge the training set. However, then we are back with an a priori specification of the invariance operation. Furthermore, these approaches do not supply a convincing explanation of the development of invariant response properties in the biological system. Following this observation the investigation of unsupervised learning of invariances became a research topic in itself.

To explain such properties Becker and Hinton [5] proposed the Imax algorithm. In this learning algorithm, the mutual information is maximized between output units which receive separate inputs. Thus, this algorithm can detect features in the input which are coherent across the eyes [5] space, time or modality [4, 12]. However, the algorithm requires backpropagation of derivatives from output units to hidden units, and the storage of several variables on each unit. A related approach for learning with multi-unit local processors with multivariate binary outputs has been proposed by Phillips et al. [37] (see also [24]). It is able to extract a number of coherent features. These learning rules are similar to the scheme described here as two separate summations occur. One of these defines the receptive fields and the other the context fields [36]. They share several properties with the basal and apical dendrite in our studies respectively (see also [44]): The context gates learning and is simultaneously used for improving processing of the receptive field input [37].

Their learning rule contains a threshold where the sign of learning changes; it depends on the conditional probabilities with regard to the context and the activity of the other output units. Thus, this signal is local to the cell but not to the synapse. The learning mechanism proposed here can lend a straightforward biological interpretation to such approaches. Although being closely related with regard to the learning goal, this study also shows a major computational difference. The context field represents the covariance matrix of post-synaptic activity with contextual input. Thus the number of weights which needs to be stored and updated is smaller, increasing learning speed at the cost of a larger statistical error.

4.1. Simplifications

Despite this increased physiological realism in this simulation, several aspects are simplified. First, and perhaps most importantly, the detailed temporal dynamics of activity is not taken into account. Indeed, it has been proposed previously that the sign of change of synaptic efficacy depends on the relative sequence of pre and post-synaptic action potentials [16]. This hypothesis is supported by recent in vitro experiments [31]. Such a dependence would lead to an automatic normalization of total afferent signals [25, 42]. Indeed, in vivo experiments provide evidence that optimally activated neurones tend to fire prior to suboptimally activated neurones on a millisecond time scale [26]. Thus, differences in firing rates may be mapped onto the relative timing of action potentials, justifying simplifications done in previous modelling studies. Furthermore, in the present simulation feed-forward input follows a Poisson statistics and its temporal dynamics contains no information. The interesting question of how tangential interactions leading to a synchronization of neuronal activity interact with the proposed learning rule has to be left for future work [30]. Second, the complex non-linear dendritic properties [23, 33] are reduced to a threshold mechanism triggering dendritic calcium spikes. The sole justification for this approach we can present is the argument of Occam's razor. In view of the large number of parameters needed for a detailed compartmental simulation, and our lack of knowledge of these, a more detailed realistic model would actually include many more estimates not solidly based on physiological and anatomical results and thus obscure the issues investigated. Third, along similar lines in our simulation the current flow from the

apical dendrite to the soma is neglected and calcium spikes have an effect on learning but not on firing rates. Introducing the effect on activity could result in information about the context not only being used for learning but also for enhancing the signal itself [37]. The effects of calcium spikes on the dynamics of the network were analysed in a previous work [44]. We could demonstrate that top-down information leads to enhanced processing of bottom-up signals. We conjecture that both effects are compatible and act simultaneously, the top-down signals improving signal processing on a short time scale and gating learning on a long time scale.

4.2. Experimental predictions

The simulations described here imply some experimental predictions.

First, cooling or lesions of higher areas should significantly reduce the frequency of bursting activity in lower areas by decreasing the amount of top-down mediated signals available. This would reduce synaptic plasticity of all synapses as could be assessed by paradigms as monocular deprivation or selective rearing influencing ocular dominance or orientation tuning of cortical neurones [14, 15, 18]. Furthermore, as these top-down mediated signals act on a spatial scale given by the receptive field size of the neurones in a higher area the loss of top-down signals should appear as a reduction of non-classical receptive field effects. In particular, this effect should be most pronounced for stimuli matching receptive field properties of neurones in the higher cortical areas [44].

Second, patching the trunk of the apical dendrite of pyramidal cortical cells would make it possible to block the induction of calcium spikes in the apical dendrites without inflicting strong changes upon the potential at the soma (except for the lack of calcium spikes). Simultaneously patching the soma of that cell in deeper layers and a pre-synaptic cell would make it possible to assess synaptic plasticity. The removal of calcium spikes should lead to strongly reduced LTP not only for synapses at the apical dendrite, but as well for synapses at the basal dendrites.

Third, along the same lines the effects of additional calcium spikes can be investigated. Patching an apical dendrite in layer I *in vivo* would make it possible to artificially induce calcium spikes in analogy to Larkum et al. [28]. Following the predictions of the proposed model it should be possible to investigate unsupervised learning *in vivo*

similar to Debanne et al. [11]. The cell should learn to represent those stimuli that are paired with excitation of the apical dendrite. The plasticity resulting from calcium spikes could thus directly be compared with the plasticity resulting from spikes not associated with calcium spikes. Should some normal spikes have less influence on plasticity than a calcium spike resulting in a volley of the same number of spikes?

Fourth, direct interference with the burst generating mechanism should have a particular strong effect on top-down mediated signals and also affect receptive field properties in the described way. Slice recordings show that the backpropagation of action potentials into the apical dendrite depends on muscarinic input, inhibitory input and the firing rate [9, 47]. Furthermore it has been demonstrated that the triggering of bursts by correlated synaptic input at the apical dendrite and back-propagating action potentials is highly sensitive to inhibitory input [28]. These findings suggest that altering the activity of inhibitory or modulatory systems could be a suitable way to interfere with the burst generating mechanism.

These experiments are demanding, but within the reach of state of the art techniques.

Acknowledgements

This work has been supported by the Swiss National Science Foundation (P.K. 31-51059.97) and the Boehringer Ingelheim Fond (K.P.K.).

References

- [1] Amitai Y., Friedman A., Connors B.W., Gutnick M.J., Regenerative activity in apical dendrites of pyramidal cells in neocortex, *Cereb. Cortex* 3 (1993) 26–38.
- [2] Artola A., Bröcher S., Singer W., Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex, *Nature* 347 (1990) 69–72.
- [3] Barnard E., Casasent D.P., Invariance and Neural Nets, *IEEE Trans. Neural Networks* 2 (1991) 489–508.
- [4] Becker S., Models of cortical self-organisation, *Network: Comput. Neural Syst.* 7 (1996) 7–31.
- [5] Becker S., Hinton G.E., Self-organizing neural network that discovers surfaces in random-dot stereograms, *Nature* 355 (1992) 161–163.
- [6] Bernander O., Koch C., Douglas R.J., Amplification and linearization of distal synaptic input to cortical pyramidal cells, *J. Neurophysiol.* 72 (1994) 2743–2753.

- [7] Bienenstock E., Cooper L.N., Munro P.W., Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex, *J. Neurosci.* 2 (1982) 32–48.
- [8] Booth M.C., Rolls E.T., View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex, *Cereb. Cortex* 8 (1998) 510–523.
- [9] Buzsáki G., Kandel A., Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat, *J. Neurophysiol.* 79 (1998) 1587–1591.
- [10] Cauller L.J., Connors B.W., Synaptic physiology of horizontal afferents to layer I in slices of rat SI neocortex, *J. Neurosci.* 14 (1994) 751–762.
- [11] Debanne D., Shulz D.E., Fregnac Y., Activity-dependent regulation of ‘on’ and ‘off’ responses in cat visual cortical receptive fields, *J. Physiol. (Lond.)* 508 (1998) 523–548.
- [12] de Sa V.R., Ballard D.H., Category learning through multimodality sensing, *Neur. Comp.* 10 (1998) 1097–1117.
- [13] Dudek S.M., Bear M.F., Homosynaptic long term depression in area CA1 of hippocampus and the effects on NMDA receptor blockade, *Proc. Natl. Acad. Sci. USA* 89 (1991) 4363–4367.
- [14] Fregnac Y., Shulz D.E., Activity-dependent regulation of receptive field properties of cat area 17 by supervised Hebbian learning, *J. Neurobiol.* 41 (1999) 69–82.
- [15] Friedlander M.J., Fregnac Y., Burke J.P., Temporal covariance of postsynaptic membrane potential and synaptic input-role in synaptic efficacy in visual cortex, *Prog. Brain Res.* 95 (1993) 207–223.
- [16] Gerstner W., Kempter R., van Hemmen J.L., Wagner H., A neuronal learning rule for sub-millisecond temporal coding, *Nature* 383 (1996) 76–78.
- [17] Graziano M.S., Gross C.G., Spatial maps for the control of movement, *Curr. Opin. Neurobiol.* 8 (1998) 195–201.
- [18] Greul J.M., Luhmann H.J., Singer W., Pharmacological induction of use-dependent receptive field modifications in visual cortex, *Science* 242 (1988) 74–77.
- [19] Helmchen F., Svoboda K., Denk W., Tank D.W., In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons, *Nat. Neurosci.* 2 (1999) 989–996.
- [20] Hinton G.E., Learning translation invariant recognition in a massively parallel network, in: Goos G., Hartmanis J. (Eds.), *PARLE: Parallel Architectures and Languages Europe*, Lecture Notes in Computer Science, Springer-Verlag, Berlin, 1987, pp. 1–13.
- [21] Hirsch J.A., Alonso J.M., Reid R.C., Visually evoked calcium action potentials in cat striate cortex, *Nature* 378 (1995) 612–616.
- [22] Hubel D.H., Wiesel T.N., Receptive Fields, binocular interaction and functional architecture in the cat’s visual cortex, *J. Physiol.* 160 (1962) 106–154.
- [23] Johnston D., Magee J.C., Colbert C.M., Christie B.R., Active properties of neuronal dendrites, *Annu. Rev. Neurosci.* 19 (1996) 165–186.
- [24] Kay J., Floreano D., Phillips W.A., Contextually guided unsupervised learning using local multivariate binary processors, *Neural Networks* 11 (1998) 117–140.
- [25] Kempter R., Gerstner W., van Hemmen J.L., Hebbian learning and spiking neurons, *Phys. Rev. E* 59 (1999) 4498–4514.
- [26] König P., Engel A.K., Roelfsema P.R., Singer W., How precise is neuronal synchronization?, *Neural Comp.* 7 (1995) 469–485.
- [27] Körding K.P., König P., Learning with two sites of synaptic integration, *Network: Comput. Neural Syst.* 11 (2000a) 1–15.
- [28] Larkum M.E., Zhu J.J., Sakmann B., A new cellular mechanism for coupling inputs arriving at different cortical layers, *Nature* 398 (1999) 338–341.
- [29] Larkum M.E., Kaiser K.M., Sakmann B., Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials, *Proc. Natl. Acad. Sci. USA* 96 (1999) 14600–14604.
- [30] Lumer E.D., A neural model of binocular integration and rivalry based on the coordination of action-potential timing in primary visual cortex, *Cereb. Cortex* 8 (1998) 553–561.
- [31] Markram H., Lübke J., Frotscher M., Sakmann B., Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs, *Science* 275 (1997) 213–215.
- [32] Martin S.J., Grimwood P.D., Morris R.G., Synaptic plasticity and memory: an evaluation of the hypothesis, *Annu. Rev. Neurosci.* 23 (2000) 649–711.
- [33] Mel B.W., Synaptic integration in an excitable dendritic tree, *J. Neurophysiol.* 70 (1993) 1086–1101.
- [34] Miller K.D., MacKay D.J.C., The role of constraints in Hebbian learning, *Neural Comput.* 6 (1994) 100–126.
- [35] Perrett D.I., Oram M.W., Harries M.H., Bevan R., Hietanen J.K., Venson P.J., Thomas S., Viewer-centred and object-centred coding of heads in the macaque temporal cortex, *Exp. Brain. Res.* 86 (1991) 159–173.
- [36] Phillips W.A., Singer W., In search of common foundations for cortical computation, *Behav. Brain Sci.* 20 (1997) 657–683.
- [37] Phillips W.A., Kay J., Smyth D., The discovery of structure by multi-stream networks of local processors with contextual guidance, *Network: Comput. Neural Syst.* 6 (1995) 225–246.
- [38] Pike F.G., Meredith R.M., Olding A.W.A., Paulsen O., Postsynaptic bursting is essential for ‘Hebbian’ induction of associative long-term potentiation at excitatory synapses in rat hippocampus, *J. Phys.* 518 (2) (1999) 571–576.
- [39] Riesenhuber M., Poggio T., Hierarchical models of object recognition in cortex, *Nat. Neurosci.* 2 (1999) 1019–1025.
- [40] Rolls E., Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas, *Phil. Trans. R. Soc. Lond. B* 335 (1992) 11–21.
- [41] Salin P.A., Bullier J., Corticocortical connections in the visual system: Structure and function, *Physiol. Rev.* 75 (1995) 107–154.

- [42] Sanchez-Montanes M.A., Verschure P.F., König P., Local and global gating of synaptic plasticity, *Neural Comput.* 12 (2000) 519–529.
- [43] Schiller J., Schiller Y., Stuart G., Sakmann B., Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons, *J. Physiol. (Lond.)* 505 (1997) 605–616.
- [44] Siegel M., Körding K.P., König P., Integrating bottom-up and top-down sensory processing by somato-dendritic interactions, *J. Comput. Neurosci.* 8 (2000) 161–173.
- [45] Stuart G.J., Sakmann B., Active propagation of somatic action potentials into neocortical pyramidal cell dendrites, *Nature* 367 (1994) 69–72.
- [46] Stuart G.J., Schiller J., Sakmann B., Action potential initiation and propagation in rat neocortical pyramidal neurons, *J. Physiol. (Lond.)* 505 (1997) 617–632.
- [47] Tsubokawa H., Ross W.N., IPSPs modulate spike back-propagation and associated $[Ca^{2+}]_i$ changes in the dendrites of hippocampal CA1 pyramidal neurons, *J. Neurophysiol.* 76 (1996) 2896–2906.
- [48] Turrigiano G.G., Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same, *Trends Neurosci.* 22 (1999) 221–227.
- [49] Turrigiano G.G., Leslie K.R., Desai N.S., Rutherford L.C., Nelson S.B., Activity-dependent scaling of quantal amplitude in neocortical, *Nature* 391 (1998) 892–896.
- [50] Zeki S., Shipp S., The functional logic of cortical connections, *Nature* 335 (1988) 311–317.