

Learning with two sites of synaptic integration

Konrad P Körding[†] and Peter König

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

E-mail: koerding@ini.phys.ethz.ch

Received 20 May 1999, in final form 24 September 1999

Abstract. Since the classical work of D O Hebb 1949 *The Organization of Behaviour* (New York: Wiley) it is assumed that synaptic plasticity solely depends on the activity of the pre- and the post-synaptic cells. Synapses influence the plasticity of other synapses exclusively via the post-synaptic activity. This confounds effects on synaptic plasticity and neuronal activation and, thus, makes it difficult to implement networks which optimize global measures of performance. Exploring solutions to this problem, inspired by recent research on the properties of apical dendrites, we examine a network of neurons with two sites of synaptic integration. These communicate in such a way that one set of synapses mainly influences the neurons' activity; the other set gates synaptic plasticity. Analysing the system with a constant set of parameters reveals: (1) the afferents that gate plasticity act as supervisors, individual to every cell. (2) While the neurons acquire specific receptive fields the net activity remains constant for different stimuli. This ensures that all stimuli are represented and, thus, contributes to information maximization. (3) Mechanisms for maximization of coherent information can easily be implemented. Neurons with non-overlapping receptive fields learn to fire correlated and preferentially transmit information that is correlated over space. (4) We demonstrate how a new measure of performance can be implemented: cells learn to represent only the part of the input that is relevant to the processing at higher stages. This criterion is termed 'relevant infomax'.

1. Introduction

The brain has been a source of inspiration for the construction of technical systems for centuries. Its architecture has motivated the development of distributed computational systems. The idea of parallel distributed processing (PDP)—the notion that intelligent behaviour emerges from the interaction of a large number of simple processing units—has been attractive for many researchers trying to understand the basics of human perception, memory, language, and thought (Rumelhart and McClelland 1986). Input stimuli are stored in the connection strengths between units. The fitness of such a system can often be expressed by a goal function (e.g. per cent of correctly classified input stimuli). Learning is then a matter of finding those connection strengths that maximize the expected value of the goal function. Usually a very simple mechanism is assumed to modulate the connection strength, which is based on information locally available at the synapse. As a consequence in connectionist systems synaptic plasticity usually depends only on the average pre- and post-synaptic activity (e.g. Hebb 1949, Stent 1973, Sejnowski 1977, Bienenstock *et al* 1982).

For the analysis of the relation of global goal functions and local learning rules two approaches have been applied. First, global goal functions may be obtained via integration of a local learning rule (e.g. Intrator and Cooper 1992). Second, learning rules can be derived

[†] Author to whom correspondence should be addressed.

via differentiation of a global goal function. This approach has been applied to complex goal functions maximizing transmission of coherent information in parallel channels (Phillips *et al* 1995, Kay *et al* 1998). The resulting learning rules rest on information that is locally available at a synapse (activity of pre- and post-synaptic cells) as well as on an additional signal conveying information about distinct parts of the network. These signals are integrated separately and have different effects on processing and on learning (cf Becker 1996, Kay *et al* 1998). Without a known basis for such an additional signal these learning rules have often been considered biologically implausible (but see Phillips and Singer 1997).

Here we investigate a learning rule which addresses this conflict. It is inspired by a remarkable asymmetry of pyramidal neurons in cerebral cortex: their basal and the apical dendrites have distinguished morphologies and functional properties. Furthermore, there is some evidence that these dendrites give rise to two different sites of synaptic integration, the soma and the apical dendrite (Bernander *et al* 1994, Larkum *et al* 1999). We assume that the communication between these sites is asymmetric: synaptic inputs at one site mainly influence the activity of the post-synaptic neuron. Inputs at the other site gate synaptic plasticity via a threshold process. This separation makes it possible to have global information locally available at the synapses and to avoid confounding effects of processing and learning. We use the terms ‘apical dendrite’ and ‘soma’ for the two sites of integration (figure 1(A)). However, we do not propose that this is the only possible implementation and other anatomical configurations may be thought of. In this paper we show how several attractive goal functions can be easily implemented.

2. Methods

We simulate a rate coding neural network where a unit’s output is a real number representing the average firing rate. Every cell is described by two main variables, corresponding to the two sites of integration (see figure 1(A)): A is referred to as the activity of the neuron and D represents the average potential at the apical dendrite.

To calculate the activity $A_i^{(j)}$ of the cell i in layer j we multiply the weight vector \mathbf{W} with the vector of the pre-synaptic activity: $A_i^{(j)} = \mathbf{A}^{(j-1)} \mathbf{W}_i$.

Three factors influence the potential D at the apical dendrite: the inhibitory signal I from local interneurons that lowers D ; the excitatory activation E that increases D ; the cell’s activity A that increases its own D . These result in the following formula: $D = A + E - I$. Thus the somatic activity increases the dendritic potential.

Whenever the potential D at the ‘apical dendrite’ integration site exceeds a threshold Θ a ‘learning event’ is induced. It leads to long-term potentiation (LTP) at active synapses and the weights W_{ik} from neuron k in the input layer to neuron i in the learning layer are changed: $\Delta W_{ik} = \text{Rate}^* A_k^{(j-1)}$. Thus, the potential at the second integration site influences somatic activity only indirectly via its effect on synaptic plasticity. This learning rule has not yet the property that the synaptic weights stay finite. We therefore normalized the learned weight vectors to unit length. The threshold $\Theta_{t,i}$ is regulated so that cells with rare ‘learning events’ increase their triggering probability and vice versa: $\Theta_{t+1,i} = \Theta_{t,i} - \tau_\Theta$ without and $\Theta_{t+1,i} = \Theta_{t,i} + \tau_\Theta^* \Delta\Theta$ with those events. The simulations started with $\Theta = 0$.

In simulations where connections to the apical dendrite are plastic, they are modified according to Hebbian learning with exponential decay of time constant 1000 iterations. In this case E is the product of the activity \mathbf{A} on the higher layer with the weight vector of the respective projections.

The inhibition I is the mean of the layer activity in all simulations. In the beginning all

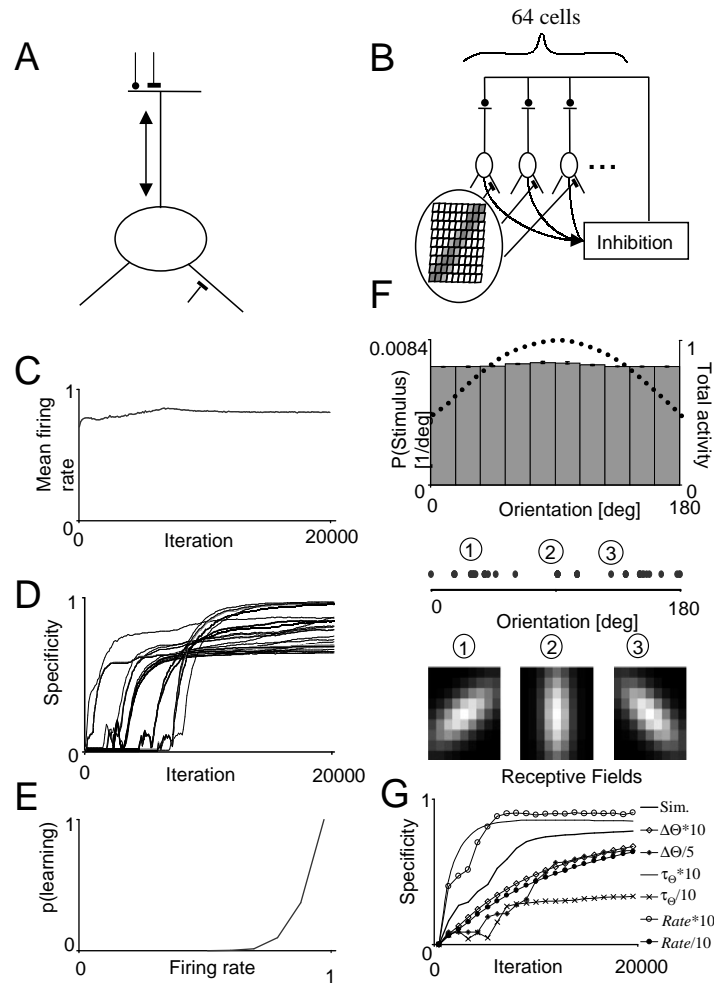


Figure 1. Learning, specificity and constant network activity. (A) The simplified cell model in our simulations is shown. The cell consists of two different sites of integration, the cell body with the basal dendrites is one compact site of integration; the apical tuft is regarded as the second site of integration. They are connected by the electrotonically long apical dendrite where action potentials propagate from the soma to the apical dendrite and calcium spikes propagate from the apical dendrite to the soma. (B) The network layout is sketched. Cells on the second layer are excited from the input layer. The weights from input to second layer are learned. All cells on the learning layer send lateral inhibition to each other's apical dendrites. (C) The total network activity averaged over 100 iterations and over all neurons is shown over time. (D) The specificity index of all neurons is plotted as a function of the iteration. (E) The probability of 'learning events' is plotted as a function of the cells firing rate. 'Learning events' are triggered if the cell activity is high. (F) The probability density for a stimulus with a given orientation is shown by the dotted curve in the upper graph. The ordinate shows the probability density in units of $1/\text{deg}$. The abscissa shows the orientation of the given bar. The bar graph shows the mean of the network activity acquired over ten simulation runs. The error bars represent the standard deviation of this activity. The (\cdot) symbols in the middle trace show the orientations preferences of the cells for one run. Please note that although the net activity is almost constant the distribution of neurons in orientation space is not uniform. Typical receptive fields are shown grey scale coded in the lower graph. (G) Systematic explorations of parameter space were performed. The mean of the specificity index is plotted against the iteration. Each line is taken for one set of parameters. The thick solid curve shows the results of the original simulation. The other curves are taken for one of the parameters $\Delta\Theta$, τ_{Θ} and $Rate$ changed by a factor as indicated in the legend.

weights are chosen to be uniform with 10% noise. The parameters chosen are: $\Delta\Theta = 10$; $\tau_{\Theta} = 0.00002$; $Rate = 0.005$; The system is simulated for 20 000 (100 000 in the last simulation) iterations, so that the system reaches a stable state in all cases.

Stimuli in some simulations resemble ‘bars’ as used in physiologic 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.

To analyse the results in some simulations a specificity index is calculated for every neuron: for each part of the receptive field the angle of the distance vector is doubled and added up with a weight according to the efficacy of the respective synapse. This yields two numbers, the specificity (corresponding to the length of that vector) and the orientation (half the angle of the vector). The specificity is normalized so that the stimulus vector with the highest specificity has the specificity of 1. This means that even if receptive fields directly represent input stimuli their specificity cannot exceed 1.

3. Results

The properties of the learning mechanism, where one site of integration defines the firing rates and the other one gates learning, is analysed in different network topologies with an identical set of parameters. Signals with increasing complexity, given to the apical dendrite, lead to increasingly interesting learning behaviours.

3.1. Specificity and constant network activity

In this simulation we only consider local inhibitory action and no specific excitation on the apical dendrites. The potential at the apical dendrite is thus negatively proportional to the network activity and proportional to the cell’s own activity.

A network (figure 1(B)) is trained with bars on the 9×9 input-plane with their orientation randomly chosen out of $[0, \pi)$ every iteration. The probability density of the stimuli is non-uniform (figure 1(F), dotted curve). The ratio of probabilities for the most likely and the least likely orientation was 2.

Analysing the behaviour of the network, we observed that the weight normalization together with the learning mechanism leads to almost constant net activity (figure 1(C)). Its almost constant high average value of about 0.8 (where 1 is the maximal possible response) shows that neurons are rather broadly tuned.

To analyse how specific the responses are, a specificity index is calculated for all neurons (see methods section). Initially all neurons have unspecific response properties (figure 1(D)). As learning proceeds, more and more neurons develop increasingly specific response properties. However, this occurs not as a continuous process. Instead, after an initial delay they quickly increase their specificity for some iterations and converge to their final value. Interestingly not all neurons do this at the same time. This is because of the following: if a neuron gets specific it strongly increases its activity for some stimuli. This leads to an even higher probability for ‘learning events’. The increased network activity prevents the other neurons from learning, until they have adjusted their learning thresholds to adapt to the increased inhibition. In the end all neurons are highly specific.

It is interesting to ask why neurons acquire specific responses: In analogy to Hebbian learning the neurons acquire specificity because learning one stimulus increases the probability of ‘learning events’ for similar stimuli. This happens because their probability of ‘learning events’ rises with increasing activity (figure 1(E)).

Because the potential at the apical dendrite is inversely proportional to the mean activity this mean activity gets almost identical for all stimuli (figure 1(F), bars). This happens although the firing probabilities were non-uniform. This shows the effectiveness of the proposed mechanism in ensuring that the net activity is constant. A slight increase in net activity can just be seen for frequent stimuli which is because τ_{\ominus} has a finite value and thus neurons tuned to frequent stimuli on the average need to wait a shorter period after the threshold gets permissive.

Figure 1(F) (middle trace) shows a typical distribution of preferred orientations for one run of the simulation. An interesting aspect can be observed: the neurons do not acquire a perfectly uniform distribution in orientation space but are partially clustered. This phenomenon could only be reduced if one introduced a threshold for activation of each neuron or super-linear inhibition (data not shown). The threshold has the effect that many different orientation preferences are necessary to achieve uniform network activity. Super-linear inhibition has the effect that clusters in orientation space tend to get unstable. Nevertheless the network has a homogeneous distribution of network activities although the neuronal distribution is clustered in stimulus space. The lower part of figure 1(F) shows typical receptive fields. We performed systematic variations of the involved parameters (figure 1(G)) to analyse the speed of convergence. Higher *Rate* leads to faster but non-uniform learning, with more neurons getting very similar receptive fields with higher specificity. Higher τ_{\ominus} leads to faster convergence because the time between different neurons starting to learn is smaller then. And finally at higher $\Delta\Theta$ neurons learn less often and thus more iterations are needed to reach convergence. At lower $\Delta\Theta$ neurons almost always learn and thus remain rather unspecific. A similar effect, where the variance of the network activity over all stimuli is minimized, could be reached with lateral inhibition on the somas of the cells. But in this case fewer neurons would fire impairing the discernability of stimuli. For recognizing stimuli the activity of many neurons gives most information. To ensure that the net activity is constant one wants only few neurons to learn a stimulus. With just one site of integration this problem is ill posed. This trade-off between learning and representing will be discussed below. There are other learning rules that also use a gating mechanism. In hard or soft competitive learning, ART as well as in self-organizing feature maps, a separate signal determines that only one or a few cells should learn. This signal is not transmitted via the post-synaptic activity and is difficult to be assigned a biological interpretation (but see Körding and König 2000). With the proposed mechanism similar effects can be realized as with the discussed learning rules but with a straightforward biological interpretation.

3.2. Pairing with teacher signal

In the preceding simulation all cells shared the inhibition I at the apical dendrite. Only their self-excitation A leads to different learning behaviour for each cell. As the next step of complexity we analyse the same system with an additional individual excitation at the apical dendrite (figure 2(A)). Depending on the orientation of the shown bar different sets of cells receive excitation E on the apical dendrite. Figures 2(B) and (C) show that the cells acquire receptive fields so that their firing is correlated with the respective teacher signal. The same effects shown in the preceding simulation are still active in this simulation. The total net activity is almost constant and the neurons acquire a high degree of specificity (data not shown). Summarizing, in this simulation the potential at the apical dendrite acts as a supervising signal defining that those stimuli paired with excitation to the apical dendrite are to be learned.

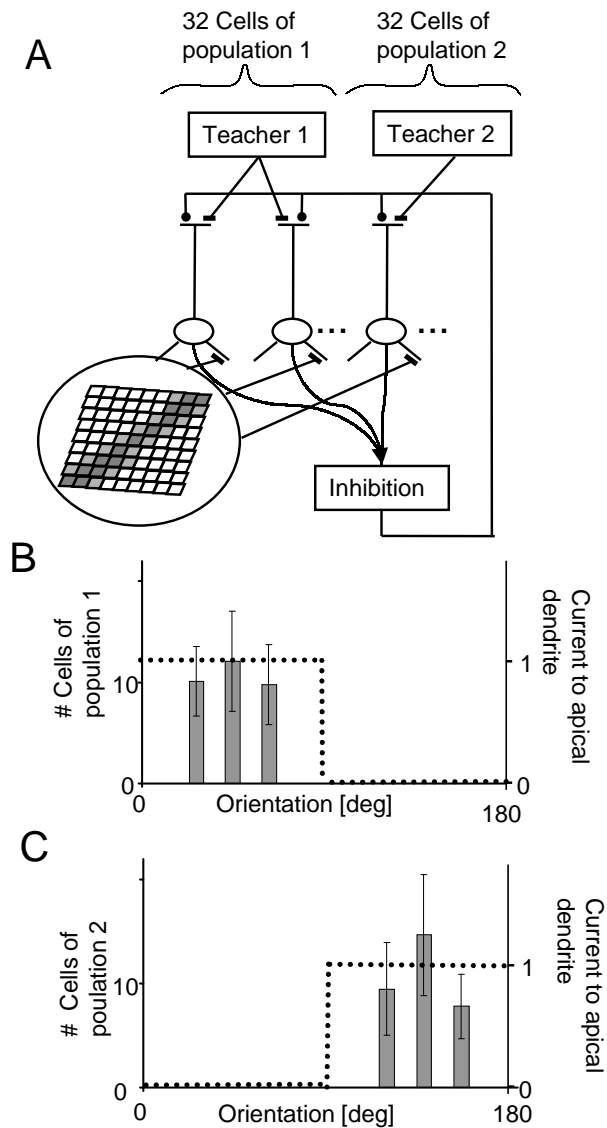


Figure 2. Pairing with teacher signal. (A) The network as shown in figure 1 is examined with additional excitatory input to the apical dendrite. This input acts as a teacher signal and depends on the stimulus. (B) The current is given onto the apical dendrites of some neurons (#1–32) only when the angle of the stimulus bar is between 90 and 180°. The dotted trace shows this current as a function of the orientation of the given bar. The bar-plot shows the distribution of orientation tuning of the first population of neurons (#1–32). In this case the neurons represent those stimuli that are paired with current at the apical dendrite. (C) This is the same plot as (B) but for the other set of neurons.

3.3. Coherent infomax

The two sites of integration make it possible that the cells not only learn when the firing rate is high but also can learn arbitrary stimuli given by the teacher. That can be exploited for maximizing the coherence between streams, a global goal which has been used to define learning rules (Becker 1996, Phillips *et al* 1995). We implement such a mechanism where the signal at the apical dendrite is set by the activity of another neuron.

We observe effects resembling this principle in the toy system similar to the one used by Phillips *et al* (1995) and shown in figure 3(A). The network consists of two separate streams. Both streams contain an input layer of 4×4 neurons and one output neuron. The output neurons of both streams project to the other one's apical dendrite. The self-excitation is turned off for this simulation, implying that there is no current flow from the soma to the apical dendrite. The inputs can be interpreted as coming from different modalities or non-overlapping parts of the same modality. For the present purpose we interpret the inputs as non-overlapping parts of the visual field. Although they do not share afferents their input can be correlated. The stimuli and their probability distribution are shown in figure 3(B). Vertical bars occurred with a probability of 0.7 with contrast not correlated across streams, and horizontal bars occurred with a probability of 0.3 with contrast correlated across streams.

In the beginning, all stimuli result in comparable firing rates (figure 3(C)). In all ten simulations performed the system learned one of the horizontal stimuli, which are correlated across the streams. This is even more remarkable as the vertical bars were by far more frequent. To analyse why the network learns the correlated stimulus we plot the difference between the potential at the apical dendrite and the threshold as a function of the stimulus and the iteration in figure 3(D). If the right neuron 'learns' a vertical pattern, this increases D of the left neuron for vertical patterns of both contrasts. Since these patterns average to zero, 'learning' at the left neuron does not favour learning of that particular pattern by the right neuron. Instead of a positive feedback, a kind of random walk is observed (seen in figure 3(D), lower trace). At some time a pair of horizontal stimuli, corresponding to correlated contrasts, leads to learning events increasing D for that stimulus for both neurons. This represents a positive feedback, the effect strengthens itself and that correlated stimulus 'wins'. Spontaneous symmetry breaking occurs when small random differences in the initial activity between both horizontal patterns progressively get larger until one wins. The pattern of weights then approaches that stimulus. Figure 3(E) shows the receptive fields after convergence. Horizontal bars are learned by the network (corresponding to correlated activity).

In contrast to this, learning with a Bienenstock–Cooper–Munro (BCM) learning rule (Bienenstock *et al* 1982), where a threshold separating regions of LTP and long-term depression is regulated by the average post-synaptic activity, results in the network learning vertical bars, the most common stimuli. Normalized Hebbian learning leads to almost flat receptive fields. Thus this example shows that exploiting the properties of the additional 'learning signal' it is possible to implement complex learning goals in a rather straightforward way. In the study of Phillips *et al* (1995) similar effects were reached with a learning rule derived from information theory. In summary, in this simulation each cell is the other one's supervisor, thus they learned correlated stimuli.

3.4. Relevant infomax

In the coherent infomax simulation the second layers consisted of only one neuron each. The mechanism makes it possible that those neurons become responsive to correlated activity. In a system of two areas, each consisting of several neurons, with strong inter-area connectivity,

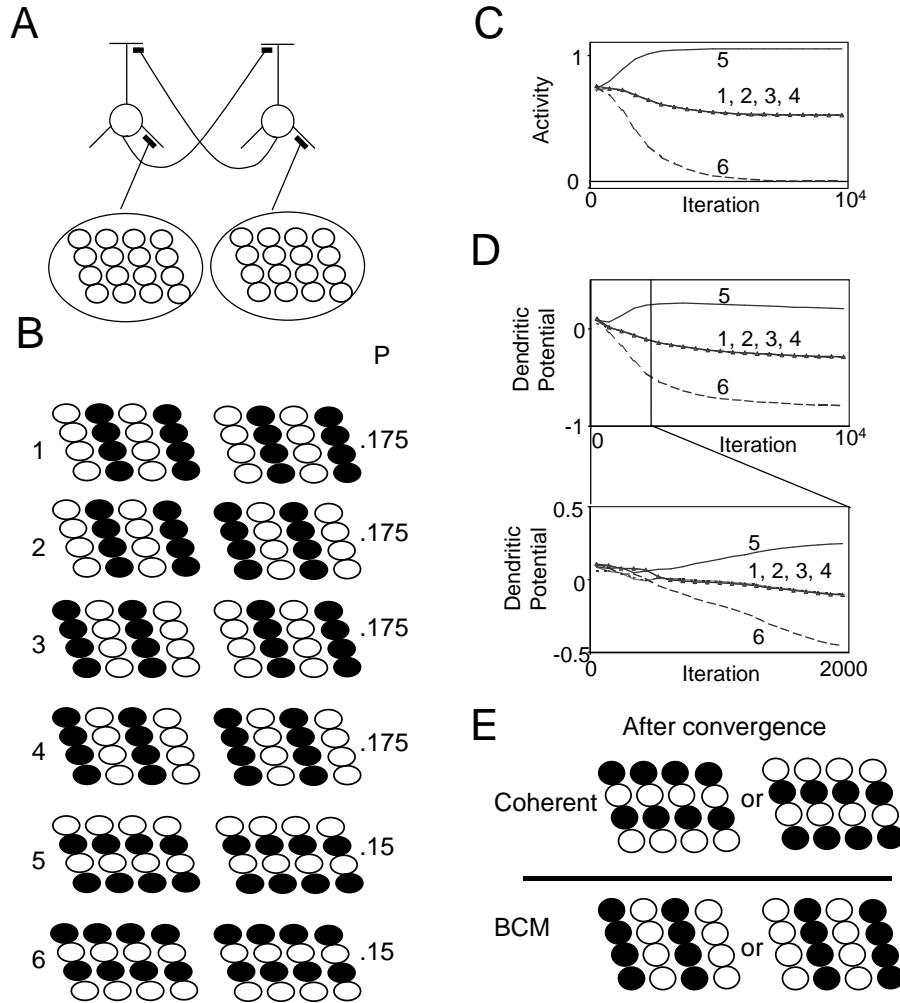


Figure 3. Coherent infomax. (A) The layout of the examined network is shown. Both streams have an input layer consisting of 4×4 neurons. Each of these cells excites one neuron on the second layer. The neurons in both streams are reciprocally coupled to their apical dendrites. (B) The six stimuli are presented in a pseudo random sequence with probabilities as shown. Please note that vertical stimuli are more likely by a factor of $\frac{7}{3}$ compared with the horizontal stimuli. The contrast of the vertical bars is not correlated across streams. For horizontal bars the contrast of the bars is correlated across streams. The numbers on the left of the figure show the number of the stimulus. (C) The response of the left neuron on the second layer to each of the stimuli is quantified. The average activity of the neuron for every stimulus, averaged over 100 stimulus presentations, is plotted over the iteration. (D) To analyse the learning behaviour further the difference of the dendritic potential to the threshold (averaged over 100 iterations for each of the six stimuli) is plotted as a function of the iteration. Each linestyle stands for a different stimulus as indicated by the stimulus number near to the trace. The lower trace shows a magnification of the time during which the potentials at the apical dendrites separate. If the displayed value is positive learning events are induced. In the beginning the cells perform a type of random walk. But whereas the non-correlated stimuli do not have an ‘advantage’ of learning the correlated stimuli have and thus the network learns them in all cases. (E) The results of the simulations are summarized in the upper part. Receptive fields are shown. In the lower part the results of the simulation with BCM learning are shown. The preferred stimulus was the same in each of ten trials (except for the contrast).

all neurons would have the tendency to learn the same ‘thing’. This can be useful and was applied to extract information that smoothly varies over space (e.g. depth in Becker and Hinton (1992)). In contrast to this, we examine a network where the connections between two areas are not uniform but learned following a standard normalized Hebbian learning rule.

We consider a higher area where information from different modalities is integrated. Let neurons in this area be translation invariant. Here we show that with the described learning rule neurons in an area, that is one step lower in the hierarchy of processing, also acquire those invariances and therefore are able to optimally and sparsely predict the responses of the higher area. Only the relevant information is transmitted to the subsequent layers. This goal function is termed relevant infomax.

On the input layer of two dimensions (space and feature, figure 4(A)), two neurons are active at each instant. Thus, the input represents pre-processed information and has a sparse representation. The two neurons transmit their activity to the middle layer of seven neurons where learning occurs. The self-excitation is set to zero for this simulation. The activity of the higher layer is also set by the stimulus: each neuron on that layer fires when a specific combination of features is shown on the input layer. These cells are therefore translation invariant.

After convergence, the neurons of the middle layer have acquired receptive fields so that they are translation invariant as well. Typical receptive fields are shown in figure 4(B). In each of the ten simulations exactly one neuron was tuned to each feature. Sometimes meta-stable states were observed where two neurons had almost identical receptive fields for some 10 000 iterations.

To characterize neuronal response properties, measures of specificity in the direction of space and feature are analysed separately. The standard deviation of synaptic efficacies, which define the receptive field, in direction of feature and of space are shown in figure 4(C) top and bottom, respectively. The neurons respond specifically to features and unspecifically to space. Thus, an invariance property of a higher area is propagated down to lower areas. In this way the dimensionality of the representation is reduced as early as possible. This is desirable because decisions based on low-dimensional data need far less training to be efficient (Edelman and Intrator 1997).

In this simulation the neurons learned to transmit that part of the information that is relevant to higher areas. Such a mechanism, where the mutual information with the higher area is maximized, can enhance the effects of information maximization in a network where for example the number of cells in the hidden layer defines constraints on the transmitted information. The network does not necessarily transmit as much information as possible, but only the relevant part of the present information.

4. Discussion

4.1. Biological implementation

Here we discuss one possible implementation of our learning algorithm and assemble evidence in favour of it. Recent research on the properties of apical dendrites suggests that the apical dendrite acts, in addition to the soma, as a second site of integration (Bernander *et al* 1994, Larkum *et al* 1999). 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 (Schiller *et al* 1997). 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 (Amitai *et al* 1993, Stuart and Sakmann 1994, Buzsaki and Kandel

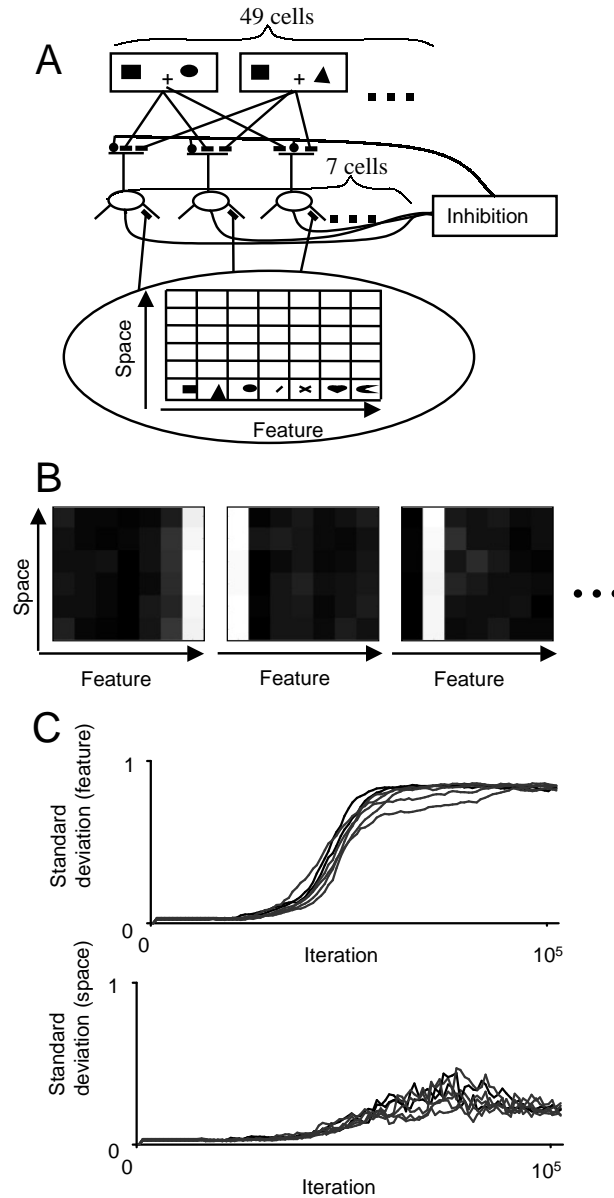


Figure 4. Relevant infomax. (A) The network architecture is shown. The input plane consists of neurons arranged in a two-dimensional array (7×6) with the axes corresponding to feature and space. In every iteration two randomly chosen cells on the input plane are activated. Their output is transmitted to the second layer of seven neurons, where the learning takes place. The third layer consists of cells that fire for a combination of two active features on the input plane. (B) Three typical receptive fields are shown in grey scale after convergence. The neurons learned that space is irrelevant from the higher layer where space is already ignored. (C) The temporal evolution of the standard deviation in space and feature averaged over the other dimension is shown. Please note that a small standard deviation in the direction of space is a sign of invariant responses and that a large standard deviation in the direction of feature indicates a high specificity.

1998) or passive current flow. Theoretical and experimental studies support the view that excitation to the apical dendrite is strongly attenuated on its way to the soma but that nonlinearities are able to amplify it again (Bernander *et al* 1994, Cauller and Connors 1994, Schiller *et al* 1997, Stuart and Spruston 1998, Larkum *et al* 1999). It is discussed below how the system changes with increasing influence of the apical dendrite on the cell activity. Slow regenerative calcium spikes have been observed *in vitro* (Schiller *et al* 1997) and *in vivo* (Hirsch *et al* 1995). They are one of the dendritic nonlinearities, are initiated in the apical dendrites, and cause a strong and prolonged depolarization (Stuart *et al* 1997, Larkum *et al* 1999). Furthermore 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 (Larkum *et al* 1999). So in conclusion, research on apical dendrites points to a system where a subset of synapses is able to induce rare discrete events of strong prolonged depolarization.

In experiments investigating the influence of post-synaptic depolarization on plasticity of synapses the synaptic weight change was measured as the function of post-synaptic activity (Artola *et al* 1990, Dudek and Bear 1991). Those experiments demonstrated that at low firing rates or depolarization long-term depression occurs whereas at higher activities a switch to LTP takes place. Thus, if there is no calcium spike and, as a consequence, a limited depolarization, the cell is in the region where synapses are depressed or only weakly potentiated. If calcium spikes are triggered, the high depolarization leads to LTP. So integrating the above two lines of research naturally leads to a plausible biological implementation for the hypothesised mechanism, where calcium spikes could correspond to the hypothesised 'learning events'. Indeed Pike *et al* (1999) gave evidence for the idea that post-synaptic bursting is essential for the induction of LTP in hippocampus giving further support for our learning mechanism.

In this proposed implementation there would definitely be an influence of the apical dendrite on the activity of the cell. But since the 'learning events' are considered to be rare, the effect on the firing rate is limited. Nevertheless, there will be an effect of the potential at the apical dendrite and thus the learning signal. But the effect can be weaker compared with the one site of the integration model. Indeed such a limited effect can be desirable and was analysed previously by Siegel *et al* (2000).

Data about cortical connectivity seem compatible with the investigated examples. There is evidence for local inhibitory projections to the apical dendrites and for associational cortico-cortical projections preferentially innervating the apical tuft branches of pyramidal neurons (Zeki and Shipp 1988, Cauller and Connors 1994). Top-down projections usually terminate in layer 1, where many apical tufts can be observed (for a review, see Salin and Bullier 1995). This supports the mechanism described in the invariance simulation.

4.2. Simplifications

In this implementation several physiological aspects are simplified:

- (1) The complex nonlinear dendritic properties (Mel 1993) are reduced to a threshold mechanism triggering dendritic calcium spikes. Current flow from the apical dendrite to the soma is neglected. In our simulation 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 (Phillips *et al* 1995). We chose to neglect it here to avoid obscuring the paper with too many details and parameters. The effects of calcium spikes on the activity of the cells and their signal processing was analysed in a previous work (Siegel *et al* 2000). There it was analysed how top-down information can lead to enhanced processing of bottom-up signals. Both effects can happen at the same time,

the top-down effect can improve signal processing on a short time scale and gate learning on a long time scale.

(2) The effect of inhibition on the cells is reduced to its effect on the potential at the apical dendrite. Introducing a more realistic inhibition, which not only acts on learning but also on firing rates could enhance the contrast and lead to enhanced decorrelation (Rumelhart and McClelland 1986). Furthermore excitatory connections within one layer were neglected; they could also result in a wider set of possibly computed functions. These assumptions were made to analyse the learning rule in a simple system to make it possible to understand its learning behaviour.

(3) In our simulations we use a normalization of the weight vectors to avoid diverging weights. There is some experimental support for such a biological mechanism. Blocking neural activity leads to increased miniature excitatory post-synaptic synaptic currents, whereas blocking GABA mediated inhibition leads to higher firing rates and to smaller synaptic currents (Turrigiano *et al* 1998). Alternatively a BCM like learning rule (Bienenstock *et al* 1982) 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.

4.3. Alternative learning mechanisms

The type of supervised learning described here may be contrasted with different types of learning:

It can be compared with a supervisor signal employed globally and unspecifically to large areas. Candidate mechanisms are triggered by neuromodulators like acetylcholine, which has been supposed to act as a ‘print now’ signal (Singer *et al* 1979). This implies that a large percentage of cells have the tendency to learn the given stimulus which was also shown in experiments (Kilgard and Merzenich 1998). The disadvantage of such methods of supervised learning is that they are unspecific, a learning goal like the one used here for the invariance simulation would not be achievable. Compared with such a signal the ‘learning signal’ in this study can transmit highly specific information, individual to every cell. It may be a complex signal deriving from cells of the same area, bottom-up inputs from lower areas and top-down information from higher areas. The cells can thus specifically change their receptive fields, so that a global network property is optimized. Nevertheless there exist stimuli, whose importance is such that many cells should change their receptive fields rapidly. Both mechanisms, specific and unspecific, could work together in optimizing receptive fields for globally defined goals.

It is possible to do the same type of learning with a classical local signal in the sense that the change of weights is defined exclusively by pre- and post-synaptic activities. If the effect of the learning signal on the post-synaptic activity increases the firing rate gets defined by the ‘learning signal’. Even with a rate-based learning rule one can then reach the same learning goal as the one proposed here. In this case, however, a dominant part of the post-synaptic activity must be defined by the inputs that set the learning goals. The activity transmitted to subsequent steps of processing is therefore dominated by the learning signal. Lowering the influence of the ‘learning signal’ on the activity can make this effect smaller. But at the same time the ‘learning signal’ is no longer defined that well, cells get increasingly biased to represent stimuli that lead to high firing rates. The cell activity in such a framework is always a mixture of the bottom up signal containing information about the actual stimulus and the ‘learning signal’. A cell can either be trained to acquire the optimal bottom-up receptive field but have an output dominated by the learning signal, or have a worse bottom-up receptive fields but only low interference by the ‘learning signal’. Thus there is a trade-off between efficient learning and efficient signal transmission. The set of possible goal functions that

can be defined for learning is limited; it is difficult for cells to learn stimuli that would not result in cell activity without the teacher. Nevertheless this approach is feasible, and is for example used by Eisele and Sejnowski (1998), where specially designed learning rules make sure that the influence of the top-down 'learning signal' on the activity of the cell stays limited. The top-down signal nevertheless always has a strong impact on cell activities leading to the problems discussed above.

The learning occurring here can be compared to the learning rules for multi-unit local processors with multivariate binary outputs for the extraction of a number of coherent features as proposed by Kay *et al* (1998). Their learning rules are derived for a class of information theoretic objective functions including the classical infomax and maximization of coherent information. The latter is related to relevant information maximization described in this paper. 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. They share several properties with the basal and apical dendrite in our studies, respectively (see also Siegel *et al* 2000). The context gates learning and is simultaneously used for improving processing of the receptive field input (Phillips and Singer 1995). Their learning rules contain 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.

4.4. Experimental predictions

The simulations described here imply some experimental predictions. 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 presynaptic 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 also for synapses at the basal dendrites.

Along the same lines one could not only analyse blocking of calcium spike but also the effects of additional calcium spikes. Patching an apical dendrite in layer 1 *in vivo* would make it possible to artificially induce calcium spikes in analogy to Larkum *et al* (1999). Following the predictions of the proposed model one should be able to conduct supervised learning *in vivo* similar to Debanne *et al* (1998). 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. Some normal spikes should have less influence on plasticity than one calcium spike resulting in a volley of the same number of spikes.

In conclusion, the presented paper analyses a learning algorithm where synaptic inputs are integrated at two sites. One site defines the activity and the other one gates learning. The resulting learning rule can cope with several common problems and leads to an interesting new algorithm.

Acknowledgments

It is a pleasure to thank our referees for constructive comments on a previous version of this manuscript. This work has been supported by the Swiss National Science Foundation (PK 31-51059.97) and the Boehringer Ingelheim Fond (KPK).

References

- Amitai Y, Friedman A, Connors B W and Gutnick M J 1993 Regenerative activity in apical dendrites of pyramidal cells in neocortex *Cereb. Cortex.* **3** 26–38
- 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
- Becker S 1996 Models of cortical self-organization *Network: Comput. Neural Syst.* **7** 7–31
- Becker S and Hinton G E 1992 Self-organizing neural network that discovers surfaces in random-dot stereograms *Nature* **355** 161–3
- Bernander O, Koch C and Douglas R J 1994 Amplification and linearization of distal synaptic input to cortical pyramidal cells *J. Neurophysiol.* **72** 2743–53
- Bienenstock E L, Cooper L N and Munro P W 1982 Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex *J. Neurosci.* **2** 32–48
- Buzsáki G and Kandel A 1998 Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat *J. Neurophysiol.* **79** 1587–91
- Caulier L J and Connors B W 1994 Synaptic physiology of horizontal afferents to layer I in slices of rat S I neocortex *J. Neurosci.* **14** 751–62
- Debanne D, Shulz D E and Fregnac Y 1998 Activity-dependent regulation of ‘on’ and ‘off’ responses in cat visual cortical receptive fields *J. Physiol.* **508** 523–48
- Dudek S M and Bear M F 1991 Homosynaptic long term depression in area CA1 of hippocampus and the effects on NMDA receptor blockade *Proc. Natl Acad. Sci. USA* **89** 4363–7
- Edelman S and Intrator N 1997 Learning as extraction of low-dimensional representations *Perceptual Learning* ed R Goldstone, D Medin and P Schyns (New York: Academic) pp 353–80
- Eisele M and Sejnowski T 1998 Unsupervised and reinforcement learning by pyramidal neurons in a layered model of neocortex *Soc. Neurosci.* **24** 1571
- Hebb D O 1949 *The Organization of Behaviour* (New York: Wiley)
- Hirsch J A, Alonso J M and Reid R C 1995 Visually evoked calcium action potentials in cat striate cortex *Nature* **378** 612–6
- Intrator N and Cooper L N 1992 Objective function formulation of the BCM theory *Neural Networks* **5** 3–17
- Kay J, Floreano D and Phillips W A 1998 Contextually guided unsupervised learning using local multivariate binary processors *Neural New.* **11** 117–40
- Kilgard M P and Merzenich M M 1998 Cortical map reorganization enabled by nucleus basalis activity *Science* **279** 1714–8
- Körding K P and König P 2000 A learning rule for dynamic recruitment and decorrelation *Neural Networks* **13** 1–9
- Larkum M E, Zhu J J and Sakmann B 1999 A new cellular mechanism for coupling inputs arriving at different cortical layers *Nature* **398** 338–41
- Mel B W 1993 Synaptic integration in an excitable dendritic tree *J. Neurophysiol.* **70** 1086–101
- Phillips W A, Kay J and Smyth D 1995 The discovery of structure by multi-stream networks of local processors with contextual guidance *Network: Comput. Neural Syst.* **6** 225–46
- Phillips W A and Singer W 1997 In search of common foundations of cortical computation *Behav. Brain Sci.* **20** 657–722
- Pike F G, Meredith R M, Olding A W A and Paulsen O 1999 Postsynaptic bursting is essential for ‘Hebbian’ induction of associative long-term potentiation at excitatory synapses in rat hippocampus *J. Physique* **518.2** 557–76
- Rumelhart D and McClelland J 1986 *Parallel Distributed Processing* (Cambridge: Bradford Books)
- Salin P A and Bullier J 1995 Corticocortical connections in the visual system: structure and function *Physiol. Rev.* **75** 107–54
- Schiller J, Schiller Y, Stuart G and Sakmann B 1997 Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons *J. Physiol.* **505** 605–16
- Sejnowski T J 1977 Storing covariance with nonlinearly interacting neurons *J. Math. Biol.* **4** 303–21
- Siegel M, Körding P K and König P 2000 Integrating bottom-up and top-down sensory processing by somato-dendritic interactions *J. Comput. Neurosci.* **8** 161–73

- Singer W, von Grünau M W and Rauschecker J P 1979 Requirements for the disruption of binocularity in the visual cortex of strabismic kittens *Brain Res.* **171** 536–40
- Stent G S 1973 A physiological mechanism for Hebb's postulate of learning *Proc. Natl Acad. Sci. USA* **70** 997–1001
- Stuart G J, Schiller J and Sakmann B 1997 Action potential initiation and propagation in rat neocortical pyramidal neurons *J. Physiol.* **505** 617–32
- Stuart G J and Spruston N 1998 Determinants of voltage attenuation in neocortical pyramidal neuron dendrites *J. Neurosci.* **18** 3501–10
- Stuart G J and Sakmann B 1994 Active propagation of somatic action potentials into neocortical pyramidal cell dendrites *Nature* **367** 69–72
- Turrigiano G G, Leslie K R, Desai N S, Rutherford L C and Nelson S B 1998 Activity-dependent scaling of quantal amplitude in neocortical *Nature* **391** 892–6
- Zeki S and Shipp S 1988 The functional logic of cortical connections *Nature* **335** 311–7