Temporal Sequence Detection with Spiking Neurons: Towards Recognizing Robot Language Instructions.

Christo Panchev and Stefan Wermter

Centre for Hybrid Intelligent Systems
School of Computing and Technology, University of Sunderland
St. Peter's Campus, Sunderland SR6 0DD, United Kingdom
www.his.sunderland.ac.uk

Abstract

We present an approach for recognition and clustering of spatio temporal patterns based on networks of spiking neurons with active dendrites and dynamic synapses. We introduce a new model of an integrate-and-fire neuron with active dendrites and dynamic synapses (ADDS) and its synaptic plasticity rule. The neuron employs the dynamics of the synapses and the active properties of the dendrites as an adaptive mechanism for maximizing its response to a specific spatio-temporal distribution of incoming action potentials. The learning algorithm follows recent biological evidence on synaptic plasticity. It goes beyond the current computational approaches which are based only on the relative timing between single pre- and post-synaptic spikes and implements a functional dependence based on the state of the dendritic and somatic membrane potentials around the pre- and post-synaptic action potentials. The learning algorithm is demonstrated to effectively train the neuron towards a selective response determined by the spatio-temporal pattern of the onsets of input spike trains. The model is used in the implementation of a part of a robotic system for natural language instructions. We test the model with a robot whose goal is to recognize and execute language instructions. The research in this article demonstrates the potential of spiking neurons for processing spatio-temporal patterns and the experiments present spiking neural networks as a paradigm which can be applied for modeling sequence detectors at word level for robot instructions.

Key words: spiking neurons, active dendrites, dynamic synapses, synaptic plasticity, temporal sequence detection, natural language, intelligent robotics

1 Introduction

The concept of exploiting the timing of spikes as an alternative of or complimentary to the mean firing rate has provided new directions for further progress in neural computing models. Different models of spiking neurons have been developed (Hodgkin and Huxley, 1952; Rall, 1989; Segev et al., 1989; Kistler et al., 1997; Panchev et al., 2002), but there is still an ongoing debate on which are the essential properties of the biological neurons necessary to be simulated in order to achieve the computational power of a real neural system. The work presented in this article extends the current modeling paradigms of spiking neurons, in particular the leaky integrate-and-fire neuron (Maass and Bishop, 1999), by introducing a computational interpretation and exploring the functionalities of active dendrites and dynamic synapses in an integrated model.

For a long time, dendrites have been thought to be the structures where complex neuronal computation takes place, but only recently have we begun to understand how they operate. The dendrites do not simply collect and pass synaptic inputs to the soma, but in most cases they actively shape and integrate these signals in complex ways (Stuart et al., 2001; Poirazi and Mel, 2001; Aradi and Holmes, 1999). With our growing knowledge of such processing in the dendrites, there is a strong argument for taking advantage of the processing power and active properties of the dendrites, and integrating their functionality into artificial neuro-computing models (Panchev et al., 2002; Horn et al., 1999; Mel et al., 1998).

Furthermore, there is a variety of dynamic processes in the axonal terminal, including paired-pulse facilitation or depression, augmentation, post-tetanus potentiation, etc. The real neurons use these short term dynamics as an additional powerful mechanism for temporal processing. Several studies have explored the mechanisms of synaptic dynamics (Tsodyks et al., 1998; Abbott et al., 1997; Zucker, 1989) as well as their computational properties (Natschläger et al., 2001; Pantic et al., 2002), highlighting the advantages of neurons and neural networks with such synapses.

The mechanisms of the active dendrites and dynamic synapses operate on different time scales, and can be complimentary to each other. The combination of their functionality is most likely to be heavily used by the real neurons and could add significant computational advantages into the artificial neural networks. However, so far these two neuronal mechanisms have been primarily modeled in isolation (Spencer and Kandel, 1961; Schutter and Bower, 1994a,b,c; Liaw and Berger, 1996; Senn et al., 2001).

The work presented here introduces a computational interpretation and integration of functional properties of neurons with active dendrites and dynamic synapses as well as a synaptic plasticity rule associated with such neurons. The active dendrites manipulate the membrane time constants and resistance of the neuron and are able to precisely shape the post-synaptic potentials within a time scale of up to a few hundred milliseconds. Complimentary to this, the dynamics of the synapse is able to manipulate the post-synaptic responses on a time scale from a few hundred milliseconds up to several seconds. The in-

tegration of active dendrites and dynamic synapses into a model of a spiking neuron adds a functionality for temporal integration, which could be particularly powerful in detecting the temporal structure of incoming action potentials. Such functionality is required in many perceptual and higher level cognitive systems in the brain, e.g. speech and language processing. The model developed here is based on the integrate-and-fire neuron with active dendrites presented in Panchev et al. (2002). The new development introduces short-term synaptic facilitation and depression and generalization of the training algorithm from input stimuli of single spikes to spike trains.

The new model of a spiking neuron is used in the development of a system for language instructions given by humans to a robot. The approach of using a robot platform for testing the performance and functionality of an artificial neural network model has many advantages (Webb, 2001). Similar to living organisms, robots can be autonomous behaving systems equipped with sensors and actuators. Furthermore, the introduction of biologically inspired models into robotics could bring significant further advances into the field of intelligent robotics (Sharkey and Ziemke, 2001, 2000; Sharkey and Heemskerk, 1997). Using existing technologies, robots can only master basic reactive or preprogrammed behaviors. They fail to follow some of the fundamental functions of living organisms: their development, learning and adaptation to the environment. The work presented in this article is a contribution towards overcoming these limitations and building intelligent and adaptable robots. In the experimental section of this article, we present a model of spiking neural network for word recognition as part of a robot understanding natural language instructions.

In section 2, we introduce the new model of a spiking neuron with active dendrites and dynamic synapses and examine some of its critical properties for temporal integration of incoming spike trains. Section 3 presents the learning algorithm for the synapses included in the model. The experiments in section 4 explore the performance of the model for phoneme sequence detection and word recognition as part of a system for language instructions of a robot.

2 Spiking neurons with active dendrites and dynamic synapses (ADDS neurons): temporal integration of input spike trains

The neuron model presented in this section introduces the advantages of the combined functionality of dynamic synapses and active dendrites. The neuron is able to maximize its response and detect a particular temporal sequence via a system of implicit "delay" mechanisms. These mechanisms are based on the modulation of the generated post-synaptic potentials resulting from the dynamics of the synapses and the active properties of the dendrites. The learning algorithm presented in the next section tunes the "delay" mechanisms such that they generate a maximum response, in terms of the membrane potential at the soma, for a particular temporal sequence of input spike trains.

2.1 The Neuron Model

A schematic presentation of the model is shown in figure 1. The neuron receives input spikes via sets of dynamic synapses $\mathcal{S}_i$, each attached to a particular active dendrite i. In addition, the neuron has a set of synapses $\overline{\mathcal{S}}$ attached close to or directly at the soma.

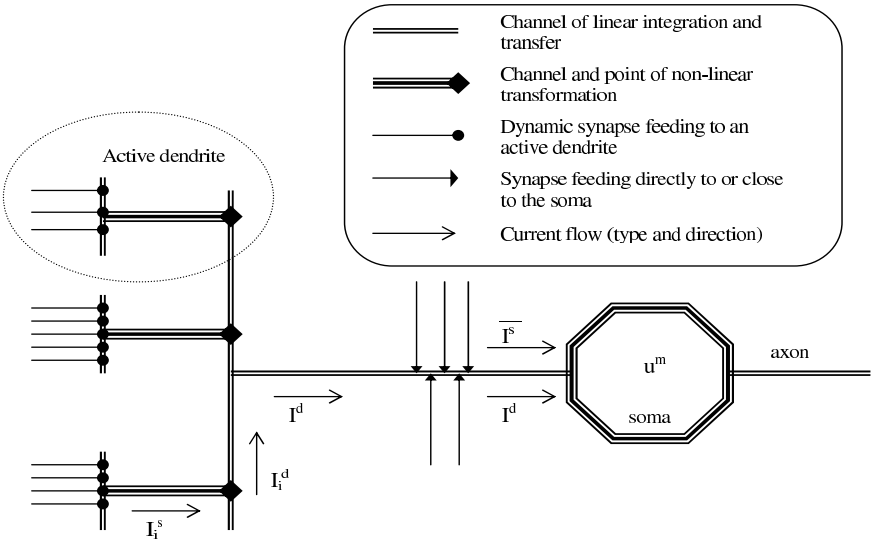

Figure 1: Model of a neuron with active dendrites and dynamic synapses.

The total post-synaptic current I_i^s generated by all dynamic synapses at dendrite i is described by:

$$\tau^s \frac{d}{dt} I_i^s = -I_i^s(t) + \sum_{j \in \mathcal{S}_i} w_{ij} \sum_{t^{(f)} \in \mathcal{F}_j} \rho(\Delta t^{(f)}) \delta(t - t^{(f)}) \tag{1}$$

where synaptic connection j at dendrite i has weight w_{ij}, $\mathcal{F}_j$ is the set of pre-synaptic spike times received at the synapse and $\delta(\cdot)$ is the Dirac δ-function.

Furthermore, the post-synaptic current passing through the dendrite into the soma is described by:

$$\tau_i^d \frac{d}{dt} I_i^d = -I_i^d(t) + R_i^d I_i^s(t) \tag{2}$$

Experimental evidence shows that the synaptic transmission properties of cortical neurons are strongly dependent on the recent pre-synaptic activity (Abbott et al., 1997; Tsodyks et al., 1998). The individual post-synaptic responses

are dynamic and can increase (in case of synaptic facilitation) or decrease (in case of of synaptic depression). Here, the synaptic dynamics is described by the function $\rho(\cdot)$ which depends on the time $\Delta t^{(f)}$ between the current and the earliest spikes in $\mathcal{F}_j$:

$$\rho(\Delta t^{(f)}) = \mu e^{-\left[\frac{\Delta t^{(f)} - \tau^{ds}}{\sigma}\right]^2} \tag{3}$$

with a time constant $\tau^{ds} = 1 - w_{ij}$, and scaling parameters σ and μ. Since the time constant depends on the weight of the synapse, an input spike train arriving at a stronger synapse will lead to a quicker short-term facilitation of the synapse, followed by a sharp depression, and generate an earlier increase of the membrane potential at the soma (figure 2 (A)). Respectively, a spike train arriving at a weaker synapse will generate a delayed increase of the membrane potential. The short-term facilitation and depression of the synapse operate on a time scale from a few hundred milliseconds up to a few seconds. The combined response of the dynamic synapses partially leads to the prolonged synaptic integration of cortical neurons presented in Beggs et al. (2000).

Real neurons show a passive response only under very limited conditions. In many brain areas, a reduction of ongoing synaptic activity has been shown to increase the membrane time constant and input resistance, suggesting that synaptic activity can reduce both parameters (Häusser and Clark, 1997; Paré et al., 1998). The computational model of active dendrites presented in this article is based on such observations. Here, the time constant τ_i^d and the resistance R_i^d are set to be dependent on the post-synaptic current into dendrite i and determine the active properties of the dendrite (see Panchev et al. (2002) and appendix A.1 for details). As shown in the next section, the effect is that a dendrite receiving strong post-synaptic input from a single spike generates a sharp earlier increase of the membrane potential at the soma, whereas the potential generated from a weaker input signal will be prolonged (figure 2 (A)).

A simpler equation holds for the total current $\overline{I}^s$ from all synapses feeding directly to the soma:

$$\tau^s \frac{d}{dt}\overline{I}^s = -\overline{I}^s(t) + \sum_{j \in \overline{\mathcal{S}}} w_j \sum_{t^{(f)} \in \mathcal{F}_j} \delta(t - t^{(f)}) \tag{4}$$

Finally the soma membrane potential u^m is:

$$\tau^m \frac{d}{dt} u^m = -u^m(t) + R^m(I^d(t) + \overline{I}^s(t)) \tag{5}$$

where $I^d(t) = \sum_i I_i^d(t)$ is the total current from the dendritic tree, and $\overline{I}^s(t)$ is the total current from synapses attached to the soma.

The current from dendrite i generates part of the potential at the soma, which we will call *partial membrane potential* and annotate as u_i^m. If pre-synaptic input arrives only at dendrite i then $u_i^m = u^m$. The *total partial membrane potential* $u_d^m = \sum_i u_i^m$ is the somatic membrane potential generated from all dendrites, i.e. excluding the synapses feeding directly to the soma.

The next section will present how the above dynamic and active mechanisms of the neuron work and allow spatio-temporal integration of incoming action potentials which facilitates the neuron's sensitivity to the temporal structure of the incoming spike trains.

2.2 Spatio-temporal integration of synaptic input in a neuron with dynamic synapses and active dendrites

As shown in the next section, two of the critical factors defining the neuron's response and adaptation are the timing and the amplitude of the maximum of the membrane potential at the soma. Figure 2 shows the response of a neuron with dynamic synapses and active dendrites to the same spike trains coming through synapses with different strength. In the case where $w = 0.8$ the spike train generates an earlier and sharp peak of the membrane potential at the soma. If the synaptic strength is smaller, e.g. $w = 0.2$, the time constants of the dynamic synapse τ^{ds} and the active dendrite τ_i^d will be longer, the resistance R_i^d will be lower and consequently the peak of the membrane potential is delayed and has a smaller amplitude.

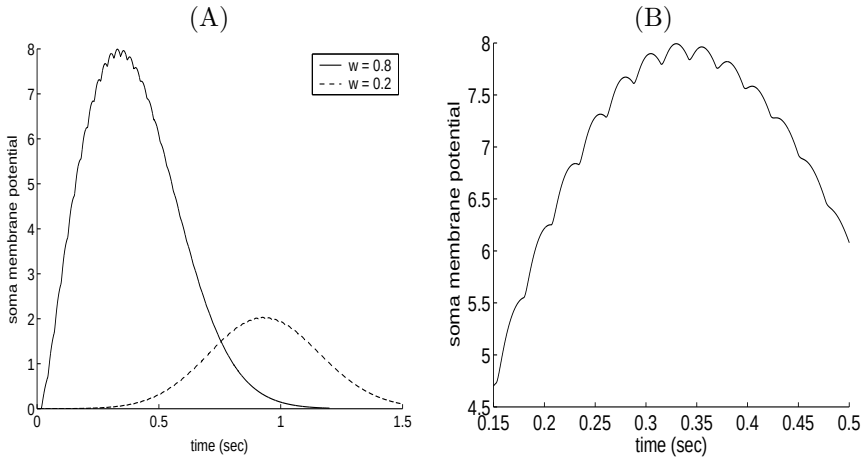

Figure 2: (A) Membrane potential at the soma generated by a single input spike train arriving at a single synapse (two cases of synapses with different weights); (B) Zoom-in around the peak of membrane potential for $w = 0.8$.

Based on such temporal integration, the neuron is able to maximize its response to spike trains with different onset times. Figure 3 shows two cases of a response of a neuron to two spike trains with different onset times. In both cases the total synaptic weight is the same, however the response at the soma

is substantially different. With appropriately adjusted weights (figure 3 top), the neuron is able to compensate for the delay of one of the spike trains, by delaying the post-synaptic current generated by the earlier one. The result is quasi-synchronous peaks of the partial membrane potentials, and a significantly higher potential generated at the soma. Such temporal integration provides the neuron with a powerful mechanism for a selective response to the temporal structure of the input stimuli.

3 Synaptic plasticity

There are two different types of synapses at the ADDS neuron, each of them having a different functional role. The dynamic synapses attached to active dendrites are part of the mechanism for spatio-temporal integration and play a role in the recognition of the temporal structure of the input signals. The neuron also includes synapses directly attached to the soma. They have a very fast and strong influence on the membrane potential at the soma and are very efficient for lateral connections between cooperative or competitive neurons. Consequently, the learning algorithms for the two types of synapses have different implementations. The specific tuning of the dynamic synapses attached to active dendrites implements the neuron's adaptation towards responding to a particular temporal sequence of input spike trains, whereas the plasticity of the synapses attached to the soma reflects the cooperation within or competition between clusters of neurons.

3.1 Plasticity in the dynamic synapses attached to the active dendrites

The task of the learning algorithm developed for the dynamic synapses attached to the active dendrites is to adjust the weights of the neuron, so that it is able to synchronize the peaks of the partial membrane potentials, and therefore maximize the response of the total somatic membrane potential for a particular temporal distribution of spike trains.

Current views on the intra-cellular mechanisms underlying synaptic plasticity postulate that the direction and magnitude of the change of the synaptic strength depend on the relative timing between the pre- and post-synaptic spikes which is expressed in the Ca^{2+} concentration modulated by pre-synaptic as well as back-propagating post-synaptic action potentials (Magee and Johnston, 1997; Markram et al., 1997; Bi and Poo, 1998; Debanne et al., 1998; Feldman, 2000; Larkum et al., 1999). The post-synaptic action potential is propagated back to the synapse by virtue of the active properties of the dendrites under complex control mechanisms. Its amplitude and duration depend on a variety of conditions, including the state of the soma, basal dendrites and the state and spatial position of the synapse's own dendritic branch (Johnston et al., 1999; Stuart et al., 1997; Buzsáki and Kandel, 1998; Spruston et al., 1995). In our approach, we view the signal carried by the back-propagating action potential

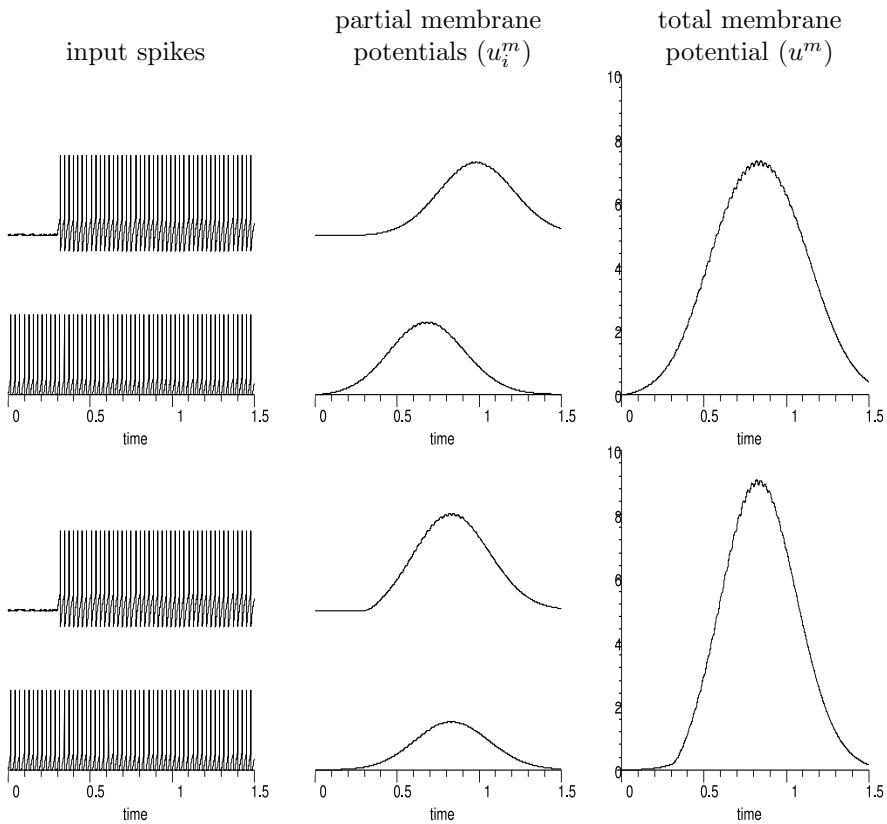

Figure 3: Integration of two spike trains through two dynamic synapses with different strength attached to separate active dendrites. The onset of one of the spike trains is delayed by 300ms. *Top*: Both spike trains arrive at two synapses with equal strength, $w_{1,2} = 0.45$ each. *Bottom*: The first spike train arrives at a synapse with strength $w_1 = 0.3$ and generates a peak of the partial membrane potential around 800ms after the onset time of the stimulus. The second spike train arrives at synapse with strength $w_2 = 0.6$ and generates a peak partial potential around 500ms after the onset time. The response of the neuron to the first spike train is delayed (prolonged) compared to the response to the second one. As a result, the neuron has achieved a quasi-synchronization of the partial membrane potentials generated from the spike trains (middle column, bottom graph), and its response (the membrane potential at the soma, bottom right graph) is much stronger.

as having two main parts: a signal depending on the state of the soma and basal dendrites, and a signal depending on the state of the dendritic branch of the synapse.

Immediately following a post-synaptic spike at time $\hat{t}$ (in a simulation with time step Δt), the synapse j at dendrite i, which has received a recent pre-synaptic spike, is sent a weight correction signal:

$$\Delta w_{ij} = \frac{\Delta u_i^m(\hat{t}) - \Delta u^m(\hat{t})}{\sqrt{\Delta t^2 + (\Delta u_i^m(\hat{t}) - \Delta u^m(\hat{t}))^2}} \tag{6}$$

where $\Delta u_i^m(\hat{t})$ and $\Delta u^m(\hat{t})$ are the changes in the partial membrane potential generated by dendrite i and the total membrane potential respectively just before the post-synaptic spike.

The weight correction signal has two main contributions: a signal from the dendrite Δu_i^m and a signal from the soma Δu^m. If we remove Δu^m, the rule will be:

$$\Delta w_i^d = \frac{\Delta u_i^m(\hat{t})}{\sqrt{\Delta t^2 + (\Delta u_i^m(\hat{t}))^2}} \tag{7}$$

and implement the following logic of the learning algorithm: if a post-synaptic spike occurs before the peak of the partial membrane potential (i.e. in the ascending phase of the membrane potential, $\Delta u_i^m > 0$), the synaptic weight will be increased, so that the next time the peak will occur earlier, i.e. closer to the post-synaptic spike time. On the other hand, if a post-synaptic spike occurs after the peak of the partial membrane potential, the synaptic weight will be decreased, and the peak will be delayed, i.e. again closer to the post-synaptic spike time.

Figure 4 (A) shows Δw_i^d for synapses with different weights as a function of the relative timing between the pre- and post-synaptic spikes. $\Delta w_i^d = 0$, i.e. no change in the weight occurs, if the post-synaptic spike coincides with the maximum of u_i^m. The synaptic plasticity learning window is different for synapses with different strength. Figure 4 (B) shows Δw_i^d as a function of the synaptic strength and the relative timing between the pre- and post-synaptic spikes. Depending on the strength of the synapse, for a pre-synaptic spike arriving shortly before the post-synaptic action potential, the synaptic weight is increased. If however, the pre-synaptic spike precedes the post-synaptic action potential by a relatively long time, the synaptic weight is reduced.

The implementation of the Δw_i^d rule does not cover the negative part of the learning window. Input spikes arriving after the post-synaptic spike are ignored. This is based on recent experiments presented in Froemke and Dan (2002), where a triplet of pre-post-pre synaptic spikes in layer II/III neurons of the visual cortex induced an LTP dominating result. For an input arriving as spike trains, which is the case modeled in this article, this means that the input spikes arriving before the post-synaptic spike are the ones that will determine the synaptic plasticity.

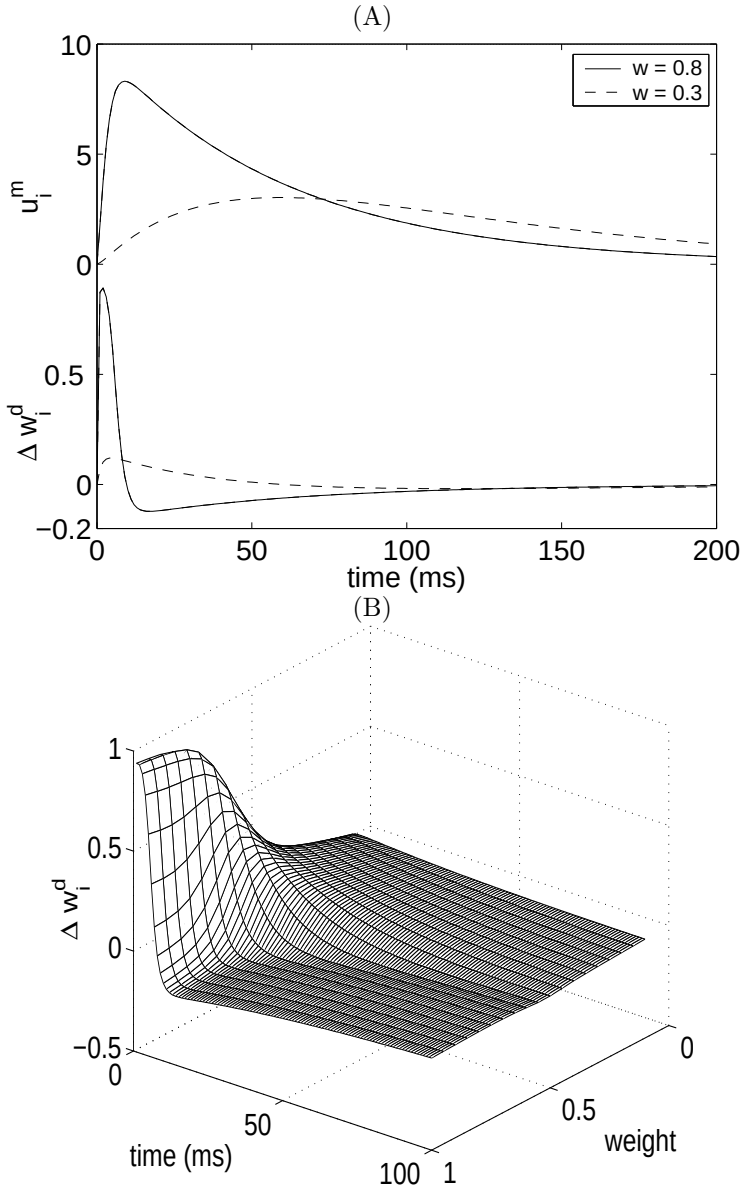

Figure 4: (A): The correction signal Δw_i^d that would be sent from the dendrite to a synapse with weight 0.8 or 0.3 in the event of a post-synaptic spike. $\Delta w_i^d = 0$, i.e. no change in the weight occurs, if the post-synaptic spike is at the point of the maximum of u_i^m. (B): Δw_i^d plotted against the relative time between the pre- and post-synaptic spikes ($t^{post} - t^{pre}$) and the synaptic weight w_i^d.

The Δu^m term implements a logic similar to Δu_i^m, but with respect to the total membrane potential and the total synaptic strength across all dendrites. It drives the neuron to fire close to the peak of the total membrane potential and has a weight normalization effect. Its role is to prevent the weights of the synapses from reaching very high values simultaneously, as well as to prevent a total decay in the synaptic strength. If the neuron is forced to always fire close to the peak of the membrane potential at the soma (coinciding with the peaks of the partial membrane potentials), the total synaptic strength will have to be relatively constant and thereby any changes of synaptic strength will lead to a redistribution (rather than a global gain/loss) of synaptic strength.

However, achieving a post-synaptic spike exactly at the peak of the total membrane potential is not always possible, and in most cases undesirable, since it will limit the noise handling capabilities of the neuron. A neuron trained to fire exactly at the peak of its membrane potential will respond only to a very specific temporal pattern without any noise. Therefore, if the post-synaptic spike is sufficiently close (for a predefined constant ε) to the peak of the membrane potential, Δu^m is ignored, i.e.:

$$\Delta u^m = 0 \text{ if } |\Delta w^m| \leq \varepsilon, \quad \text{where } \Delta w^m = \frac{\Delta u^m(\hat{t})}{\sqrt{\Delta t^2 + (\Delta u^m(\hat{t}))^2}}. \tag{8}$$

The value of ε allows control over the noise tolerance of the neuron.

Following the weight correction signal, the weights of the synapses are changed according to:

$$w_{ij}^{new} = \begin{cases} w_{ij}^{old} + \eta \Delta w_{ij}(1 - w_{ij}^{old}) & \text{if } \Delta w_{ij} \geq 0 \text{ and } \mathcal{F}_j \neq \emptyset; \\ w_{ij}^{old} + \eta \Delta w_{ij} w_{ij}^{old} & \text{if } \Delta w_{ij} < 0 \text{ and } \mathcal{F}_j \neq \emptyset; \\ w_{ij}^{old} - \eta^{decay} w_{ij}^{old} & \text{if } \mathcal{F}_j = \emptyset. \end{cases} \tag{9}$$

with a learning rate parameter η. The weight correction signal Δw_{ij} depends on the term Δu_i^m which can be non-zero only if a pre-synaptic spike has recently arrived at the synapse, i.e. $\mathcal{F}_j \neq \emptyset$. If the synapse has not received a pre-synaptic spike, the weight decays with a rate proportional to η^{decay}.

Following recent advances in the experimental evidence on synaptic plasticity in the biological neurons, several algorithms for learning with spiking neurons have been developed as functions of the relative timing of the pre- and post-synaptic spikes (Song et al., 2000; Natschläger and Ruf, 1999; Kempter et al., 2001; Panchev and Wermter, 2001; Rao and Sejnowski, 2001). However, these algorithms are explicitly based on a single pair of pre- and post-synaptic spike events and cannot be applied to more complex input stimuli involving multiple spikes arriving at the same synapse, e.g. spike trains. The aim of the algorithm presented here is to achieve synaptic plasticity which exceeds the applicability of the algorithms explicitly based on the relative timing between single pre-synaptic spikes and a post-synaptic spike, while still being consistent with the biological evidence for synaptic plasticity of the real neurons (see also the discussion in section 5). The learning algorithm presented in this section applies a

local synaptic plasticity rule, but goes beyond the simple relative spike timing, and incorporates functions of the membrane potential at the dendrite and the soma, as well as the synaptic strength. An earlier version of the algorithm, based on the same principles, has been shown to effectively train the neuron to an arbitrary precision, when responding to a temporal sequence of single pre-synaptic spikes within an interval of less than a hundred milliseconds (Panchev et al., 2002) as well as achieving weight normalization and an even distribution of the synaptic strength. The new version of the algorithm generalizes into a neural adaptation for input spike trains and temporal sequences spanning from a few hundred milliseconds up to several seconds.

One main difference in receiving an input as a spike train, in contrast to a single spike, is that the membrane potential is not a smooth curve, but it contains many local peaks. If the learning algorithm described by equations (6) - (9) is directly applied to such an input, it will converge to one of these local peaks and not drive the neuron towards firing close to the global maximum. However, a closer look into the membrane potential curve (figure 2 (B)) reveals that: (1) during the ascending phase, the local increase of the membrane potential is either steeper, or longer compared to the local decrease; and (2) during the global descending phase, the local decrease is either steeper or longer. Consequently, if during the global ascending phase, the neuron's firing time fluctuates moderately around a local peak, on average over several presentations of the same stimuli, it has a higher probability of coinciding with the local ascending too. Similarly, if the neuron's firing time fluctuates moderately around a local peak in the global descending phase, it has a higher probability of coinciding with a local decrease of the membrane potential. Similar arguments apply if, instead of fluctuating the post-synaptic spike times, the timing of the single spikes in the input contain noise, and therefore causing a fluctuation of the timing of the local peaks.

The above arguments provide the basis for the generalization of the synaptic plasticity rule for neurons receiving input as spike trains. They outline the necessary condition for consistency between the ascending and descending phases of membrane potentials generated by single spikes and by spike trains. The logic implemented by the Δw_i^d rule can be applied for post-synaptic potentials generated by spike trains. The additional condition required is a moderate fluctuation of the timings of the pre- or post-synaptic spikes. Similar to the real neural systems, there are many possible sources of such fluctuations, such as unreliable synapses, noise in the input spike trains, etc. It is well known that most real neurons receive input spike trains containing irregular spike timings and a level of noise (Softky and Koch, 1993; Shadlen and Newsome, 1994). Such a mechanism is employed in the experiments presented later in this article.

3.2 Plasticity in the synapses attached to the soma

The function performed by the synapses attached to the soma is relatively simple. They are used mainly for lateral connections between competitive or cooperative neurons. These synapses are trained using a simpler rule implementing

an asymmetric spike-timing dependent learning window with length $t^{lwin}ms$. Here the change in the synaptic strength depends only on the normalized relative timing between the pre- and post-synaptic spikes: $\Delta t_i = \frac{t_i^{pre} - t_i^{post}}{t^{lwin}}$. After a pre- or post-synaptic spike, the weight change for the synapses is calculated as:

$$\Delta w_i = \begin{cases} A\Delta t_i \exp(-\Delta t_i) & \text{if } \Delta t_i > 0, \\ -AB\Delta t_i \exp(-\Delta t_i) & \text{if } \Delta t_i < 0. \end{cases} \tag{10}$$

and the weights are changed according to:

$$w_i^{new} = \begin{cases} w_i^{old} + \eta^s \Delta w_i (1 - w_i^{old}) & \text{if } \Delta w_i > 0, \\ w_i^{old} + \eta^s \Delta w_i w_i^{old} & \text{if } \Delta w_i < 0. \end{cases} \tag{11}$$

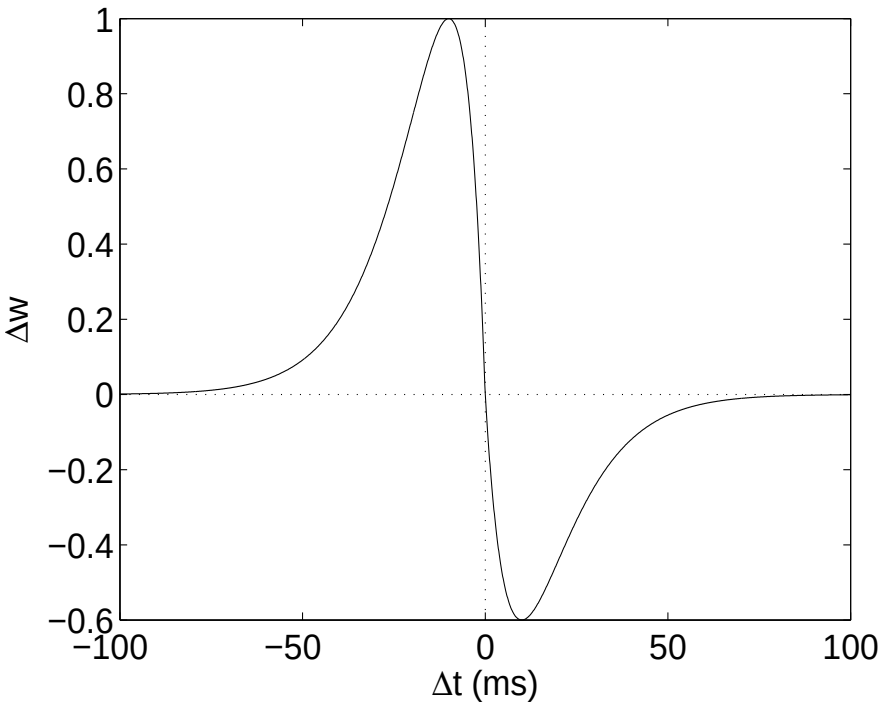

Figure 5: The asymmetric learning window for synapses attached to the soma. $A = e, B = 0.6$ and $t^{lwin} = 100ms$.

Figure 5 shows the asymmetric weight change learning window described by the above equations. If the pre-synaptic spike precedes the post-synaptic spike, the weights are increased. Respectively, if the pre-synaptic spike arrives after the post-synaptic one, the weights are reduced.

4 Experiments

There are many different systems in the brain where recognition of temporal information encoded in incoming spike trains is required. For instance, the spike trains generated in the auditory system are used in the higher cortical auditory regions for the identification of phonemes and words (Hopfield and Brody, 2000). Here, we will present some experiments where we used the new model of a neuron with dynamic synapses and active dendrites to recognize words based on input of phoneme sequences.

4.1 Short words with repeating phonemes

One aim in the first experiment is to explore how the model will perform the recognition of words which have the same phonemes, but in a different order, as well as words with repeating phonemes. In order to test this, we developed a neural network for the recognition of the words "bat", "tab", "babat" and "tatab". Each word is represented as a sequence of spike trains generated by input neurons representing the phonemes "æ", "$æ_2$", "b", "b_2", "t" and "t_2". The active neurons and the order of the spike trains defines the word in the input. Here we assume that neurons representing the same phoneme (e.g. "æ") will fire with different probability. Consequently, there are different neurons (e.g. for "æ", "$æ_2$", etc.) responding to the first, second, etc. occurrences of the same phoneme.

Figure 6 shows the architecture of the neural network. The six input units are leaky integrate-and-fire neurons, and are driven by a decaying supra-threshold current with random fluctuation within a predefined range. This random fluctuation provides the noise in the timings of the single input spikes that is necessary for the training of the neurons on the output layer (as discussed in section 3.1). The decay rate of the current driving all input neurons is the same, so all input spike trains have approximately the same mean firing rate. The role of the input neurons is to represent the incoming word as a sequence of spike trains (see input phonemes column in figure 8). The length of each phoneme (i.e. the delay of the onset of the next phoneme in the word) is 100ms. In all experiments presented in this article, the input spike trains were generated using neurons with persistent dendrites which model the activity of a type of cortical neuron presented in Egorov et al. (2002).

The output units are the leaky integrate-and-fire neurons with dynamic synapses and active dendrites presented in section 2.1. Their task is to recognize the temporal pattern of input spike trains, i.e. to recognize the words given their particular sequences of phonemes. As presented in the previous sections these neurons have an effective mechanism and learning algorithm for selective response of the temporal structure encoded in the onset times of a set of spike trains. The output neurons form a four by four map with all to all lateral connections via synapses attached to the soma. Each output neuron receives connections from all phoneme input neurons via dynamic synapses attached to different active dendrites.

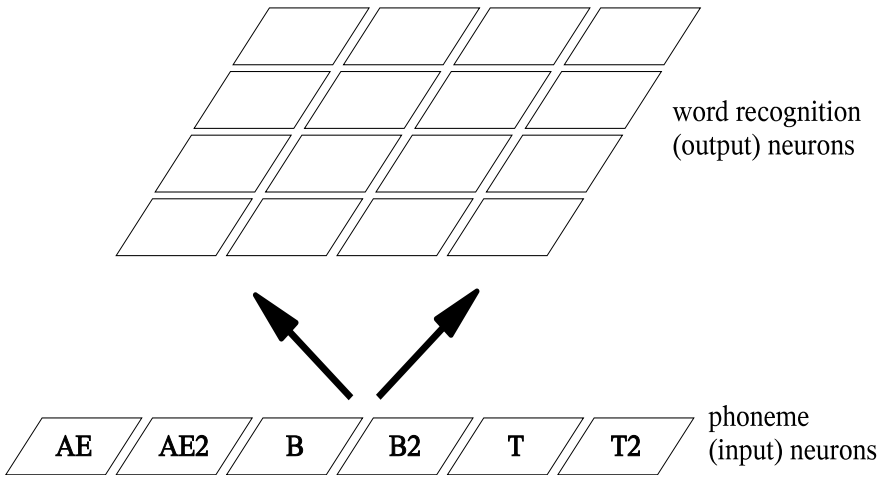

Figure 6: The network architecture for short word phoneme sequence recognition. All word recognition neurons receive connections from the phoneme neurons via dynamic synapses attached to active dendrites and are connected to each other via synapses attached to the soma.

After training (see appendix A.2 for the procedure and parameters used), the network developed a well formed tonotopic word map of small clusters of neurons recognizing each word (figure 7). Clusters representing words that sound similar are close on the map. The neighborhood was determined by the number of shared phonemes and partial phoneme sequence overlap between the words.

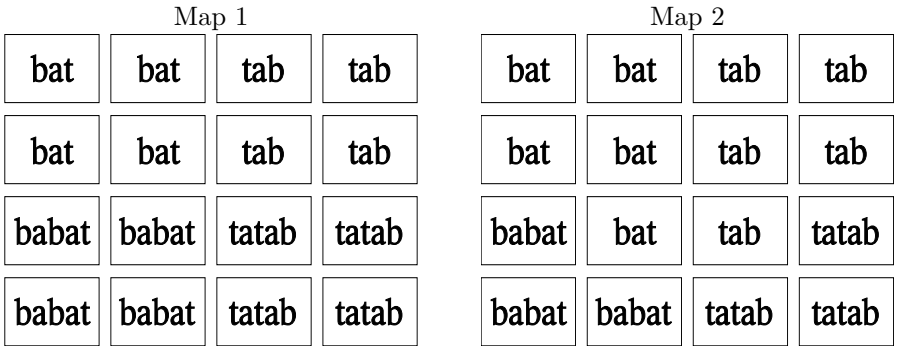

Figure 7: Two typical map formations of the words "bat", "tab", "babat" and "tatab". Each neuron responds only to one particular word. Words that sound similar, i.e. have similar phoneme sequences are recognized by neurons in neighboring clusters.

Figure 8 shows the input and output spikes for the words "bat" [b æ t], "tab" [t æ b] and "babat" [b æ b_2 $æ_2$ t]. In the former two words, the same input neurons are active, but based on the temporal order of the spike trains, the output neurons are able to distinguish between the different words. The word "babat" contains all the phonemes of the word "bat". However, the temporal distribution of the phonemes is different (a longer delay between the æ and t phonemes), and there are two additional phoneme inputs. As a result, the neurons representing "bat" do not respond to this phoneme sequence. It is recognized only by the sequence detectors for "babat".

There are different sources and types of noise in the biological neural systems. Although the neurons in our model were trained only with noise in the timing of the single spikes within the train, the model exhibits robust behavior for several different types of noise. The output neurons have reliable responses when the onset times of the input spike trains vary by up to 40ms. Further, the output neurons reliably detected temporal sequences in the presence of relatively high levels of noise from the non-active inputs. Figure 9 shows such examples for

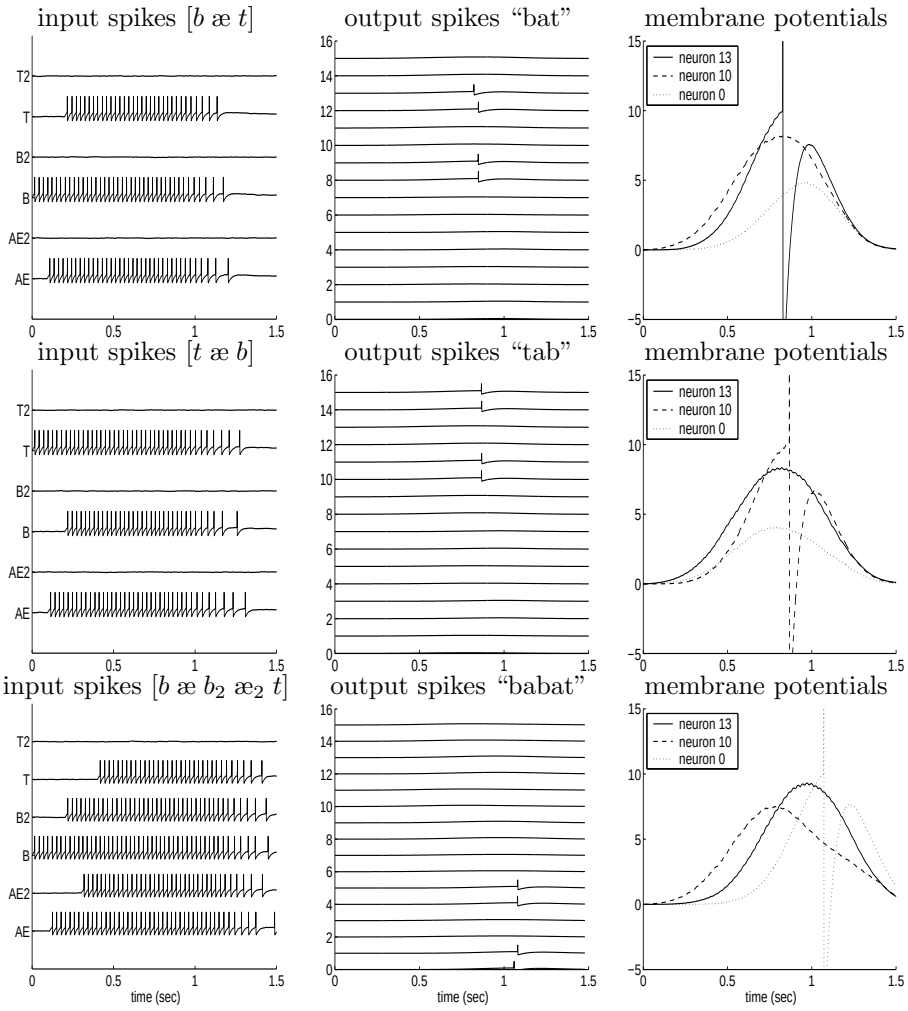

Figure 8: Recognizing the words "bat" (*top*), "tab" (*middle*) and "babat" (*bottom*). *Left column*: The input spike trains for the phonemes representing the three words. *Middle column*: Output spikes of the small clusters of neurons recognizing the particular word. Each plot line represents the activity of the 16 neurons from map 1 on figure 7. Neuron number 0 is the bottom left unit from the map, neuron 3 is the bottom right unit and neuron 15 is the top right unit. *Right column*: Total membrane potentials at the soma for three neurons recognizing the three words. Neuron 13 responds only to the word "bat", neuron 10 responds only to the word "tab" and neuron 0 responds only to the word "babat".

the words “tatab” and “bat”. The words are correctly recognized even in the presence of additional noise in the onset times and random spikes from the non-active phonemes.

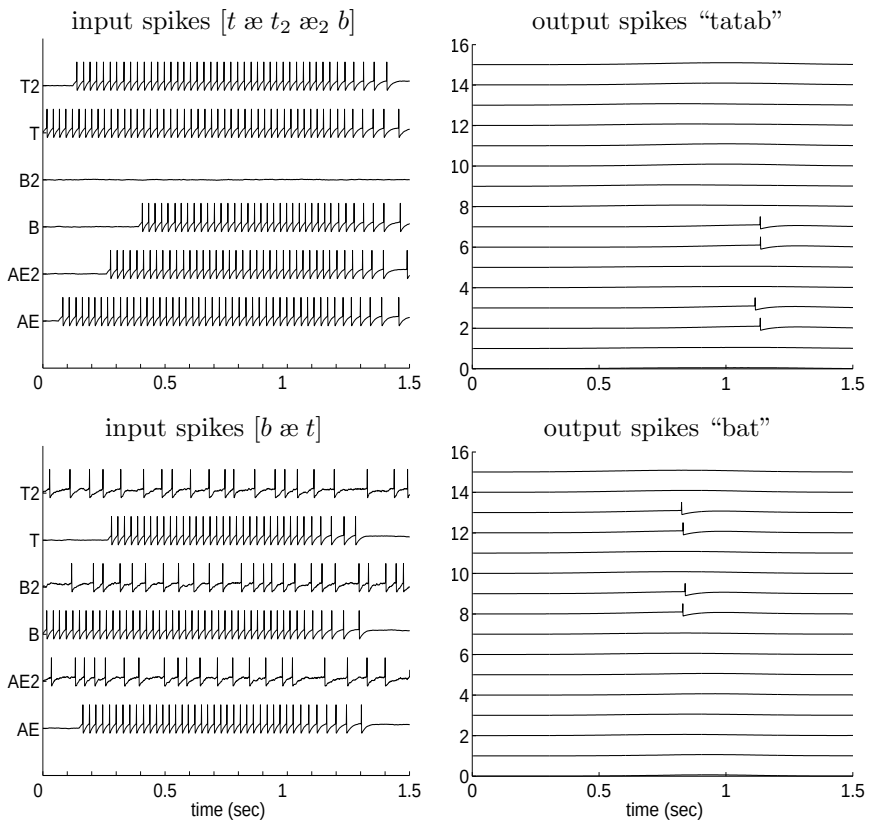

Figure 9: Processing “tatab” with noise in the onset times(up to 40%); and “bat” with noise in the onset times and lower frequency random spikes from the non-active phonemes.

4.2 Recognizing language instructions to a robot

Increasingly more attention in intelligent robotics has been paid to robots that are capable of interacting with people, responding to voice commands or deriving an internal representation from a language description (Wermter et al., 2003; Lauria et al., 2002; Yoshizaki et al., 2002; Kyriacou et al., 2002; Bugmann

et al., 2001; Crangle, 1997). Such robots exhibit learning and acquire adaptive behavior which cannot be completely preprogrammed in advance. Natural language can be used for the description of relatively specific rules or action sequences, and could be the primary means of communication between the user (sometimes computer-language-naive) and the robot.

In this section we present part of a system which is being implemented for language instructions to a robot. The goal of the system is to instruct a robot for navigation and grasping tasks. The language corpus consists of words forming instruction phrases like: "Bot go", "Bot stop", "Bot turn left" or "Bot lift". The words used in this experiment, together with their phonemic representation are shown in table 1. The overall architecture is shown in figure 10. Phoneme sequences are used as an input to the neural network module which recognizes the words. The input phoneme sequences can be generated from real speech using a speech-to-phoneme recognition module, e.g. parts of Sphinx (Lee et al., 1990). The output words as a sequence are sent to the robot, and the formed instruction interpretation is executed.

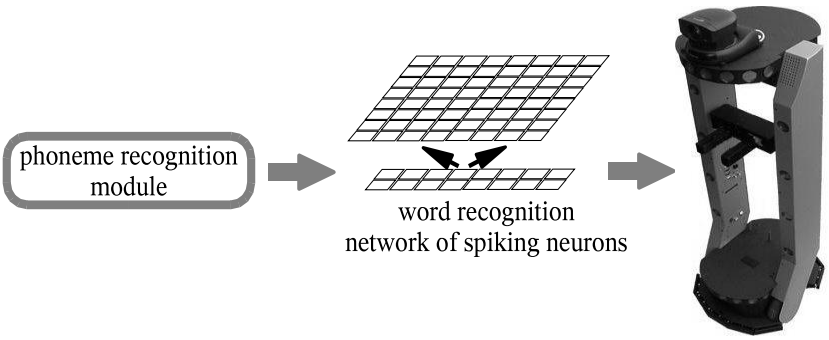

Figure 10: Robot's language instruction system.

The architecture of the neural network module is shown in figure 11. The input layer contains one neuron for each phoneme. As in the previous experiment, the input units are integrate-and-fire neurons driven by a decaying supra-threshold current injection containing noise in the amplitude. The output units on the next layer are integrate-and-fire neurons with dynamic synapses and active dendrites. They are organized in an eight by eight map with all-to-all lateral connections via synapses attached to the soma. Each output neuron receives connections from all phoneme neurons via dynamic synapses attached to different active dendrites.

The network was trained using correct, noise-free phoneme sequences as produced by a speech-to-phonemes recognition module. Based on the observed average length of the phonemes produced by this module, the length of each

Table 1: Phoneme sequences of the words.

word	phoneme sequence
GO	[*g ow*]
BOT	[*b ao t*]
TURN	[*t er n*]
RIGHT	[*r ay t*]
LEFT	[*l eh f t*]
LIFT	[*l ih f t*]
DROP	[*d r ao p*]
STOP	[*s t ao p*]

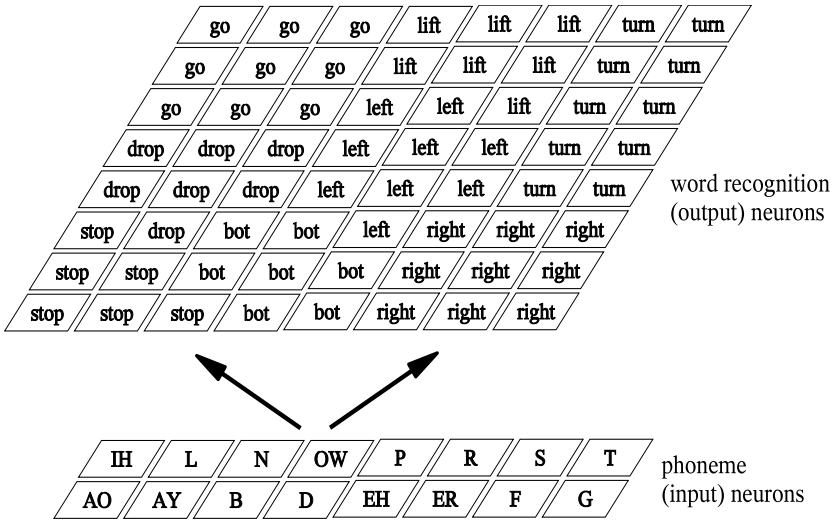

Figure 11: Network architecture and self-organized map of the word recognition neurons after training.

consonant was set to 50ms and each vowel to 100ms. As in the previous experiment, the length of a phoneme determines the delay of the onset time of the spike train representing the next phoneme in the word.

After training (see section A.2 for details on the training procedure and parameters) each of the output neurons adapted towards recognizing a particular phoneme sequence, i.e. a particular word. The word recognition neurons formed small localized clusters of units recognizing the same word (see figure 11). Typically a subset of the neurons within a cluster responded depending on the input noise level. Altogether, the neurons within a cluster covered a range of noise-dependent fluctuations in the input phoneme sequence for the particular word. Furthermore, the clusters were organized in a tonotopic word map, i.e. neurons recognizing words that have similar sounds, such as "drop" and "stop" or "left" and "lift", occupied neighboring clusters.

There are three different types of noise produced by a phoneme recognition module: (1) the lengths of the same phonemes are not exactly the same in each utterance; (2) in most sequences there are additional phonemes which do not belong to the word being represented; and (3) in some sequences there are missing phonemes. The network was tested on all three types of noise.

Figure 12 shows the processing of the words "drop" [$draop$] and "stop" [$staop$] with noise in the phoneme lengths and additional phonemes being activated. The actual network inputs are [$d\,r\,ao\,f\,l\,p$] and [$s\,t\,f\,aol\,p$] respectively, with noise in the onset time of the phoneme spike trains. The network responds reliably if some extra phonemes are added to the input. Due to the weight decay, a neuron which recognizes a particular word has synapses with a significant strength only for phoneme inputs that belong to that word. The synaptic strength for input phonemes which do not belong to the word is negligible. As a result an additional input phoneme does not have a significant influence on the selective response of the neuron recognizing its word. However, if the phonemes, which are added as noise, constitute a valid phoneme sequence, it might be recognized by other neurons. Furthermore, the output neurons have been found to respond reliably if the phoneme lengths varied by up to 40%.

If however, the input phoneme sequence representing a particular word has missing phonemes, it will most likely not be detected as a valid sequence and the neurons recognizing that word will not respond. Figure 12 shows one such example for the word "right" [$r\ ay\ t$], where the "*ay*" phoneme has been substituted with "*ih*". No output neuron from the network responded to this phoneme sequence. Although the neurons which recognize the word "right" are still significantly activated, the input stimulus can not trigger a post-synaptic spike.

Figure 13 shows the processing of the instruction "Bot turn left". Each of the words has been recognized by a particular set of neurons. The input is a set of three noisy phoneme sequences [$b\ ao\ t$], [$t\ er\ n$] and [$l\ eh\ f\ t$]. Since the model was trained to perform recognition only on a single word level, there is a delay between the consecutive words. The output is a sequence of active clusters representing each of the three words. The recognized words will be sent to the robot which will execute the given instruction.

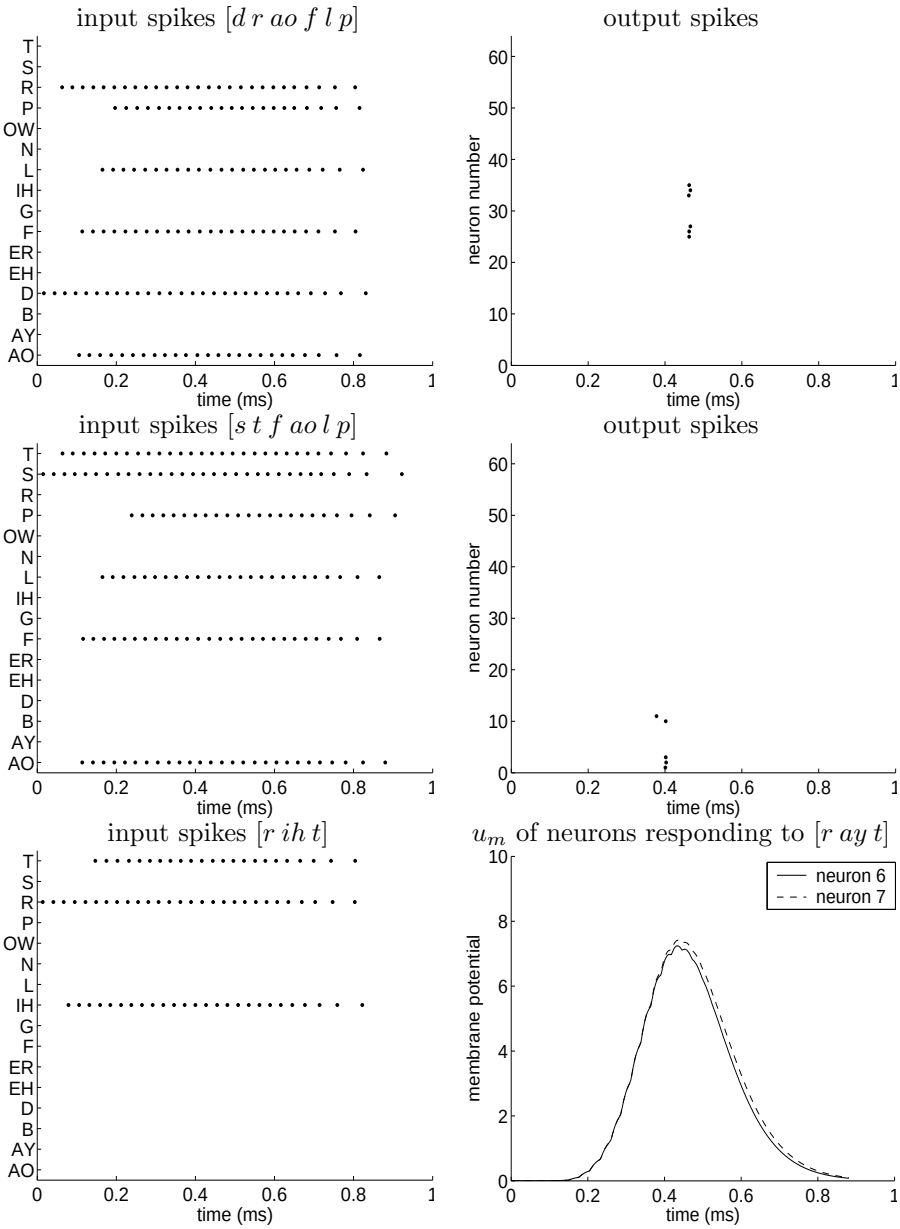

Figure 12: Processing the words "drop", "stop", and "right". *Top and middle*: Small clusters of neurons responding to noisy input for the words "drop" and "stop". *Bottom*: No output neurons from the network responded to the [*r ih t*] sequence. Although the neurons recognizing the word "right" [*r ay t*] are still significantly potentiated, the input stimulus is not sufficient to trigger a post-synaptic spike.

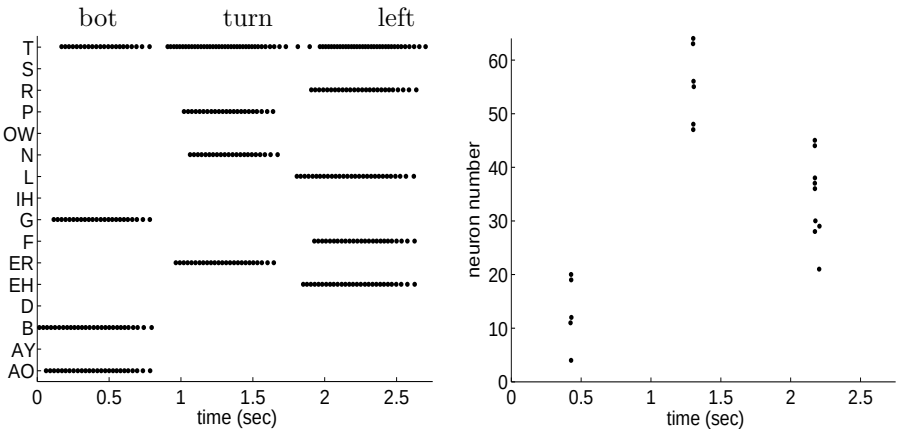

Figure 13: Processing the words "Bot turn left".

5 Discussion and Conclusion

We presented a novel model of a spiking neuron for detecting the temporal structure of spike trains and applied it in a network of spiking neurons for recognizing words for robot language instructions. The neuron exploits the dynamics of the synapse and the active properties of the dendrites in order to implement an efficient "delay" mechanism which will maximize its response to a particular input sequence. The spatio-temporal structure of the "delay" mechanism is encoded in the synaptic strengths. We have developed and presented a synaptic plasticity rule which is capable of efficiently tuning the neuron. After training, the neuron is sensitive to a specific temporal structure of the input spikes, and is selectively responsive to stimulus with spatio-temporal patterns that matches its "delay" structure.

The synaptic plasticity rules employed in this work follows recent neuro-physiological experimental results (reviewed in Bi and Poo (2001); Kepecs et al. (2002); Roberts and Bell (2002)). The implementation of plasticity of the synapses attached to the soma, section 3.2, is a direct approximation of the asymmetric learning window as a function of the relative timing of the pre- and post-synaptic spikes, as presented in several neuro-physiological studies (Markram et al., 1997; Bi and Poo, 1998; Feldman, 2000). In general, if the pre-synaptic spike precedes the post-synaptic spike, the weights are increased, and respectively, if the pre-synaptic spike arrives after the post-synaptic one, the weights are reduced. The magnitude of the change depends on the time difference between the two spikes.

The plasticity rule for the synapses attached to the active dendrites is some-

what more complicated. Studies presented in Nishiyama et al. (2000) have revealed a second negative part of the learning window. Negative changes in the synaptic strength have been observed if the pre-synaptic spike precedes the post-synaptic spike by a relatively long time. Nishiyama and colleagues discussed possible causes for the appearance of this so called "paradoxical zone" of the learning window as the different spatio-temporal pattern of calcium elevation or the activation of local inhibitory circuits. A possible cellular mechanism for this part of the learning window has been discussed in Bi (2002). The synaptic plasticity rule presented in section 3.1 implements and takes advantage of the "paradoxical" part of the learning window. Our work shows that in fact, if the active properties of the dendrites are taken into account, this part of the learning window has a possible computational interpretation and plays a critical role in the synaptic plasticity process. It allows the neuron to adapt and maximize its response to a specific spatio-temporal pattern of incoming action potentials.

Furthermore, the new model of spiking neuron was tested in the development of a self-organizing tonotopic word map of cooperative and competitive neurons. We have achieved a self-organization of the neurons representing the words which is based on the phonetic properties of the words, i.e. the spatio-temporal structure of the input stimuli. The same word is recognized by a small cluster of neighboring neurons, and the clusters representing words that sound similar are close on the map. This is a possible intermediate representation and link between the frequency-based tonotopic maps in the auditory cortex (Kaas and Hackett, 2000) and the distributed semantic-based representations of words in the brain (Pulvermüller, 1999).

The performance of the model was tested under different regimes and levels of noise. The results suggest that it can be successfully applied in processing of real data. An example of such application is the presented network for word recognition as part of a robot language instruction system. Each word was represented as noisy temporal sequence of phonemes generated as an intermediate output by a speech-to-phoneme recognition module. The output neurons were able to learn the spatio-temporal structure of the phonemic representation and recognize words from the input under noisy conditions.

Further experiments and applications of the model will involve generating the input sequence directly from the cochleagrams of the real speech input (Slaney and Lyon, 1993; Wermter and Panchev, 2002), instead of phoneme representations, and thereby achieving a model of neural speech recognition with spiking neurons. The results from the experiments presented here, i.e. the output of the self-organizing map, will be used as an input for the future development of a language processing network of spiking neurons for representation of whole phrases and sentences (Pulvermüller, 2002). As part of this, the model of a spiking neuron with dynamic synapses and active dendrites is also being further developed towards achieving more complex neural structures such as cell assemblies and synfire chains.

Acknowledgments

This research was partially supported by the MirrorBot project, EU FET-IST program grant IST-2001-35282.

A Appendix: Methods

This appendix gives details on the parameters and training procedures used in the experiments presented in the article.

A.1 Active dendrites

The time constant τ_i^d and the resistance R_i^d of the active dendrites are defined as functions of I_i^{s*} which is the maximum of $I_i^s(t)$ since the last pre-synaptic spike. Following equation 1, for a single spike arriving at the synapse, $I_i^{s*} = w_{ij}\rho(\Delta t^{(f)})/tau^s$ (see also figure 14). τ_i^d is defined as:

$$\tau_i^d \triangleq \tau_i^d(I_i^{s*}) = \tau^m - \frac{\tau^m - \tau^s}{1 + e^{-(\tau^s I_i^{s*} + \frac{1}{\tau^s})}} \tag{12}$$

For low synaptic input, this leads to values of τ_i^d approaching the time constant of the soma τ^m, and for high inputs τ_i^d approaches the time constant of the synapse τ^s which is usually much faster than τ^m.

Furthermore, R_i^d is defined such that for a single spike at a synapse with strength w_{ij}, the value of the maximum of the membrane potential at the soma is directly proportional to the neuron's firing threshold θ , i.e. equals $w_{ij}\theta$. Such a choice for R_i^d facilitates the control over the neuron in simulation and improves its adaptation during learning. Thus, R_i^d yields the equation:

$$R_i^d \triangleq R_i^d(\tau_i^d) \triangleq R_i^d(I_{i*}^s) = \frac{\theta}{BE} \tag{13}$$

$$\text{with } A = \frac{1}{\tau^s - \tau_i^d},\ B = A\frac{R^m}{\tau^m},\ C = \frac{\tau^m\tau^s}{\tau^m - \tau^s},\ D = \frac{\tau^m\tau_i^d}{\tau^m - \tau_i^d} \text{ and} \tag{14}$$

$$E = min\left(-Ce^{-\frac{t}{\tau^s}} + De^{-\frac{t}{\tau_i^d}} + (C-D)e^{-\frac{t}{\tau^m}}\right),\ t > 0 \tag{15}$$

A.2 Running parameters and training procedure

The following parameters were used for the simulations presented in this article:

For all neurons: $\tau^s = 2ms$, $\tau^m = 60ms$ and $R^m = 15$.

For the dynamic synapse (equation 3): $\sigma = 0.3$ and $\mu = 0.18$. The normalized time $\Delta t^{(f)}$ between the current and the earliest spikes in $\mathcal{F}_j$ is: $\Delta t^{(f)} = \frac{t^{current} - t^{first}}{t^{(f)}}$, with $t^{(f)} = 1sec$ in the experiment presented in section 4.1 and $t^{(f)} = 0.4sec$ in the experiment presented in section 4.2.

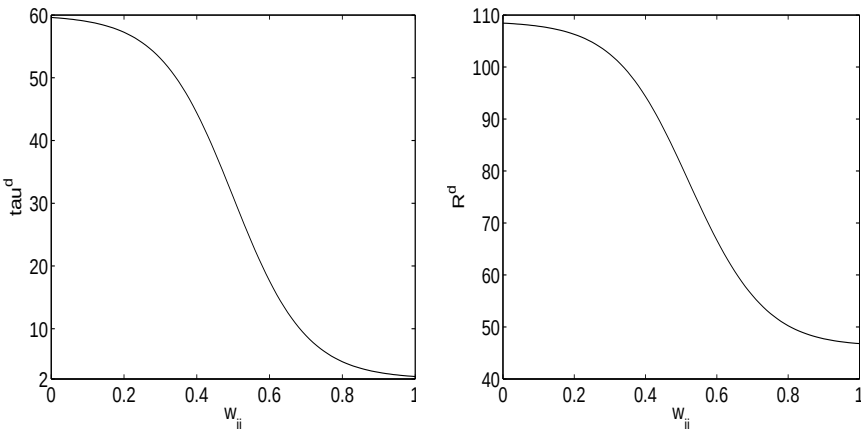

Figure 14: τ^d and R^d plotted as a function of the weight of the synapse at which a single spike has arrived (with $\tau^s = 2ms$ and $\tau^m = 60ms$).

The following parameters were found to be optimal for the training of the neurons:

For the synapses attached to active dendrites: $\epsilon = 0.03$, $\eta = 0.03$ and $\eta^{decay} = 0.001$.

For the synapses attached to the soma: $A = 0.005e$, $B = 0.6$, $tau_l = 40ms$, $t^{lwin} = 100ms$ and $\eta^s = 0.03$.

The network presented in section 4.1 was trained for 10000 epochs. During one epoch correct noise-free sequences were presented once for each of the four words. After each epoch, the weights of the lateral connections were decreased by -0.00001. In order to examine to exact responsiveness of each neuron to a particular spatio-temporal pattern, the lateral connections were removed during the tests.

The network presented in section 4.2 was trained for 15000 epochs. During one epoch correct noise-free sequences were presented for each of the eight words. Within each epoch of the first 3000 epochs, the shorter words were represented more frequently. For the remaining 12000 epochs, each word was represented once within an epoch. After each epoch, the weight of the lateral connections were decreased by -0.00001.

References

Abbott, L., Varela, J., Sen, K., and Nelson, S. (1997). Synaptic depresion and cortical gain function. *Science*, 275:220–224.

Aradi, I. and Holmes, W. (1999). Active dendrites regulate spatio-temporal synaptic integration in hippocampal dentate granule cells. *Neurocomputing*, 26-27:45–51.

Beggs, J. M., Moyer, J. R., McGann, J. P., and Brown, T. H. (2000). Prolonged synaptic integration in perirhinal cortical neurons. *Journal of Neurophysiology*, 83:3294–3298.

Bi, G. and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike-timing, synaptic strength and postsynaptic cell type. *Journal of Neuroscience*, 18:10464–10472.

Bi, G. and Poo, M. (2001). Synaptic modification by correlated activity: Hebb's postulate revised. *Annual Reviews in Neuronscience*, 24:139–166.

Bi, G.-Q. (2002). Spatio-temporal specificity of synaptic plasticity: cellular rules and mechanisms. *Biological Cybernetics*, 87:319–332.

Bugmann, G., Lauria, S., Kyriacou, T., Klein, E., Bos, J., and Coventry, K. (2001). Using verbal instructions for route learning: Instruction analysis. Proc. TIMR 2001 UMC-01-4-1, Department of Computer Science, Manchester University.

Buzsáki, G. and Kandel, A. (1998). Somadendritic backpropagation of action potentials in cortical pyramidal cells of awake rat. *Journal of Neurophysiology*, 79:1587–1591.

Crangle, C. (1997). conversational interfaces to robots. *Robotica*, 15:117–127.

Debanne, D., Ghwiler, B. H., and Thompson, S. M. (1998). Long-term synaptic plasticity between pairs of individual ca3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology*, 507.1:237–247.

Egorov, A., Hamam, B., Fransen, E., Hasselmo, M., and Alonso, A. (2002). Graded peristent activity in ethorhinal cortex neurons. *Nature*, 420.

Feldman, D. E. (2000). Timing-based ltp and ltd at vertical inputs to layer ii/iii pyramidal cells in rat barrel cortex. *Neuron*, 27:45–56.

Froemke, R. and Dan, Y. (2002). Timing-based ltp and ltd at vertical inputs to layer ii/iii pyramidal cells in rat barrel cortex. *Nature*, 416:433–438.

Häusser, M. and Clark, B. (1997). Tonic synaptic inhibition modulates neuronal output pattern and spatiotemporal synaptic integration. *Neuron*, 19:665–678.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117:500–544.

Hopfield, J. and Brody, C. (2000). What is a moment? "cortical" sensory integration over a brief interval. *Proceedings of National Academy of Sciences*, 97:13919–13924.

Horn, D., Levy, N., and Ruppin, E. (1999). The importance of nonlinear dendritic processing in multimodular memory networks. *Neurocomputing*, 26-27:389–394.

Johnston, D., Hoffman, D., Colbert, C., and Magee, J. (1999). Regulation of back-propagating action potentials in hippocampal neurons. *Current Opinion in Neurobiology*, 9:288–292.

Kaas, J. and Hackett, T. (2000). Subdivisions of the auditory cortex and processing streams in primates. *Proceedings of National Academy of Sciences*, 97(22):11793–11799.

Kempter, R., Gerstner, W., and van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Computation*, 13(12):2709–2741.

Kepecs, A., van Rossum, M., Song, S., and Tegner, J. (2002). Spike-timing-dependent palsticity: common themes and divergent vistas. *Biological Cybernetics*, 87:446–458.

Kistler, W. M., Gerstner, W., and van Hemmen, J. L. (1997). Reduction of the hodgkin-huxley equations to a single-variable threshold model. *Neural Computation*, 9(5):1015–1045.

Kyriacou, T., Bugmann, G., and Lauria, S. (2002). Vision-based urban navigation procedures for verbally instructed robots. In *Proceedings of the 2002 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, pages 1326–1331.

Larkum, M., Zhu, J., and Sakmann, B. (1999). A new cellular mechanism for coupling input arriving at different cortical layers. *Nature*, 398:338–341.

Lauria, S., Bugmann, G., Kyriacou, T., and Klein, E. (2002). Mobile robot programming using natural language. *Robotics and Autonomous Systems*.

Lee, K., Hon, H., and Reddy, R. (1990). An overview of the SPHINX speech recognition system. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(1):35–45.

Liaw, J. and Berger, T. (1996). Dynamic synapse: A new concept of neural representation and computation. *Hippocampus*, 6:591–600.

Maass, W. and Bishop, C. M., editors (1999). *Pulsed Neural Networks*. MIT-Press.

Magee, J. and Johnston, D. (1997). A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. *Science*, 275:209–213.

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275:213–215.

Mel, B. W., Ruderman, D. L., and Archie, K. A. (1998). Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. *The Journal of Neuroscience*, 18(11):4325–4334.

Natschläger, T., Maass, W., and Zador, A. M. (2001). Efficient temporal processing with biologically realistic dynamic synapses. *Network: Computation in Neural Systems*, 12:75–87.

Natschläger, T. and Ruf, B. (1999). Pattern analysis with spiking neurons using delay coding. *Neurocomputing*, 26-27(1-3):463–469.

Nishiyama, M., Hong, K., Mikoshiba, K., Poo, M.-M., and Kato, K. (2000). Calcium release from internal stores regulates polarity and input specificity of synaptic modification. *Nature*, 408:584–588.

Panchev, C. and Wermter, S. (2001). Hebbian spike-timing dependent self-organization in pulsed neural networks. In *Proceedings of World Congress on Neuroinformatics.*, pages 378–385, Austria.

Panchev, C., Wermter, S., and Chen, H. (2002). Spike-timing dependent competitive learning of integrate-and-fire neurons with active dendrites. In *Lecture Notes in Computer Science. Proceedings of the International Conference on Artificial Neural Networks*, pages 896–901, Madrid, Spain. Springer.

Pantic, L., Torres, J. J., Kappen, H. J., and Gielen, S. C. (2002). Associative memory with dynamic synapses. *Neural Computation.*

Paré, D., Shink, E., Gaudreau, H., Destexhe, A., and Lang, E. (1998). Impact of spontaneous synaptic activity on the resting properties of cat neocortial pyramidal neurons in vivo. *Journal of Neurophysiology*, 79:1450–1460.

Poirazi, P. and Mel, B. (2001). Impact of active dendrites and structural plasticity on the memory capacity of neural tissue.

Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22:253–336.

Pulvermüller, F. (2002). A brain perspective on language mechanisms: from discrete neuronal ensembles to serial order. *Progress in Neurobiology*, 67(2):85–111.

Rall, W. (1989). Cable theory for dendritic neurons. In Koch, C. and Segev, I., editors, *Methods in Neuronal Modeling*, Cambridge, Massachusetts. MIT Press.

Rao, R. P. N. and Sejnowski, T. J. (2001). Spike-timing-dependent hebbian plasticity as temporal difference learning. *Neural Computation*, 13(10):2221–2237.

Roberts, P. and Bell, C. (2002). Spike timing dependent synaptic plasticity in biological systems. *Biological Cybernetics*, 87:392–403.

Schutter, E. D. and Bower, J. M. (1994a). An active membrane model of the cerebellar purkinje cell. i. simulation of current clamps in slice. *Journal of Neurophysiology*, 71:375–400.

Schutter, E. D. and Bower, J. M. (1994b). An active membrane model of the cerebellar purkinje cell: Ii. simulation of synaptic responses. *Journal of Neurophysiology*, 71:401–419.

Schutter, E. D. and Bower, J. M. (1994c). Simulated responses of cerebellar purkinje cell are independent of the dendritic location of granule cell synaptic inputs. In *Proceedings of National Academy of Science, USA*, volume 91, pages 4736–4740.

Segev, I., Fleshman, J. W., and Burke, R. E. (1989). Compartmental models of complex neurons. In Koch, C. and Segev, I., editors, *Methods in Neuronal Modeling*, chapter 3, pages 63–96. MIT Press, Cambridge, Massachusetts.

Senn, W., Markram, H., and Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation*, 13(1):35–67.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4):569–79.

Sharkey, N. and Heemskerk, J. (1997). The neural mind and the robot. In Browne, A., editor, *Neural Network Perspectives on Cognition and Adaptive Robotics*, pages 169–194. IOP Press, London.

Sharkey, N. and Ziemke, T. (2000). Life, mind and robots: the ins and outs of embodiment. In Wermter, S. and Sun, R., editors, *Hybrid Neural Systems*, pages 313–332. Springer Verlag, Heidelberg.

Sharkey, N. and Ziemke, T. (2001). Mechanistic vs. phenomenal embodiment: Can robot embodiment lead to strong ai. *Cognitive Systems Research*, 2(4):251–262.

Slaney, M. and Lyon, R. (1993). On the importance of time - a temporal representation of sound. In Cooke, M., Beet, S., and Crawford, M., editors, *Visual representations of speech signals*, pages 95–116. John Wiley & Sons.

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *Journal of Neuroscience*, 13:334–350.

Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive hebbian learning though spike-timing dependent synaptic plasticity. *Nature Neuroscince*, 3:919–926.

Spencer, W. A. and Kandel, E. R. (1961). Electrophysiology of hippocampal neurons. iv. fast pre-potentials. *Journal of Neurophysiology*, 24:272–285.

Spruston, N., Schiller, Y., Stuart, G., and Skmann, B. (1995). Activity dependent action potential invasion and calcium influx into hippocampal ca1 dendrites. *Science*, 268:297–300.

Stuart, G., Spruston, N., and Häusser, M., editors (2001). *Dendrites*. Oxford University Press.

Stuart, G., Spruston, N., Sakmann, B., and Häusser, M. (1997). Action potential initiation and back-propagation in neurons of the mammalian cns. *Trends in Neurosciences*, 20(3):125–131.

Tsodyks, M., Pawelzik, K., and Markram, H. (1998). Neural networks with dynamic synapses. *Neural Computation*, 10(4):821–835.

Webb, B. (2001). Can robots make good models of biological behaviour. *Behavioral and Brain Sciences*, 24:1033–1050.

Wermter, S., Elshaw, M., and Farrand, S. (2003). A modular approach to self-organisation of robot control based on language instruction. *Connection Science*.

Wermter, S. and Panchev, C. (2002). Hybrid preference machines based on inspiration from neuroscience. *Cognitive Systems Research*, 3(2):255–270.

Yoshizaki, M., Nakamura, A., and Kuno, Y. (2002). Mutual assistance between speech and vision for human-robot interface. In *Proceedings of the 2002 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, pages 1308–1313.

Zucker, R. (1989). Short-term synaptic plasticity. *Annual Reviews in Neuroscience*, 12:13–31.