Journal of Computational Neuroscience 11, 207–215, 2001

Supervised and Unsupervised Learning with Two Sites of Synaptic Integration

KONRAD P. KÖRDING AND PETER KÖNIG
Institute of Neuroinformatics, ETH/UNI Zürich, Winterthurerstr. 190, 8057 Zürich, Switzerland
koerding@ini.phys.ethz.ch

Received June 1, 2000; Revised August 31, 2001; Accepted August 31, 2001

Action Editor: Terrence Sejnowski

Abstract. Many learning rules for neural networks derive from abstract objective functions. The weights in those networks are typically optimized utilizing gradient ascent on the objective function. In those networks each neuron needs to store two variables. One variable, called activity, contains the bottom-up sensory-fugal information involved in the core signal processing. The other variable typically describes the derivative of the objective function with respect to the cell's activity and is exclusively used for learning. This variable allows the objective function's derivative to be calculated with respect to each weight and thus the weight update. Although this approach is widely used, the mapping of such two variables onto physiology is unclear, and these learning algorithms are often considered biologically unrealistic. However, recent research on the properties of cortical pyramidal neurons shows that these cells have at least two sites of synaptic integration, the basal and the apical dendrite, and are thus appropriately described by at least two variables. Here we discuss whether these results could constitute a physiological basis for the described abstract learning rules. As examples we demonstrate an implementation of the backpropagation of error algorithm and a specific self-supervised learning algorithm using these principles. Thus, compared to standard, one-integration-site neurons, it is possible to incorporate interesting properties in neural networks that are inspired by physiology with a modest increase of complexity.

Keywords: calcium spike, apical dendrite, backpropagation, self-supervised learning, bursts of action potentials

1. Introduction

Cortical neurons exhibit amazingly complex dendrites. Their morphologies have spurred a lot of interest, and many studies investigate processing of synaptic input within the dendritic tree (Zador et al., 1995; Mainen et al., 1996; Mainen and Sejnowski, 1996; Carnevale et al., 1997; Agmon-Snir et al., 1998; Mel et al., 1998). Some evidence argues for a sublinear interaction of synaptic inputs due to dendritic saturation (Mel, 1993). Other studies suggest a supralinear interaction due to voltage-dependent conductances (Softky, 1994; Schiller et al., 2000). Indeed, it has been proposed that different nonlinear effects cancel each other and result in a linear superposition of synaptic inputs at the soma (Johnstone, 1996; Cash and Yuste, 1998; Cook and Johnston, 1999). These studies are based on data obtained in different cortical regions using acute or cultured brain slices. Thus, they cannot be directly compared, and current results do not allow the properties of dendritic interactions in pyramidal neurons to be unambiguously resolved under *in vivo* conditions.

As a consequence, two main lines of investigation emerged. Some studies concentrate on the dynamics within single neurons and take into account the detailed structure of the dendritic tree. Other studies investigate dynamic interactions in larger networks and use integrate-and-fire neurons or other highly simplified

model neurons as basic units. In such neurons all synaptic input is summed and only then subjected to a nonlinear transfer function. This single integration site is usually assigned to the soma resulting in a point-like neuron without dendritic structure. This is a gross simplification, which nonetheless seemed necessary because simulations of large networks considering all neurons in great detail are currently not possible. Here we argue that considering a modest increase in complexity, a second site of integration, offers a second independent variable, which allows a larger class of neural network algorithms to be implemented. Indeed, many successfully applied neural networks use two variables per cell. The first variable, called activity, transmits information in the hierarchy of processing. The second variable typically stores the derivative of the objective function with respect to the cell's input and is exclusively used for learning.

A prominent example from the domain of supervised learning is the "backpropagation of error" algorithm (Werbos, 1974/1994; Rumelhart and McCleland, 1986; LeCun et al., 1990). For each neuron it calculates activity as well as the derivative of the globally defined error function with regard to the neuron's activity. These two variables together allow computing updates of synaptic weights. In a recent survey of applications of neural networks (Arbib, 1998), this lgorithm features prominently. Among 26 different methods the backpropagation algorithm and its closely related cousin, backpropagation in recurrent networks, are applied in 13 areas of a total of 18 areas. With applications in only five areas the runner up trails far behind. Thus, it is fair to say (Arbib, 1998) that "Backpropagation is the most diversely used adaptive architecture." This has led to some proposals addressing possible implementations, and a number of mechanisms have been suggested: a second network propagating back the errors (Zipser and Rumelhart, 1990; Tesauro, 1990); global reinforcement (Mazzoni et al., 1991); bidirectional recirculation in a recurrent network (Hinton and McClelland, 1988); and contrastive Hebbian Learning (O'Reilly, 1996). In the latter two approaches early signals carry the activity, whereas late signals also contain the error signal. Nevertheless, the common notion is that backpropagation is a successful algorithm but has no obvious link to processes of the brain. Indeed, according to Arbib (1998, p.), "This architecture is an example of 'neurally inspired' modeling, not modeling of actual brain structures: there is no evidence that backpropagation represents actual brain mechanisms of learning."

As a second example, since in many technical applications as well as in most biological systems labeled data are scarce, self-supervised algorithms have been developed. They do not need an external supervisor but locally generate a supervision signal and have thus been proposed to be linked more directly to cortical function. As an example, the Imax algorithms (Becker, 1996) are based on the idea of maximizing the mutual information between outputs of different network modules. Variants of these algorithms are based on spatial (Becker and Hinton, 1992; Stone and Bray, 1995), cross-modal (de Sa and Ballard, 1998), or temporal (Földiak, 1991) smoothness criteria and also use a second variable for each cell containing the derivative of a local objective function with respect to the cell's activity. With respect to self-supervised learning, a number of possible mechanisms have been proposed as well: nondriving synapses (Becker, 1996), a nonlocal threshold mechanism (Kay et al., 1998), and a second layer of neurons comparing two activities (de Sa and Ballard, 1998). A characteristic property of these solutions is the idea pertinent in the neural network community that neurons exhibit just one site of synaptic integration and thus can add presynaptic influences at one site of synaptic integration only.

Here we explore how known characteristics of pyramidal neurons can interact to form a system comparable to the previously described algorithms and give an alternative proposal how the above principles could be mapped on physiology. We do not assume that cortex actually works as a backpropagation machine. Neither do we try to improve the convergence behavior or memory overhead of existing algorithms. Instead, we investigate a possible implementation of an interesting and widely used algorithm in physiologically plausible network architectures and like to point out new aspects when thinking about dynamics and function of cortical networks.

2. Methods and Relevant Biological Experiments

2.1. Integration in the Basal and Apical Dendritic Tree

The most abundant type of neuron in cerebral cortex, the pyramidal neuron, is characterized by its prominent apical dendrite. Recent research on the properties of layer V pyramidal neurons suggests that the apical dendrite acts, in addition to the soma, as a second site

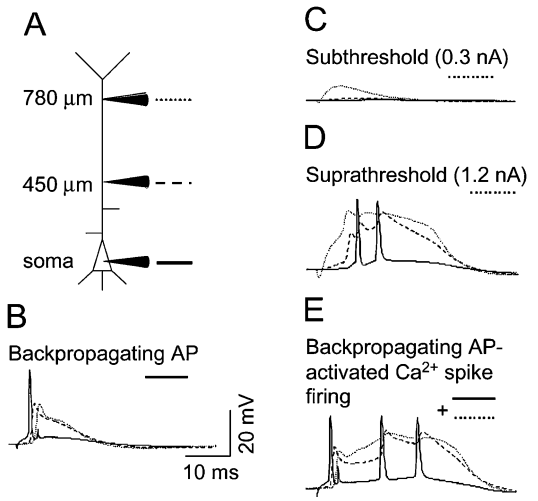

Figure 1. The relevant physiology. In vitro data kindly supplied by M. Larkum (MPI Heidelberg) modified after a figure in (Larkum, 1999b) **A:** The experimental setup. A neuron is patched simultaneously at three positions: the soma (recordings in solid line), 0.450 mm (recordings in dashed lines), and 0.780 mm (recordings in dotted lines) out on the apical dendrite. **B:** The effect of inducing an action potential by current injection into the soma. It travels anterogradly into the axon and retrogradly into the apical dendrite, where it leads to a slow and delayed depolarization. **C:** Injection of a small current at the apical dendrite does not trigger regenerative potentials and induces only a barely noticeable depolarization at the soma. **D:** Injection of a larger current into the apical dendrite elicits a regenerative event with a long depolarization (dotted line). This slow potential is less attenuated and reaches the soma with a delay of only a few milliseconds (solid line). There, a series of action potentials is triggered, riding on the slow depolarization. **E:** Combining a subthreshold current injection into the apical dendrite with a somatic action potential leads to a calcium spike and a burst of action potentials. Source: Based on Larkum et al. (1999b).

of synaptic integration (Larkum et al., 1999b; Körding and König, 2000). Each site integrates input from a set of synapses defined by their anatomical position and is able to generate regenerative potentials (Schiller et al., 1997). The two sites exchange information in the following way (see Fig. 1A, modified after a figure in Larkum et al., 1999b): (1) signals originating at the soma are transmitted to the apical dendrite by actively backpropagating dendritic action potentials (Fig. 1A, B, Amitai et al., 1993; Stuart and Sakmann, 1994; Buzsaki and Kandel, 1998); (2) signals from the apical dendrite to the soma are sent via slow regenerative calcium spikes (Fig. 1D, E), which have been observed in vitro (Schiller et al., 1997) and in vivo (Hirsch et al., 1995; Helmchen et al., 1999). These calcium spikes are initiated in the apical dendrites and cause a strong and prolonged depolarisation, typically leading to bursts of action potentials (Fig. 1D, E, Stuart et al., 1997; Larkum et al., 1999a, 1999b). 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 (Fig. 1B) (Schiller et al., 1997; Larkum et al., 1999b). In conclusion, a subset of synapses on the apical dendrite is able to induce discrete events of strong prolonged depolarization combined with bursts.

Synaptic efficacy is a function of the temporal structure of afferent action potentials. They typically either facilitate or depress (Markram et al., 1997). On presynaptic activity facilitating synapses increase their efficacy for a short interval, whereas depressing synapses react more weakly to presynaptic events within a short time following presynaptic activity. Thus, highly facilitating synapses lead to a strong signal only for bursts; highly depressing synapses strongly respond only to trains of single spikes.

To accommodate these results we consider the rate of action potentials (A) and the rate of dendritic bursts (D). Both are in general a function of synaptic efficacies for afferent action potentials ($\boldsymbol{W}_{basal}$, $\boldsymbol{W}_{apical}$) and bursts ($\boldsymbol{V}_{basal}$, $\boldsymbol{V}_{apical}$) depending on the type of synapses involved (bold letters are used to denote vectors, here the vector of all synapses afferent to the basal and apical dendritic tree of one postsynaptic neuron), and the respective rates of the presynaptic neurons ($\boldsymbol{A}_{pre}$, $\boldsymbol{D}_{pre}$):

$$\boldsymbol{A}_{post} = f(\boldsymbol{W}_{basal}\boldsymbol{A}_{pre}, \boldsymbol{V}_{basal}\boldsymbol{D}_{pre}, \boldsymbol{D}_{post}),$$

where $\mathbf{XY}$ denotes the scalar product of $\mathbf{X}$ and $\mathbf{Y}$ and

$$\boldsymbol{D}_{post} = g(\boldsymbol{W}_{apical}\boldsymbol{A}_{pre}, \boldsymbol{V}_{apical}\boldsymbol{D}_{pre}, \boldsymbol{A}_{post}),$$

where f and g are two transfer functions. Their choice is given further down in the respective context.

2.2. *Input to the Basal and Apical Dendritic Tree*

To complete the picture we have to consider which afferents target the apical and basal dendritic tree, respectively. Although the anatomy of a cortical column is complicated, some regular patterns can be discerned. The apical dendrites of the considered layer 5 pyramidal neurons receive local inhibitory projections and long-range cortico-cortical projections (Zeki

and Shipp, 1988; Cauller and Connors, 1994). Top-down projections from areas higher in the hierarchy of the sensory systems and long-range connections within the same cortical area strongly terminate in layer 1, where many apical tufts can be observed (cf. Salin and Bullier, 1995). This supports the idea that top-down connections from higher to lower areas, and long-range lateral connections 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.

We use the approximation that the bottom-up input targets the basal dendritic tree ($\boldsymbol{W}_{top\text{-}down,basal} = 0$, $\boldsymbol{V}_{top\text{-}down,basal} = 0$), whereas the apical dendrite integrates top-down information from higher areas and the same area ($\boldsymbol{W}_{bottom\text{-}up,apical} = 0$, $\boldsymbol{V}_{bottom\text{-}up,apical} = 0$). We must note though that this approximation is only partially supported by biological data. It is true that top-down projections massively target the upper layers of cortex (Salin and Bullier, 1995), but there are some connections to the basal dendrites as well. Similarly bottom-up connections do not exclusively target basal dendrites but to a certain degree also target the apical dendrites.

2.3. The Synaptic Weight Update

Experiments on hippocampal slices by Pike et al. (1999) support the idea that postsynaptic bursting is essential for the induction of LTP. Furthermore, independent experiments support the idea that strong postsynaptic activity also is necessary in cortex for induction of Hebbian learning (Artola et al., 1990; Dudek and Bear, 1993). We infer that bursts and thus calcium spikes can trigger Hebbian plasticity at active synapses.

At the site of the synapse, bursts and action potentials could trigger synaptic plasticity. Together with the time course of the presynaptic activity the synapse can change its parameters. Here we consider V and W and only changes that are proportional to the presynaptic A and D. This yields the following weight update: $\Delta\mathbf{W} = F_1(A, D)\ F_2(\mathbf{A_{pre}}, \mathbf{D_{pre}})$, $\Delta\mathbf{V} = G_1(A, D)\ G_2(\mathbf{A_{pre}}, \mathbf{D_{pre}})$, where the F and G are linear functions. Learning is proportional to presynaptic activity and burst-rate.

3. Results

Within the framework outlined above, we investigate possible implementations of supervised and unsupervised learning rules.

3.1. How to Emulate the Backpropagation of Error Algorithm

In addition to the formalization described above—which we think is compatible with experimental results—we have to make additional assumptions:

- Basal synapses are modified only when calcium spikes are triggered: $\Delta\mathbf{W_{basal}} = \beta\, D_{post}\, \mathbf{A_{pre}}$. In contrast, apical synapses learn proportional to the postsynaptic activity according to a conventional Hebbian learning rule: $\Delta\mathbf{V_{apical}} = \beta\, A_{post}\, \mathbf{A_{pre}}$. This results in symmetric synaptic weighs.
- All basal synapses are strongly depressing ($\mathbf{V_{basal}} = 0$), and all apical synapses strongly facilitating ($\mathbf{W_{apical}} = 0$).
- The transfer function f of the activity is $A_{post} = f(\mathbf{W}_{basal}\mathbf{A}_{pre},\ \mathbf{V}_{basal}\mathbf{D}_{pre},\ \mathbf{D}_{post}) = 1/(1 + \exp(-\mathbf{WA_{pre}}))$, so that A_{post} is of [0,1] and the transfer function g of the dendrite is $D_{post} = g(\mathbf{W}_{apical}\ \mathbf{A}_{pre}, \mathbf{V}_{apical}\mathbf{D}_{pre}, A_{post}) = (\mathbf{V_{apical}}\ \mathbf{D_{pre}})A_{post}(1 - A_{post})$.

Given these assumptions, the resulting dynamics are identical to those of backpropagation, the error δ is just exchanged by the burstrate D, so the system implements $\Delta\mathbf{W} \sim \delta\mathbf{A_{pre}}$ and $\delta = A(1 - A)\mathbf{W}\delta$. In essence, errors are transmitted by bursts via the top-down connections and processed first in the apical dendrite. The core signal processing is performed in the soma using action potentials transmitted in the bottom-up direction. Facilitating and depressing synapses allow multiplexing these signals via the same axons.

It has to be noted that these additional assumptions are speculative. Learning of apical synapses, for example, is more likely to be proportional to the burst rate as well, and the assumptions that synapses depress or facilitate are clearly not met. The transfer function for D_{post} could be realistic since it has been shown that increasing the firing rate increases the readiness of the cell to have calcium spikes. The saturation of D_{post} has not yet been shown but could be implemented with channels that are inactivated by backpropagating action potentials. D_{post} is continuous, which means that we assume the system is noisy enough so

that the burstrate is proportional to the input; the cell is a stochastic burster. This is also a weakness of the proposed system since it implies that the system must see the stimulus long or often enough for the statistics of the stochastic burster to even out. Another issue is that the values of D_{post} in biology are constrained to be positive; no negative burst-rates would seem possible. This problem can be overcome by having a base bursting activity; values below lead to decreasing weights, and values above to increasing weights. But nevertheless, it has to be noted that our knowledge of properties of synapses connecting particular pairs of neurons is very limited. Indeed, only recently insight has been gained with respect to specific properties of synaptic connections of different classes of interneurons (Gupta et al., 2000). Thus, time will have to show whether some of the assumptions above, which are questionable, are outright wrong. On the other hand, if the above assumptions are met, the system performs backpropagation of error.

3.2. *Learning the X-Or Function*

The backpropagation of error algorithm received widespread popularity, as it allows multilayer perceptrons to be trained, which as opposed to the linear perceptron (Minsky and Papert, 1969) can approximate nonlinear functions and solve linearly inseparable problems. Results of training a network with this algorithm have even been compared with cortical receptive fields (Zipser and Andersen, 1988). The simplest example of such a nonlinear function is the two-bit parity or X-OR function. For our simulations we choose $\beta = 2$, all initial weights randomly from (0..1(and simulate the system for 10,000 iterations. In Fig. 2A, B we demonstrate how neurons with two sites of synaptic integration emulating backpropagation learn this function. The residual error decreases steeply over time, indicating optimization of the network with respect to the globally defined error function.

Classical Hebbian learning rules applied to such a network composed of units with one integration site cannot learn this task since they approximately search for the first principal components. In the present framework such a system that is composed of units with only one site of integration unit is equivalent to mixing the signals targeting basal and apical dendritic tree for the determination of activity as well as of learning. As a stricter control we mix signals targeting

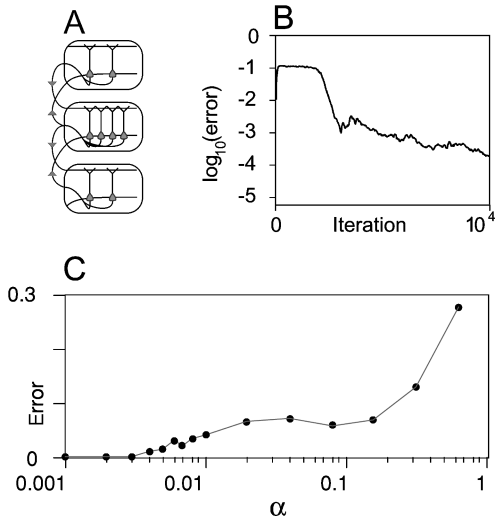

Figure 2. Backpropagation algorithm. **A:** The network layout is sketched. We consider a three-layer network with a set of input neurons, a set of hidden neurons, and a set of output neurons. **B:** The average error of the output neurons averaged is shown as a function of the iteration. **C:** The average error averaged over the last 1000 iterations is plotted as a function of the influence of the postsynaptic firing rate onto learning as parameterized by α. The standard error of the mean determined in 10 runs is about the same size as the circles.

the two sites for determination of the learning only: $\Delta \mathbf{W_{basal}} = \beta \mathbf{A_{pre}}(D_{post} + \alpha A_{post})$. The constant α determines the relative influence of the activity of the postsynaptic neuron relative to the burst rate. Here, however, the relative size of the two variables has to be considered. Initially, they are of the same order of magnitude. During the simulation the error drops (Fig. 2B) and is quickly an order of magnitude or more smaller than the average activity. Thus, a mixing parameter of 0.01 (for example) is not as small as it might seem on first sight. Nevertheless, Fig. 2C shows that performance decays for increasing influences of the firing rate of the postsynaptic neuron onto synaptic plasticity.

The error in the cases where the network does not converge to a correct solution is about 0.05, which is smaller then the variance of activity by a factor of about 20. The objective function of Hebbian learning, the detection of the principle components, is mixed with the objective function of backpropagation learning, the minimization of the global error. If the cell had only one site of integration, larger mixing parameters than those used here seem inevitable. Thus, in the investigated network the emulation of backpropagation learning relies on the second site of synaptic integration.

3.3. Self-Supervised Learning

Neurons with two sites of integration also allow implementing self-supervised learning algorithms (Körding and König, 2000; compare Becker and Hinton, 1992; Stewart and Sejnowski, 1998). Approaches for self-supervised learning use two variables per cell to integrate a learning signal from a spatiotemporal area larger then the one supplying activating input. So here we analyze such a system and test whether the same solutions can be obtained using just one site of synaptic integration.

We use a physiologically plausible transfer-function: $A_{post} = H(\mathbf{A_{pre}W}_i)/\underline{A}_i^2$, where H is the heavyside function ($H(x) = 0$ for $x < 0$, $H(x) = x$ otherwise) and $\underline{A}_i$ is the running average of the activity with a time constant of 1000. Based on the experimental result that already weak inhibition abolishes calcium spikes (Larkum et al., 1999b), we assume strong local competition on the level of the generation of calcium spikes. Within each module only in the neuron with the highest D_{post} calcium spikes are triggered. $D_{post} = (\mathbf{WA_{pre}} + \alpha A) = \max(\mathbf{WA_{pre}} + \alpha A_{post})$. Synaptic plasticity is triggered only by bursts. Cells that have not learned for an interval T increase their weights to stabilize the net activity: $\Delta W = \beta D_{post}(\mathbf{A_{pre}} - W) + T_\chi$. We choose $\beta = 0.002$, $\chi = 0.00001$, select all initial weights randomly from (0..1(and simulate the system for 10,000 iterations. The network is organized into streams (Fig. 3A). Streams are connected on the second layer. Each stream receives input from a distinct set of cells (4 by 3). The input position on the abscissa is randomly chosen and identical in both streams. The input position on the ordinate is randomly chosen for each stream, thus there is a correlation in the abscissa position but not in the ordinate position.

The goal of the network is to extract the coherent variable and generalize over the incoherent variable. Thus, in the chosen example the synaptic weights have to be adjusted such that the correlations between the higher areas of each stream are maximized. Figure 3B shows the normalized correlation against the iteration, as a function of the iteration for one run of $\alpha = 0.08$. The network learns to extract the coherent information.

To control for the relevance of two separate integration sites, we study a model with mixing signals targeting the apical and basal dendritic tree. Along the lines of the control above, a mixing parameter α determines the relative weight of the input to the basal dendritic tree in the function describing the synaptic updates. When α is large, this is equivalent to a one integration site model for learning but not for the activation of the neurons. In this sense it is an even stricter control. Figure 3C shows that increasing the effect of the postsynaptic activity on learning tampers with the performance of the system. Learning with just one site of synaptic integration leads to a tradeoff. If α is small, the cells are able to correctly learn according to their locally defined objective. But then they mostly transmit information about nonlocal properties of stimuli because their activity is dominated by the learning signal. If α is large, they do not learn correctly but represent information about local stimulus properties. In the example considered here, the majority of the cell's activity must be determined by the learning signal for the system to converge correctly. And in this case the

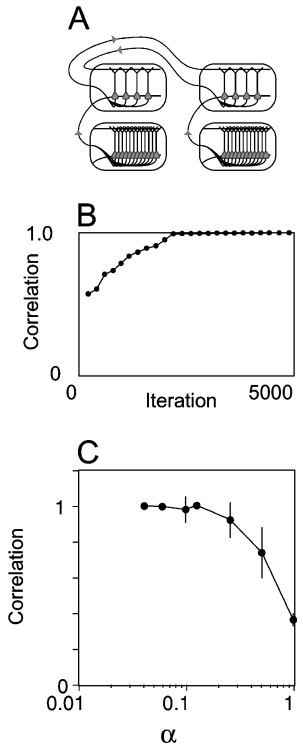

Figure 3. Self-supervised learning. **A:** The used network setup is shown. **B:** The normalized correlation $(\sum(A_{i1}A_{i2})^2/\text{sqrt}((\sum(A_{i1}A_{j1})^2)(\sum(A_{i2}A_{j2})^2)$, where numbers denote streams and i and j individual neurons, averaged over 200 iterations is shown for one run of $\alpha = 0.08$. A normalized correlation of one means that both streams exactly extract the same information. **C:** The normalized correlation (see text) averaged over the last 200 iterations is shown as a function of the influence of the firing rate onto learning as parameterized by α.

neuron hardly represents local stimulus properties. So also in the case of self-supervised learning a two-site model allows the implementation of some algorithms that seem unavailable to one-site models.

4. Discussion

In this study we demonstrate that taking into account a second integration site allows, at a modest increase in complexity, interesting properties to be incorporated in physiologically realistic neuronal networks. To our knowledge, this is the first demonstration that a network composed of units with two integration sites allows the implementation of such algorithms.

We proceed to discuss the concept of an integration site in some more detail. In the widely used integrate-and-fire model neuron all inputs are summed before a nonlinear transfer function is applied. Thus, the relative positions of synapses on the dendritic tree and potential interactions of inputs in the dendrite are neglected. The neuron may then be described as point-like with a single integration site. In contrast, a number of integration sites is taken into account by a $\Sigma\Pi$-unit. Here subsets of inputs interact in a multiplicative way (the Π), before the results are summed (the Σ). Thus the intermediate results depend only on disjoint subsets of afferent signals. In this sense such a model neuron has as many integration sites as multiplications are performed. The "real" cortical neuron is supposedly much more complex, and a description by a varying number of integration sites, albeit better then the one integration site model, must be considered a rough approximation.

An important question is whether the tradeoff between representing a stimulus and learning correctly generally holds for neurons with just one site of synaptic integration. Whenever the objective is not exclusively a function of the neuron's activity (as it is for ICA and PCA), the learning signal cannot be identical to the desired neuronal activity. It can only influence the weights via the one variable, the somatic activity, and therefore needs to distort this activity. That is why the system cannot correctly learn and respond at the same time. This tradeoff is thus a general property of neurons with just one site of integration.

The theory of graphical systems can help explain why some algorithmic problems that cannot be solved by a set number of one site neurons can be solved by the same number of two site neurons. Systems in which information flows can be described by directed graphs (cf. Frey, 1998). Every node of this graph corresponds to a variable in the system. The variables in our network are the weights, the level of activity, and in the two sites case also the dendritic potential. Any connection from one node to another indicates that the first variable directly influences the second. The graph for two sites of synaptic integration contains more connections over which information flows and thus allows a larger class of functions to be computed. The weight change, for example, can depend on somatic activity as well as dendritic potential. At the same time it is straightforward to have the two sites system emulate the one site system by appropriately setting the mixing parameter.

Here we argue that taking into account the dynamics within the apical dendritic tree results in a description with two integration sites. Thus, we have to investigate the quality of such an approximation. Recently a compartmental model was used to analyze the initiation of calcium spikes in the apical dendrite of pyramidal neurons (Schäfer et al., 2000). Considering recent results on channel dynamics, channel densities, and dendritic morphologies, they could identify factors that influence the generation of calcium spikes. The dendritic morphology and the distribution of calcium channels turned out to be key features in this process. Relating this with recent experimental results by Lüscher and his colleagues (Berger et al., 2001) suggests that the influence of synapses targeting the apical dendritic tree onto the soma are mediated mainly by regenerative potentials. Thus given the current state of our knowledge a description by two integration sites seems to be a reasonable description.

Since we are interested in analyzing effects on a network level, we concentrated on new aspects gained by considering a second integration site. Combining the results described above with additional assumptions we show here that neuronal networks might implement previously proposed unsupervised as well as supervised learning rules. Neurons in this framework do not only integrate both afferent driving activity and an error signal. They even transmit both variables through the same axon (Markram et al., 1998). Here we assume that single spikes transmit the activity and that bursts transmit error and learning signals.

We do not want to imply that the mammalian cerebral cortex literally works like a backpropagation machine. But we argue that the sometimes-perceived wide gap between physiologically realistic simulation and neural networks that successfully solve technical problems might not be that large after all. Thus, looking

for inspiration in biological networks and caring about physiological realism still are possible. In particular, neurons might transmit information in efferent spikes above and beyond the instantaneous firing rate. Synchronization relative to the activity of other neurons (Singer and Gray, 1995; König and Engel, 1995) and the temporal structure of individual spike trains as expressed by the rate of bursts (Livingstone et al., 1996; Siegel et al., 2000) are promising candidates for further research.

In this article, we explored how the mammalian cortex could learn under defined constraints. We argued that two sites of synaptic integration within each pyramidal neuron could allow the implementation of supervised as well as unsupervised learning rules. We pointed out which type of computations could be performed, with the most interesting aspect being that two sites of integration allow combining two classes of algorithms—supervised and unsupervised learning rules—in a homogeneous architecture. This might be an important contribution to the superior performance of biological systems in learning tasks.

References

Agmon-Snir H, Carr CE, Rinzel J (1998) The role of dendrites in auditory coincidence detection. Nature 393: 268–272.

Amitai Y, Friedman A, Connors BW, Gutnick MJ (1993) Regenerative activity in apical dendrites of pyramidal cells in neocortex. Cereb. Cortex 3: 26–38.

Arbib MA (1998) The Handbook of Brain Theory and Neural Networks. MIT Press: Cambridge, MA.

Artola A, Bröcher S, 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-organisation. Network: Comput. Neural Syst. 7: 7–31.

Becker S, Hinton GE (1992) Self-organizing neural network that discovers surfaces in random-dot stereograms. Nature 355: 161–163.

Berger T, Larkum ME, Lüscher H-R (2001) A high I_h channel density in the distal apical dendrite of layer 5 neocortical pyramidal cells increases bidirectional attenuation of EPSPs. J. Neurophysiol. 85: 855–868.

Buzsaki G, Kandel A (1998) Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat. J. Neurophysiol. 79: 1587–1591.

Carnevale NT, Tsai KY, Claiborne BJ, Brown TH (1997) Comparative electrotonic analysis of three classes of rat hippocampal neurons. J. Neurophysiol. 78: 703–720.

Cash S, Yuste R (1998) Input summation by cultured pyramidal neurons is linear and position-independent. J. Neurosci. 18: 10–15.

Cauller LJ, Connors BW (1994) Synaptic physiology of horizontal afferents to layer I in slices of rat SI neocortex. J. Neurosci. 14: 751–762.

Cook EP, Johnston D (1999) Voltage-dependent properties of dendrites that eliminate location-dependent variability of synaptic input. J. Neurophysiol. 81: 535–543.

Crick FHC (1989) The recent excitement about neural networks. Nature 337: 129–132.

De Sa VR, Ballard DH (1998) Category learning through multimodality sensing. Neur. Comput. 10: 1097–1117.

Dudek SM, Bear MF (1993) Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. J. Neurosci. 13: 2910–2918.

Földiak P (1991) Learning invariance from transformation sequences. Neur. Comput. 3: 194–200.

Frey, B. (1998) Graphical Models for Machine Learning and Digital Communication. MIT Press: Cambridge, MA.

Gupta A, Wang Y, Markram H (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. Science 287: 273–278.

Helmchen F, Svoboda K, Denk W, Tank DW (1999) In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons. Nat. Neurosci. 2: 989–996.

Hinton GE, McClelland JL (1988) Learning representations by recirculation. In: Anderson DZ, ed. Neural Information Processing Systems. American Institute of Physics, New York. pp. 358–366.

Hirsch JA, Alonso JM, Reid RC (1995) Visually evoked calcium action potentials in cat striate cortex. Nature 378: 612–616.

Johnston D, Magee JC, Colbert CM, Cristie BR (1996) Active properties of neuronal dendrites. Ann. Rev. Neurosci. 19: 165–186.

Kay J, Floreano D, Phillips WA (1998) Contextually guided unsupervised learning using local multivariate binary processors. Neural Networks 11: 117–140.

König P, Engel AK (1995) Correlated firing in sensory-motor systems. Curr. Opin. Neurobiol. 5: 511–519.

Körding KP, König P (2000) Learning with two sites of synaptic integration. Network: Comput. Neural Syst. 11: 1–15.

Larkum ME, Kaiser KM, Sakmann B (1999a) Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of backpropagating action potentials. Proc. Natl. Acad. Sci. USA 96: 14600–14604.

Larkum ME, Zhu JJ, Sakmann B (1999b) A new cellular mechanism for coupling inputs arriving at different cortical layers. Nature 398: 338–341.

LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, Jackel LD (1990) Handwritten digit recognition with a back-propagation network. In: Touretzky D, ed. Advances in Neural Information Processing Systems 2. Morgan Kaufman. Au: Pls. add city

Mainen ZF, Carnevale NT, Zador AM, Claiborne BJ, Brown TH (1996) Electrotonic architecture of hippocampal CA1 pyramidal neurons based on three-dimensional reconstructions. J. Neurophysiol. 76: 1904–1923.

Mainen ZF, Sejnowski TJ (1996) Influence of dendritic structure on firing pattern in model neocortical neurons. Nature 382: 363–366.

Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275: 213–215.

Markram H, Wang Y, Tsodyks M (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proc. Natl. Acad. Sci. USA 98: 5323–5328.

Mazzoni P, Andersen RA, Jordan MI (1991) A more biologically plausible learning rule for neural networks. Proc. Natl. Acad. Sci. USA 88: 4433–4437.

Mel BW (1993) Synaptic integration in an excitable dendritic tree. J. Neurophysiol. 70: 1086–1101.

Mel BW, Ruderman DL, Archie KA (1998) Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. J. Neurosci. 18: 4325–4334.

Minsky M, Papert S (1969) Perceptrons: An Introduction to Computational Geometry. MIT Press: Cambridge, MA.

O'Reilly RC (1996) Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. Neur. Comput. 8: 895–938.

Pike FG, Meredith RM, Olding AWA, Paulsen O (1999) Postsynaptic bursting is essential for "Hebbian" induction of associative long-term potentiation at excitatory synapses in rat hippocampus. J. Phys. 518: 571–576.

Au: Cambridge MA ? or Eng ?

Rumelhart D, McCleland J (1986) Parallel Distributed Processing. Bradford Books, Cambridge.

Salin PA, Bullier J (1995) Corticocortical connections in the visual system: Structure and function. Physiol. Rev 75: 107–154.

Schäfer AT, Roth A, Sakmann B (2000) Morphological correlates of BAC-firing threshold in model layer 5 pyramidal neurons. Forum of European Neuroscience Brighton. Eur. J. Neurosci. Supp. 167.08.

Schiller J, Major G, Koester HJ, Schiller Y (2000) NMDA spikes in basal dendrites of cortical pyramidal neurons. Nature 404: 285–289.

Schiller J, Schiller Y, Stuart G, Sakmann B (1997) Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. J. Physiol. (Lond.) 505: 605–616.

Siegel M, Körding KP, König P (2000) Integrating bottom-up and top-down sensory processing by somato-dendritic interactions. J. Comput. Neurosci. 8: 161–173.

Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Ann. Rev. Neurosci. 18: 555–586.

Softky W (1994) Submillisecond coincidence detection in active dendritic trees. Neuroscience 58: 13–41.

Stewart BM, Sejnowski TJ (1998) Viewpoint invariant face representation from visual experience by temporal association. In: Wechsler et al., eds. Face Recognition: From Theory to Applications. Springer, Berlin.

Au:Kindly check

Stone JV, Bray AJ (1995). A learning rule for extracting spatio-temporal invariances. Network: Comput. Neural Syst. 6: 429–436.

Stuart GJ, Sakmann B (1994) Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. Nature 367: 69–72.

Stuart GJ, Schiller J, Sakmann B (1997) Action potential initiation and propagation in rat neocortical pyramidal neurons. J. Physiol. (Lond.) 505: 617–632.

Tesauro G (1990) Neural models of classical conditioning: A theoretical viewpoint. In: SJ, Hanson CR Olson, eds. Connectionist Modeling and Brain Function. MIT Press: Cambridge, MA.

Tsodyks MV, Markram H (1997) The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc. Natl. Acad. Sci. USA 94: 719–723.

Werbos PJ (1974/1994) The Roots of Backpropagation . Wiley, New York (includes Ph.D. thesis of Werbos P from 1974).

Zador AM, Agmon-Snir H, Segev I (1995) The morphoelectrotonic transform: A graphical approach to dendritic function. J. Neurosci. 15: 1669–1682.

Zeki S, Shipp S (1988) The functional logic of cortical connections. Nature 335: 311–317.

Zipser D, Andersen RA (1988) A backpropagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature 331: 679–684.

Zipser D, Rumelhart DE (1990) Neurobiological significance of new learning models. In: E Schwartz, ed. Computational Neuroscience MIT Press: Cambridge, MA, pp. 192–200.