

Learning Sensory Maps With Real-World Stimuli in Real Time Using a Biophysically Realistic Learning Rule

Manuel A. Sánchez-Montañés, Peter König, and Paul F. M. J. Verschure

Abstract—We present a real-time model of learning in the auditory cortex that is trained using real-world stimuli. The system consists of a peripheral and a central cortical network of spiking neurons. The synapses formed by peripheral neurons on the central ones are subject to synaptic plasticity. We implemented a biophysically realistic learning rule that depends on the precise temporal relation of pre- and postsynaptic action potentials. We demonstrate that this biologically realistic real-time neuronal system forms stable receptive fields that accurately reflect the spectral content of the input signals and that the size of these representations can be biased by global signals acting on the local learning mechanism. In addition, we show that this learning mechanism shows fast acquisition and is robust in the presence of large imbalances in the probability of occurrence of individual stimuli and noise.

Index Terms—Auditory system, learning, natural stimuli, real time, real world, spiking neurons.

I. INTRODUCTION

OVER the past years neuroscientists have gained insight in the neural mechanisms responsible for the ability of learning and adaptation in biological systems [1], [11]. The substrate of learning in these systems is thought to be provided by the mechanisms which regulate the change of synaptic efficacies of the connections among neurons [31], [43]. In his seminal work Hebb proposed that neurons which are consistently coactivated strengthen their coupling [20] and form associative networks. Since then many experiments have addressed different mechanisms which regulate changes in synaptic efficacies dependent on specific properties of pre- and postsynaptic activity [9], [11]. Based on these experiments, a number of Hebbian learning rules have been proposed with different desirable properties [8], [10], [14], [36], [39].

These learning rules have been considered physiologically realistic when they only rely on signals which are available to

the synapse local in time and space. However, recent physiological results on neurons in mammalian cortex give a richer picture. These studies demonstrate, first, that an action potential triggered at the axon hillock propagates not only anterogradely along the axon, but also retrogradely through the dendrites [40], [12]. Second, on its way into the dendrite the action potential may be attenuated or blocked by inhibitory input from other neurons [38], [45]. Third, it has been demonstrated that these backpropagating action potentials directly affect mechanisms regulating synaptic plasticity [30] which depends on postsynaptic calcium dynamics [27]. In addition, the dramatic effect of even single inhibitory inputs on the calcium dynamics in the dendritic tree, in particular in its apical compartments, suggests that regulation of synaptic plasticity can be strongly influenced by inhibitory inputs [28]. Thus, the backpropagating action potential can make information on the output of the neuron available locally at each of its afferent synapses, and inhibitory inputs onto a neuron can in turn regulate the effectiveness of this signal.

The above described mechanism makes a change in synaptic efficacy dependent on the temporal relation between pre- and postsynaptic activity. In particular, it will be strongly affected by the temporal relation between the inhibition and excitation a neuron receives and its own activity. Neurons which fire with the shortest latency to a stimulus will receive inhibition after they have generated backpropagating action potentials. In this case active synapses can be potentiated [28]. In contrast, neurons which fire late to a stimulus would receive inhibition before they have generated a spike. Their backpropagating action potentials are modulated by this inhibition preventing potentiation of their active synapses. This dynamic seems to be reflected in the physiology of the visual system where the optimality of the tuning of a neuron seems to be directly reflected in its response latency to a stimulus [24]. Given the above mechanism this would imply that the optimally tuned neurons prevent further learning by other neurons in the map.

Synaptic plasticity, however, is not only dependent on the dynamics of the local network but also on modulatory signals [2] arising from subcortical structures. For instance, it has been shown that cholinergic and gabaergic neurons in the basal forebrain, which project to many areas including the cerebral cortex, can strongly regulate mechanisms of synaptic plasticity [50]. These results were obtained in classical conditioning experiments where tones were paired with aversive stimuli such as a footshock [51]. Subsequently it was shown that the aversive stimulus could be replaced by direct stimulation of the basal

Manuscript received September 20, 2000; revised July 23, 2001 and January 3, 2002. This work was supported by SPP Neuroinformatics and SNF (Grant 31-51059.97, awarded to P. König). The work of M. A. Sánchez-Montañés was supported by a FPU grant from MEC and Grant BFI2000-0157 from MCYT (Spain). This work was done in part at the Telluride Workshop on Neuromorphic Engineering (1999).

M. A. Sánchez-Montañés is with the Institute of Neuroinformatics, ETH/University, 8057 Zürich, Switzerland, and also with the Grupo de Neurocomputación Biológica, ETS de Informática, Universidad Autónoma de Madrid, 28049 Madrid, Spain.

P. König and P. F. M. J. Verschure are with the Institute of Neuroinformatics, ETH/University, 8057 Zürich, Switzerland.

Publisher Item Identifier S 1045-9227(02)02405-0.

forebrain [23]. In these latter experiments it was shown that more neurons in the primary auditory cortex would respond to the reinforced frequency while the representation of the others was not increased.

These changes are retained indefinitely [50] due to long-term changes in the neural circuit. However, synaptic efficacies are also regulated by mechanisms with short-term dynamics, operating on a time scale of hundreds of milliseconds [46], [41], [42]. Two different types of such a dynamics can be distinguished: facilitation and depression. These mechanisms enhance or depress the efficacy of a synapse following the occurrence of a presynaptic spike. Both effects can occur simultaneously at a synapse affecting its band pass properties [44]. For instance, short-term depression of a synapse can act as a high-pass filter blocking the dc component of a presynaptic signal [3].

In previous work we have shown, using computer simulations, how some of these biological mechanisms can support learning allowing extremely high learning rates and robustness to inhomogeneities of the stimulus set [26], [34]. In addition, our models allow for the combination of local learning mechanisms, supporting the development of sensory maps, with global signals that convey information of the behavioral significance of the stimuli [34].

However, the brain works in the real-world, properties of which are difficult to capture in simulations. In addition, in simulation studies models are not exposed to the real-time constraints imposed on the brain. Hence, our models should be evaluated under similar constraints. Here we investigate the properties of a biophysically realistic real-time model of learning in the auditory cortex which is tested using real-world stimuli. In particular we address the ability of this learning mechanism to develop sensory maps of auditory stimuli.

II. METHODS

A. Hardware Setup

All experiments are conducted in a standard office environment with a room size of about 30 m². The analog audio signals are sampled using a microphone (ME64, Sennheiser, Wedemark, Germany) 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 fast Fourier transform (DFT) is computed [17]. Input to the model is provided by the absolute values of the first 128 FFT coefficients. The whole system for the control of the setup, the stimulus generation protocol, the simulation, and data acquisition is defined within the distributed neural simulation environment IQR421 [47] using three Pentium III 450 MHz PCs [Fig. 1(a)].

B. The Network

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 [Fig. 1(b)]. All neurons are simu-

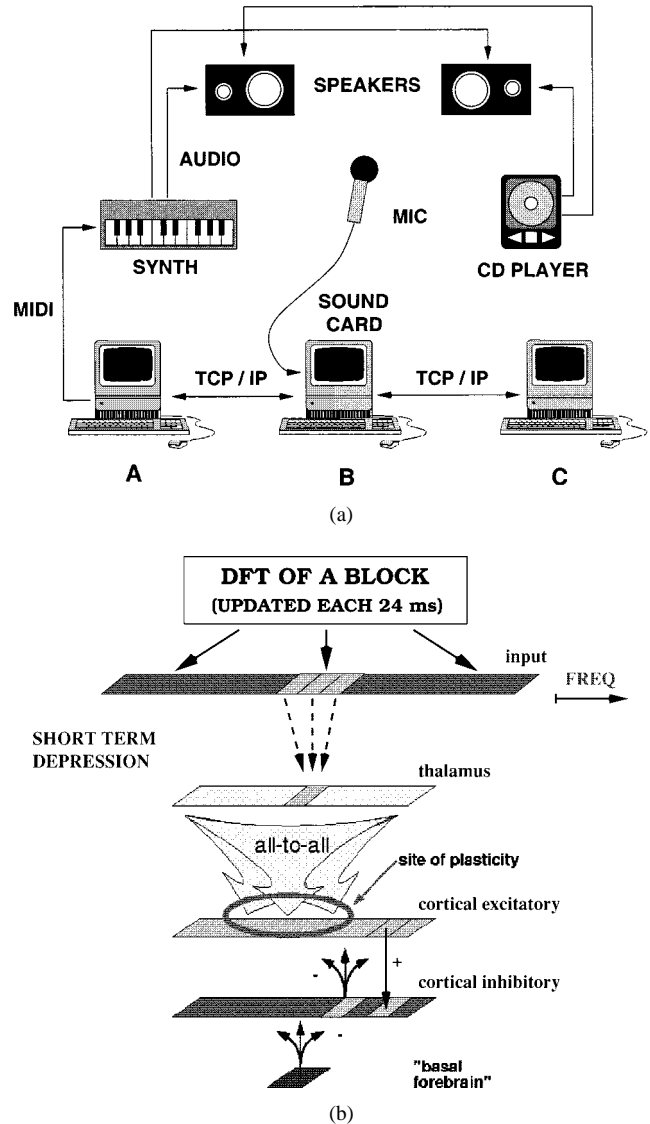


Fig. 1. (a) Schema of the hardware. The sounds are generated by a pair of speakers that receive audio signals either from a synthesizer which is controlled by a computer using the MIDI protocol (A) or from a CD player. The microphone sends the audio signal to the soundcard installed in computer B; this computer calculates the FFT and sends this data to C, where the neural model is implemented. Each computer communicates to the others through TCP/IP (connections between computers A and C are omitted for clarity). The whole system is controlled by the distributed neural simulation environment IQR421. (b) Schema of the neural model. Each one of the first 128 FFT coefficients excites its corresponding input neuron. The output of each group of three of these neurons converges into a thalamic neuron through short-term depressing synapses. Each cortical excitatory neuron receives excitation from the whole thalamic population. These synapses are subject to long-term synaptic plasticity. Each cortical excitatory neuron is connected to a cortical inhibitory neuron which sends back inhibition to the whole excitatory population. Finally, the unit representing the basal forebrain activity sends inhibition to the cortical inhibitory population. See text for further details.

lated in strict real time, i.e., simulated biological time matches 1 : 1 spent physical time. The dynamics of the membrane potential of neuron j in population a , $V_j^a(t)$, is defined as

$$\dot{V}_j^a(t) = -\tau^{-1}V_j^a(t) + C^{-1} \sum_{b,i} I_{i \rightarrow j}^{b \rightarrow a}(t) - C^{-1} \Delta Q_{sp} \delta(t - t_{spike_j}) \quad (1)$$

TABLE I
PARAMETERS OF THE DIFFERENT POPULATIONS. “BF”: BASAL FOREBRAIN.
“ARP”: ABSOLUTE REFRACTORY PERIOD

| | Input | Thalamic | Cortical excitatory | Cortical inhibitory | BF |
|-------------|-------|----------|---------------------|---------------------|----|
| size | 128 | 43 | 36 | 36 | 1 |
| τ (ms) | 19 | 19 | 19 | 19 | 19 |
| ARP (ms) | 6 | 10 | 10 | 6 | 8 |

where τ is the time constant; C is the membrane capacitance; $I_{i \rightarrow j}^{b \rightarrow a}(t)$ represents the current injected by the synapse from neuron i in population b . In case V is greater than the threshold (V_{Th}) the neuron emits a spike (t_{spike} is current time) and the membrane potential is reset to zero by injecting the charge ΔQ_{sp} instantaneously (represented by the Dirac delta function, δ). The dynamics of the membrane potential also includes an absolute refractory period (see Table I for concrete values for each population).

The dynamics of the synapses formed by neurons of population b with neurons of population a are modeled using a first-order approximation

$$\dot{I}_{i \rightarrow j}^{b \rightarrow a}(t) = -\tau^{-1} I_{i \rightarrow j}^{b \rightarrow a}(t) + \gamma^{b \rightarrow a} w_{i \rightarrow j}^{b \rightarrow a} \Gamma_i^{b \rightarrow a} \delta(t - t_{spike_i}) \quad (2)$$

where $\gamma^{b \rightarrow a}$ is a constant gain factor that defines the type of connection (positive for excitatory, negative for inhibitory) and its maximum gain; $w_{i \rightarrow j}^{b \rightarrow a}$ and $\Gamma_i^{b \rightarrow a}$ are two variables ranging from zero to one that express the long-term and short-term efficacies of the synapse, respectively. There are three types of connections in the model: 1) nonplastic (w and Γ are both constant and equal to one); 2) subject to short-term plasticity; and 3) subject to long-term plasticity. A transmission delay of 2 ms is taken into account in all the connections. In Table II the numerical details for each type of connection are given.

Each input neuron receives an excitation proportional to the absolute value of its respective analog Fourier coefficient (frequencies up to a quarter of the Nyquist frequency, 5.5 kHz). Thus we simulate in a first approximation the spectral decomposition achieved by the cochlea and subcortical nuclei [22]. The spectra of all the sounds used in the experiments are kept in this range. Because it takes 23 ms to sample a block of audio signals, the input to these neurons is updated every 24 ms and kept constant until the next update. The time spent in DFT computation is very low compared to the sampling time.

Each thalamic neuron receives excitation from three input neurons in a tonotopic manner Fig. 1(b). This convergence of information allows to process a broad frequency band with a reduced number of neurons, making real-time processing possible. The details of this connectivity, however, are not critical to the performance of the model.

The synapses from input neurons are subject to short-term depression [46], making the efficacy of the synapse dependent on previous presynaptic activity

$$\dot{\Gamma}_i = \tau_d^{-1}(1 - \Gamma_i) - f \Gamma_i \delta(t - t_{spike_i}). \quad (3)$$

τ_d defines the recuperation time of the synapse (4 s). f defines the speed of adaptation, being 0.1.

Each cortical excitatory neuron receives excitatory input from all thalamic neurons and in turn projects to one cortical inhibitory neuron. All cortical inhibitory neurons project to all

TABLE II
PARAMETERS OF THE CONNECTIONS BETWEEN POPULATIONS. “C. EXCITATORY”: CORTICAL EXCITATORY. “C. INHIBITORY”: CORTICAL INHIBITORY. “BF”: BASAL FOREBRAIN. THE CONNECTION STRENGTH, $C^{-1}\gamma$, IS GIVEN IN UNITS OF THE POSTSYNAPTIC THRESHOLD V_{Th}

| Connection | Plasticity | Connections | $C^{-1}\gamma$ | τ (ms) |
|---|------------|-------------|----------------|-------------|
| Input \rightarrow Thalamus | Short-term | 3 to 1 | .15 | 19 |
| Thalamus \rightarrow C. excitatory | Long-term | all to all | .1 | 19 |
| C. excitatory \rightarrow C. inhibitory | No | 1 to 1 | 1 | 19 |
| C. inhibitory \rightarrow C. excitatory | No | 1 to all | -.0025 | 19 |
| BF \rightarrow C. inhibitory | No | 1 to all | -.06 | 19 |

cortical excitatory neurons. The synaptic strengths $w_{i \rightarrow j}$ of the connections from thalamic neurons to cortical excitatory neurons are initially random, with values between 0.7 and 0.8 (homogeneous distribution); therefore, the receptive fields of the excitatory cortical neurons are initially diffuse. These synapses are subject to long-term synaptic plasticity (see Section II-C). To model the context of a larger network, we added an independent excitatory input to each cortical neuron which is firing at 10 Hz following a Poisson distribution. Finally, the one neuron representing basal forebrain activity sends inhibitory connections to the cortical inhibitory neurons [15], [16].

C. Learning Dynamics

The synaptic strength of the thalamic projections to the cortical excitatory neurons evolves according to a modification of a recently proposed learning rule [26], [34].

- 1) When the backpropagating action potential and the presynaptic action potential arrive within a temporal association window \bar{W} (i.e., the absolute value of the time difference between the two events is smaller than $\bar{W} = 20$ ms), the efficacy of the respective synapse is increased [7], [18], [30], [29]

$$\Delta w_{ij} = \alpha \frac{t_0}{t_0 + |t_i - t_j|}. \quad (4)$$

t_i is the time when the postsynaptic cell fires and t_j is the time when the action potential of the presynaptic cell arrives at the synapse.

- 2) If the backpropagating action potential and the afferent action potential occur within the temporal association window \bar{W} , but the inhibitory input attenuates the backpropagating action potential [38], [45], the efficacy of the respective excitatory synapse is decreased

$$\Delta w_{ij} = -\beta \frac{t_0}{t_0 + |t_i - t_j|}. \quad (5)$$

- 3) In case of nonattenuated backpropagating action potentials which do not coincide with presynaptic activity, synaptic efficiency decreases with a constant amount

$$\Delta w_{ij} = -\eta. \quad (6)$$

Thus, in this learning rule the changes of synaptic efficacy are crucially dependent on the temporal dynamics in the neuronal network. In our model we used the values $\alpha = 0.02$, $\beta = 0.005$, $\eta = 0.01$, $t_0 = 10$ ms. The weights are kept by saturation in the 0–1 range.

D. Training Protocol and Analysis

The network is trained with different types of acoustic stimuli. First, we use a commercial CD that is continuously played for 2.5 h. In this experiment the synaptic weights are sampled at intervals of 20 s for further analysis. In the second set of experiments we use a music synthesizer (QS8, Alesis, Santa Monica, CA) for generating the stimuli. Simple sinusoids are played for a few minutes either together with continuous low-band noise or without. The noise is obtained by passing white noise through a low-pass linear filter with a cutoff frequency of 600 Hz. Network activity, synaptic weights, sound frequency, and amplitude are continuously recorded for further analysis. All the parameters of the model are kept constant over all experiments and the learning mechanism is continuously active. Data analysis is performed using a commercial software package (MatLab, MathWorks, Natick, MA).

III. RESULTS

A. Development of Specific Receptive Fields Presenting Natural Stimuli

In the real world events do not occur in isolation but are combined in a variety of ways. In the first experiment we assess whether our model is able to develop specific and stable representations under these circumstances. The initial weights of the synapses from thalamic neurons to cortical excitatory neurons are randomly chosen in the range of 0.7–0.8 [Fig. 2(a)]; this makes the initial receptive field of all the cortical excitatory neurons diffuse and no knowledge about the stimuli is put into the network. The network is exposed for 2.5 h to the music from the CD (“Cabo do Mundo” by Luar na Lubre, Warner Music Spain, 1999). The CD style is celtic music played with traditional instruments, vocals, drums and synthesizers. The CD is available worldwide by music stores such as Amazon.

In this period the learning mechanism continuously acts on the synaptic efficacies of the thalamo–cortical projections shaping the receptive fields of the cortical neurons.

Due to the short-term depression in the projection from the input neurons to the thalamic neurons, not the absolute intensity but the fast dynamics of the different frequency components is transmitted to the cortical neurons. However, due to the initial homogeneous connections from thalamic neurons to cortical excitatory neurons, most of these excitatory neurons are active, resulting in a high level of inhibition in the network. This inhibition leads to an attenuation of most backpropagating action potentials within the excitatory neurons and, thus, to a depression of thalamo–cortical synapses [Fig. 2(b), 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 [Fig. 2(b), 200–500 s]. After 30 min most neurons have highly specific and stable receptive fields which practically cover the full frequency spectrum presented to the system. In addition, the different receptive fields provide a practically homogeneous coverage of the stimulus space [Fig. 2(c)].

The ability of the network to develop receptive fields which cover the full range of presented frequencies is the result of a competitive process. Neurons with a receptive field which

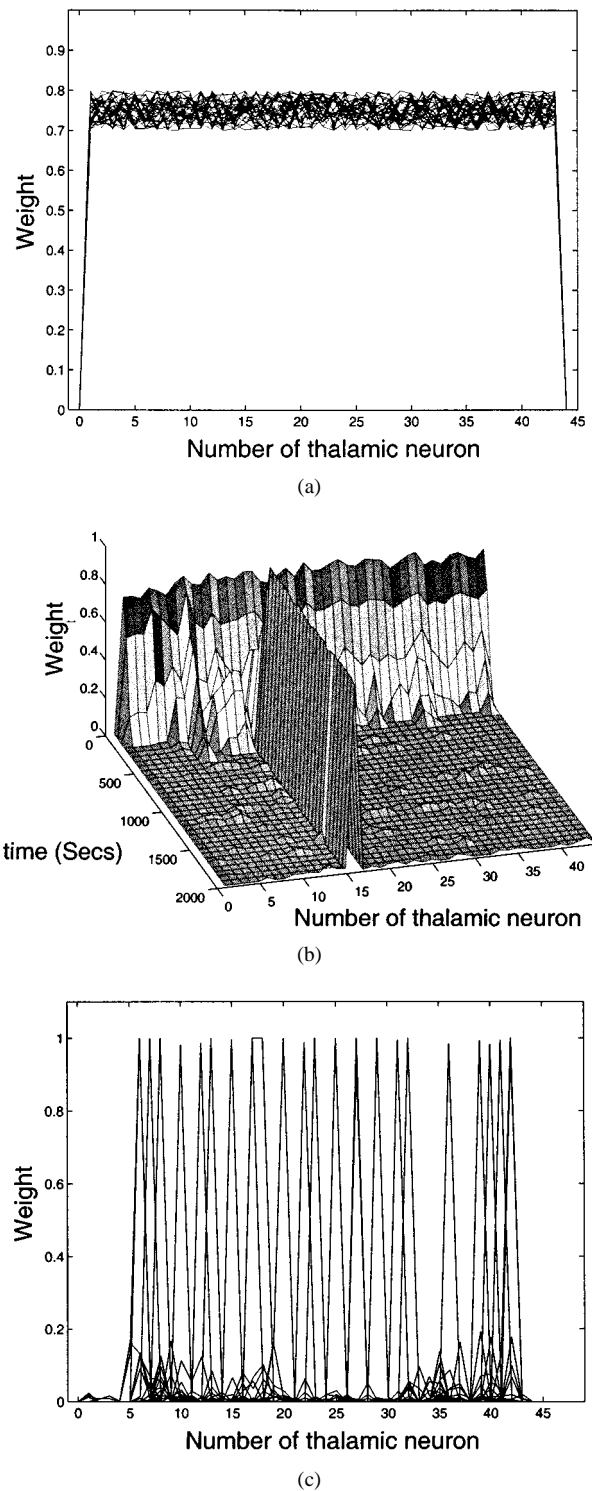


Fig. 2. Receptive field dynamics under continuous stimulation with music. (a) Superposition of the initial receptive fields of every cortical excitatory neuron. (b) Evolution of the receptive field of one of the cortical excitatory neurons. (c) Superposition of the final receptive fields of every cortical excitatory neuron after 2.5 h of stimulation.

is specific to the provided input respond with a short latency after stimulus onset. This in turn drives the inhibitory population rapidly, shunting the backpropagating actions potentials in those neurons which are not effectively representing the input, preventing a change in synaptic efficacy to occur in their afferents.

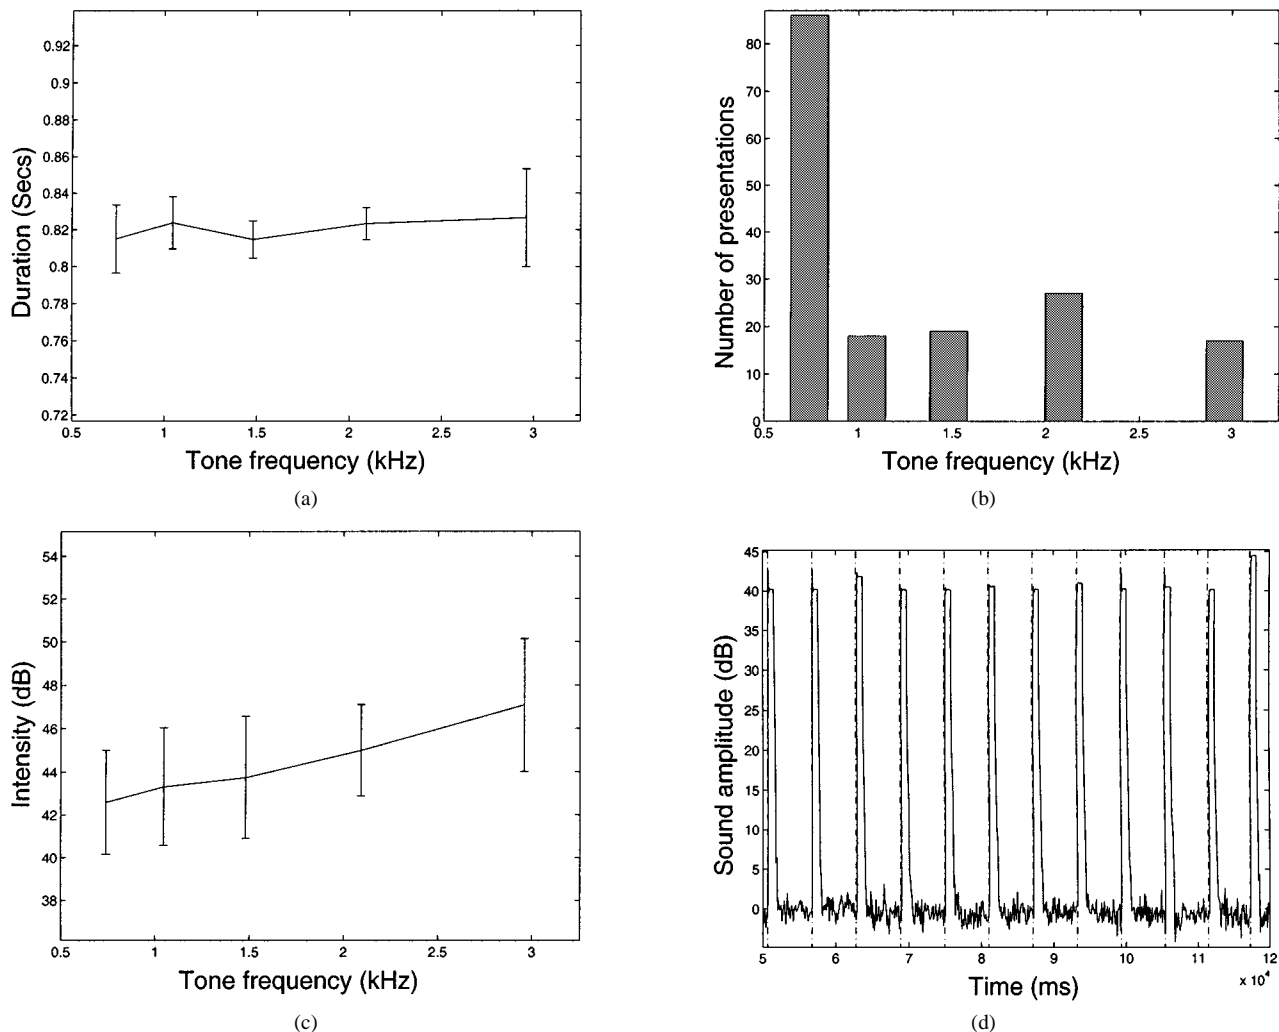


Fig. 3. Stimulus statistics using a pseudorandom sequence of five different tones (0.74, 1.05, 1.48, 2.09 and 2.96 kHz). The probability of occurrence is 1/2, 1/8, 1/8, 1/8, and 1/8, respectively. (a) Mean duration of each stimulus. (b) Number of presentations of each stimulus. (c) Mean intensity of each stimulus. The 0-dB level is chosen as the averaged level of noise in the room. (d) Sound amplitude over time.

These results demonstrate that this local learning mechanism allows single neurons to develop specific receptive fields within minutes, which are for realistic input conditions stable over hours. In addition, at the level of the network it allows the full range of inputs to be represented.

B. Dynamic Modulation of Representation Size

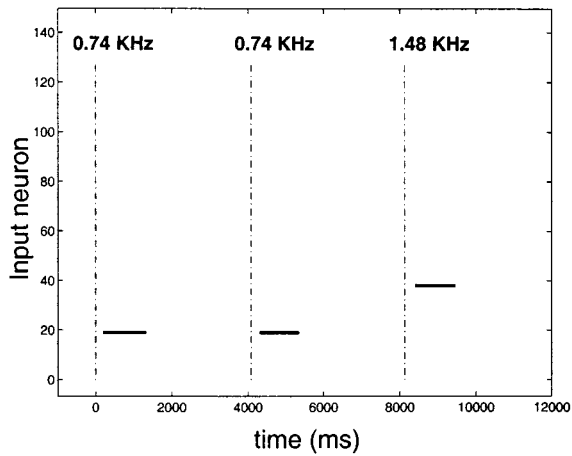
The brain uses global signals to provide information on the behavioral relevance of events. These signals can affect local mechanisms which govern changes in synaptic plasticity. An example of such a system is the basal forebrain, mentioned in the introduction. It was recently shown that the paired activation of this structure with a particular tone induces an enlargement of the representation of this tone in the primary auditory cortex [23]. This change in representation size does, however, not affect the size of other representations in the cortical map and the presentation of unpaired tones does not seem to affect the organization of this cortical area.

We investigate our model using an equivalent stimulation protocol. Sinusoidal tones with frequencies of 0.74, 1.05, 1.48, 2.09, and 2.96 kHz are generated on a digital synthesizer. These

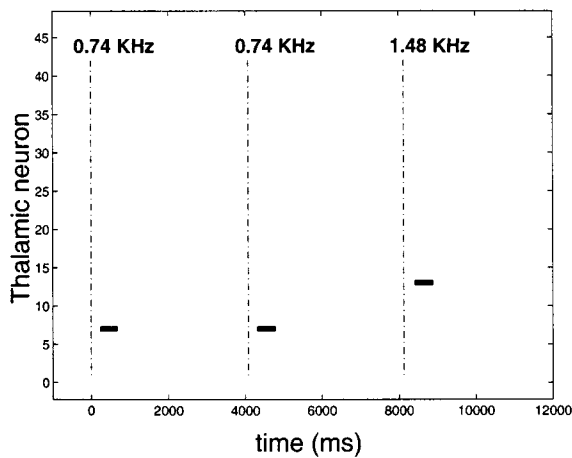
frequencies are presented in a pseudorandom order with an average duration of 0.8 s [Fig. 3(a)] and a probability of occurrence of 1/2, 1/8, 1/8, 1/8, and 1/8, respectively [Fig. 3(b)]. In these experiments the signal-to-noise ratio is above 30 dB [Fig. 3(c) and (d)]. High learning rates are used, $\alpha = 0.05$, $\beta = 0.4$, $\eta = 0.1$, in order to demonstrate the ability of the learning mechanism to learn with few stimulus presentations.

Fig. 4 shows a typical example of the responses in the network after the presentation of the sequence 0.74, 0.74, 1.48 kHz. When a tone is presented, typically 1–3 neurons fire in the input population and 1–2 neurons in the thalamic population, depending on the intensity of the sound. As observed in the first experiment, nearly all cortical excitatory neurons respond initially [Fig. 4(c), 0–1000 ms; Fig. 6(a)] to a novel stimulus. However, after a few presentations, the number of neurons which respond to this stimulus stabilizes [Fig. 6(a)].

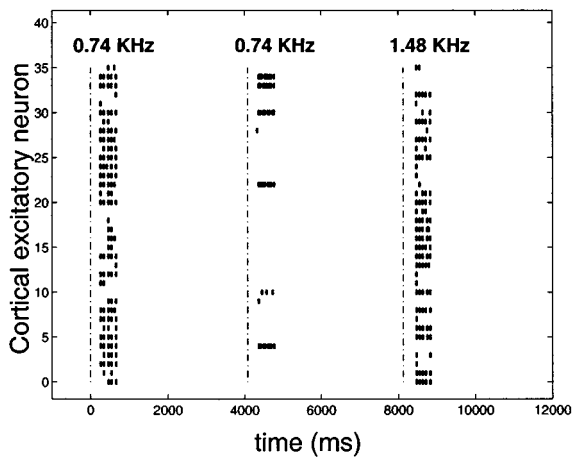
The developed receptive fields are specific: a neuron that responds to one tone does not respond to any of the others [Fig. 5(a) and (b)]. Furthermore, the size of the representation of each tone, i.e., the number of neurons responding to it, does not depend on its probability of occurrence [Fig. 6(a)]



(a)



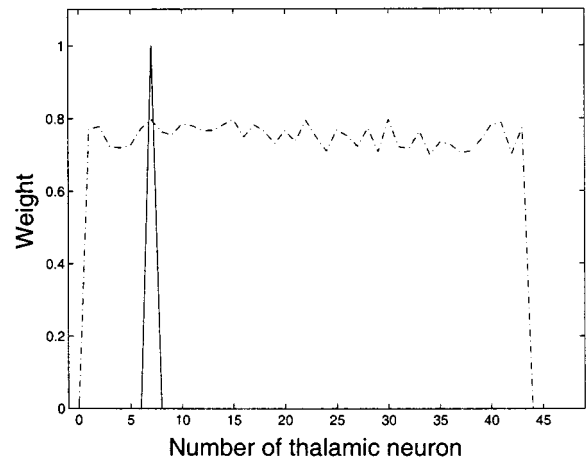
(b)



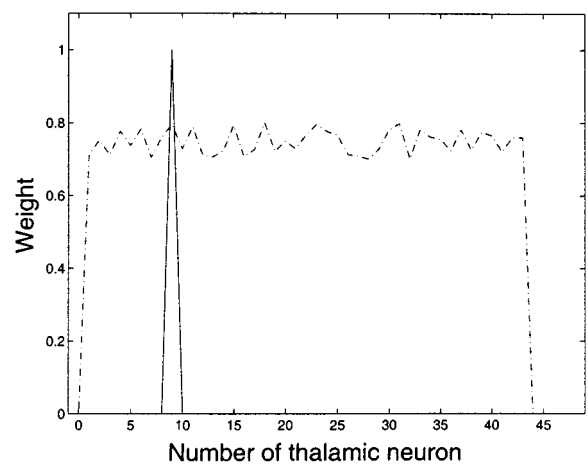
(c)

Fig. 4. Raster of network activity responding to three tones in sequence (0.74 kHz, 0.74 kHz, 1.48 kHz). Time zero corresponds to the onset of the first tone. Vertical dashed lines represent the onset of each tone. (a) Input population. (b) Thalamic population. (c) Cortical excitatory population.

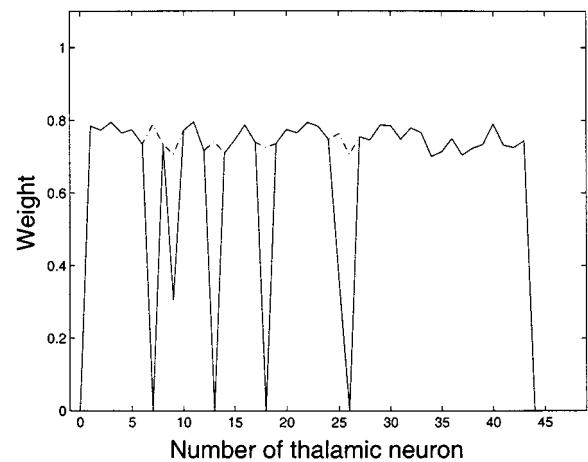
and (b)]. These results demonstrate that the learning rule is robust and can handle inhomogeneities in the occurrence of different stimuli. In addition, it shows the ability of the network for dynamic recruitment [26]. That is, those neurons that do not develop specific receptive fields remain “unspecific,” while



(a)



(b)



(c)

Fig. 5. Initial (dashed line) and final (solid line) receptive fields of some neurons. (a) Neuron finally selective to the 0.74 kHz tone. (b) Neuron finally selective to the 1.05 kHz tone. (c) Neuron that finally does not respond to any of the five tones used in the training. The final receptive field of each neuron is either selective to one tone (a), (b), or insensitive to any tone used in the training (c).

losing any sensitivity to frequencies represented by other neurons in the population [Fig. 5(c)]. These unspecific neurons can be activated by novel tones and develop receptive fields specific

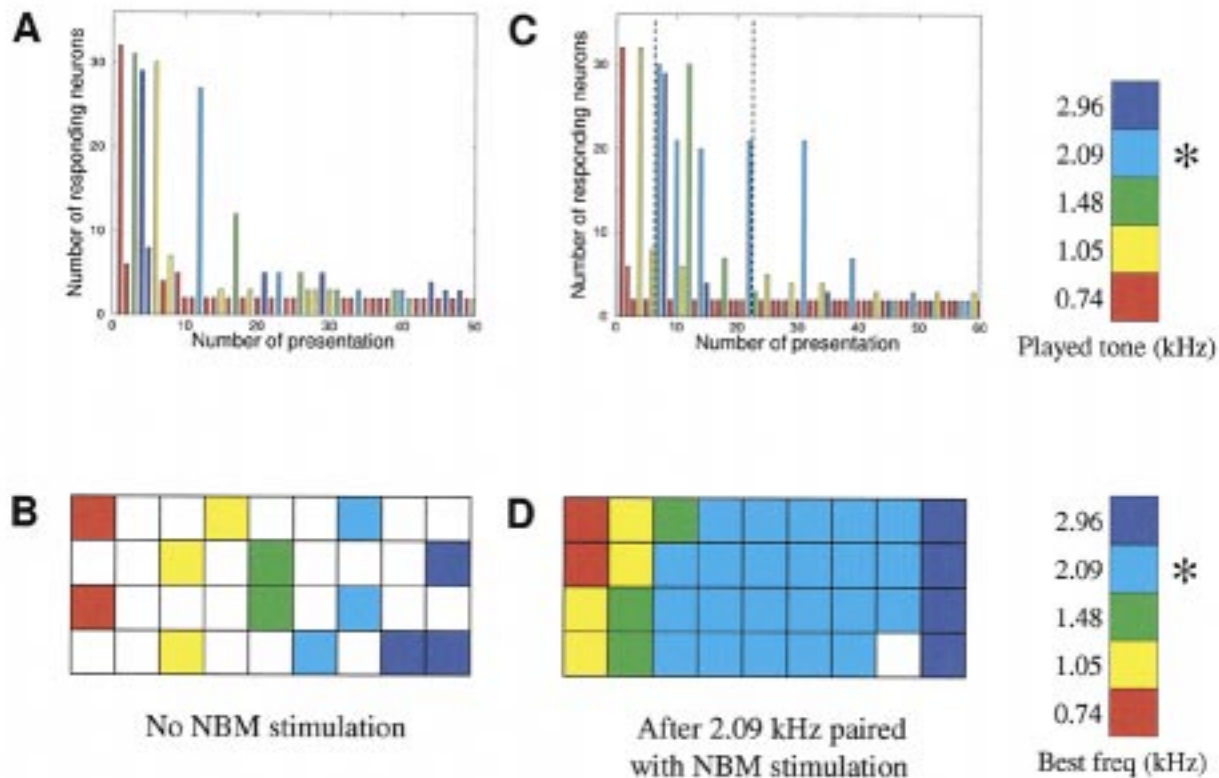


Fig. 6. Response of the cortical excitatory neurons during training. (A) Number of neurons responding to each tone in pseudorandom sequence of consecutive presentations. In each presentation a tone is randomly chosen from the set (0.74, 1.05, 1.48, 2.09, and 2.96 kHz) with a probability of (1/2, 1/8, 1/8, 1/8, 1/8), respectively. The color of a bar indicates which tone is presented and its height represents the number of neurons which respond to it. (B) Distribution of the preferred frequency of the 36 cortical excitatory neurons after training, without basal forebrain activity. Each square corresponds to a neuron. Color indicates the preferred stimulus frequency. For better visibility, the neurons are arranged in four rows in order of increasing preferred stimulus frequency. Neurons marked in white are not selective to any of the used tones. Thus, this representation might be compared to a top view onto the primary auditory cortex as used by Kilgard and Merzenich (1998). (c) Same as (a), but now one of the rare stimuli (2.09 kHz) is paired with basal forebrain activity (both the start and ending of the pairing phase are indicated by vertical dashed lines). (d) Same as (b), for the experiment described in panel (c). The receptive fields are measured after the last paired presentation (presentation 22).

to them. Hence, the network has the ability to “reserve” neurons for representing future novel stimuli.

As a next step, comparable to recent physiological experiments done by Kilgard and Merzenich [23], we pair one of the rare stimuli (2.09 kHz) with the activation of the basal forebrain unit. Basal forebrain stimulation occurs simultaneously with stimulus onset. After a paired presentation 22 neurons develop specific receptive fields to this tone [their receptive field are similar to those in Fig. 5(a) and (b), data not shown]. The number of neurons responding to this tone is stable since in the following paired presentations this number does not decrease [Fig. 6(c)]. Therefore we see that the size of representation of this rare tone is much increased compared to the previous experiment where the basal forebrain remained inactive [Fig. 6(d)]. In that experiment the representation size of this stimulus stabilizes after several presentations in four neurons [Fig. 6(b)].

The basal forebrain input hyperpolarizes the cortical inhibitory neurons, delaying their activity with respect to the cortical excitatory neurons by about 6 ms and, thus, enlarging the temporal window for the backpropagating action potential to induce the potentiation of synaptic efficacies. This results in an increase in the representation size of this stimulus. This effect is independent of the presentation frequency of the

stimulus (data not shown) and does not affect the size of the representation of the other stimuli. With no basal forebrain activation, the final number of specific neurons responding to the tones 0.74, 1.05, 1.48, 2.09, and 2.96 kHz is 2, 3, 2, 3, and 3 neurons, respectively [Fig. 6(b)]. In the experiment where the basal forebrain is paired with the 2.09 kHz tone, the number of specific neurons responding to the tones is 2, 4, 3, 22, and 4 neurons, respectively [Fig. 6(d)].

When pairing is discontinued after presentation 22, the size of the representation of the previously paired tone is reduced and reaches a size comparable with the representation of the other tones (2, 3, 2, 2, and 2 neurons, respectively). Thus, the learning rule dynamically modifies the size of representation of the stimuli according to their behavioral importance, represented by the level of activity in the basal forebrain. This effect is independent of the probability of occurrence of the stimuli [Fig. 6(d)]. In addition, the dynamic modification does not affect the representations of other stimuli [Fig. 6(d)].

C. Learning With Acoustic Noise

As an additional control we investigate the properties of the proposed learning rule using stimuli with acoustic noise of

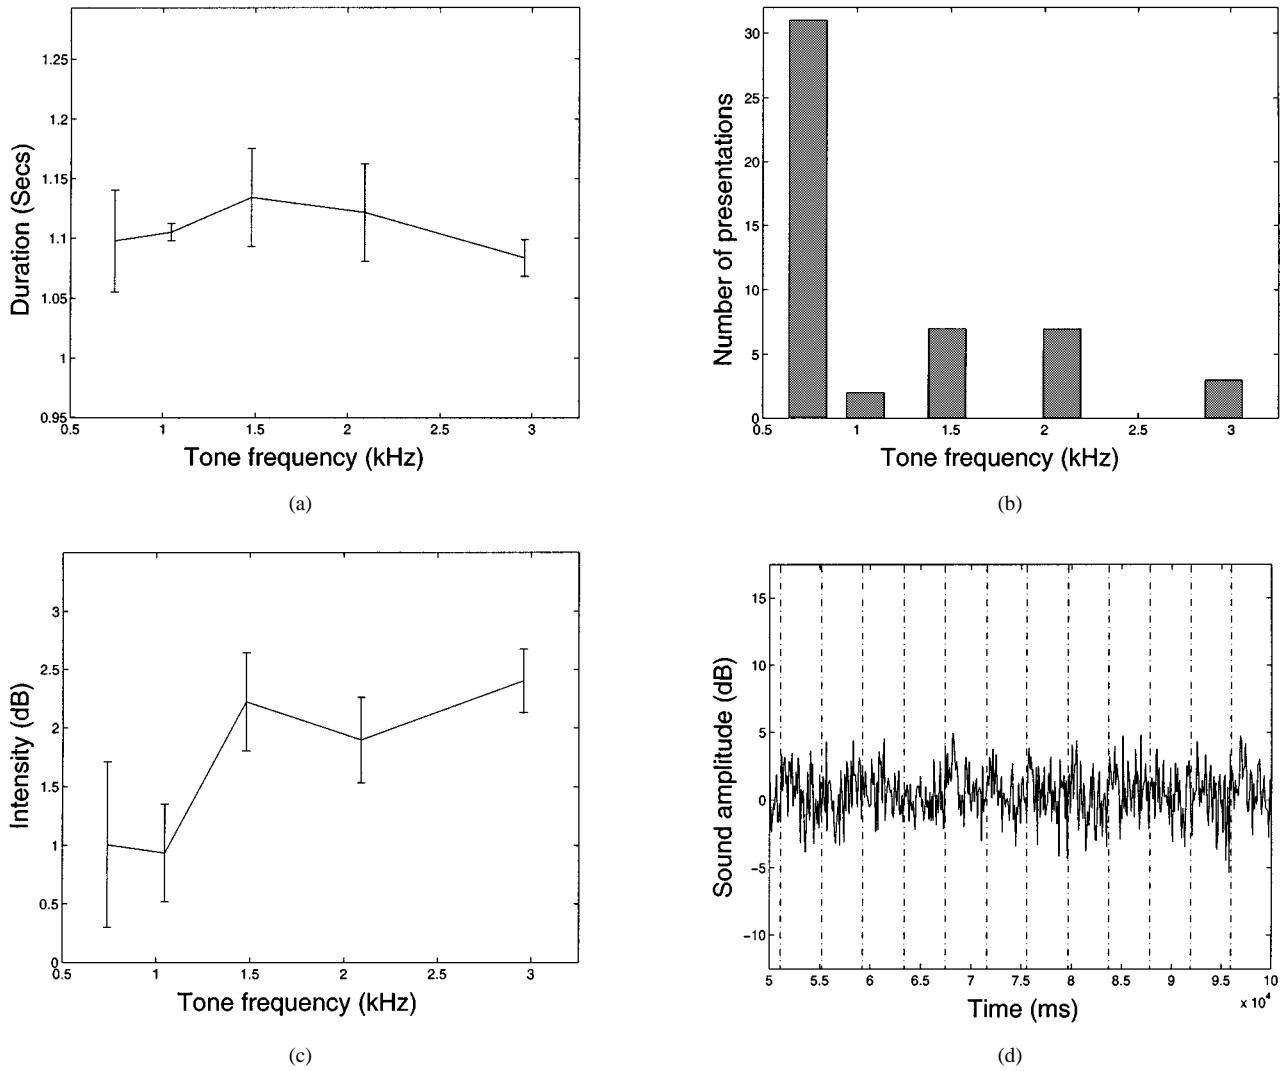


Fig. 7. Stimulus statistics using a pseudorandom sequence of five different tones (0.74, 1.05, 1.48, 2.09, and 2.96 kHz) and very loud noise as background. The probability of occurrence is 1/2, 1/8, 1/8, 1/8, and 1/8, respectively. (a) Mean duration of each stimulus. (b) Number of presentations of each stimulus. (c) Mean intensity of each stimulus. The 0 dB level is chosen as the averaged level of noise in the room. (d) Sound amplitude over time.

greater amplitude in a nonoverlapping frequency band. We use the same protocol as in the previous experiment [see Fig. 7(a) and (b)], while a continuous low-band noise is played by the synthesizer. The global signal-to-noise ratio of all the stimuli is close to 1 [Fig. 7(c) and (d)].

The noise continuously excites the input neurons corresponding to the lowest frequencies [Fig. 8(a)]. This in turn drives the thalamic neurons tuned to low frequencies leading to a response of all cortical excitatory neurons at the first presentation due to their initially diffuse receptive fields. After a few seconds, however, the efficacy of the synapses from the input population to the thalamic population, which transduce the presented frequencies, diminishes due to their short-term depression. This prevents continuously present harmonics from further activating the thalamic and cortical populations. This can be analyzed by calculating the expected value of the changes of the short-term depressing synapses (3)

$$\langle \dot{\Gamma}_i(t) \rangle = \tau_d^{-1}(1 - \langle \Gamma_i(t) \rangle) - f \langle \Gamma_i(t) \delta(t - t_{\text{spike}_i}) \rangle. \quad (7)$$

The temporal dynamics of Γ (time constant of 4 s) is much slower than the temporal dynamics of the input neuron (time constant of 19 ms) so we can write

$$\langle \Gamma_i(t) \delta(t - t_{\text{spike}_i}) \rangle = \langle \Gamma_i(t) \rangle \langle \delta(t - t_{\text{spike}_i}) \rangle = \langle \Gamma_i(t) \rangle F_N \quad (8)$$

where F_N is the firing rate of the input neuron which responds to a continuous stimulus. Using this in (7) we have that $\langle \Gamma_i(t) \rangle$ converges to

$$\langle \Gamma_i \rangle = \frac{1}{1 + \tau_d f F_N}. \quad (9)$$

Thus $\langle \Gamma_i \rangle$ goes to zero as the firing rate of the input neuron increases. Hence, continuously presented audio signals are filtered out through rapid synaptic depression. The main harmonics of the continuously presented signal induce high activation in the corresponding input neuron leading to a continuous excitation

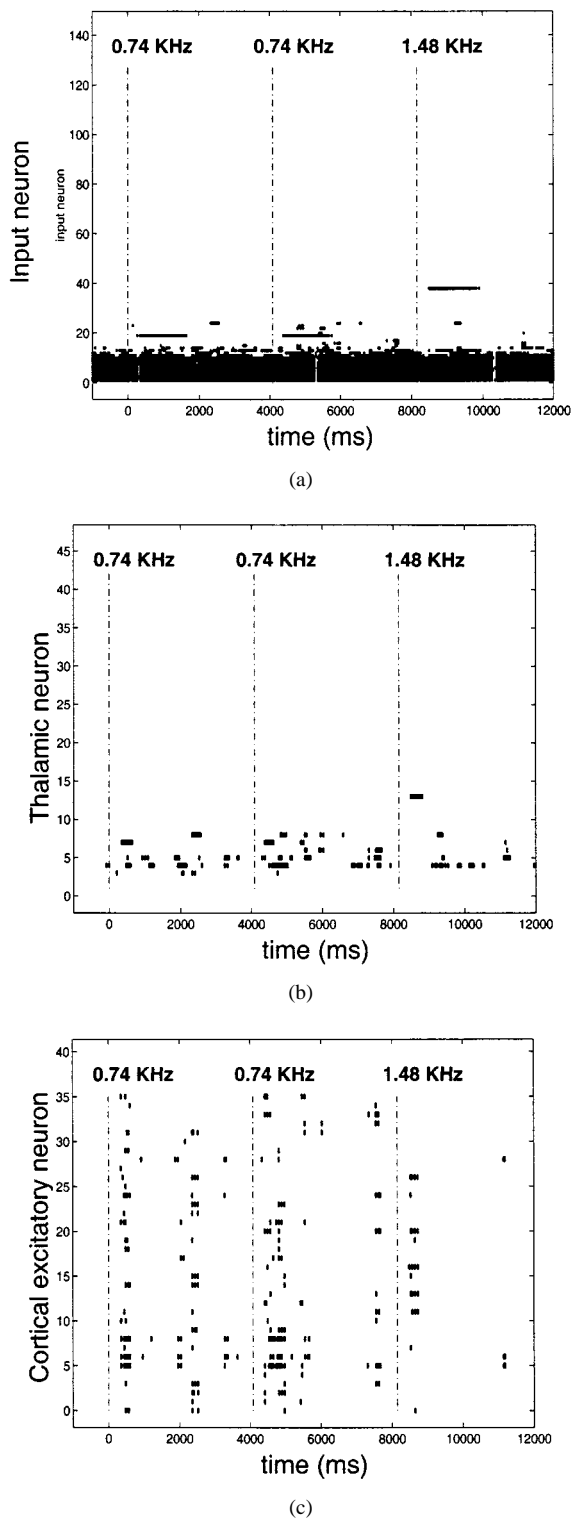


Fig. 8. Raster of network activity following the presentation of three tones (0.74 kHz, 0.74 kHz, 1.48 kHz) with very loud noise as background. Time zero corresponds to the onset of the first tone. Vertical dashed lines represent the onset of each tone. (a) Input population. (b) Thalamic population. (c) Cortical excitatory population.

of the thalamic neurons they project to. The efficacies of the synapses connecting these input neurons with the thalamic neurons would have a very low $\langle \Gamma_i \rangle$ strongly attenuating the signal they transduce.

However, as shown in Fig. 8(b) not all aspects of the continuously presented stimulus are filtered out. This is due to fluctuations in harmonics which have a small contribution to the signal and are not filtered out by the short-term depressing synapses. The weak contribution of these harmonics makes the corresponding input neuron fire at a low firing rate. From (9) we see that its connection to the thalamic neuron is not strongly affected. However, in our system a thalamic neuron needs to receive 2–3 effective spikes in a short period of time in order to fire. This means that those input neurons firing at a low frequency are not able to trigger a spike in thalamic neurons, even when the connection has a high $\langle \Gamma_i \rangle$. However, a momentary increment in the harmonic contribution would increase the firing rate of the input neuron, thus having the possibility to fire two to three spikes in a short period of time with a high $\langle \Gamma_i \rangle$, making the corresponding thalamic neuron fire. Therefore, we see that short-term depressing synapses are not able to completely filter out the continuous noise. As a result fluctuations in the harmonics of the noise are processed by the cortical network, mixed with the information about the tones presented to the system. Furthermore, these fluctuations in the noise can activate those input neurons that are activated when the 0.74 kHz tone is presented (Fig. 8). Therefore we see that the noise overlaps with the signal both temporally and spatially.

Hence, one would expect that the continuously presented noise would interfere with the development of receptive fields specific to the tones. However, those thalamo–cortical synapses that transduce information about the noise tend to get weaker. The learning rule decorrelates signals that are independent, in this case the fluctuations in the spectrum of the noise and the tones played by the synthesizer. This effect can be understood by calculating the expected increment of synaptic strength per postsynaptic spike, $\Delta W_{i,j}$, from (4)–(6). First, we introduce the notation

$$\Pi_{i,j} \equiv \left\langle \frac{t_0}{t_0 + ||t_i - t_j||} \middle| BPAP_j \wedge (||t_i - t_j|| < W) \right\rangle \quad (10)$$

that is, the expected value of the quantity $t_0/(t_0 + |t_i - t_j|)$ given that the backpropagating action potential in the postsynaptic neuron is not attenuated by the inhibition (“BPAP_j”), and that this event and the presynaptic spike fall within the temporal association window W (see Section II-C); t_i and t_j are the times when the presynaptic and postsynaptic neurons spike, respectively. Analogously

$$\Phi_{i,j} \equiv \left\langle \frac{t_0}{t_0 + ||t_i - t_j||} \middle| BPAP_m_j \wedge (||t_i - t_j|| < W) \right\rangle \quad (11)$$

that is, the expected value of the quantity $t_0/(t_0 + |t_i - t_j|)$ given that the backpropagating action potential in the postsynaptic neuron is attenuated by the inhibition (“BPAP_{m_j}”), and that this event and the presynaptic spike occur within the temporal association window W . Given these definitions the ex-

pected increment of synaptic strength per postsynaptic spike can be calculated from (4)–(6) resulting in

$$\begin{aligned} \langle \Delta W_{i,j} \rangle &= \alpha \Pi_{i,j} p(BPAP_j \wedge (\|t_i - t_j\| < W) | sp_j) \\ &\quad - \beta \Phi_{i,j} p(BPAP_m_j \wedge (\|t_i - t_j\| < W) | sp_j) \\ &\quad - \eta p(BPAP_j \wedge (\|t_i - t_j\| \geq W) | sp_j) \end{aligned} \quad (12)$$

where “ sp_j ” means “there is a spike in the j th postsynaptic neuron.” In case the i th thalamic cell encodes just fluctuations in the noise while another signal is making the cortical excitatory population fire, the activity of this cell is uncorrelated with activity of cortical excitatory cell j , that is

$$\begin{aligned} p(BPAP_j \wedge (\|t_i - t_j\| < W) | sp_j) \\ = p(BPAP_j | sp_j) p(\|t_i - t_j\| < W | sp_j) \end{aligned} \quad (13)$$

$$\begin{aligned} p(BPAP_j \wedge (\|t_i - t_j\| \geq W) | sp_j) \\ = p(BPAP_j | sp_j) p(\|t_i - t_j\| \geq W | sp_j). \end{aligned} \quad (14)$$

Because both β and $\Phi_{i,j}$ are positive

$$\begin{aligned} \langle \Delta W_{i,j} \rangle \leq \alpha \Pi_{i,j} p(BPAP_j \wedge (\|t_i - t_j\| < W) | sp_j) \\ - \eta p(BPAP_j \wedge (\|t_i - t_j\| \geq W) | sp_j). \end{aligned} \quad (15)$$

Noting that $\Pi_{i,j} \leq 1$, and using (13) and (14)

$$\begin{aligned} \langle \Delta W_{i,j} \rangle \leq \alpha p(BPAP_j | sp_j) p(\|t_i - t_j\| < W | sp_j) \\ - \eta p(BPAP_j | sp_j) p(\|t_i - t_j\| \geq W | sp_j). \end{aligned} \quad (16)$$

Therefore, $\Delta \langle W_{i,j} \rangle$ is guaranteed to be negative (forcing the final value of the synaptic strength toward zero) if

$$\eta > \alpha \frac{p(\|t_i - t_j\| < W | sp_j)}{1 - p(\|t_i - t_j\| < W | sp_j)}. \quad (17)$$

If the noisy output of thalamic neuron i can be described as a Poisson process with rate F , then

$$p(\|t_i - t_j\| < W | sp_j) = \int_0^{2W} f e^{-tF} = 1 - e^{-2FW}. \quad (18)$$

Using this in (17) we obtain

$$\eta > \alpha (e^{2FW} - 1). \quad (19)$$

This equation shows that the smaller the association window W is, the easier it is for the learning mechanism to prune the synapses that carry noisy information. In addition, the smaller the Poisson noise rate is, the easier it is for the learning mechanism to prune the synapse which transduces it.

Therefore, this learning mechanism decorrelates the noise from the receptive fields of the cortical excitatory cells sensitive to tones. Effectively, we see in [Fig. 9(a)] that the receptive fields of the neurons that fire to the tones are decorrelated from noise. A few neurons develop receptive fields specific to frequencies that are part of the noise: two are finally selective to frequencies lower than 0.7 kHz and one is selective to 0.90 kHz (Fig. 10). These neurons, however, do not respond to any of the tones [Fig. 9(b)]. Finally, the remaining neurons do not respond to either the tones or the noise, remaining “unspecific”

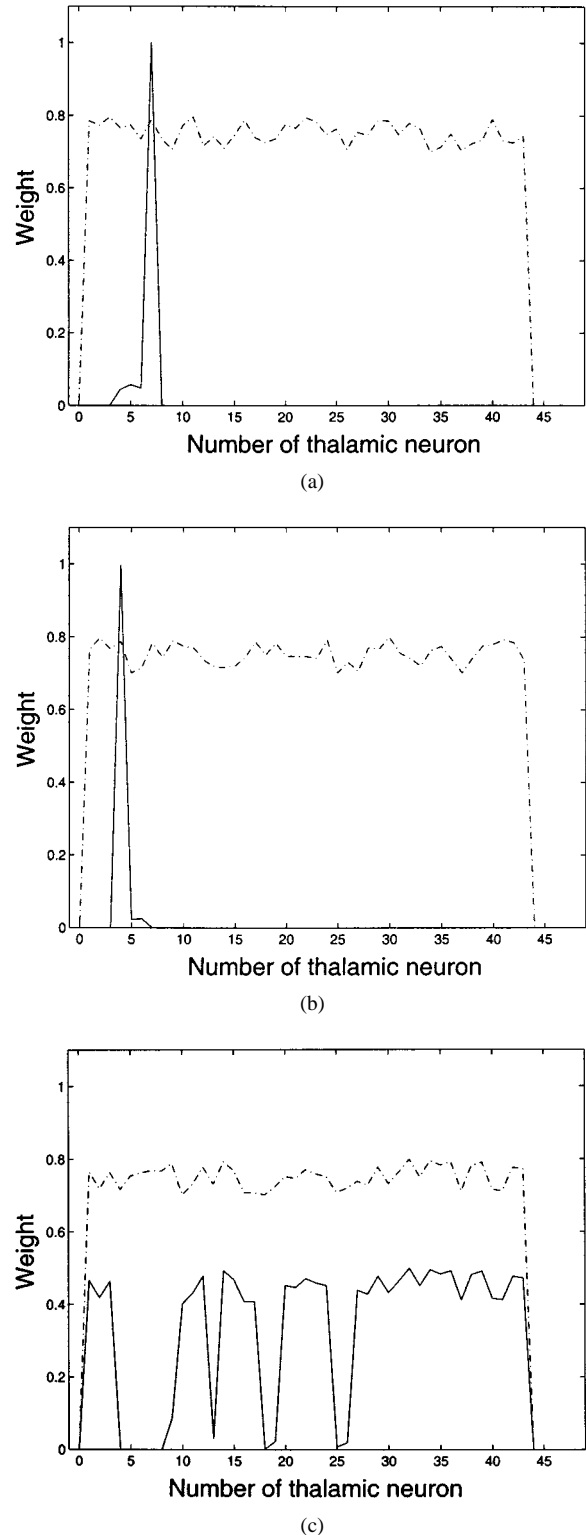


Fig. 9. Initial (dashed line) and final (solid line) receptive fields of selected neurons in the experiment with very loud noise as background. (a) Neuron finally selective to the 0.74 kHz tone. (b) Neuron finally selective to the 0.40 kHz component of the noise. (c) Neuron that finally does not respond to any of the five tones used in the training.

[Fig. 9(c)]. This ensures the ability of the network to learn future tones. In conclusion, the system proved to be robust against high noise levels and the results obtained are similar to those without noise (Fig. 10).

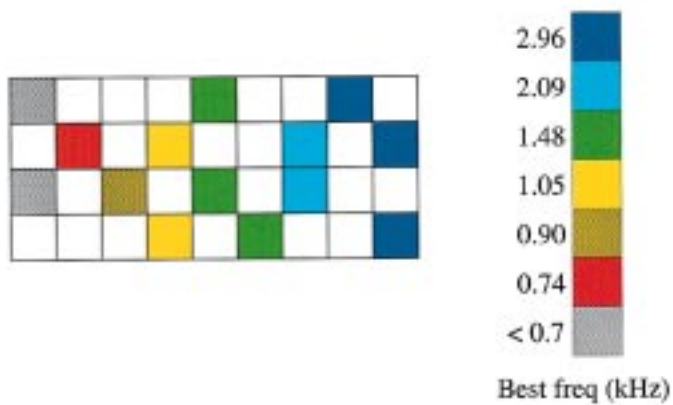


Fig. 10. Distribution of the preferred frequency of the 36 cortical excitatory neurons after training with a sequence of tones with very loud noise ad background using the same convention as in Fig. 6(b) and (d). Each tone in the sequence is randomly chosen from the set (0.74, 1.05, 1.48, 2.09, and 2.96 kHz) with a probability of (1/2, 1/8, 1/8, 1/8, 1/8), respectively. The displayed receptive fields were stable and resulted after 50 presentations. Neurons marked in gray and gold are selective to frequencies which are part of the noise: gray indicates a preferred frequency lower than 0.7 kHz, and gold indicates a preferred frequency of 0.90 kHz.

IV. DISCUSSION

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 which can be affected by a global mechanism. We show that this model supports continuous and fast learning, provides an even coverage of stimulus space, 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 [26], [34]. In addition we have shown that this biophysically realistic learning method is robust against noise and strong imbalances in the statistics of the input events.

In implementing our model we made some simplifications which are not critical to the presented results but could be important in further extensions. For example, in our model the auditory input neurons are band-pass filters with constant bandwidth. However, in biology the bandwidth of these neurons changes accordingly to their preferred frequency (narrower at low frequencies and wider for high frequencies). This would change slightly the activity statistics of the input neurons, diminishing the activity of input neurons tuned to lower frequencies (since now they have a reduced receptive field) while increasing the activity of input neurons tuned to higher frequencies, which now have a broader receptive field. As it is shown in the paper, the learning rule is robust against inhomogeneities in the input statistics. Therefore, the use of input neurons with different bandwidths is expected to lead to similar results to those we present in the paper.

Another aspect is the omission of temporal cues to frequency coding for the model. The input neurons of our model code the stimulus spectrum using a place representation (each neuron responds to a characteristic frequency). However, in the biology

the frequency of low-frequency tones is also coded by timing information (the spikes in the neuron are locked to the phase of the sound). We did not include this mechanism for simplicity. This mechanism could be easily integrated by using the phase of the DFT coefficients to modulate the firing of the input neurons in order to produce phase locked spikes. However, in the present context, the cortical network would not be able to differentiate between locked or not locked action potentials, and performance would not be affected.

Another simplification is that we are using chunked sampling (24 ms), keeping stimulation of the input neurons constant during that time. This is not crucial in our system, as long as the envelope of the sounds we used in the experiments changes slowly. This chunking of the input signals, however, limits the phase information available. Although not relevant for our present experiments this type of information is crucial for sound source localization. Hence, the present system could not support such forms of processing due to the loss of sound onset information as a result of the input chunking. One solution to this problem could be to shorten the length of the sample data. This loss of resolution, however, would increase the sensitivity of the DFT to noise. Other solutions could be found in using overlapping samples or wavelet based analysis.

Another simplification is that we do not try to replicate the huge dynamic range of the auditory system (about 100 dB). For this reason the network does not adapt to strong changes in the amplitude of the sound (e.g., it does not respond to very weak sounds). A possible solution to this problem is to use adapting gains in the neurons.

Biological systems learn from the real-world in real-time in order to enhance their ability to survive. Learning on one hand allows the identification of stimuli which are relevant for behavior and on the other to shape the behavioral output of the system. The learning abilities of biological systems has so far not been paralleled by those of artificial systems. For instance, one important area of research is in the domain of extraction of basic features of sensory stimuli using learning methods [5], [6], [13], [21], [33]. We would like to discuss critical aspects of our work in the context of the problems which pertain to this area.

- In the real world sensory input is a long continuous stream of unlabeled events. Potentially relevant and to be learned stimuli may occur at any time. Hence, a learning system needs to be able to adaptively identify and construct representations of relevant events or stimulus identification [49]. In our experiments there is no distinction between a training phase and an application or test phase. The learning mechanisms of the system are continuously active, allowing the system to learn new stimuli at any time.
- In contrast to simulation studies, in the real world it is not possible to tightly control stimulus statistics. Since online learning has no late or early phase and needs to store information at any time/continuously, the learning mechanism must be able to deal with high inhomogeneities in the coverage of stimulus space. Our model can deal with strong inhomogeneities and we show that the size of the representation of a stimulus does not depend on its frequency of presentation. This is a highly desirable feature

for a learning mechanism since neutral stimuli might be very frequent while those which are behaviorally relevant only occur sporadically. A learning mechanism which is sensitive to the presentation frequency of a stimulus would be saturated by neutral stimuli and insensitive to relevant ones. Hence, a learning system which interacts with the real world should dedicate more resources for processing relevant or potentially relevant stimuli. This is also suggested by the experimental data on auditory learning [50], [23] and captured by our model.

- Individual stimuli should be learned on the basis of a single or few presentations (“one-shot-learning”), requiring high learning rates. However, this could destabilize the system leading to a loss of previously learned stimuli (“speed-stability problem,” [19]). Hence, a learning mechanism working in the real world should be able to use high learning rates while keeping the learned representations stable. The learning mechanism we propose satisfies this requirement. It allows the learning of new stimuli while conserving the previously learned representations.

Furthermore, it allows one-shot learning, due to the ability of the learning mechanism to work with high learning rates: neurons that respond faster to the stimulus show rapid acquisition, while neurons responding late will suffer strong depression of their activated sensory synapses extinguishing their response to a future presentation of the stimulus. This strong competition permits the model to be stable while allowing one-shot learning. Another consequence of this competition is that the receptive fields of the neurons tend to be nonoverlapping. Using low learning rates, however, would diminish the “average competition” allowing receptive fields to overlap. The details of this process depend on the details of the stimulus statistics. One-shot learning, however, needs to be traded off against the signal to noise ratio, that is, for a learning system it is impossible to distinguish signal from noise in just one trial. If the signal to noise ratio is close to one, the learning rule needs more trials to prune the dynamic noise, achieving a noise free representation of the signal.

- Real-world stimuli are inherently noisy. This includes dynamic noise due to imperfections of sensors, but also static noise, due to cross-talk of different individual stimuli present simultaneously. Furthermore, some variations (e.g., distance of an object and thus its retinal size) are not important for one aspect (e.g., recognizing its identity) and thus represent noise but can be all decisive for other aspects (e.g., grasping it). Thus, the system must be able to cope with different kinds of noise appropriately depending on the task at hand.

In our model we included a short-term depressing mechanism in the synapses formed by the peripheral neurons on the thalamic ones [46]. We show that this mechanism can filter out the continuous part of the dynamic noise. The remaining part is decorrelated from the receptive-fields through the learning mechanism. In this case the information provided by the temporal relations between presynaptic and postsynaptic spikes is

crucial. However, in dealing with static noise without a high level memory system, or other means to describe larger statistics, noise can not be differentiated from signal in a single trial. Our system can partly deal with this problem if the learning rate used is not too high. In this case, if stimuli are not presented alone but mixed in different combinations (e.g., by using a typical music CD) the system achieves a “sparse” representation of the environment that minimizes the redundancy while covering the complete stimulus space. Interestingly, this is the type of representation that the visual cortex seems to use [32], having the advantages of minimizing the energy consumed [4] while minimizing the reconstruction error [32]. In addition, it is important to obtain low-redundancy codes (“minimum entropy codes”) in order to make the processing by higher stages as simple as possible [5].

- Due to limitations in hardware and/or software, simulations of learning in neural networks are often far from real time, and thus not suitable to be used in real world tasks. In this work we show that a real-time implementation of a biophysically detailed model, with some reasonable simplifications, is possible. The hardware used is based on a Pentium III processor running at 500 MHz; the model includes 244 neurons, 1548 plastic and 1496 nonplastic connections. In this type of study the critical element affecting computational load is not the number of neurons, but the number of connections. We estimate that with a Pentium IV at 2 GHz it is possible to increase the size of the two cortical populations up to 110 neurons each. This renders an overall system size of 392 neurons, 4730 plastic connections and 12 448 nonplastic connections in total.

Further optimizations of the implementation could be achieved by recoding parts of the simulation and reducing the number of connections. For instance, the number of connections can be drastically reduced if we use a rough topology in the connections between thalamic neurons and cortical excitatory neurons and in the connections between cortical inhibitory neurons and cortical excitatory neurons instead of the presently applied “all to all” topology. We gained in the computational efficiency by relying on a distributed implementation of the overall system, running each component of the system on a separate computer, which is a prerequisite for the real-time simulation of large scale networks. This is provided by the neural simulator we use IQR421 [47]. Further optimizations can be provided by using FPGAs or analog VLSI (aVLSI) devices. Currently we are experimenting with replacing the digital FFT with an aVLSI model of the cochlea.

In our model a local learning rule interacts with a global modulatory system. This represents the substantial gabaergic projection originating in the basal forebrain, which terminates on cortical inhibitory neurons [15], [16]. Its activation increases the proportion of successfully backpropagating action potentials in the cortical excitatory neurons [34]. Interestingly, the much better investigated basal forebrain cholinergic projection [37] increases the fraction of backpropagating action potentials in cortical neurons [45]. Thus, these two subcortical projections may act synergistically, enhancing each other’s effect. Of

course these findings do not preclude additional actions of acetylcholine, e.g., the modulation of effective intracortical connectivity [48]. In addition we include a short-term depressing mechanism in the synapses formed by the peripheral neurons on the thalamic ones [46]. Experiments on this property do not investigate the synapses which we assume to adapt, but the abundance of data about these mechanisms in different types of synapses [52], [41], [46] supports our hypothesis. The integration of these biologically realistic mechanisms in the model renders a number of interesting abilities which did not interfere with each other.

In conclusion, we have presented a biologically realistic neural network that performs in real time, learning sensory maps of the environment using real-world stimuli. However, the interaction of learning sensory representations and overt behavior is an essential aspect of an autonomous agent. Ultimately, only this interaction allows an assignment of relevance to a subset of stimuli. Indeed, experimental evidence demonstrates that compared to passive reception of stimuli, active sensing qualitatively influences the development of sensory representations [35]. Hence, the implementation of the model on a robot interacting with the environment [25], [49] will be an exciting topic for future work.

ACKNOWLEDGMENT

The authors would like to thank P. Paschke for his help in developing the data acquisition system and the FFT analysis module. They are also grateful to A. van Schaik, D. Klein, and F. Corbacho for helpful discussions.

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, and C. Wible, "Learning and memory," *Brain Res. Rev.*, vol. 16, pp. 193–220, 1991.
- [2] L. F. Abbott, "Learning in neural network memories," *Network*, vol. 1, pp. 105–122, 1990.
- [3] L. F. Abbott, J. A. Varela, K. Sen, and S. B. Nelson, "Synaptic depression and cortical gain control," *Science*, vol. 275, pp. 220–224, 1997.
- [4] R. Baddeley, "Visual perception. An efficient code in V1?," *Nature*, vol. 381, pp. 560–561, 1996.
- [5] H. B. Barlow, "Unsupervised learning," *Neural Comput.*, vol. 1, pp. 295–311, 1989.
- [6] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Comput.*, vol. 7, no. 6, pp. 1129–1159, 1995.
- [7] G. Q. Bi and M. M. Poo, "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type," *J. Neurosci.*, vol. 18, pp. 10464–10472, 1998.
- [8] E. L. Bienenstock, L. N. Cooper, and P. W. Munro, "Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex," *J. Neurosci.*, vol. 2, no. 1, pp. 32–48, 1982.
- [9] T. V. Bliss and G. L. Collingridge, "A synaptic model of memory: Long-term potentiation in the hippocampus," *Nature*, vol. 361, pp. 31–39, 1993.
- [10] T. H. Brown and S. Chattarji, "Hebbian synaptic plasticity," in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. Cambridge, MA: MIT Press, 1998.
- [11] D. V. Buonomano and M. M. Merzenich, "Cortical plasticity: From synapses to maps," *Annu. Rev. Neurosci.*, vol. 21, pp. 149–186, 1998.
- [12] X. Buzsáki and E. Kandel, "Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat," *J. Neurophysiol.*, vol. 79, no. 3, pp. 1587–1591, 1998.
- [13] P. Comon, "Independent component analysis—A new concept?," *Signal Processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [14] Y. Fregnac, "Hebbian synaptic plasticity: Comparative and developmental aspects," in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. Cambridge, MA: MIT Press, 1998.
- [15] T. F. Freund and A. I. Gulyás, "GABAergic interneurons containing calbindin D28K or somatostatin are major targets of GABAergic basal forebrain afferents in the rat neocortex," *J. Comp. Neurol.*, vol. 314, pp. 187–199, 1991.
- [16] T. F. Freund and V. Meskenaite, "Gamma-aminobutyric acid-containing basal forebrain neurons innervate inhibitory interneurons in the neocortex," *Proc. Nat. Academy Sci. USA*, vol. 89, pp. 738–742, 1992.
- [17] M. Frigo and S. G. Johnson, "FFTW: An adaptive software architecture for the FFT," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing*, vol. 3, Seattle, WA, 1998, pp. 1381–1384.
- [18] W. Gerstner, R. Ritz, and J. L. van Hemmen, "Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns," *Biol. Cybern.*, vol. 69, pp. 503–515, 1993.
- [19] S. Grossberg, "How does the brain build a cognitive code?," *Psych. Rev.*, vol. 87, pp. 1–15, 1980.
- [20] D. O. Hebb, *The Organization of Behavior*. New York: Wiley, 1949.
- [21] I. T. Jolliffe, *Principal Component Analysis*. New York: Springer-Verlag, 1986.
- [22] J. P. Kelly, "Auditory system," in *Principles of Neural Science*, E. Kandel and J. Schwartz, Eds. Amsterdam, The Netherlands: Elsevier, 1985, pp. 396–408.
- [23] M. P. Kilgard and M. M. Merzenich, "Cortical map reorganization enabled by nucleus basalis activity," *Science*, vol. 279, pp. 1714–1718, 1998.
- [24] P. König, A. K. Engel, P. R. Roelfsema, and W. Singer, "How precise is neuronal synchronization?," *Neural Comput.*, vol. 7, pp. 469–485, 1995.
- [25] P. König, E. Bizzi, N. Burgess, N. Franceschini, M. P. Kilgard, M. Oram, G. Sagerer, and C. Scheier, "Group report: Representations in natural and artificial systems," *Z. Naturforsch [C]*, vol. 53, no. 7–8, pp. 738–751, 1998.
- [26] K. P. Körding and P. König, "A learning rule for dynamic recruitment and decorrelation," *Neural Networks*, vol. 13, pp. 1–9, 2000.
- [27] H. J. Köster and B. Sakmann, "Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials," *Proc. Nat. Academy Sci. USA*, vol. 95, no. 16, pp. 9596–9601, 1998.
- [28] M. E. Larkum, J. J. Zhu, and B. Sakmann, "A new cellular mechanism for coupling inputs arriving at different cortical layers," *Nature*, vol. 398, pp. 338–341, 1999.
- [29] J. Magee, D. Hoffman, C. Colbert, and D. Johnston, "Electrical and calcium signaling in dendrites of hippocampal pyramidal neurons," *Annu. Rev. Physiol.*, vol. 60, pp. 327–346, 1998.
- [30] H. Markram, J. Lubke, M. Frotscher, and B. Sakmann, "Regulation of synaptic efficacy by coincidence of postsynaptic AP's and EPSPs," *Science*, vol. 275, pp. 213–215, 1997.
- [31] S. J. Martin, P. D. Grimwood, and R. G. Morris, "Synaptic plasticity and memory: An evaluation of the hypothesis," *Annu. Rev. Neurosci.*, vol. 23, pp. 649–711, 2000.
- [32] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, pp. 607–609, 1996.
- [33] R. P. Rao, "An optimal estimation approach to visual perception and learning," *Vision Res.*, vol. 39, no. 11, pp. 1963–1989, 1999.
- [34] M. A. Sánchez-Montañés, P. F. M. J. Verschure, and P. König, "Local and global gating of synaptic plasticity," *Neural Comput.*, vol. 12, no. 3, pp. 519–529, 2000.
- [35] W. Schultz and A. Dickinson, "Neuronal coding of prediction errors," *Annu. Rev. Neurosci.*, vol. 23, pp. 473–500, 2000.
- [36] T. Sejnowski, "Storing covariance with nonlinearly interacting neurons," *J. Math. Biol.*, vol. 4, pp. 303–321, 1977.
- [37] W. Singer and J. P. Rauschecker, "Central core control of development plasticity in the kitten visual cortex. II: Electrical activation of mesencephalic and diencephalic projections," *Exp. Brain Res.*, vol. 47, pp. 223–233, 1982.
- [38] N. Spruston, Y. Schiller, G. Stuart, and B. Sakmann, "Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites," *Science*, vol. 268, pp. 297–300, 1995.
- [39] G. S. Stent, "A physiological mechanism for Hebb's postulate of learning," *Proc. Nat. Academy Sci. USA*, vol. 70, pp. 997–1001, 1973.

- [40] G. J. Stuart and B. Sakmann, "Active propagation of somatic action potentials into neocortical pyramidal cell dendrites," *Nature*, vol. 367, pp. 69–72, 1994.
- [41] A. M. Thomson, J. Deuchars, and D. C. West, "Large, deep layer pyramid–pyramid single axon EPSP's in slices of rat motor cortex display paired pulse and frequency-dependent depression, mediated presynaptically and self-facilitation, mediated postsynaptically," *J. Neurophysiol.*, vol. 70, no. 6, pp. 2354–2369, 1993.
- [42] A. M. Thomson and J. Deuchars, "Temporal and spatial properties of local circuits in neocortex," *Trends Neurosci.*, vol. 17, no. 3, pp. 119–126, 1994.
- [43] J. Z. Tsien, "Linking Hebb's coincidence-detection to memory formation," *Curr. Opin. Neurobiol.*, vol. 10, pp. 266–273, 2000.
- [44] M. Tsodyks, K. Pawelzik, and H. Markram, "Neural networks with dynamic synapses," *Neural Comput.*, vol. 10, no. 4, pp. 821–835, 1998.
- [45] H. Tsubokawa and W. N. Ross, "IPSP's modulate spike backpropagation and associated Ca^{2+} changes in the dendrites of hippocampal CA1 pyramidal neurons," *J. Neurophysiol.*, vol. 76, pp. 2896–2906, 1996.
- [46] J. A. Varela, K. Sen, J. Gibson, J. Fost, L. F. Abbott, and S. B. Nelson, "A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex," *J. Neurosci.*, vol. 17, pp. 7926–7940, 1997.
- [47] P. F. M. J. Verschure, "Xmorph," Int. Rep., Inst. Neuroinform., ETH-UZ, 1997.
- [48] P. F. M. J. Verschure and P. König, "On the role of biophysical properties of cortical neurons in binding and segmentation of visual scenes," *Neural Comp.*, vol. 11, pp. 1113–1138, 1999.
- [49] T. Voegtlin and P. F. M. J. Verschure, "What can robots tell us about brains? A synthetic approach toward the study of learning and problem solving," *Rev. Neurosci.*, vol. 10, no. 3–4, pp. 291–310, 1999.
- [50] N. M. Weinberger, "Learning induced changes of auditory receptive fields," *Cur. Opin. Neurobiol.*, vol. 3, pp. 570–577, 1993.
- [51] N. M. Weinberger, R. Javid, and B. Lapan, "Long-term retention of learning-induced receptive-field plasticity in the auditory cortex," *Proc. Nat. Academy Sci. USA*, vol. 90, no. 6, pp. 2394–2398, 1993.
- [52] R. S. Zucker, "Short-term synaptic plasticity," *Annu. Rev. Neurosci.*, vol. 12, pp. 13–31, 1989.



Manuel A. Sánchez-Montañés received the B.Sc. degree (with honors) in physics from the Universidad Complutense de Madrid, Spain, in 1997. Currently, he is pursuing the Ph.D. degree at the Universidad Autónoma de Madrid.

In 1998 and 1999, he visited the Institute of Neuroinformatics in Zürich. His main interests include 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 of Bonn, Bonn, Germany. He was with the Department of Neurophysiology at the Max-Planck-Institute for Brain Research, Frankfurt, Germany, where he received the Habilitation degree in 1990.

After working as a Senior Fellow at the Neurosciences Institute in San Diego, CA, he joined the Institute of Neuroinformatics, Zürich, Switzerland, 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.



Paul F. M. J. Verschure received both the M.A. and Ph.D. degrees in psychology.

He is a group leader at the Institute of Neuroinformatics ETH-University, Zurich, Switzerland. He works on biologically realistic models of perception, learning, and problem solving, which are applied to robots and on the tools and methods that support this research.