ELSEVIER
Neurocomputing 38–40 (2001) 859–865
NEUROCOMPUTING
www.elsevier.com/locate/neucom

Learning in a neural network model in real time using real world stimuli

Manuel A. Sánchez-Montañés[a,b,*], Peter König[a], Paul F.M.J. Verschure[a]

[a]Institute of Neuroinformatics, ETH/University Zürich, Winterthurerstr. 190, 8057 Zürich, Switzerland
[b]E.T.S. de Informática, Universidad Autónoma de Madrid, Madrid 28049, Spain

Abstract

In this paper we present a model of the auditory system that is trained using real-world stimuli and running in real-time. The system consists of different sound sources, a microphone, an A/D board, a peripheral auditory system implemented in software and a central network of spiking neurons. The synapses formed by peripheral neurons on the central ones are subject to synaptic plasticity. We implemented a learning rule that depends on the precise temporal relation of pre- and post-synaptic action potentials. We demonstrate that this mechanism allows the development of receptive fields combining learning in real-time, learning with few stimulus presentations and robust learning in the presence of large imbalances in the probability of occurrence of individual stimuli. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Learning; Spiking neurons; Real time; Natural stimuli; Auditory system

1. Introduction

Research on neuronal networks has been very much motivated by the ability of these systems to learn from experience [1,3]. In existing biological systems, this property is paired with the ability to interact with the real world in real time. Although such a combination is most attractive for any type of application, their success has been limited. For this, several reasons can be identified. First, the bulk of the work on learning in neural networks is done using artificial stimuli. These are cleanly defined and often more symbolic than natural signals. Second, due to

* Corresponding author. E.T.S. de Informática, Universidad Autónoma de Madrid, Madrid 28049, Spain.
E-mail address: manuel.sanchez-montanes@ii.uam.es (M.A. Sánchez-Montañés).

PII: S0925-2312(01)00357-5

limitations in hardware and/or software, simulations of neural networks are often far from real-time, and thus not suitable to be used in real world tasks. Furthermore, learning rules tend to be slow, often needing thousands of stimulus presentations. This combination makes real-time applications very difficult. Third, learning in neuronal networks in general use well balanced stimulus sets to avoid any instability of the dynamics of learning. This property cannot be guaranteed in real-world applications. These three issues have to be addressed if neural networks shall live up to the expectations put into the field.

Here we investigate a learning rule inspired by recent biological results [13,10]. It allows extremely high learning rates and is simultaneously very robust to inhomogeneities of the stimulus set [8,11]. It is implemented in a model of the auditory system using a high performance distributed simulation environment, IQR421 [16], computed in real time and trained with real world stimuli on-line. Furthermore, we demonstrate that a global mechanism modeled after the action of the basal forebrain couples seamlessly into the local learning rule and adds flexibility in emphasizing important stimuli.

2. Methods

Sounds are generated either by a computer controlled synthesizer or a CD-player. Using a microphone (ME64, Sennheiser, Wedemark, Germany) the analog signals are sampled at 44.1 kHz and digitized with 16 bit resolution on an interface card (Soundblaster, Creative Technology Ltd, Singapore, Singapore). On each block of 1024 sampled signals a digital FFT is computed. Input to the model is provided by the absolute values of the first 128 FFT coefficients.

The neural network is a very rough sketch of the mammalian auditory system and includes five sets of integrate and fire neurons: an input population, a thalamic population, cortical excitatory and inhibitory neurons and an additional neuron representing the basal forebrain. All neurons are simulated in strict real time, i.e. simulated biological time matches 1:1 spent physical compute time. The excitation of neurons in the input population is directly determined by the absolute values of the respective Fourier coefficients. Projections of 3 units of this set converge on each neuron in the thalamic population. These synapses are subject to short-term depression [15]. Our model of the cortical circuit consists of populations of excitatory principal neurons and inhibitory interneurons where the excitatory ones project one-to-one to the inhibitory neurons. The excitatory neurons receive inhibitory input from the interneurons and excitatory input from all neurons of the thalamic population. To model the context of a larger network, the population of cortical excitatory neurons receives an additional excitatory input from a noise source firing at 10 Hz with a Poisson distribution. The one neuron representing basal forebrain activity injects a hyperpolarizing current into the inhibitory neurons, effectively delaying their activity relative to the excitatory neurons by a few milliseconds [4,11].

The synaptic strength of the thalamic projections to the excitatory cortical neurons evolves according to a modification of a recently proposed learning rule [8,11], which utilizes a backpropagating action potential. First, when this backpropagating action potential arrives at a synapse nearly simultaneously with pre-synaptic activity the efficacy of the synapse is increased [5,10,9,2]. Second, if the backpropagating action potential coincides with an afferent action potential, but is attenuated by inhibitory input [12,14], the efficacy of the respective excitatory synapse is decreased. Third, in case of non-attenuated backpropagating action potentials which do not coincide with pre-synaptic activity the synaptic efficacy is also decreased. Thus, in this learning rule the changes of synaptic efficacy are crucially dependent on the temporal dynamics in the neuronal network. Those neurons which are activated strongest and fire earliest [7] prevent other neurons from learning the same stimuli due to the recurrent inhibition they cause.

3. Results

In the first experiment we used a commercial CD ("Cabo do Mundo" by Luar na Lubre, Warner Music Spain, 1999) for training the network. Due to the short-term depression in the connections from the input neurons to the thalamic neurons, not the absolute loudness but the fast dynamics of the different frequency components dominates the dynamics of the cortical neurons. The effects of the learning rule onto the thalamo–cortical synapses can be described in several stages. Due to the initial homogeneous connectivity between thalamic and excitatory cortical neurons, most excitatory neurons are active, resulting in a high level of inhibition in the network. This inhibition, however, leads to an attenuation of most backpropagating action potentials within the excitatory neurons and, thus, to a depression of thalamo–cortical synapses (Fig. 1, left, 0–200 s). With the decrease of the activity level, inhibition is reduced as well, and some synapses are potentiated, leading to the formation of well defined receptive fields. After 30 min most neurons have stable receptive fields, which do not change anymore after prolonged stimulation of 2.5 h (Fig. 1, right).

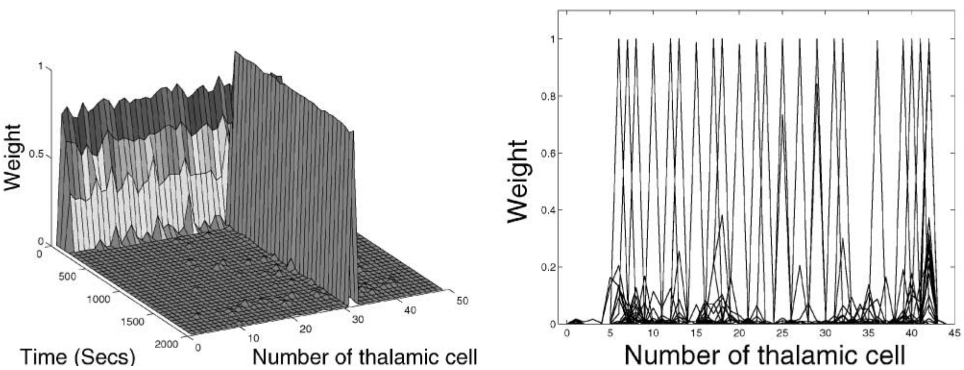

Fig. 1. (Left), evolution of the receptive field of one of the cortical excitatory cells when the CD was played. (Right), superposition of the final receptive fields of every cortical excitatory neuron.

Due to the competition between the excitatory neurons via the inhibitory neurons on the level of synaptic plasticity, most of the frequency spectrum is covered. Thus the described learning rule does not only lead to a formation of stable receptive fields, but also to an even coverage of the whole range of input stimuli.

In a second experiment, we investigated the effect of basal forebrain activation. For this purpose we used a digital synthesizer (QS8, Alesis, Santa Monica, USA). The stimuli were sinusoidal tones with frequencies of 0.74, 1.05, 1.48, 2.09 and 2.96 kHz. The stimulation consisted of a pseudo-random sequence of these tones, each presented for 800 ms with a probability of occurrence of $\frac{1}{2}$, $\frac{1}{8}$, $\frac{1}{8}$, $\frac{1}{8}$, and $\frac{1}{8}$, respectively. As observed in the first experiment, nearly all excitatory neurons respond initially. However, after a few presentations the number of neurons which respond to a particular stimulus stabilizes (Fig. 2, left). Furthermore, the size of the representation of each tone does not depend on its probability of occurrence (Fig. 2, left). Thus, the learning rule is robust and can handle inhomogeneities in the occurrence of different stimuli.

As a next step, comparable to recent physiological experiments [6] we paired one of the rare stimuli (2.09 kHz) with the activation of the basal forebrain unit. The basal forebrain input hyperpolarizes the inhibitory cortical neurons, delaying their activity with respect to the excitatory neurons by about 6 ms and, thus, effectively enlarging the temporal window for the backpropagating action potential to induce the potentiation of synaptic efficacies. The representation of this stimulus is now much increased (Fig. 2, right). This effect is independent of the presentation frequency of the stimulus and does not affect the size of the representation of the other stimuli. If pairing is discontinued after presentation 22, the size of the representation of the previously paired tone is reduced and reaches a size comparable to the representation of the other tones (Fig 2, right). Thus, the learning rule allows to dynamically modify the "importance" of stimuli independently of their probability of occurrence.

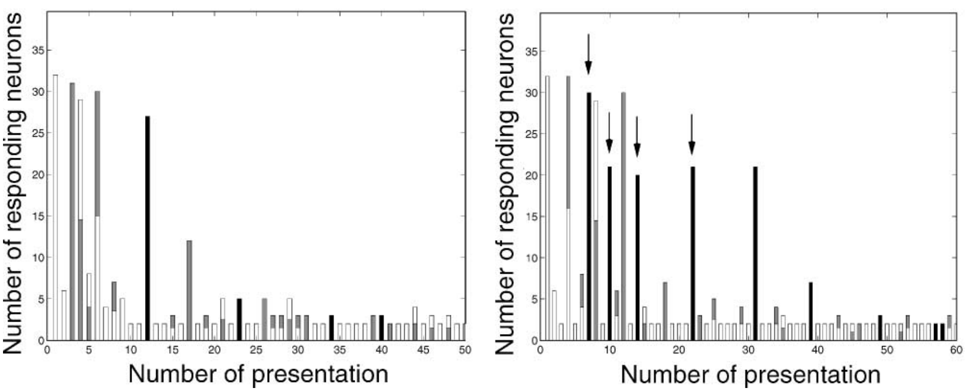

Fig. 2. Number of cortical excitatory neurons responding to the tones (each bar style corresponds to one of the 5 different tones used). (Left), no stimulus is paired with basal forebrain. (Right), stimuli paired with basal forebrain are labeled by arrows.

4. Conclusions

In this study we investigate the properties of a real-time implementation of a biophysically realistic learning rule using real world stimuli. Within the framework of a model of the mammalian auditory system we investigate a single-integrated learning mechanism which combines a local learning rule with a global gating mechanism. We show that this model supports continuous and fast learning, provides an even coverage of stimulus space, and generates stable representations combined with the flexibility to change representations in relation to task requirements. This is in good accord with our previous results using computer simulations of biological neural networks [8,11].

In implementing our model we made some simplifications which are not critical to the presented results, but could become important in further extensions. First, we are using chunked sampling (23 ms), keeping stimulation of the input neurons constant during that time. This limits the phase information available, and hence does not allow usual approaches to localization of sound sources. This problem may be addressed by shortening the length of data samples, or using partly overlapping samples or a wavelet based analysis. A second simplification is that we do not try to replicate the huge dynamic range of the auditory system (about 100 dB). A possible solution to this problem is the use of adapting gains in the neurons. Due to limitations in hardware and/or software, simulations of learning in neural networks are often far from real-time, and thus not suitable for real-world tasks. Here we use “standard” hardware based on Pentium III processors (500 MHz) connected by a TCP/IP network to demonstrate a real-time implementation of a biophysically realistic neural network model. The computational efficiency is provided by the IQR421 software package [16] allowing easy parallelization of large simulations in such a cluster of computers. Larger systems might be implemented by a coarse grained parallel implementation on a larger number of processors, as well as reducing computational and communication load by omitting full connectivity and imposing a rough topology of connections between the different pools of neurons. Further optimizations can be provided by using FPGAs or aVLSI devices. Currently we are experimenting with replacing the digital FFT with an aVLSI model of the cochlea.

Acknowledgements

We would like to thank Peter Paschke for his help in developing the data acquisition system and the FFT analysis module. We are also grateful to Andre van Schaik, David Klein and Fernando Corbacho for helpful discussions. This work was supported by SNF (grant 31-51059.97, awarded to P.K.), SPP Neuroinformatics (grants 5002-44888/2&3 to P.F.M.J.V.) and MCyT (grant BFI2000-0157 to M.A.S.-M.). The work was partly done at the Telluride Workshop on Neuromorphic Engineering (1999).

References

[1] D.L. Alkon, D.G. Amaral, M.F. Bear, J. Black, T.J. Carew, N.J. Cohen, J.F. Disterhoft, H. Eichenbaum, S. Golski, L.K. Gorman, G. Lynch, B.L. Mcnaughton, M. Mishkin, J.R. Moyer, J.L. Olds, D.S. Olton, T. Otto, L.R. Squire, U. Staubli, L.T. Thompson, C. Wible, Learning and memory, Brain Res. Rev. (1991) 193–220.

[2] G.Q. Bi, M.M. Poo, Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type, J. Neurosci. 18 (1998) 10 464–10 472.

[3] D.V. Buonomano, M.M. Merzenich, Cortical plasticity: From synapses to maps, Annu. Rev. Neurosci. 21 (1998) 149–186.

[4] T.F. Freund, A.I. Gulyas, GABAergic interneurons containing calbindin D28K or somatostatin are major targets of GABAergic basal forebrain afferents in the rat neocortex, J. Comp. Neurol. 314 (1) (1991) 187–199.

[5] W. Gerstner, R. Ritz, J.L. van Hemmen, Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns, Biol. Cybern. 69 (1993) 503–515.

[6] M.P. Kilgard, M.M. Merzenich, Cortical map reorganization enabled by nucleus basalis activity, Science 279 (1998) 1714–1718.

[7] P. König, A.K. Engel, P.R. Roelfsema, W. Singer, How precise is neuronal synchronization? Neural Comput. 7 (1995) 469–485.

[8] K.P. Körding, P. König, A learning rule for dynamic recruitment and decorrelation, Neural Networks 13 (2000) 1–9.

[9] J. Magee, D. Hoffman, C. Colbert, D. Johnston, Electrical and calcium signaling in dendrites of hippocampal pyramidal neurons, Annu. Rev. Physiol. 60 (1998) 327–346.

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

[11] M.A. Sánchez-Montañés, P.F.M.J. Verschure, P. König, Local and global gating of synaptic plasticity, Neural Comput. 12 (3) (2000) 519–529.

[12] N. Spruston, Y. Schiller, G. Stuart, B. Sakmann, Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites, Science 268 (1995) 297–300.

[13] G.J. Stuart, B. Sakmann, Active propagation of somatic action potentials into neocortical pyramidal cell dendrites, Nature 367 (1994) 69–72.

[14] H. Tsubokawa, W.N. Ross, IPSPs modulate spike backpropagation and associated Ca_i^{2+} changes in the dendrites of hippocampal CA1 pyramidal neurons, J. Neurophysiol. 76 (5) (1996) 2896–2906.

[15] J.A. Varela, K. Sen, J. Gibson, J. Fost, L.F. Abbott, S.B. Nelson, A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex, J. Neurosci. 17 (1997) 7926–7940.

[16] P.F.M.J. Verschure, Xmorph, Internal Report, Institute of Neuroinformatics, ETH-UZ, 1997.

Manuel A. Sánchez-Montañés received a B.Sc. degree (with honors) in Physics from the Universidad Complutense de Madrid, Spain (1997). Currently he is a Ph.D. student at the Universidad Autónoma de Madrid. In 1998 and 1999 he visited the Institute of Neuroinformatics in Zürich joining Dr. Verschure's group. His main interests are learning and adaptation in biological systems, and the validation of these concepts in autonomous artificial systems.

Peter König studied physics and medicine at the University in Bonn/Germany. Working from 1987 to 1990 in the Department of Neurophysiology at the Max-Planck-Institute for Brain Research in Frankfurt/Germany he received the habilitation at the medical faculty of the University. After working as a senior fellow at the Neurosciences Institute in San Diego, he joined the Institute of Neuroinformatics in Zürich in 1997. Here he is using experimental and theoretical approaches to study the mammalian visual system, with a particular interest in synchronization neuronal activity, the role of top–down signals and their interaction with learning and synaptic plasticity.

Dr. Paul F.M.J. Verschure is group leader at the Institute of Neuroinformatics ETH-University Zurich, Switzerland. He received both his M.A. and Ph.D. in psychology. He works on biologically realistic models of perception, learning, and problem solving which are applied to robots and on the software tools which support this research, IQR421. In addition, he applies the methods and models he, and his collaborators, develop in the domain of art and technology.