Processing of temporal structured information by spiking neural networks

A dissertation presented
by
Cătălin V. Rusu
under supervision
of
prof. Leon Ţâmbulea

– summary –

Submitted to the
Babeş-Bolyai University
in partial fulfilment of the
requirements for the degree of
Doctor of Computer Science

2012
Cluj Napoca

Key words: spiking neural networks, synaptic plasticity, robotic control, evolutionary techniques, magnetic stimulation

CONTENTS

CHAPTER 1

INTRODUCTION

How artificial neural systems can reproduce the rich information processing and self-organizational capabilities displayed by realistic cortical microcircuits remains a major unanswered question in the world of computational neuroscience. Moreover, despite several interesting results (Kempter et al., 1999; Del Giudice et al., 2003; Vogels et al., 2005; Cessac et al., 2009; Bialek and Rieke, 1992; Victor and Purpura, 1997; Mazor and Laurent, 2005) the precise effects of synaptic plasticity on the dynamics and computational performance of spiking neural networks together with how is input represented by the spatio-temporal patterns of activity exchanged by neurons still remain poorly understood. The present Thesis constitutes a small step towards achieving these goals, by providing among other contributions, novel spike train measures and new insights regarding some of the effects of synaptic plasticity and inhibitory cells on the activity of spiking neural networks. Our results are particularly interesting in the context of designing better controllers of artificial intelligent systems and in the analysis of the information contained in the signals exchanged by neurons.

The purpose of Chapter 2 is to make the thesis self contained. It briefly presents the most basic biological structures that lie at the core of state of the art computational models. Several types of artificial neural networks are also discussed with emphasis on their biological resemblance and computational power. At the end of the chapter, we present several plasticity mechanisms that underlie and shape the activity of realistic cortical microcircuits and their computational models.

Among artificial neural networks, spiking networks remain the most biological relevant as they represent information in the shape of spiking times not unlike the brain does. In addition, such networks have been shown to be computationally superior and more robust to noise than previous generations of networks (Maass, 1996, 1997). In Chapter 3 we present a biologically plausible, universal computational paradigm featuring spiking neurons suitable for real-time computing separately introduced as the Echo State Network (ESN) (Jaeger, 2001b) and the Liquid State Machine (LSM) (Maass et al., 2002b). Due to their intrinsic computational properties and high biological resemblance, these special types of spiking networks make suitable controllers of artificial intelligent systems. Evolutionary computational techniques such as Particle Swarm Optimization (PSO) have been applied to solve a number of optimization problems from finding the shortest path in a graph, designing optimal communication networks to developing learning algorithms like supervised or reinforcement learning for neural networks, by optimizing the synaptic connections between neu-

rons. In Chapter 3 we show how PSO can be used in such an application and propose how it could be employed to increase LSM performance.

Several studies (Florian, 2010b; Brickhard, 1993; Chiel and Beer, 1997; Steels and Brooks, 1995) suggest that intelligent behaviour can emerge only in an embodied system emerged in an environment, through a process of continuous interaction. This constant interaction between the cognitive system and its surrounding environment creates a closed loop in which the actions of the agent shape the environment which in turn affect the sensory information and the way it is perceived by the agent. Through such a continuous process of interaction and discovery the system self-organizes, develops its own conceptualization of the environment and is eventually able to learn. However, additional studies (Oka et al., 2001) claim that embodiment could not necessarily be given by materiality and that the physical interaction between the cognitive system and the environment are arguably not needed. This essentially suggests that the relation between the environment and cognitive system emerges computationally through a surrogate. A major disadvantage of such an approach is that these computational interactions and virtual measurements lack the intrinsic noise found in ones from the material world. Thus noise should be artificially added and it might introduce in the process unnecessary artefacts which would make experiments either too difficult or too simple. In this context neural simulators which are able to simulate large-scale neural networks efficiently and robotic frameworks allowing them to control robotic devices are highly desirable. Such frameworks allow the facile control of physical cognitive agents and enable one to spend less time on programming details and more on detailing experiments. In Chapter 3 we introduce a flexible distributed control framework for robotic interaction with spiking neural networks ideal for large-scale simulations. The framework enables the control of different robotic platforms by multiple types of neural networks featuring different synaptic plasticity mechanisms. It has been successfully used at the Center for Cognitive Studies (Coneural) and we think it might be relevant for the scientific community interested in robotic control.

In the context of constructing artificial neural networks for robotic control, analyzing the neural code together with determining how sensorial information is represented in patterns of neural discharge becomes crucial. The distance between two spike trains reflects their similarity. Spike train distances (or measures) were successfully used in classification of neural recordings in response to different stimuli in an attempt to predict the presented stimulus, to measure the variability of neural responses to same stimulus across multiple recordings and to quantify the degree of synchrony between neurons. To this end, in Chapter 4 we derive several novel spike train measures which enable the analysis of neural variability and the information content of a spiking sequence. The newly introduced spike train measures are inspired by the Pompeiu-Hausdorff distance between two non-empty compact sets. They compute the distance between pairs of spike trains and yield a value that is dependent on the precise timing of the differences across the two spike trains. The principle which underlies their functioning is that a single spike may become as important as the spike train itself (Rieke et al., 1997).

A fundamental issue in neuroscience is understanding how learning and memory emerge out of neural plasticity. There has been a considerable interest in the past decade in Spike-timing Dependent Plasticity (STDP), a phenomenon where synaptic changes depend on the relative timing of pre- and postsynaptic action potentials. Such plasticity rules are considered to be a basis for learning (Hebb, 1949). Temporal difference (TD) methods are incremental techniques which enable a system to predict its future behavior based on the difference between two successive predictions (Sutton and Barto, 1998). Using a complex biophysical model of a cortical neuron and a simple

setup (a single presynaptic spike followed by a single current pulse) Rao and Sejnowski (2001) have shown that TD learning reproduces a Hebbian window of plasticity similar to those observed experimentally. However, it is not clear whether the result is a consequence of their complex model and holds in the case of simpler neurons, like the ones that are commonly used in large scale computer simulations; nor whether the phenomenon holds for more complex setups that are likely to appear both in simulations and in the brain. In Chapter 5 we show that in general TD learning in spiking neurons does not lead to Hebbian STDP. Using simple neural models, we verify that such a spike-timining based TD learning mechanism enables the prediction of input sequences by spiking neurons moments before their expected arrival. Additionally, we show that the same predictive capabilities can be obtained using a plasticity rule that reproduces only the causal part of Hebbian STDP used together with a homeostatic regulatory mechanism. Moreover, we show that the synaptic modifications are achieved in an optimal way when they are proportional to the value of the postsynaptic potential.

Spatio-temporal patterns of activity have been observed in the hippocampus and cortex, and were associated to memory traces. The coding of information in the phases of spikes relative to a background oscillation has been observed in many brain regions such as auditory or visual with such patterns found to be stimulus dependent and convey more information than for example firing rate alone (Gerstner and Kistler, 2002). In a simple simple setup, consisting of a spiking neuron receiving input from a number of presynaptic neurons, STDP was shown (through computer simulations) to enable the detection of spatio-temporal patterns of activity embedded in the input spike trains (Masquelier et al., 2008, 2009). In Chapter 6 we analytically derive a set of values for the input synaptic weights which facilitates such a detection. In addition to that, we study the effects of having multiple input patterns in the case of a single output neuron (the case of multiple neurons was treated in Masquelier et al. (2009)). We show that in the presence of Intrinsic Plasticity, a homeostatic regulatory mechanism, the neuron is able to respond to more than one input pattern. Such a mechanism is extremely desirable because, in contrast to other supervised learning approaches (Guetig and Sompolinsky, 2006; Florian, 2010a), it is simple, computationally cheap and biologically plausible and in addition to that also allows a fast, online implementation.

Non-invasive brain stimulation techniques such as Transcranial Magnetic Stimulation (TMS) have been hypothesized to improve learning, facilitate stroke rehabilitation, treat depression, schizophrenia, chronic pain, or addictions such as alcoholism. In a standard TMS paradigm, single-pulse stimulation over motor cortex produces high-frequency repetitive responses of around 600Hz in descending motor pathways called I-waves. Although this paradigm is well-established experimentally, the detailed mechanisms of I-wave generation have remained unclear. In Chapter 7 we introduce a model that reproduces I-waves similar to those observed in epidural responses during in vivo recordings of conscious humans. The model consists of a detailed layer 5 (L5) pyramidal cell and a population of layer 2 and 3 (L2/3) neurons projecting to it. The model parsimoniously explains the mechanisms underlying I-wave generation together with some of their basic properties such as frequency and timing. We argue that I-waves are a product of both extrinsic and intrinsic factors. By depolarizing large populations of L2/3 cells, magnetic stimulation causes a synchronized volley of postsynaptic potentials to impinge onto the dendritic trees of L5 cells. The intrinsic membrane properties and spiking mechanism of the L5 cells are then responsible for generating trains of action potentials at the characteristic I-wave frequency. Our model is shown to reproduce the effects of pharmacological interventions with drugs affecting GABA-ergic transmission on I-waves. By incorporating short-term synaptic depression of synapses from L2/3 onto L5 cells, our model

also accounts for facilitation and inhibition effects observed in paired-pulse stimulation protocols. Overall, our model parsimoniously explains findings from a range of experiments and brings us one step closer to designing optimized protocols for specific clinical purposes. Such a model is relevant not only because it can uncover the biophysical mechanisms behind magnetic stimulation but also enables the study of the functional roles of inhibitory neurons and short-term plasticity in simple cortical circuits containing complex compartmental cells – a central theme of the present Thesis.

CHAPTER 2

BIOLOGICAL FOUNDATIONS

In this chapter we present the biological structures and biophysical mechanisms that underlie the activity of cortical microcircuits and how they are reflected in state of the art computational models.

2.1 The neuron

2.1.1 Biological neurons

We discuss the anatomy together with the most basic features of biological neurons.

2.1.2 Neural models

We present a realistic model of a biological neuron.

2.1.2.1 Formal models

We discuss the need and advantages of introducing formal neuron models.

2.1.2.2 Modelling noise

We discuss methods of introducing noise into neural models.

2.2 The synapse

We present the anatomy of a synapse.

2.2.1 Synaptic plasticity

We discuss typical synaptic plasticity mechanisms present in cortical microcircuits.

2.2.1.1 Spike-Timing Dependent Plasticity

We present a basic model of Spike-timing dependent plasticity.

2.3 From neural microcircuits to artificial neural networks

We present types of cortical microcircuits and how they are reflected in computational models.

2.3.1 Classes of neural networks

We present several classes of artificial neural networks.

2.3.1.1 Threshold gates

We discuss the basic features of threshold gates.

2.3.1.2 Analog networks

We discuss the basic features of analog networks.

2.3.1.3 Spiking networks

We discuss the basic features of spiking networks.

2.3.2 Biological resemblance of artificial models

We discuss the biological relevance of artificial neural network models.

2.4 Models of spiking neurons

We present several neural models which provide a reasonable balance between biological relevance and computational efficiency.

2.4.1 The spike response model

We introduce the Spike Response Model (Gerstner and Kistler, 2002).

2.4.2 The integrate-and-fire neuron

We introduce the integrate-and-fire neuron (Gerstner and Kistler, 2002).

2.4.2.1 Pulsed input currents

2.4.2.2 Exponential input currents

2.4.3 The Izhikevich neuron

We introduce the Izhikevich (2003) neuron.

2.4.4 The stochastic Poisson neuron

We introduce a probabilistic neuron model based on a Poisson process (Gerstner and Kistler, 2002; Kempter et al., 1999).

2.4.5 The BMS neuron

We introduce a simplified discrete model (Soula et al., 2006; Cessac, 2008).

2.4.6 Measurements and typical parameters

2.5 Conclusion

We have briefly reviewed here the basic computational units – neurons and synapses – of neural microcircuits and how their characteristics are reflected into formal models of artificial neurons and networks.

CHAPTER 3

RESERVOIR COMPUTING

In this chapter we present a biologically plausible, universal computational paradigm suitable for real-time computing separately introduced as the Echo State Network (ESN) (Jaeger, 2001b) and the Liquid State Machine (LSM) (Maass et al., 2002b). In this context we introduce a novel robotic platform for distributed control of robotic agents with spiking neural networks.

The work presented in this chapter has been published and submitted as (Rusu, 2011; Rusu and Ahn, 2011).

3.1 Introduction

The LSM and ESN universal online computational frameworks were introduced with the goal to shift the focus away from training the entire set of connection within a pool of randomly connected neurons to training instead the weights of a single set of neurons.

3.2 Model for robotic control

We present the basic building blocks of LSM.

3.2.1 Robby: a robotic platform

We present "Robby", a robotic platform for distributed control, which enables an efficient and easy manipulation of different devices by spiking neural networks.

3.2.1.1 Purpose

"Robby" aims to provide a distributed control interface that allows multiple controllers to access and manipulate robotic devices located anywhere on a network. In the framework controllers are primarily neural networks, but in principle they can be any custom user-defined controller as long as it adheres to a specified control interface. Additional support for joystick controllers is provided to allow direct manipulation of devices.

3.2.1.2 Architecture and implementation details

The architecture of "Robby" is modular. It consists of a control structure (the server), a behavioural component (the client) and a common component. Figure 3.1 depicts the architecture in the form of a component diagram (Cheesman and Daniels, 2000). The server is in strict relation with the devices through an instantiation of corresponding drivers. It forwards commands received from clients, and awaits and reads replies from devices. Besides providing the communication functionality it also provides an interface to plot the raw sensory data received from the devices together with basic communication cycle parameters. The client reads data from the controller and maps it into robot commands which are later sent to the server where they are processed and forwarded. After it sends them to the server it awaits a reply before notifying the controller that its communication cycle is over. A controller implements the *ControllerInterface* and runs in a separate thread. Currently the available controller types are: a *NeuralController* which is a spiking neural network and a *JoystickController* which is useful for direct manipulation of devices. The common component contains the neural simulator which facilitates the creation and simulation of spiking neural networks (Gerstner and Kistler, 2002), device drivers and various image processing algorithms like Laplacian of Gaussian and log-polar filters (Haralick and Shapiro, 1992; Wolberg and Zokai, 2000) used to process device video data.

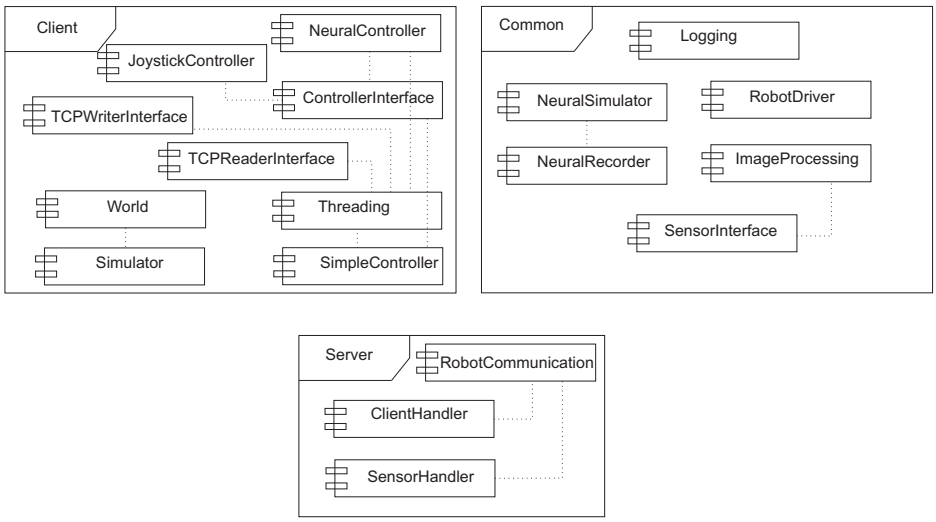

Figure 3.1: "Robby" architecture.

This server/client strategy acts as proxy between the client controller and server driver entities to increase flexibility in control and allow the controller and robot to be located in different locations with server and client communication mediated through Ethernet. Figure 3.2 depicts the communication processes inside "Robby".

3.2.1.3 Robby as a framework for robot learning

We present a simple experiment to demonstrate some of the features of "Robby" and their application.

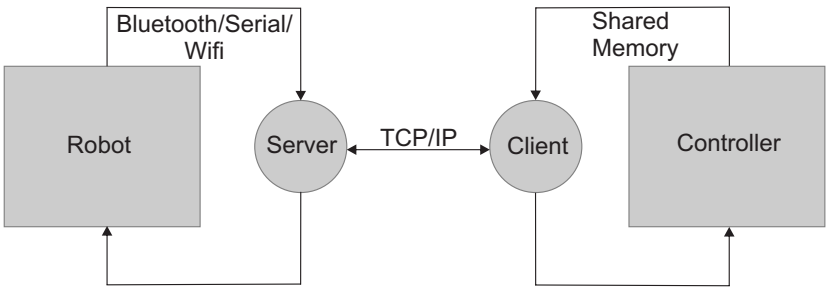

Figure 3.2: "Robby" communication. Communication between the server and client is implemented using the TCP/IP protocol.

3.3 Model for particle swarm optimization

The Particle Swarm Optimization (PSO) computational paradigm introduced by Kennedy and Eberhart (1995) has at its core a strong social interaction component.

3.3.1 Particle swarm optimization

We briefly present PSO.

3.3.2 Optimal network localization

We briefly present the problem of optimal network localization and how PSO could be used to solve it.

3.3.2.1 Particle encoding

3.3.2.2 The complete algorithm

3.3.2.3 Equivalent problem and non-optimal solutions

We briefly consider the alternative problem of constructing a localizable network from any given non-localizable one.

3.3.2.4 Simulation results

We investigate the convergence properties of the proposed PSO algorithm when applied to the problem of constructing minimal localizable networks.

3.3.3 PSO in reservoir computing

As previously outlined structured reservoirs have been shown to have better performances when compared to unstructured ones. Using the technique outlined here PSO could be used to optimize the weights and delays of the liquid with the aim of maximizing some of its computational properties (such as its separation property (Huang et al., 2009)). PSO could additionally be employed to

optimize the topology of the liquid in order to maximize the mutual information passed to the readout layer. More precisely, the topology of the liquid could be modeled as a directed graph and PSO could optimize its structure by adding and removing connections while maximizing the estimated mutual information at each moment in time.

3.4 Conclusion

In this chapter we have briefly presented the LSM (ESN) as a paradigm for online computation on continuous input streams. Because such systems allow computation in real time and are extremely robust to noise (Goodman and Ventura, 2005b) they are suitable for the control of embodied agents. In the literature they have been successfully used to control a robotic arm (Joshi and Maass, 2005), to model an existing robot controller (Burgsteiner, 2005), to perform object tracking and motion prediction (Maass et al., 2002a; Burgsteiner et al., 2007), or path planning (Zhang and Wang, 2009). Reservoirs were also been successfully used in signal processing task such as speech recognition (Verstraeten et al., 2005; Schrauwen et al., 2007). Additional applications include dynamic pattern classification (Jaeger, 2001a) and detection (Goodman and Ventura, 2005a, 2006), autonomous sine generation (Jaeger, 2001b), the computation of highly nonlinear functions on the instantaneous rates of spike trains (Maass et al., 2004), or chaotic time series generation and prediction (Jaeger, 2003; Jaeger and Haas, 2004; Steil, 2005, 2006). Particle Swarm Optimization was also presented and hypothesized to be a viable technique for the optimization of reservoirs. We have also introduced a flexible distributed control framework for robotic interaction with spiking neural networks ideal for large-scale simulations. Such control frameworks and computational paradigms could provide a basis to easily explore the theoretical principles discussed in following chapters in the context of real computational tasks involving physical autonomous agents.

CHAPTER 4

NOVEL SPIKE TRAIN MEASURES

In this chapter we introduce a new class of spike train metrics inspired by the Pompeiu-Hausdorff distance between two non-empty compact sets. They compute the distance between pairs of spike trains and yield a result that is dependent on the precise timing of the differences across the two spike trains.

The work presented in this chapter has been published as (Rusu and Florian, 2010).

4.1 Introduction

Here we introduce a new class of spike train metrics inspired by the Pompeiu-Hausdorff distance between two non-empty compact sets. The new spike train metrics yield a result dependent of the exact timing of differences among two spike trains. In the context of the information exchanged by two neurons, each spike may be as important as the spike train itself (Rieke et al., 1997). Therefore, such metrics, based on the specific timing of differences within spike trains, become desirable.

4.2 A new class of spike metrics

We consider bounded, nonempty spike trains of the form

$$T = \{t^{(1)}, \ldots, t^{(n)}\}, \tag{4.1}$$

where $t^{(i)} \in \mathbb{R}$ are the ordered spike times and $n \in \mathbb{N}^*$ is the number of spikes in the spike train. We denote by a and b the bounds of the considered spike trains, i.e. $a \leq t^{(i)} \leq b$, $\forall t^{(i)}$, with $a, b \in \mathbb{R}$, finite, and $a < b$. We denote by $\mathscr{S}_{[a,b]}$ the set of all possible such spike trains. We study metrics that compute the distances between two spike trains T and $\bar{T}$ from $\mathscr{S}_{[a,b]}$.

The new metrics that we introduce are inspired by the Pompeiu-Hausdorff distance Pompeiu (1905); Hausdorff (1914). When applied to a pair of spike trains, the Pompeiu-Hausdorff distance h returns the largest difference, in absolute value, between the timings of a spike in one train and of

the closest spike in the other spike train:

$$h(T,\bar{T}) = \max\left\{\sup_{t\in T}\inf_{\bar{t}\in\bar{T}}|t-\bar{t}|, \sup_{\bar{t}\in\bar{T}}\inf_{t\in T}|t-\bar{t}|\right\}, \tag{4.2}$$

or, equivalently, the minimal number ϵ such that the closed ϵ-neighborhood of T includes $\bar{T}$ and the closed ϵ-neighborhood of $\bar{T}$ includes T:

$$h(T,\bar{T}) = \inf\left\{\epsilon \geq 0 \text{ such that } |t-\bar{t}| \leq \epsilon,\ \forall\, t\in T,\ \forall\, \bar{t}\in\bar{T}\right\}. \tag{4.3}$$

Another equivalent form of the Pompeiu-Hausdorff distance is (Ponulak, 2005, pp. 105–110; Rockafellar and Wets, 2009, pp. 117–118; Deza and Deza, 2009, pp. 47–48)

$$h(T,\bar{T}) = \sup_{x\in\mathbb{R}}\left|\inf_{t\in T}|t-x| - \inf_{\bar{t}\in\bar{T}}|\bar{t}-x|\right|. \tag{4.4}$$

We introduce a distance d between an arbitrary timing $x\in\mathbb{R}$ and a spike train T:

$$d(x,T) = \inf_{t\in T}|t-x|. \tag{4.5}$$

Eq. 4.2 can then be rewritten as

$$h(T,\bar{T}) = \max\left\{\sup_{t\in T} d(t,\bar{T}), \sup_{\bar{t}\in\bar{T}} d(\bar{t},T)\right\} \tag{4.6}$$

and Eq. 4.4 as

$$h(T,\bar{T}) = \sup_{x\in\mathbb{R}}\left|d(x,T) - d(x,\bar{T})\right|. \tag{4.7}$$

We also have (Section 4.6):

$$h(T,\bar{T}) = \sup_{x\in[a,b]}\left|d(x,T) - d(x,\bar{T})\right|. \tag{4.8}$$

The Pompeiu-Hausdorff metric has a quite poor discriminating power, as for many variations of the spike trains the distances will be equal and any spike train space endowed with this metric would be highly clusterized. Our new metrics generalize the form of the Pompeiu-Hausdorff distance given in Eq. 4.8, by introducing features that are more sensitive to spike timings.

We consider $\mathbb{B}$ to be the space of arbitrary continuous, strictly positive functions $\mathscr{H}\colon \mathbb{R}\to\mathbb{R}^+$. On compact sets such functions are bounded (Protter, 1998, p. 56). We denote by m the upper bound of $\mathscr{H}$ on the interval $[0, b-a]$, i.e.

$$0 < \mathscr{H}(x) < m < \infty,\ \forall\ x\in[0,b-a]. \tag{4.9}$$

By $\mathbb{B}^+$ we denote $\mathbb{B}$ class functions with domains restricted to $\mathbb{R}^+$.

4.2.1 The max-metric

Consider an arbitrary function $\mathscr{H}\in\mathbb{B}^+$. Typically, $\mathscr{H}(x)$ has a maximum for $x=0$ and is a decreasing function of x, for example an exponential,

$$\mathscr{H}_E(x) = \frac{1}{\tau}\exp\left(-\frac{x}{\tau}\right), \tag{4.10}$$

or a Gaussian,

$$\mathscr{H}_G(x) = \frac{1}{\tau\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\tau^2}\right), \tag{4.11}$$

with τ a positive parameter.

We introduce the max-metric as

$$d_m(T,\bar{T}) = \int_a^b \sup_{x\in[a,b]} \left\{|d(x,T) - d(x,\bar{T})|\,\mathscr{H}(|s-x|)\right\} \mathrm{d}s. \tag{4.12}$$

The max-metric integrates, through the variation of s along the interval $[a,b]$ that contains the two spike trains, the maximum difference, in absolute value, between the distances from a point x in that interval to the two spike trains, weighted by the kernel $\mathscr{H}(|s-x|)$ which focuses locally around s. Figure 4.1 shows how the distance d_m between two spike trains is computed.

The max-metric is a generalization of the Pompeiu-Hausdorff distance, since in the particular case that $\mathscr{H}(\cdot) = 1/(b-a)$ we have $d_m(T,\bar{T}) = h(T,\bar{T})$. In Section 4.7 we show that d_m is finite and that it satisfies the properties of a metric. We also show that regardless of the kernel $\mathscr{H}$ all the max-metrics are topologically equivalent to each other (O'Searcoid, 2007, p. 229) because they are equivalent to the Pompeiu-Hausdorff distance. Each metric will generate the same topology and thus any topological property is invariant under an homeomorphism. This means that the metrics generate the same convergent sequences in the space of spike trains $\mathscr{S}_{[a,b]}$. The implication of this for learning is that learning rules derived from these metrics will converge in the same way, independently of the choice of $\mathscr{H}$.

4.2.2 The modulus-metric

We define the modulus-metric as

$$d_o(T,\bar{T}) = \int_a^b |d(s,T) - d(s,\bar{T})|\,\mathrm{d}s. \tag{4.13}$$

The modulus-metric is a particular case of the max-metric in the limit that $\mathscr{H}$ is

$$\mathscr{H}(x) = \begin{cases} 1, & \text{if} \quad x = 0, \\ 0, & \text{otherwise.} \end{cases} \tag{4.14}$$

The modulus-metric uses the d distance like the max-metric does, but it does not depend on any kernels or parameters and it also allows a fast computer implementation in linear complexity. Algorithm 1 presents a simple implementation of the d_o metric in pseudo-code. The algorithm first builds P as an ordered set that contains all spikes in the two spike trains T and $\bar{T}$, the bounds a and b, as well as the time moments that lie at the middle of the interval between two spikes from the same spike train, for both spike trains. As exemplified by Fig. 4.3 D, the graph of the function $f(s) = |d(s,T) - d(s,\bar{T})|$ is made out of line segments that join in points from P. In order to compute the integral of this function $d_o = \int_a^b f(s)\,\mathrm{d}s$, it is sufficient to compute the function at the joining points. Since between these joining points the function is linear, the integral can be then computed exactly. The algorithm's duration depends linearly on the number of spikes in the two spike trains, $n + \bar{n}$.

It can be shown that the distance d_o is finite and that it satisfies the properties of a metric by particularizing the proofs in Section 4.10 with $\mathscr{L}(x) = 1$, $\forall x \in \mathbb{R}$.

Input: The pair of spike trains T_1, T_2 and the bounds a and b.
Output: The distance d_o between the spike trains.
$n_1 = \texttt{length}(T_1)$; $n_2 = \texttt{length}(T_2)$;
T_1, T_2 and P are ordered sets of real numbers, indexed starting from 0.
$P := T_1 \bigcup T_2 \bigcup \{a, b\}$;
for $i := 1 \ldots n_1 - 1$ **do**
  $P := P \bigcup \{(T_1[i] - T_1[i-1])/2\}$;
for $i := 1 \ldots n_2 - 1$ **do**
  $P := P \bigcup \{(T_2[i] - T_2[i-1])/2\}$;
If P is not automatically sorted, it should be explicitly sorted:
$P := \text{sort}(P)$;
$s_p := a$; $f_p := |T_1[0] - T_2[0]|$;
$i_1 := 1$; $i_2 := 1$;
for $i := 1 \ldots \texttt{length}(P) - 1$ **do**
  $s := P[i]$;
  if $s \geq T_1[i_1]$ `and` $i_1 < n_1 - 1$ **then**
    $i_1 := i_1 + 1$;
  if $s \geq T_2[i_2]$ `and` $i_2 < n_2 - 1$ **then**
    $i_2 := i_2 + 1$;
  $d_1 := 0$; $d_2 := 0$;
  if $i_1 > 1$ **then**
    $d_1 := s - T_1[i_1 - 1]$;
  $d_1' = |T_1[i_1] - s|$;
  if $d_1' < d_1$ **then**
    $d_1 := d_1'$;
  if $i_2 > 1$ **then**
    $d_2 := s - T_2[i_2 - 1]$;
  $d_2' = |T_2[i_2] - s|$;
  if $d_2' < d_2$ **then**
    $d_2 := d_2'$;
  We can now compute the value of f at s:
  $f := |d_1 - d_2|$;
  The integration is performed here:
  $d_o := d_o + (s - s_p)(f + f_p)/2$;
  $s_p := s$; $f_p := f$;
return d_o;

Algorithm 1: The algorithm for computing the distance d_o between two spike trains T_1 and T_2. The text in italic represents comments.

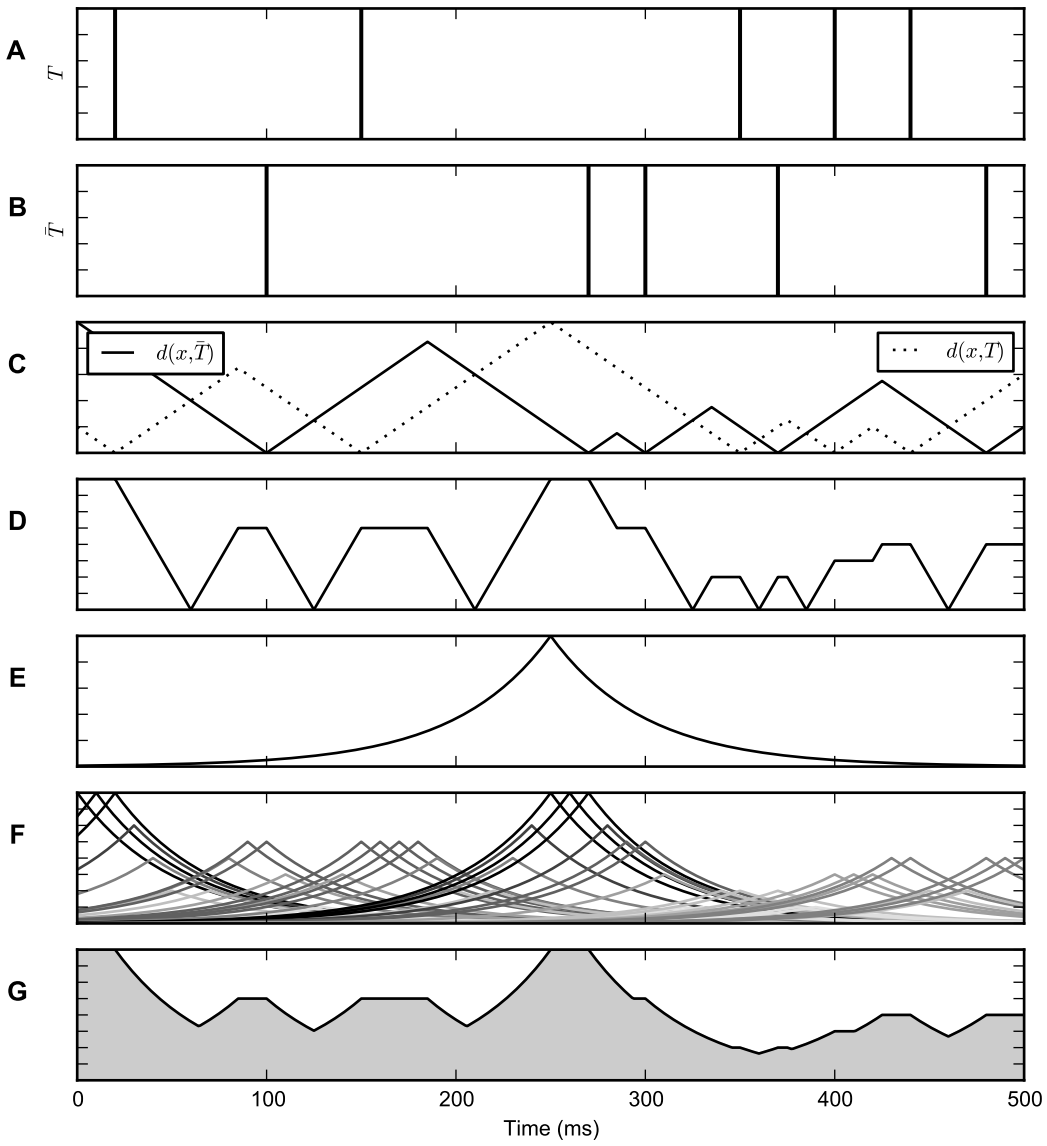

Figure 4.1: The max-metric modus operandi. (A) Spike train $T = \{20, 150, 350, 400, 440\}$ ms. Each spike time is represented as a vertical bar. (B) Spike train $\bar{T} = \{100, 270, 300, 370, 480\}$ ms. (C) The distance between a point x and the timings of spikes in the spike trains, $d(x,T)$ and $d(x,\bar{T})$ as a function of x. (D) The difference $\left|d(x,T) - d(x,\bar{T})\right|$ as a function of x. (E) The kernel $\mathscr{H}(|s-x|)$ as function of s with a fixed x and a 50 ms decay constant. (F) The weighted difference $\left|d(x,T) - d(x,\bar{T})\right| \mathscr{H}(|s-x|)$ as function of s for discrete values of x. (G) The supremum of the difference weighted by the kernel, $\sup_{x\in[a,b]} \left\{\left|d(x,T) - d(x,\bar{T})\right| \mathscr{H}(|s-x|)\right\}$. The distance d_m is the area under this curve.

4.2.3 The convolution max-metric

The max-metric can also be given in a convolution form. To construct this form of the metric we consider an arbitrary continuous, positive kernel $\mathscr{K}: \mathbb{R} \to \mathbb{R}^+$, with the property that

$$0 \leq \mathscr{K}(x) \leq 1 \text{ for every } x \in \mathbb{R}, \text{ and } \mathscr{K}(0) > 0, \tag{4.15}$$

and which is strictly increasing for $x < 0$ and strictly decreasing for $x > 0$. We convolve the two spike trains T and $\bar{T}$ with the filtering kernel $\mathscr{K}$ to obtain

$$f(x) = \sum_{i=1}^{n} \mathscr{K}(t - t^{(i)}) \tag{4.16}$$

$$\bar{f}(x) = \sum_{i=1}^{\bar{n}} \mathscr{K}(t - \bar{t}^{(i)}). \tag{4.17}$$

We denote by $\mathscr{F}_{[a,b]}$ the set of all possible filtered spike trains from $\mathscr{S}_{[a,b]}$.

For the convolution max-metric, we request that $\mathscr{H} \in \mathbb{B}^+$ is derivable on $(0, b-a)$ and that it has bounded derivative.

The convolution max-metric is defined as

$$d_c(T, \bar{T}) = \int_a^b \sup_{x \in [a,b]} \left\{ |f(x) - \bar{f}(x)| \mathscr{H}(|s-x|) \right\} \mathrm{d}s. \tag{4.18}$$

The kernels influence the performance and response of the metric so they should be chosen according to the task at hand. Figure 4.2 shows how the distance d_c between two spike trains is computed. In Section 4.8 we show that d_c is finite and that it satisfies the properties of a metric.

4.3 Localized metrics

In the case of the max-metric, with or without convolution, the use of the kernel $\mathscr{H}$ served the purpose of providing a local perspective, around each point within $[a, b]$, of the distance between the spike trains. These local perspectives were then integrated in the final distance. In this section we introduce different metrics that also depend on a kernel $\mathscr{L} \in \mathbb{B}^+$, but for which it has a different purpose. More precisely, it may be regarded as a magnifying glass to be used to focus on one specific area of the spike trains.

Such a metric is biologically relevant if, for example, we take into consideration how a neuron responds to input spikes. Recent spikes influence more the neuron than old ones. If we would like to measure the distance between two spike trains according to how the differences between them influence the activity of a neuron at a particular moment of time, recent differences should account more than differences in the distant past. For the localized metrics, $\mathscr{L}$ could thus model the shape of postsynaptic potentials (PSP) that reflects the dynamics of the effect of one presynaptic spike on the studied neuron. Thus, $\mathscr{L}$ could typically be an exponential, $\mathscr{L}_E = \mathscr{H}_E$ (Eq. 4.10), an alpha function,

$$\mathscr{L}_\alpha(x) = \frac{x}{\tau^2} \exp\left(-\frac{x}{\tau}\right), \tag{4.19}$$

a double exponential,

$$\mathscr{L}_D(x) = \frac{\tau}{\tau - \tau_s} \left[\exp\left(-\frac{x}{\tau}\right) - \exp\left(-\frac{x}{\tau_s}\right) \right], \tag{4.20}$$

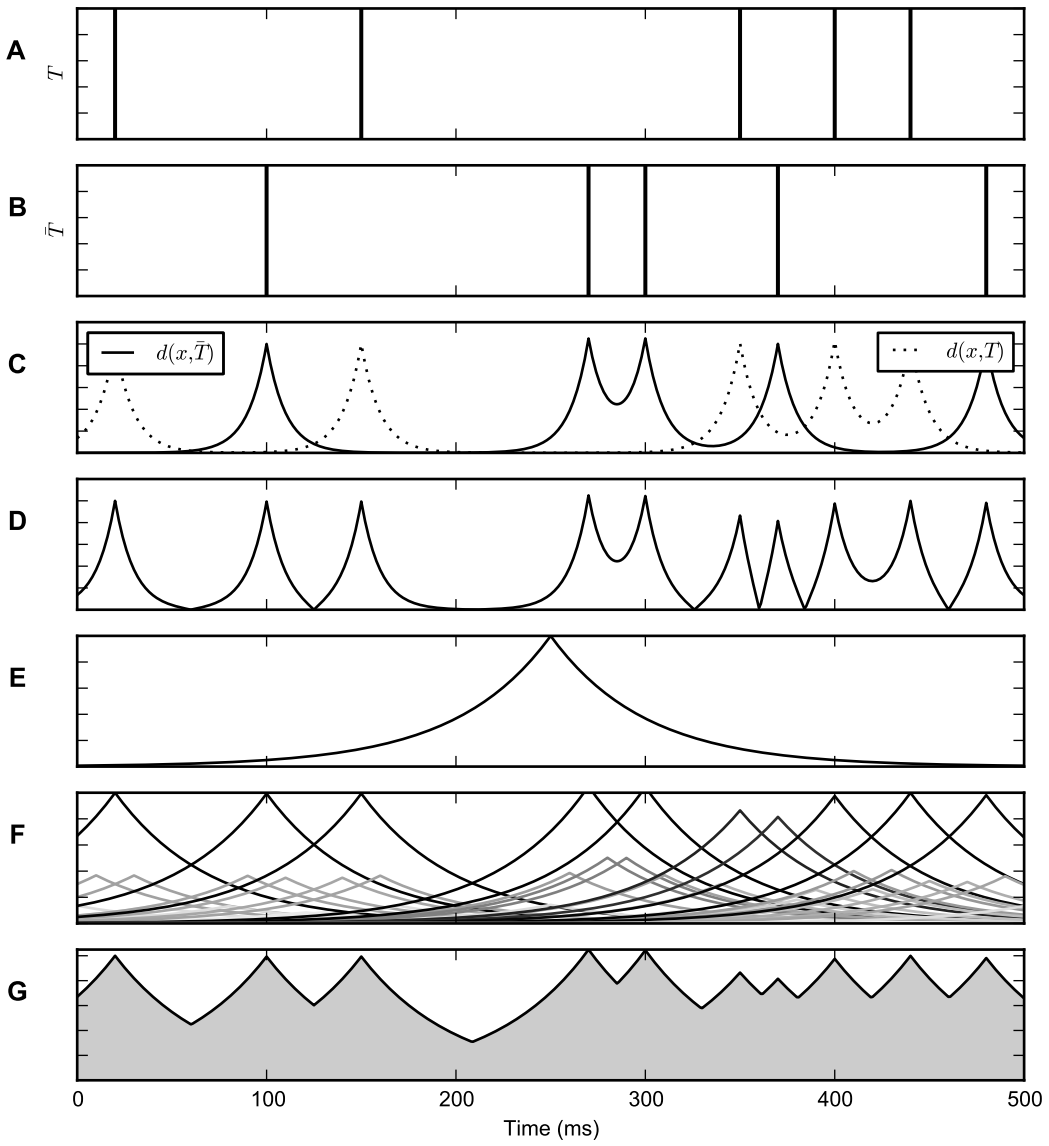

Figure 4.2: The convolution max-metric modus operandi. (A) Spike train $T = \{20, 150, 350, 400, 440\}$ ms. Each spike time is represented as a vertical bar. (B) Spike train $\bar{T} = \{100, 270, 300, 370, 480\}$ ms. (C) The spike trains T and $\bar{T}$ filtered with an exponential kernel with a 10 ms decay constant. (D) The difference $\left|f(x) - \bar{f}(x)\right|$ as a function of x. (E) The kernel $\mathscr{H}(|s-x|)$ as function of s with a fixed x and a 50 ms decay constant. (F) The weighted difference $\left|f(x) - \bar{f}(x)\right| \mathscr{H}(|s-x|)$ as function of s for discrete values of x. (G) The supremum of the difference weighted by the kernel, $\sup_{x\in[a,b]} \left\{\left|f(x) - \bar{f}(x)\right| \mathscr{H}(|s-x|)\right\}$. The distance d_c is the area under this curve.

or, if we model the PSP generated by a double exponential synaptic current for an integrate-and-fire

neuron,

$$\mathscr{L}_I(x) = \frac{\tau}{\tau_s - \tau_r} \left\{ \frac{\tau_s}{\tau - \tau_s} \left[\exp\left(-\frac{x}{\tau}\right) - \exp\left(-\frac{x}{\tau_s}\right) \right] - \frac{\tau_r}{\tau - \tau_r} \left[\exp\left(-\frac{x}{\tau}\right) - \exp\left(-\frac{x}{\tau_r}\right) \right] \right\}, \tag{4.21}$$

where τ, τ_s, and τ_r are positive parameters.

4.3.1 Localized max-metric

We introduce the localized max-metric as

$$d_l(T, \bar{T}) = \int_a^b \mathscr{L}(b-s) \sup_{x \in [s,b]} |d(x, T) - d(x, \bar{T})| \, \mathrm{d}s. \tag{4.22}$$

Fig. 4.3 shows how the distance d_l between two spike trains is computed. The differences between the spike trains that account the most for the distance are those that are close to b. The shape of $\mathscr{L}$ has a high impact on the behavior of the metric. In Section 4.9 we show that the distance d_l is finite and that it satisfies the properties of a metric.

4.3.2 Localized modulus-metric

The modulus-metric can also be given in a localized form. This form does not require the restriction of the kernel to $\mathbb{R}^+$. Thus, we define it to depend on a kernel $\mathscr{L} \in \mathbb{B}$,

$$d_n(T, \bar{T}) = \int_a^b |d(s, T) - d(s, \bar{T})| \, \mathscr{L}(b-s) \, \mathrm{d}s. \tag{4.23}$$

In Section 4.10 we show that d_n is finite and that it satisfies the properties of a metric.

4.3.3 Localizing existing metrics

A property of the distances introduced here in Eq. 4.22 and 4.23 is their intrinsic sensitivity to the timings of differences across the spike trains. A similar localization by a kernel can be applied to existing metrics. Let $f, \bar{f} \in \mathscr{F}_{[a,b]}$ be two filtered spike trains obtained by convolution of the spike trains $T, \bar{T} \in \mathscr{S}_{[a,b]}$. Consider the van Rossum (2001) distance defined as

$$d_R(T, \bar{T}) = \int_{-\infty}^{\infty} (f(s) - \bar{f}(s))^2 \, \mathrm{d}s. \tag{4.24}$$

When localized with $\mathscr{L}$ the distance becomes

$$d_R l(T, \bar{T}) = \int_{-\infty}^{\infty} (f(s) - \bar{f}(s))^2 \, \mathscr{L}(b-s) \, \mathrm{d}s. \tag{4.25}$$

Here $\mathscr{L}$ may be chosen to have the same qualitative properties as the kernel used in Eqs. 4.19–4.21.

4.4 Application

We analyzed the behavior of the introduced metrics through computer simulations using simple setups.

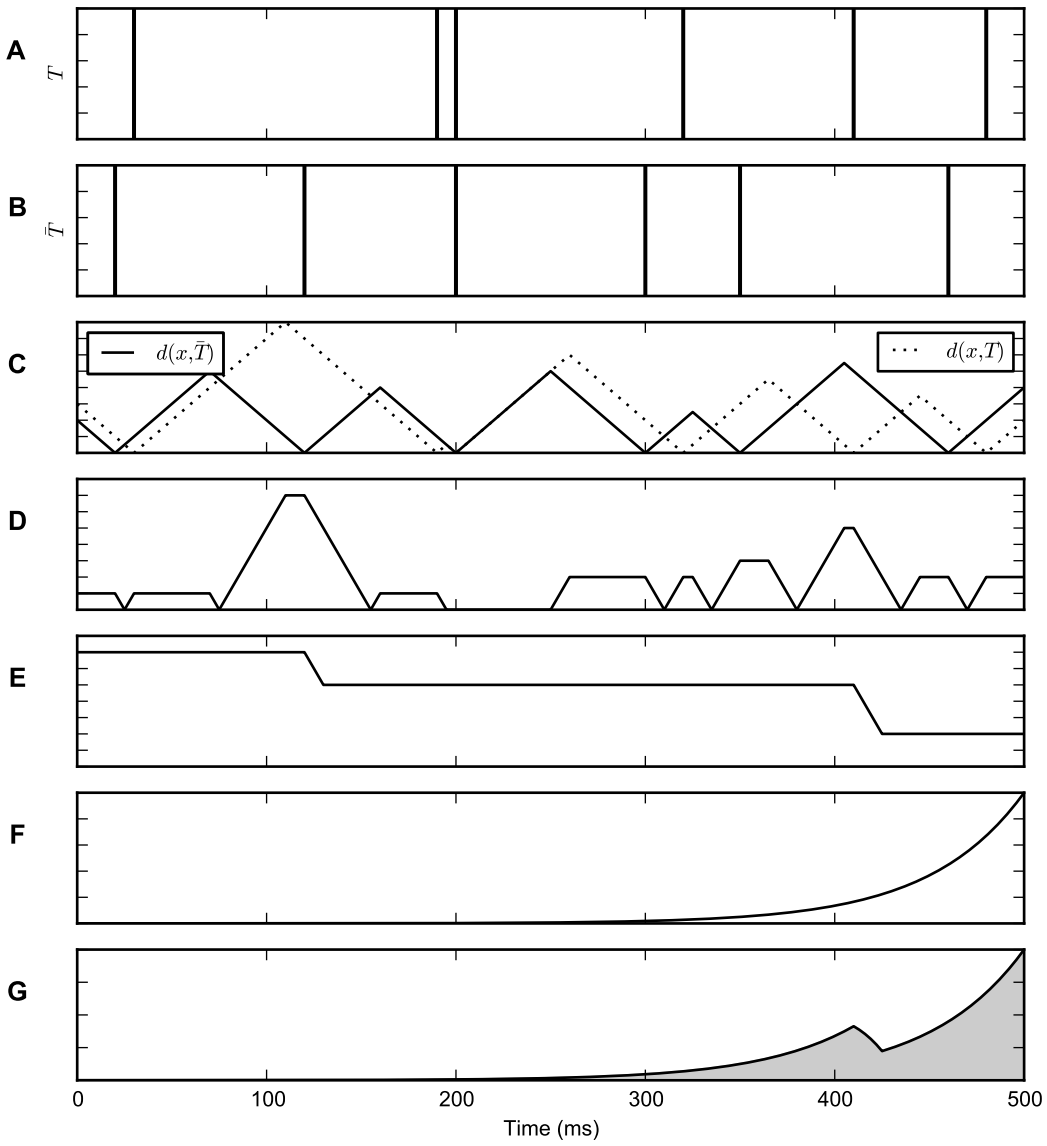

Figure 4.3: The localized max-metric modus operandi. (A) Spike train $T = \{50, 180, 200, 300, 400, 480\}$ ms. Each spike time is represented as a vertical bar. (B) Spike train $\bar{T} = \{20, 120, 200, 300, 350, 470\}$ ms. (C) The distance between a point x and the timings of spikes in the spike trains, $d(x,T)$ and $d(x,\bar{T})$ as function of x. (D) The difference $\left|d(x,T) - d(x,\bar{T})\right|$ as a function of x. (E) The kernel $\mathscr{H}(b-s)$ as function of s with a 50 ms decay constant. (F) The supremum, $\sup_{x\in[s,b]} \left|d(x,T) - d(x,\bar{T})\right|$. (G) The weighted difference $\left|d(x,T) - d(x,\bar{T})\right| \mathscr{H}(b-s)$ as function of s for discrete values of x. (H) The supremum weighted by the kernel, $\mathscr{H}(b-s)\sup_{x\in[s,b]} |d(x,T) - d(x,\bar{T})|$, as a function of s. The distance d_l is the area under this curve.

4.5 Conclusion

We have introduced a new class of spike train metrics inspired by the Pompeiu-Hausdorff distance. Their underlying principle of functioning is that the precise timing of each individual spike in a spike train is important. Each metric is given dependent on a kernel $\mathscr{H}$ which can be particularized to cause distinct behaviors. On one side the kernel can be used to provide a local perspective around individual spikes while on the other it can be used as a magnifying glass to focus on specific parts of spikes trains. From a mathematical point of view the kernel $\mathscr{H}$ and can be just about any function because the generated metrics are commensurable. Some, however, will have a lesser physiological interpretation than others. Because the metrics generate the same topologies regardless of the choice of kernels, topological properties are the same across all spike train spaces. Since they are not dependent on a filtering kernel and are able to use the raw spike times directly they do not introduce additional time constants and therefore have the advantage of being less restrictive. An additional, convolution based distance, was also presented and shown to have a behaviour similar to the van Rossum distance. A simple metric which uses the spike trains directly and is not dependent on $\mathscr{H}$ was also presented. It was shown to have similar properties to the other introduced metrics while allowing a much faster computer implementation. An optimal algorithm to compute the metric, operating in linear time was also introduced.

All distances with the exception of the convolution based max-metric were shown to be negatively correlated to the rate of the spike trains. In contrast, the convolution based max-metric together with the Victor & Purpura and van Rossum distances were shown to be correlated with spike rate. Regardless of their type all introduced metrics exhibited a behaviour dependent on the precise timing of the differences across the two spike trains. In contrast to other metrics such as Victor & Purpura and van Rossum they were shown to depend when single spikes were either inserted or moved depending on their precise timing. This essentially suggests that the location of a spike within a spike train may become just as important as the inter-spike-interval and that a single spike might be as important as the spike train itself.

4.6 Equivalent Pompeiu-Hausdorff form

Proposition 4.6.1. *The Pompeiu-Hausdorff metric*

$$h(T,\bar{T}) = \sup_{x\in\mathbb{R}} \left| \inf_{t\in T} |t-x| - \inf_{\bar{t}\in\bar{T}} |\bar{t}-x| \right| \tag{4.26}$$

can be equivalently expressed as

$$h(T,\bar{T}) = \sup_{x\in[a,b]} \left| d(x,T) - d(x,\bar{T}) \right|. \tag{4.27}$$

4.7 Analysis of the max-metric

Proposition 4.7.1. $d_m(T,\bar{T}) < \infty$.

Proposition 4.7.2. $d_m : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ *is a metric.*

Proposition 4.7.3. *The metric* $d_m : \mathscr{S}_{[a,b]} \times \mathscr{S}_{[a,b]} \to \mathbb{R}$ *is topologically equivalent to the Pompeiu-Hausdorff distance.*

4.8 Analysis of the convolution max-metric

Lemma 4.8.1. *Let $g\colon [a,b] \to \mathbb{R}$ be a continuous function and $h\colon [0,b-a] \to \mathbb{R}$ be a continuous function which is derivable on $(0,b-a)$ and has bounded derivative. Then the function $q\colon [a,b] \to \mathbb{R}$,*

$$q(s) = \sup_{x\in[a,b]} \left[g(x)\, h(|s-x|) \right] \tag{4.28}$$

is continuous on $[a,b]$.

Proposition 4.8.1. $d_c(T,\bar{T}) < \infty$.

Proposition 4.8.2. $d_c : \mathcal{S}_{[a,b]} \times \mathcal{S}_{[a,b]} \to \mathbb{R}$ *is a metric.*

4.9 Analysis of the localized max-metric

Proposition 4.9.1. $d_l(T,\bar{T}) < \infty$.

Proposition 4.9.2. $d_l : \mathcal{S} \times \mathcal{S} \to \mathbb{R}$ *is a metric.*

4.10 Analysis of the localized modulus-metric

Proposition 4.10.1. $d_n(T,\bar{T}) < \infty$.

Proposition 4.10.2. $d_n : \mathcal{S}_{[a,b]} \times \mathcal{S}_{[a,b]} \to \mathbb{R}$ *is a metric.*

CHAPTER 5

EXPLORING THE LINK BETWEEN STDP AND TD LEARNING

In this chapter we show that in general Temporal Difference (TD) learning in spiking neurons does not lead to Hebbian Spike-timing Dependent Plasticity (STDP). Using simple neural models, we verify that such a spike-timining based TD learning mechanism enables the prediction of input sequences by spiking neurons moments before their expected arrival. Additionally, we show that the same predictive capabilities can be obtained using a plasticity rule that reproduces only the causal part of hebbian STDP used in conjunction with a homeostatic regulatory mechanism.

The work presented in this chapter has been published and communicated as (Rusu, 2008, 2009; Florian and Rusu, 2009; Rusu and Florian, 2009).

5.1 Introduction

In an early study that tried to derive the experimentally-observed Spike-timing Dependent Plasticity (STDP) from first principles, Rao and Sejnowski (2000, 2001, 2003) have shown that temporal difference (TD) learning of the value of the membrane potential of a neuron, at a fixed delay after the neuron received a presynaptic spike, leads to a plasticity rule that is very similar to STDP. The results have been obtained using a relatively complex, but biologically plausible, biophysical model and a simple setup (a single presynaptic spike followed by a single current pulse). It is thus not clear whether TD learning also leads to STDP in networks of simpler neurons, like the ones that are commonly used in large scale computer simulations; nor whether the phenomenon holds for more complex setups that are likely to appear both in simulations and in the brain. It is important to know this in order to establish whether STDP leads to TD learning in these networks and setups.

We thus study here the same phenomenon using integrate-and-fire (IAF) and Izhikevich (2003) neurons, which are commonly used in large scale computer simulations of spiking neural networks. We used setups that are more complex than the ones in the original study (Rao and Sejnowski, 2001), with postsynaptic spikes being generated by the irregular firing of synaptic afferents or by a constant input current. The plasticity curves obtained from TD learning were determined not only

through computer simulations, but also analytically, in the case of the IAF neuron. We also derive and analyze a more general and biologically plausible form of such a plasticity rule. Using a simple example we show how a neuron can learn to predict input sequences through recurrent excitation by using both a TD based plasticity rule and a plasticity rule that reproduces only the causal part of Hebbian STDP used in conjunction with a homeostatic plasticity mechanism.

5.2 TD learning

We briefly present TD learning.

5.3 TD learning for spiking neurons

Rao and Sejnowski (2000, 2001, 2003) implemented the prediction of the membrane potential through TD learning by incrementing or decrementing the value of the synaptic strength by an amount proportional to the TD in the postsynaptic membrane potential at time instants $t+\Delta t$ and t for presynaptic activation at time t. This can be written as

$$\dot{w}_j = \alpha \cdot \Delta u \cdot \Phi_j(t), \tag{5.1}$$

where

$$\Phi_j(t) = \sum_f \delta(t - t_j^{(f)}) \tag{5.2}$$

is the spike train of presynaptic neuron j represented as a sum of Dirac functions.

In the following we check whether the TD based learning rule (Eq. 5.1) reproduces an STDP-like asymmetric window of plasticity (WOP) in the case of IAF and Izhikevich neurons.

5.3.1 TD learning in integrate-and-fire neurons

This section is divided into two parts. On one side we analytically derive the shape of the WOP while on the other, using computer simulations we numerically check the assumption that STDP can be expressed as a form of TD learning.

5.3.1.1 The model

Neurons are modeled as IAF (for a detailed introduction, see Section 2.4.2 and Gerstner and Kistler (2002)).

5.3.1.2 Analytical results

In the analytical computations throughout this section we used setups with postsynaptic spikes being generated by a single input spike, a transient pulse of current, or by a constant input current and using the IAF neuron in two distinct cases. In the first, we treated the IAF model as it is described in Gerstner and Kistler (2002), while in the second, we added an action potential of non-zero duration to the membrane potential in order to model a shape for the postsynaptic action potential. The added action potential was modeled as a function ζ

$$\zeta\colon [t_\theta, t_\theta + K] \longrightarrow \mathbb{R} \qquad \zeta(t) = a e^{b\frac{t-t_\theta}{K}}. \tag{5.3}$$

Single input spike A neuron was stimulated by a single spike from some presynaptic neuron. Postsynaptic action potentials were paired to this single input spike by some unspecified input. The variation in membrane potential caused by the pairing is summarised in the following

$$\Delta u = \begin{cases} \eta_0 \left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right) & \\ \quad + w\epsilon_0 \left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < 0; \\ w\epsilon_0 \left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } 0 < \Delta t < \tau; \\ \eta_0 \exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) & \\ \quad + w\epsilon_0 \left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } 0 < \tau < \Delta t. \end{cases} \tag{5.4}$$

When an action potential is added to the IAF model the variation in membrane potential changes correspondingly

$$\Delta u = \begin{cases} \eta_0 \left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right) & \\ \quad + w\epsilon_0 \left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < 0; \\ w\epsilon_0 \left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } K < K + \Delta t < \tau; \\ \zeta(K - (\tau - \Delta t)) & \text{if } K < \tau < \Delta t + K < \tau + K; \\ \eta_0 \exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) & \\ \quad + w\epsilon_0 \left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } K < \tau < \Delta t; \\ \eta_0 \exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \zeta(K - \tau) & \\ \quad + w\epsilon_0 \left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < K < \Delta t. \end{cases} \tag{5.5}$$

Figure 5.1 depicts the WOP obtained from an IAF stimulated by a single input spike with and without an added action potential.

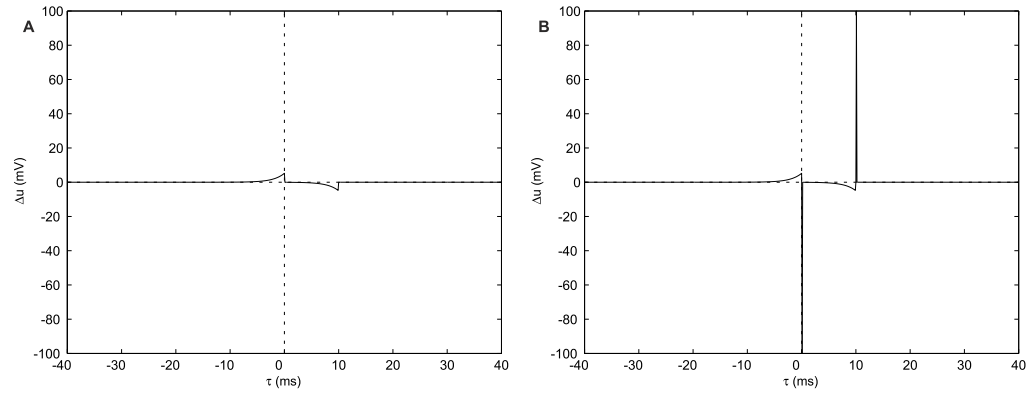

Figure 5.1: The WOP. An IAF neuron stimulated by a single input spike. The added action potential was modeled by a function $\zeta(x) = ae^{bx}$ with $K = 1$ ms. The reset value of the potential was negative and $\Delta t = 10$ ms. (A) without the added action potential. (B) with the added action potential.

Constant input current A neuron received input on a single excitatory synapse. The input consisted of a constant current $I(t) = I_0$ which caused the neuron to fire output spikes periodically. Postsynaptic action potentials were paired to presynaptic activation by test presynaptic spikes. The test action potential was chosen such that the total charge q that it delivers to the neuron was much smaller than I_0. The variation in membrane potential caused by the pairing is summarised in the following

$$\Delta u = \begin{cases} (\eta_0 - RI_0)\left[\exp\left(-\frac{-\tau+\Delta t}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t}{\tau_m}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < \Delta t < \tau + T,\ \tau < 0; \\ (\eta_0 - RI_0)\left(\exp\left[-\frac{-T-\tau+\Delta t}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < \Delta t - T,\ \tau < 0; \\ (\eta_0 - RI_0)\left[\exp\left(-\frac{T-\tau+\Delta t}{\tau_m}\right) - \exp\left(-\frac{T-\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left(\exp\left[-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } 0 < \Delta t < \tau < T; \\ (\eta_0 - RI_0)\left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \exp\left(-\frac{T-\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < \Delta t < T. \end{cases} \tag{5.6}$$

When an action potential is added to the IAF model the variation in membrane potential changes correspondingly

$$\Delta u = \begin{cases} (\eta_0 - RI_0)\left[\exp\left(-\frac{-\tau+\Delta t}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t}{\tau_m}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < \Delta t < \tau + T,\ \tau < 0; \\ (\eta_0 - RI_0)\left[\exp\left(-\frac{-T-\tau+\Delta t}{\tau_m}\right) - \exp\left(\frac{\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } \tau < \Delta t - T,\ \tau < 0; \\ (\eta_0 - RI_0)\left[\exp\left(-\frac{T-\tau+\Delta t}{\tau_m}\right) - \exp\left(-\frac{T-\tau}{\tau_m}\right)\right] & \\ \quad + \epsilon_0\left[\exp\left(-\frac{\Delta t}{\tau_m}\right) + \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } K < K + \Delta t < \tau; \\ \zeta(K - (\tau - \Delta t)) - RI_0 - (\eta_0 - RI_0)\exp\left(-\frac{T-\tau}{\tau_m}\right) & \text{if } K < \tau < \Delta t + K < \tau + K; \\ (\eta_0 - RI_0)\left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \exp\left(-\frac{T-\tau}{\tau_m}\right)\right] & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } K < \tau < \Delta t; \\ RI_0 + (\eta_0 - RI_0)\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right) - \zeta(K - \tau) & \\ \quad + w\epsilon_0\left[\exp\left(-\frac{\Delta t-\tau}{\tau_m}\right)\exp\left(-\frac{\tau}{\tau_s}\right) - \exp\left(-\frac{\Delta t}{\tau_s}\right)\right] & \text{if } 0 < \tau < K < \tau + K. \end{cases} \tag{5.7}$$

Figure 5.2 depicts the WOP obtained from an IAF stimulated by constant input current with and without an added action potential.

Pulsed input current Finally, a neuron was stimulated by a presynaptic spike which did not cause the generation of a postsynaptic action potential. The presynaptic activation was paired with the postsynaptic action potential by a rectangular pulse of current I. Figure 5.3 depicts the WOP obtained from an IAF stimulated by a rectangular pulse of current with and without an added action potential.

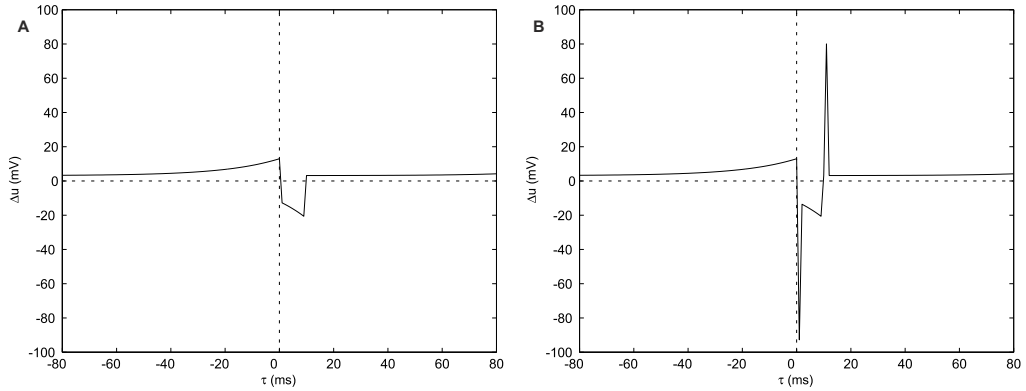

Figure 5.2: The WOP. An IAF neuron stimulated by a constant current. The added action potential was modelled by a function $\zeta(x) = ae^{bx}$ with $K = 1$ ms. The reset value of the potential was negative and $\Delta t = 10$ ms. (A) without the added action potential. (B) with the added action potential. Note that the plasticity window is smaller than the period of postsynaptic firing.

5.3.1.3 Simulations

We numerically checked the assumption that STDP can be expressed as a form of TD learning. We used setups with postsynaptic spikes being generated by a transient pulse of current, a constant input current or by presynaptic activity.

Pulsed input current A neuron which received input on a single excitatory synapse. Postsynaptic action potentials were paired to presynaptic activation by an externally injected rectangular pulse of current. Figure 5.4 shows such pre-postsynaptic pairings with a postsynaptic action potential triggered $\Delta t = 30$ ms before and after the presynaptic spike.

Figure 5.5 shows the changes in postsynaptic activity for different values of Δt and u_r. It can be seen that the values of Δt had a clear effect on the shape of the WOP. Similar to hebbian STDP, for delays between presynaptic activation and postsynaptic spiking much larger than Δt the change in amplitude tends to zero.

Input caused by presynaptic activity A neuron was connected to a number N of presynaptic inputs and was stimulated by randomly generated spike trains which were generated using a random-walk-like algorithm. The weights of the synapses were chosen randomly from the exponential distribution, with 80 % excitatory and 20% inhibitory. We chose the first synapse to monitor.

In the pairing between presynaptic spiking and postsynaptic activation some presynaptic spikes can be paired with more than one postsynaptic activation. We therefore paired an input spike with the first postsynaptic spike before –this corresponds to a negative τ, and with the first postsynaptic spike after –which corresponds to positive values of τ (Figure 5.6).

Constant input current A neuron received input on a single excitatory synapse. The input consisted of a constant current $I(t) = I_0$ which caused the neuron to fire output spikes periodically. Postsynaptic action potentials were paired to the presynaptic activation by test presynaptic spikes. The test action potential was chosen so that the total charge q that it delivers to the neuron was

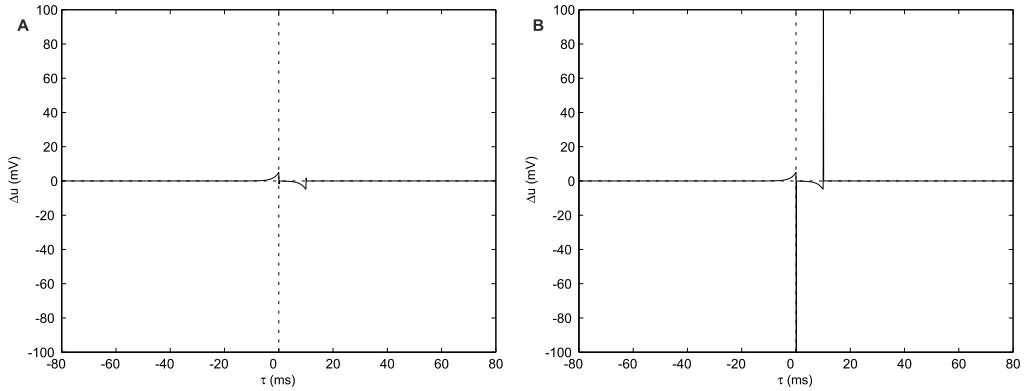

Figure 5.3: The WOP. An IAF neuron stimulated by a rectangular pulse of current. The added action potential was modelled by a function $\zeta(x) = ae^{bx}$ with $K = 1$ ms. The reset value of the potential was negative and $\Delta t = 10$ ms. (A) without the added action potential. (B) with the added action potential.

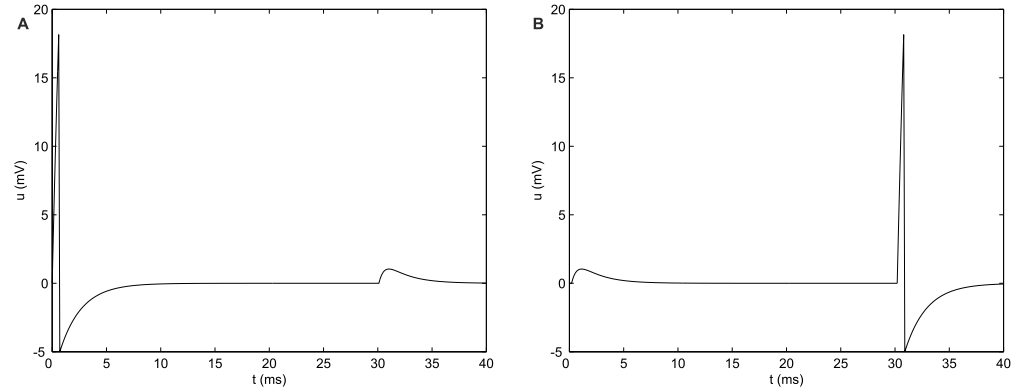

Figure 5.4: The membrane potential of an IAF neuron. (A) The neuron stimulated at time $t = 0$ ms by a pulse of current paired at $t = 30$ ms with a presynaptic spike. (B) The neuron stimulated at time $t = 0$ ms by a presynaptic spike paired at time $t = 30$ ms with a pulse of current.

much smaller than I_0. We only constructed the WOP for values of the delay within the interval $]0, T[$, where T is the period of postsynaptic firing. Within an interval of length T there is only one postsynaptic spike and we considered interactions between the presynaptic spike and the closest postsynaptic spike before –corresponding to a negative τ and after –corresponding to a positive τ.

Figure 5.7 shows the changes for different values of Δt and u_r. It can be seen that the values of Δt have a clear effect on the shape of the plasticity window while u_r has no direct effect. For negative τ much larger than Δt the change in amplitude tends to zero (similar to Hebbian STDP), but for positive $\tau > \Delta t$ the change is positive and ascending.

5.3.2 TD learning in Izhikevich neurons

Using the same test setups as in the case of the IAF neuron we numerically checked if STDP can implement a form of TD learning using a regular spiking Izhikevich (2003) neuron.

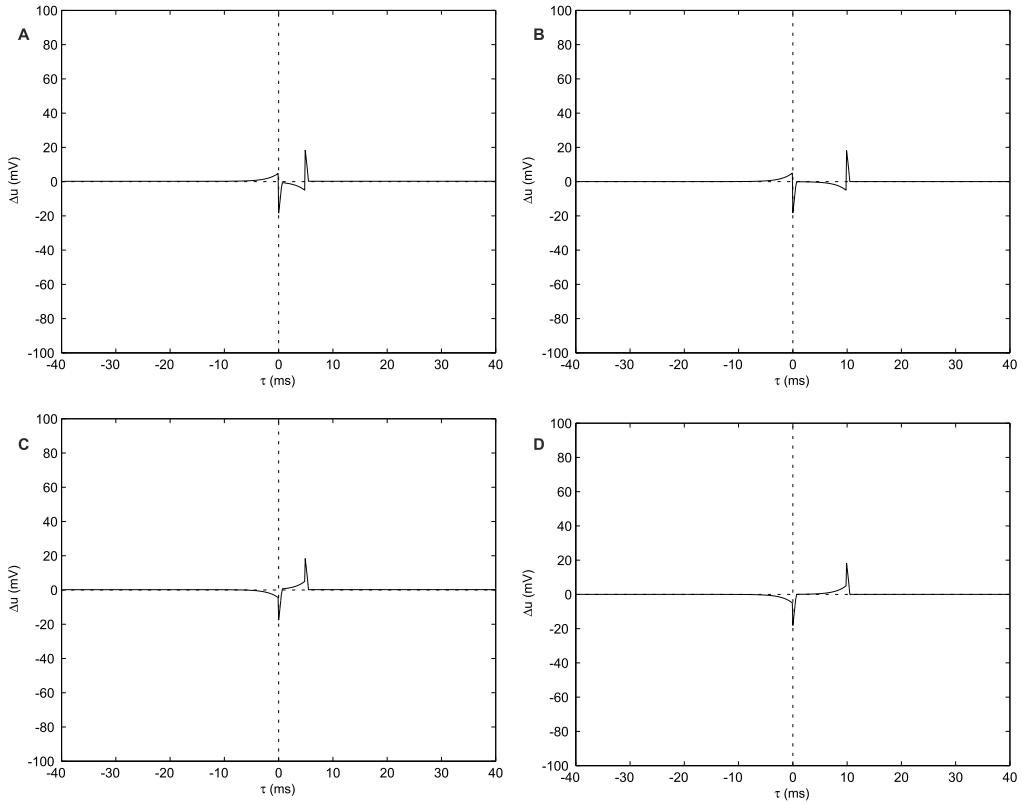

Figure 5.5: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An IAF neuron stimulated by a transient pulse of current. The reset potential u_r and Δt were modified. (A) $\Delta t = 5$ ms and $u_r < 0$. (B) $\Delta t = 10$ ms and $u_r < 0$. (C) $\Delta t = 5$ ms and $u_r > 0$. (D) $\Delta t = 10$ ms and $u_r > 0$.

5.3.2.1 The model

Izhikevich (2003) introduced a model that has the advantage of being able to reproduce the rich dynamics (bursting, chattering, adaptation, resonance) displayed by complex neurophysiological models, like the Hodgkin-Huxley model, while at the same time being far less computationally expensive (for a detailed introduction, see Section 2.4.3 and Izhikevich (2003)).

5.3.2.2 Simulations

Pulsed input current Postsynaptic action potentials were paired to the presynaptic activation by a rectangular pulse of current. Figure 5.8 shows the change in postsynaptic activity caused by the pairing.

Constant input current Like in the case of the IAF neuron, when stimulated by a constant current $I = I_0$, a neuron responded by emitting spikes at regular intervals. The WOP was constructed only for τ contained in the interval $]0, T[$ with T the period of postsynaptic firing. Figure 5.9 shows the changes in the postsynaptic activity caused by the pairing.

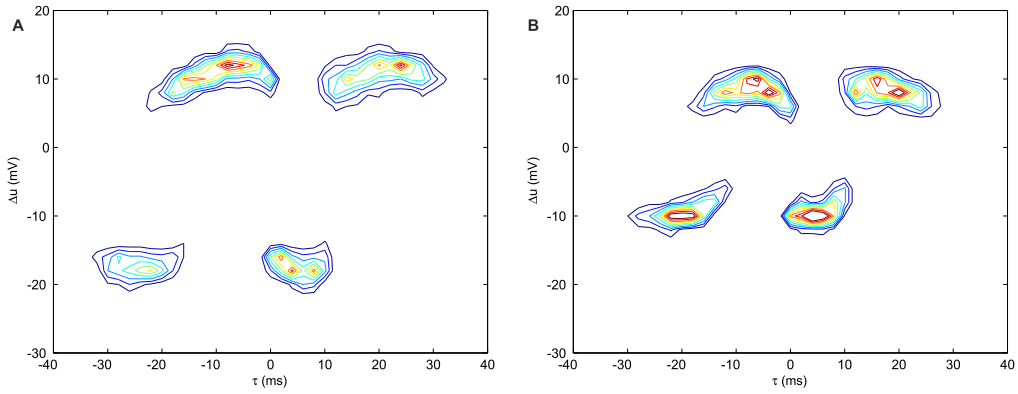

Figure 5.6: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An IAF neuron connected to $N = 1000$ presynaptic neurons. The reset potential u_r was modified. (A) $u_r < 0$. (B) $u_r > 0$.

Input caused by presynaptic activity Like in the case of the IAF, a neuron received inputs on N distinct afferents. The neuron was stimulated by randomly generated spike trains which were generated using a random-walk-like algorithm. The weights of the synapses were generated randomly from the exponential distribution with 80 % excitatory whilst only 20% were chosen arbitrarily as inhibitory. We chose the first synapse to monitor.

In the pairing between presynaptic spiking and postsynaptic activation some presynaptic spikes can be paired with more than one postsynaptic activation. We therefore paired an input spike with the first postsynaptic spike before –this corresponds to a negative τ, and with the first postsynaptic spike after –which corresponds to positive values of τ (Figure 5.10).

5.4 A new implementation of TD learning

In this section we derive a more general and biologically plausible TD rule.

5.4.1 Analysis of the learning rule

We show how the synaptic weights of a neuron evolve under such a plasticity rule.

5.5 Predictive learning

In this section we address the problem of training a network of neurons by modifying the plastic connections of the output neurons in order to bring their generated spikes closer to a given target.

5.5.1 Results

5.5.2 Simulation parameters

5.6 Analytical results for the integrate-and-fire neuron

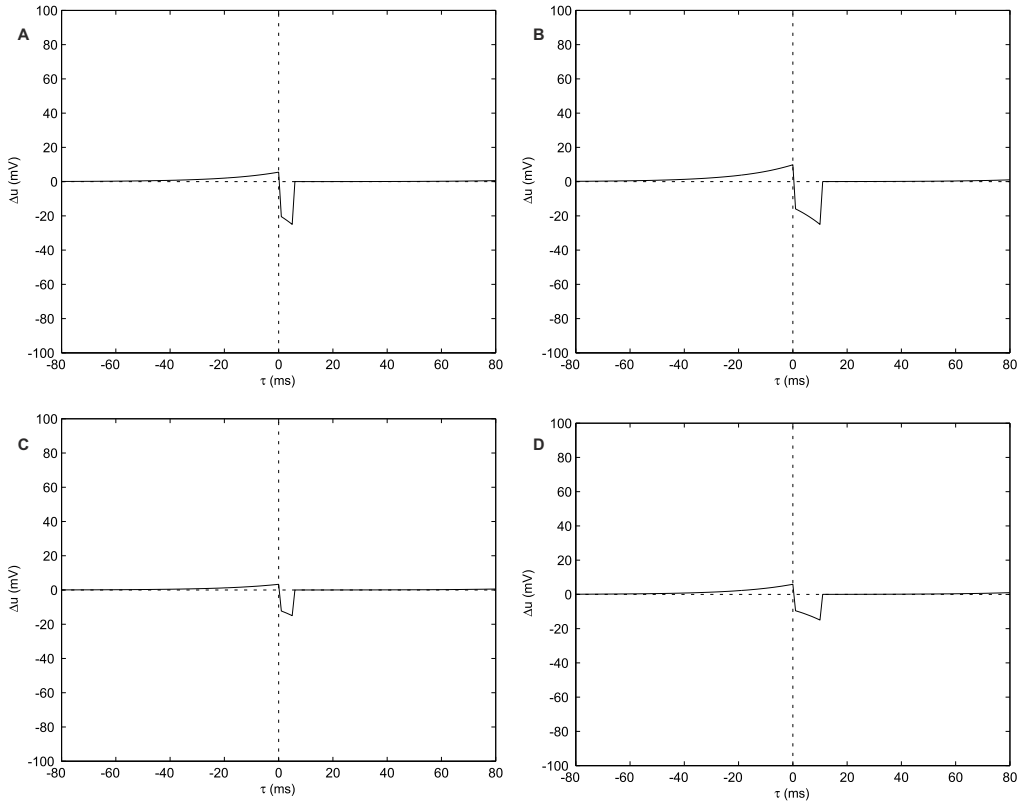

Figure 5.7: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An IAF neuron stimulated by constant input current. The reset potential u_r and Δt were modified. (A) $\Delta t = 5$ ms and $u_r < 0$. (B) $\Delta t = 10$ ms and $u_r < 0$. (C) $\Delta t = 5$ ms and $u_r > 0$. (D) $\Delta t = 10$ ms and $u_r > 0$.

5.6.1 Single input spike

Consider an IAF neuron stimulated by a single spike from some presynaptic neuron.

5.6.2 Constant input current

Consider an IAF neuron stimulated by a constant current of intensity I_0, $I(t) = I_0$.

5.6.3 Pulsed input current

Consider an IAF neuron stimulated by a presynaptic spike which does not cause the membrane potential to cross the threshold and emit a postsynaptic action potential. We therefore inject the neuron with a rectangular pulse of current I which depolarises the membrane and causes the neuron to fire an action potential.

5.7 Conclusion

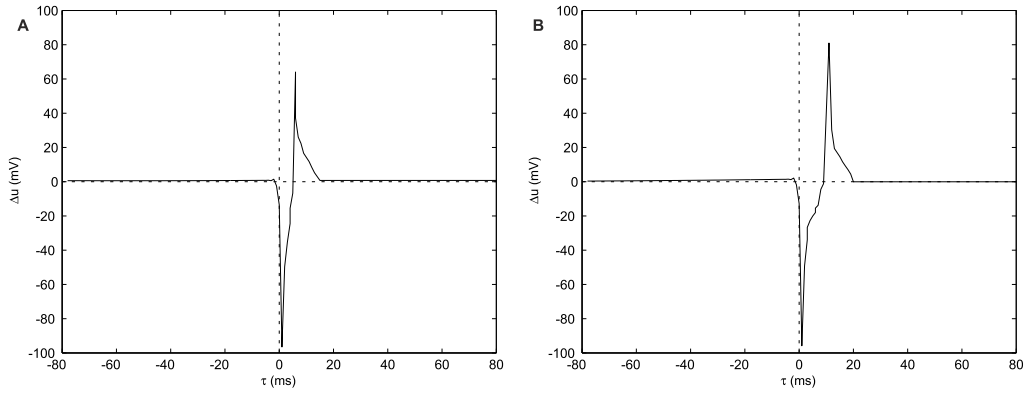

Figure 5.8: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An Izhikevich neuron stimulated by a transient pulse of current. The value of the parameter Δt was modified. (A) $\Delta t = 5$ ms. (B) $\Delta t = 10$ ms.

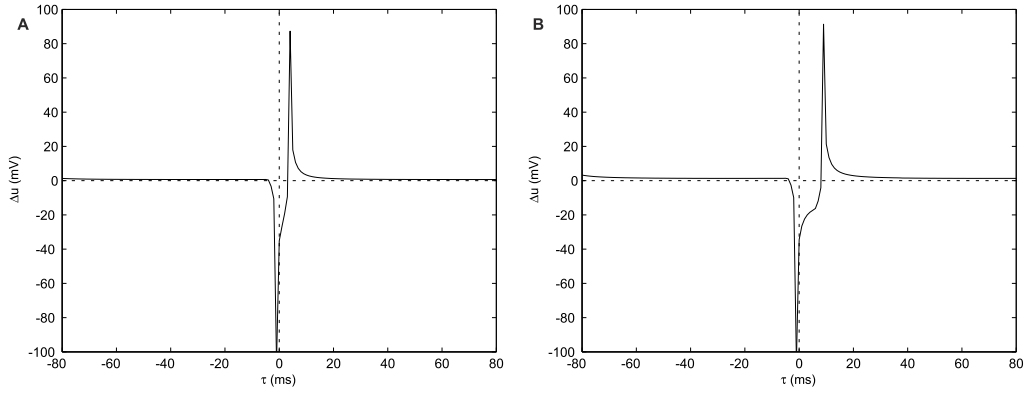

Figure 5.9: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An Izhikevich neuron stimulated by constant input current. The value of the parameter Δt was modified. (A) $\Delta t = 5$ ms. (B) $\Delta t = 10$ ms.

In the case of the IAF neuron, through both analytical derivations and computer simulations we have found that a TD learning rule reproduces a Hebbian STDP-like window of plasticity when postsynaptic spikes are paired to presynaptic activation by a transient pulse of current and with a positive reset potential. For the same pairing but negative reset potential, as well as for pairings caused by constant input current (regardless of the reset potential), the resulting window of plasticity was anti-Hebbian.

For the Izhikevich (2003) neuron we obtained Hebbian STDP-like windows of plasticity for pairings caused by both constant and pulsed input current. There is a qualitative difference with respect to the IAF case because the Izhikevich neuron incorporates the dynamics of the membrane potential during the onset of the action potential. By adding an action potential of non-zero duration to the IAF model, the shape of the plasticity function changes significantly and becomes similar to Hebbian STDP. This shows that the plasticity function resulted from TD learning depends critically on whether the neuron adapts its synapses to learn the shape of its action potential or not. However, the shape of the action potential is commonly considered not to carry information. When we

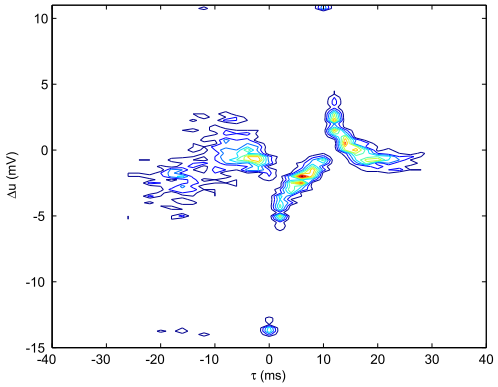

Figure 5.10: The WOP obtained by varying the delay between pre- and post-synaptic spiking. An Izhikevich neuron connected to $N = 1000$ presynaptic neurons.

consider just the TD learning of the sub-threshold dynamics of the membrane potential, the shape of the resulted learning function can loose its similarity with Hebbian STDP.

For both neural models, in the case of irregular synaptic input there was no clear relationship between the plastic changes predicted by TD learning and the temporal delay between the pre- and postsynaptic spikes. Moreover, the sign of these plastic changes did not depend uniquely on the sign of the temporal delay.

We have also derived an alternative form of the TD based plasticity rule.

Using a very simple setup and IAF neurons we have verified that Hebbian STDP implemented as a form of TD learning can be used to predict input sequences moments prior to their expected arrival. Because Hebbian STDP provides a mixture of homeostatic equilibrium and synaptic competition through its LTP and LTD components we have obtained a similar result by using a plasticity rule that reproduces a plasticity window with only a causal LTP part used together with a homeostatic plasticity mechanism in place of the anti-causal LTD part. In this context we have also shown that the synaptic modifications are achieved in an optimal way when the plasticity functions are correlated to the EPSPs caused by input spikes. Essentially, this result suggests that the prediction capabilities of neurons as expressed in the present work and in Rao and Sejnowski (2001) are not the specific consequence of a TD based plasticity rule but rather any rule that provides and STDP-like causality and some intrinsic, homeostatic regulatory mechanism.

CHAPTER 6

PATTERN DETECTION WITH STDP

Given a spiking neuron receiving input on a number of afferents, it has been shown (through computer simulations) that Spike-timing Dependent Plasticity (STDP) enables the detection of spatio-temporal patterns of activity embedded in its synaptic input (Masquelier et al., 2008, 2009). We analytically derive a set of synaptic weights for the input connections which facilitate such a detection mechanism. We also discuss the case of multiple input patterns.

6.1 Introduction

Previous studies (Masquelier et al., 2008, 2009) have shown that Spike-timing Dependent Plasticity (STDP) facilitates the detection of spatio-temporal patterns of activity by spiking neurons. Such precise spatio-temporal patterns of activity lasting from a few ms to several seconds have been found both in vivo and in vitro (Frostig et al., 1990; Prut et al., 1998; Fellous et al., 2004) and were connected to different behavioral states. To verify whether STDP provides a suitable framework for the neuron to learn to respond to such patterns in an unsupervised manner, a given spatio-temporal pattern was embedded into the synaptic input of a single spiking neuron. The neuron received input on 1000 afferents over the time course of 14 s. The input consisted of randomly generated Poisson spike trains at varying frequencies. At random times, instead of the stochastically firing patterns, a precise firing pattern present on about 50% of afferents was delivered to the spiking neuron. The input was thus divided into two parts: a deterministic part given by the embedded pattern and a stochastic distractor part (Figure 6.1). The repeated pattern has the same spike density as the stochastic distactor parts making it invisible in terms of firing rates (Figure 6.1 bottom). This suggests that to enable the neuron to detect such repeating input patterns a mechanism which takes into account spike times is required. The synaptic weights of the connections between the input and the output neuron evolved according to a Hebbian STDP rule.

We essentially show that STDP enables the neuron to act as a coincidence detector and derive analytically a set of weights that allows the neuron to respond selectively to embedded spatio-temporal patterns. We also investigate what happens when multiple patterns are presented to the neuron. In such a situation the neuron becomes randomly selective to one of the patterns. By endowing the neuron with an Intrinsic Plasticity (IP) mechanism (Desai et al., 1999; Daoudal and

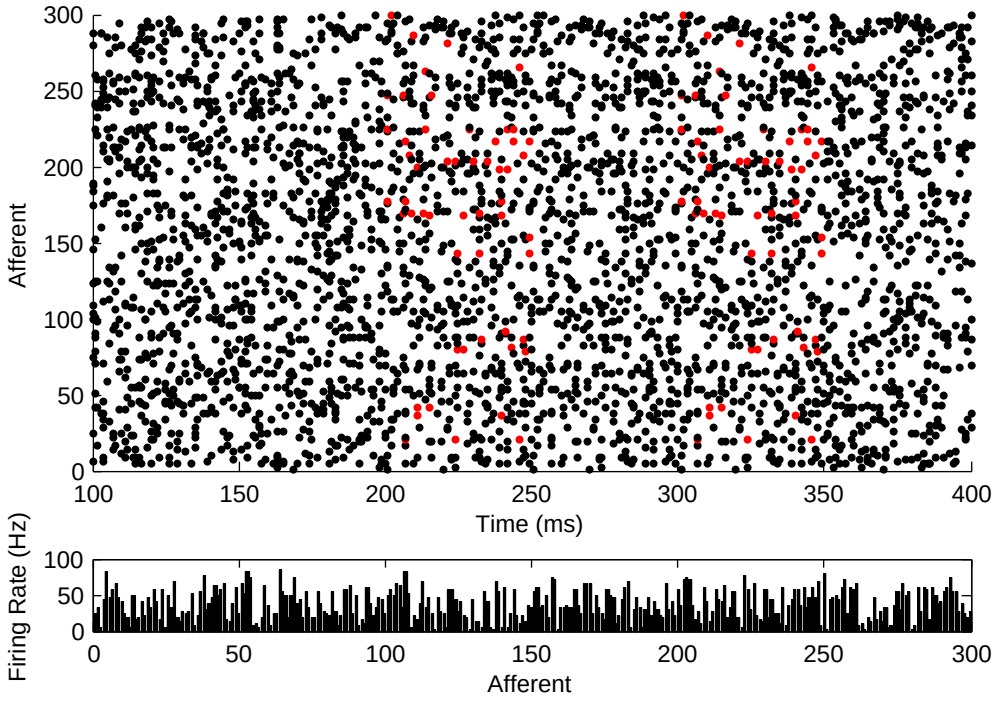

Figure 6.1: Spatio-temporal spike pattern of activity. (top) A repeating 50 ms long firing pattern (shown in red) which affects 150 of the 300 afferents shown. (bottom) The average firing rate of each afferent.

Debanne, 2003; Turrigiano and Nelson, 2004; Lazar et al., 2007) we obtain that the neuron is able to respond to more than one pattern.

6.2 Network description

We present the setup used.

6.3 Results

Initially, the neuron was not selective to the pattern and discharged almost periodically with a period dependent on the strength of the connections between input and output. From this initial behavior, the neuron learned to selectively spike only when the pattern was presented. STDP reinforced the connections which took part in the firing of the neuron and thus postsynaptic firing became correlated with the input. Thus, SDTP detected correlations in the input and reinforced the causal links while decreasing the ones who do not contribute to the postsynaptic firing. Figure 6.2 shows the behaviour of the neuron during the first and last 4 seconds of simulation.

When the neuron was presented with more than one pattern it became selective to only one of them. The pattern was chosen randomly among the two (Masquelier et al., 2008). In Masquelier et al. (2009) the results were extended such that multiple patterns are detected by using more than one neuron. In a simple case in which the input featured only two patterns IP and STDP enabled a

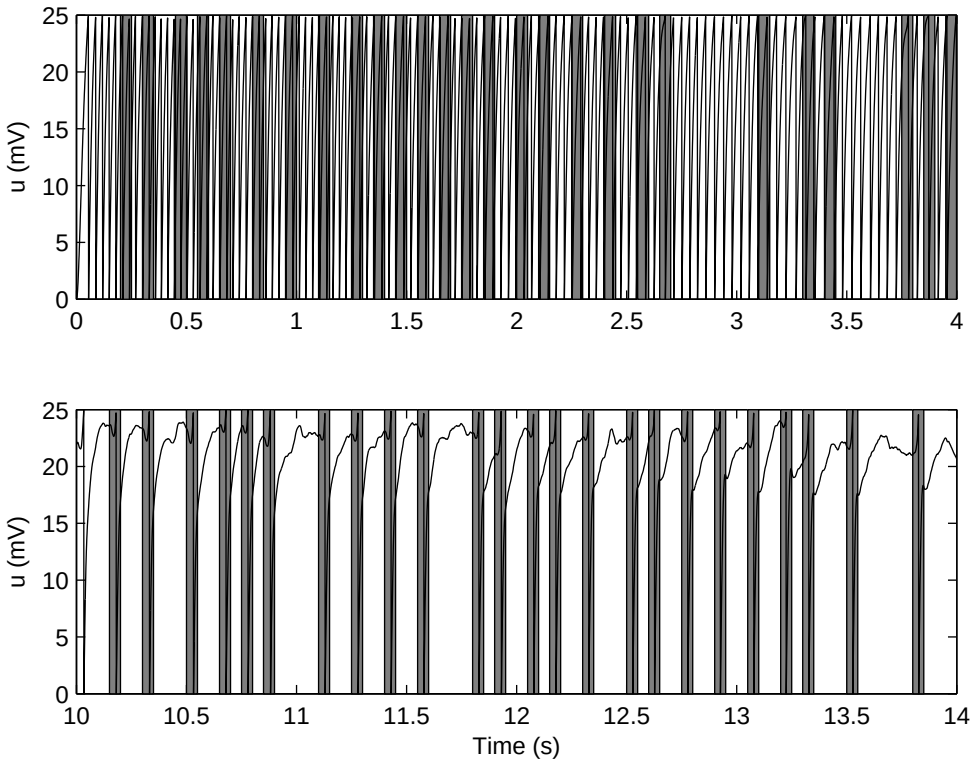

Figure 6.2: Detection of a single pattern by STDP. Initially the neuron fired regardless of pattern presentation. After a period of learning the neuron responded only when the patterns was presented. During the last 4 s of simulation no spikes occurred outside of a pattern presentation.

single neuron to become selective to both of them (Figure 6.3).

6.3.1 Theoretical analysis

It is possible to analytically derive a set of synaptic weights which allows the neuron to detected such embedded spatio-temporal patterns of activity. Using a discretized model of a spiking neuron (for a detailed introduction, see Section 2.4.5 and Cessac (2008)) we show that such a configuration exists. The dynamics of the output neuron are given by

$$u_o(t+1) = \sum_{j=1}^{N} W_j \sum_{l=s}^{t} \gamma^{t-l} W_j(l) := \sum_{j=1}^{N} W_j \, I_j(s,t,\boldsymbol{\eta}). \tag{6.1}$$

Fix a spatio-temporal pattern of activity defined as a raster $[\boldsymbol{\eta}]_{t-R}^{t}$ of width R. The input to the neuron can be divided into two –one deterministic and one stochastic– parts

$$\begin{aligned} I(s,t,\boldsymbol{\eta}) &= I(s,t-R,\boldsymbol{\eta}) + I(t-R,t,\boldsymbol{\eta}) \\ &= \sum_{j=1}^{N} W_j \sum_{l=s}^{t-R} \gamma^{t-l} \eta_{j;l} + \sum_{j=1}^{N} W_j \sum_{l=t-R}^{t} \gamma^{t-l} \eta_{j;l}. \end{aligned} \tag{6.2}$$

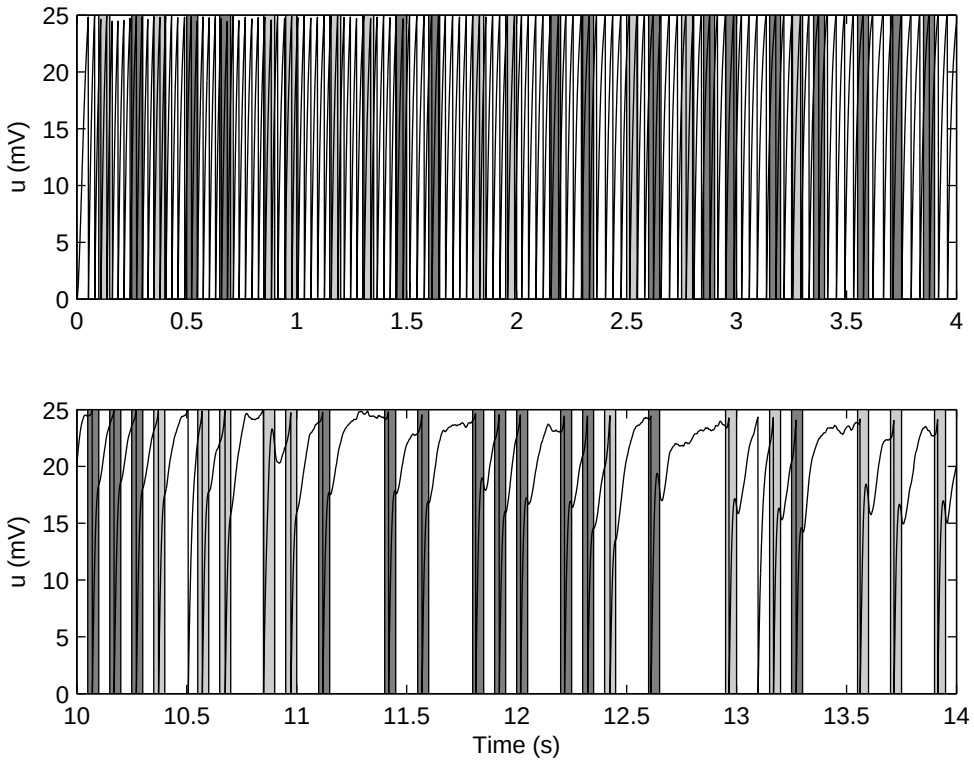

Figure 6.3: Detection of two patterns by STDP and IP. Initially the neuron fired regardless of pattern presentation. After a period of learning the neuron responded only when the two patterns were presented.

We denote the stochastic and deterministic parts by

$$S(j) := I(s, t-R, \boldsymbol{\eta}) \tag{6.3}$$

$$\Delta(j) := I(t-R, t, \boldsymbol{\eta}) \tag{6.4}$$

respectively. To completely determine the distribution of the input $S(j)$ we compute the characteristic function

$$\varphi_S(u) = E\left[e^{i\,u\,S(j)}\right] = E\left[\prod_{j=1}^{N}\prod_{l=s}^{t-R} e^{i\,u\,W_j\,\gamma^{t-l}\,\eta_{j;l}}\right], \tag{6.5}$$

where E is the expected value of a random variable. Because the input spike patterns satisfy Poisson statistics we obtain

$$\begin{aligned}\varphi_S(u) &= \prod_{j=1}^{N}\prod_{l=s}^{t-R}\left[\rho\, e^{i\,uW_j\,\gamma^{t-l}} + 1 - q\right] \\ &= \prod_{j=1}^{N}\prod_{l=s}^{t-R}\left[\rho\left(e^{i\,W_j\,\gamma^{t-l}} - 1\right) + 1\right].\end{aligned} \tag{6.6}$$

There is a one-to-one correspondence between cumulative distribution functions and characteristic functions. Knowing any of these allows the computation the probability density

$$f_S(u) = F'_S(u) = \frac{1}{2\,\pi}\int_{-\infty}^{\infty} e^{-i\,t\,u}\varphi_S(t)\mathrm{d}t. \tag{6.7}$$

From the probability density function of $S(j)$ the expected value

$$E[S] = \rho \sum_{j=1}^{N} W_j \left[\sum_{l=s}^{t-R} \gamma_{t-l} + \sum_{l=t-R}^{t} \gamma_{t-l}\,\eta_{j;l}\right] = \theta, \tag{6.8}$$

and variance

$$\mathrm{Var}(S) = \rho\,(1-\rho)\sum_{j=1}^{N} W_j^2 \sum_{l=s}^{t-R} \gamma^{2(t-l)} = \epsilon. \tag{6.9}$$

can be easily obtained. In order to force the output neuron to generate spikes during the presentation of the pattern and be silent during the stochastic distractor part the variance of the distribution $S(j)$ needs to be minimal while the mean close to the threshold. We obtain the following Lagrange problem

$$\mathrm{var}(S) + \lambda\,E(S) = G(S). \tag{6.10}$$

It follows that

$$\frac{\partial G}{\partial W_j} = 0, \quad \forall\, j \le N \tag{6.11}$$

$$\frac{\partial^2 G}{\partial W_i\,\partial W_j} = 0, \quad \forall\, i,j \le N. \tag{6.12}$$

Eq. 6.11 is equivalent to

$$\lambda\,\rho\sum_{l=s}^{t-R}\gamma^{t-l} + 2\rho\,(1-\rho)\,W_i\sum_{l=s}^{t-R}\gamma^{2(t-l)} + \rho\sum_{l=t-R}^{t}\gamma^{t-l}\,\eta_{j;l}. \tag{6.13}$$

Finally, we obtain an expression for the synaptic weights

$$W_j = -\frac{\lambda\left[\sum_{l=s}^{t-R}\gamma^{t-l} + \sum_{l=t-R}^{t}\gamma^{t-l}\,\eta_{j;l}\right]}{\rho\,(1-\rho)\sum_{l=s}^{t-R}\gamma^{2(t-l)}}. \tag{6.14}$$

6.4 Conclusion

A neuron equipped with STDP was shown to find spatio-temporal patterns of activity embedded into equally dense distractor parts of the input. In a Hebbian-like fashion, the learning mechanism works by strengthening the connections that take part in the firing of the neuron and weakening connections which feature uncorrelated spikes with the output. This enables the neuron to gradually respond only to pattern presentations. We have analytically derived a set of weights which enables such a detection of spiking patterns. The weights have been obtained by minimizing the variance of the output spike firing distribution while keeping the mean close to the threshold during the presentation of the patterns. Additionally, we extended the STDP-based learning mechanism to incorporate additional input patterns by including a homeostatic plasticity mechanism. The presence of IP caused the neuron to respond to more than one pattern.

Such a STDP-based, unsupervised detection mechanism is extremely desirable because, in contrast to other supervised learning approaches (Guetig and Sompolinsky, 2006; Florian, 2010a), it is simple, computationally cheap and biologically plausible and in addition to that also allows a fast, online implementation.

CHAPTER 7

A MODEL OF TMS-INDUCED I-WAVES IN MOTOR CORTEX

Transcranial Magnetic Stimulation (TMS) allows to manipulate neural activity non-invasively and much research is trying to exploit this ability in clinical settings. But the details of how TMS induces the high-frequency repetitive responses (I-waves) observed during epidural recordings remain poorly understood, which hampers targeted clinical application. In this chapter we present a model that reproduces I-waves similar to those observed in epidural responses during in vivo recordings of conscious humans. The model parsimoniously explains the mechanisms underlying I-wave generation together with some of their basic properties such as frequency and timing.

The work presented in this chapter has been published and submitted as (Rusu et al., 2011a,b).

7.1 Introduction

Non-invasive brain stimulation techniques such as Transcranial Magnetic Stimulation (TMS) have gained much attention in recent years after promising results in treating several neurological disorders such as depression or stroke (Liepert et al., 2000; Loo and Mitchell, 2005). The ability to influence brain activitiy non-invasively is very appealing, but it has been difficult to establish how exactly TMS activates different types of neurons in cortical circuits. In a standard single-pulse paradigm a TMS coil is placed over the motor cortex. The fluctuating magnetic field induces an electric field which affects the excitability of central motor pathways and causes strong depolarization of large neuronal populations. As a result, high-frequency ($\sim$ 600 Hz) descending volleys of activity can be observed by placing electrodes in the epidural space (Di Lazzaro et al., 1998b,a, 2000, 2001). The earliest wave that persists after cortical depression or after cortical ablation is thought to be generated by direct stimulation of pyramidal tract neurons and is therefore termed D-wave. The later waves are considered to have an indirect origin and are thought to be the result of action potentials (APs) from presynaptic fibres impinging on the dendritic tree of pyramidal tract neurons. They are therefore termed I-waves. In order to gain a better understanding of the biophysical basis underlying the magnetic stimulation and how cortical circuits give rise to D- and I-waves it is

useful to develop a sufficiently detailed computational model to account for the effects of TMS at the cellular level. Several theoretical mechanisms believed to be responsible for the generation of

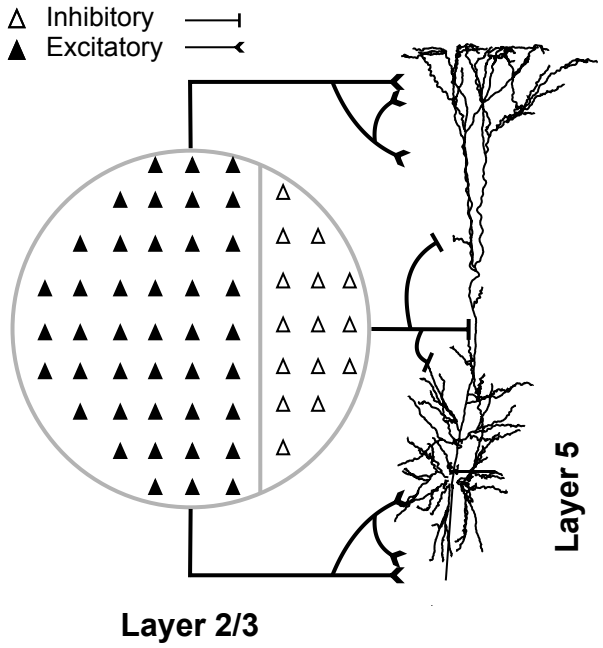

Figure 7.1: The model used including a reconstructed dendritic tree of a L5 pyramidal cell. A total of 300 excitatory and inhibitory L2/3 cells (ratio 4:1) project synapses on to the L5 cell.

I-waves have been proposed, but none of them has gained widespread acceptance (Ziemann and Rothwell, 2000; Esser et al., 2005). Here we aim at uncovering mechanisms behind D- and I-wave generation by investigating the effects of magnetic stimulation at the cellular level. We constructed a model consisting of a complex compartmental layer 5 (L5) cell stimulated by a pool of layer 2 and 3 (L2/3) excitatory and inhibitory cells in a 4:1 ratio (Beaulieu and Colonnier, 1985) (see Figure 7.1). These cells project randomly onto the basal and apical dendrites of the L5 cell. Ion channel kinetics were modeled using a Hodgkin-Huxley formalism while neurotransmission is mediated by excitatory voltage-independent and voltage-dependent channels together with inhibitory channels. The model explains the generation of D- and I-waves and accounts for their frequency and timing. Specifically, the generation of I-waves is shown to be the product of intrinsic and extrinsic factors. Sychronous volleys of excitatory and inhibitory post-synpatic potentials (EPSPs and IPSPs) from L2/3 cells interact on the complex dendritic tree of the L5 cell. In response, the L5 cell's spiking mechanism generates brief trains of action potentials at typical I-wave frequencies. As shown in the following, our model reproduces findings from a range of experiments including pharmacological or behavioural modulation of I-waves and facilitation and inhibition effects observed in paired-pulse stimulation protocols.

7.2 Results

7.2.1 L5 cell spiking at I-wave frequency in response to direct current injection

The simulated L5 cell is capable of firing spikes at I-wave frequencies. We simulated the direct injection of rectangular current pulses into the soma of the L5 cell. Figure 7.2A shows a simulated voltage trace recorded at the L5 cell axon exhibiting fast APs in response to a somatic current injection of 1 nA for 50 ms. Figure 7.2B plots the number of APs generated during the first 10 ms of stimulation as a function of the injected current, which was varied in steps of 0.1 nA. Indeed, with sufficient stimulation, the neuron generates firing rates up to 600 Hz (6 spikes in 10 ms), matching typical I-wave frequencies.

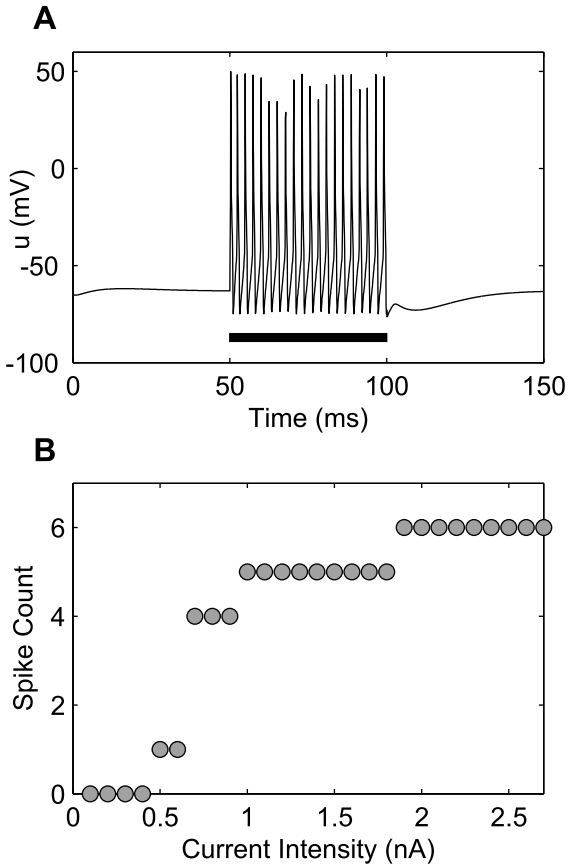

Figure 7.2: High-frequency firing of L5 cell in response to direct current injection. (A) Repetitive axonal response evoked by injection of a rectangular current of 1 nA (horizontal bar). (B) Spike count (measured in the first 10 ms after stimulation) vs. current amplitude.

7.2.2 L5 cell spiking at I-wave frequency in response to synchronous L2/3 input

After establishing that our model L5 cell can in principle generate AP trains at I-wave frequencies, we next tested whether synchronous TMS-induced volleys of synaptic input from L2/3 cells could also generate spike trains with I-wave frequencies in the L5 cell. The electric field induced by a TMS

pulse causes the generation of APs across populations of cortical neurons. In the present model, depending on the number of excited L2/3 neurons the L5 cell is capable of generating multiple output spikes at I-wave frequencies. Specifically, we show how the number of such spikes depends on the percentage of active presynaptic neurons and the balance between inhibition and excitation. We assume that depending on the strength and orientation of the TMS pulse, different percentages of inhibitory and excitatory L2/3 fibers get activated. Inhibitory fast spiking (FS) cells have smaller somata compared to pyramidal regular spiking (RS) cells. Due to this and their different intrinsic membrane properties (see Materials and methods and Pospischil et al. (2008) for more details), inhibitory cells are more excitable (Markram et al., 2004). Figure 7.3 shows output spikes recorded at the level of the L5 cell axon for different fractions of active excitatory and inhibitory inputs. Depending on the balance between excitation and inhibition, many or few APs are generated. Interestingly, identical numbers of APs can be produced by quite different stimulation conditions (compare, e.g., the trace corresponding to 10% excitation and 25% inhibition to the one for 30% excitation and 100% inhibition).

7.2.3 Modeling TMS-induced D- and I-waves

In vivo recordings of epidural responses typically show a D-wave followed by three or four larger I-waves at peak intervals of about 1.4 ms with the D-wave clearly separated from subsequent I-waves. Recordings from single neurons have shown that the delays between D- and I-waves form a continuous distribution between 0.75 and 1.34 ms (Rosenthal et al., 1967). We model the measured D- and I-waves as the superposition of spike trains from many instantiations of the L5 cell model. Large L5 pyramidal cells feature dendritic trees which span all layers of the cortex. This makes L5 cells targets for direct TMS activation (Silva et al., 2008). In the present model TMS elicits responses from an L5 cell through both direct and indirect activation. Direct activation is modeled as a brief direct current injection of variable strength into the L5 cell wich sometimes causes the appearance of a single spike contributing to the D-wave. Indirect activation is caused by stimulation of presynaptic L2/3 neurons which lead to a series of subsequent spikes in the L5 neuron contributing to I-waves.

To simulate D- and I-waves we pool the spiking responses of 1000 model instantiations with different random synaptic connections drawn from log-normal distributions, different direct activation strengths drawn from a normal distribution, and different conductions delays between the L5 cell soma and the site of the epidural recording also drawn from a normal distribution (see Materials and methods for details). Figure 7.4A shows the simulated D- and I-wave response to a TMS pulse activating 100% of the excitatory and inhibitory L2/3 inputs. A smaller D-wave is followed by 4 strong I-waves. Figure 7.4B shows an epidurally measured response adapted from Di Lazzaro et al. (1998b) for comparison. Figure 7.4C shows how the number and size of I-waves changes when different percentages of L2/3 fibers are activated by the TMS pulse. When the percentage of active inhibitory and excitatory L2/3 fibers was increased from 25% and 10% to 75% and 30%, respectively, the second I-wave gained in amplitude.

7.2.4 Pharmacological interventions

The administration of CNS active drugs with known modes of actions has been shown to enable the use of TMS as a measure of cortical excitability (for a review see Ziemann (2004); Paulus et al. (2008)). We set out to test if simulated changes to cortical excitability induced by pharmacological interventions produce alterations to I-waves similar to those observed during in vivo experiments.

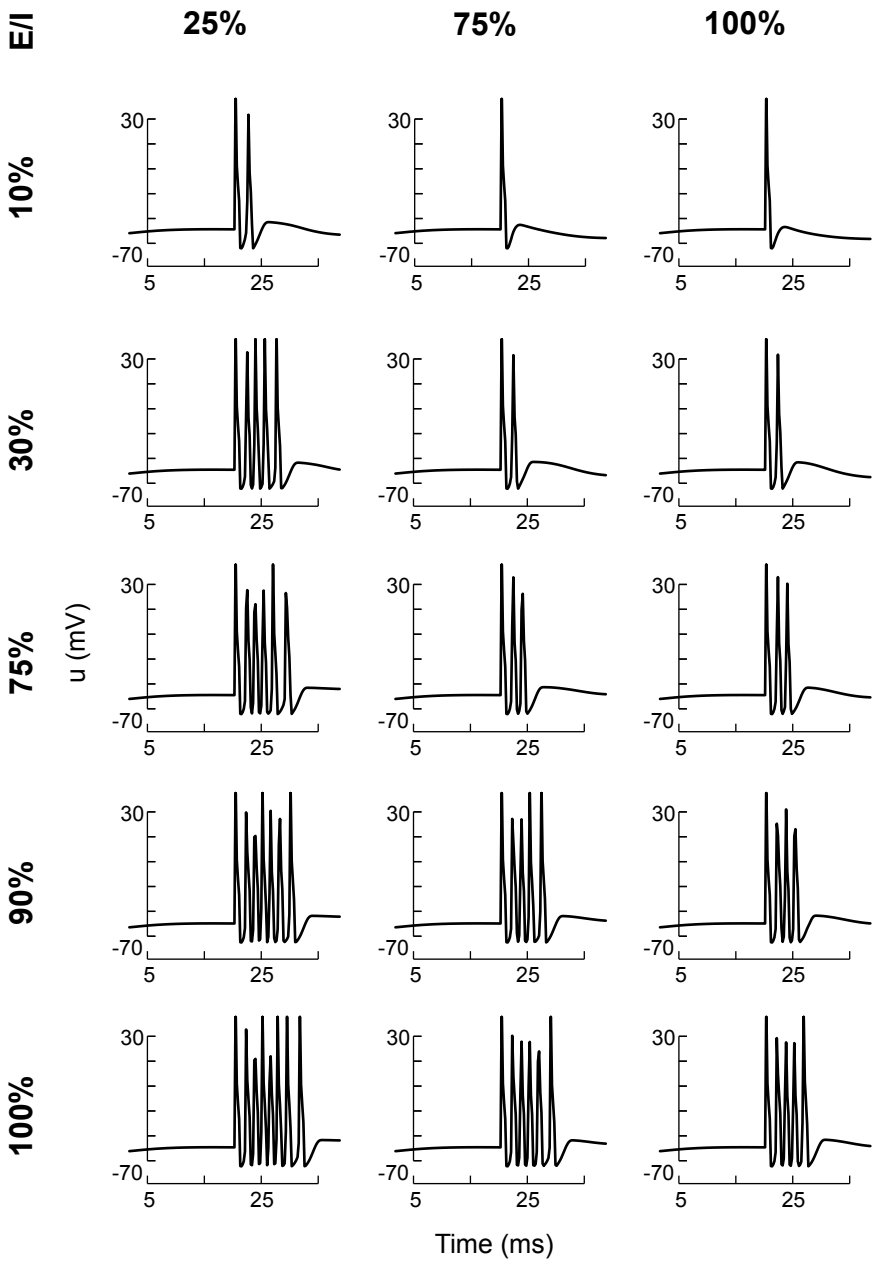

Figure 7.3: L5 cell response to different amounts of excitatory and inhibitory drive. TMS pulses may activate different percentages of inhibitory (I) and excitatory (E) L2/3 fibers. The timing of the TMS pulse is 20 ms. Spikes were measured at the level of the L5 cell axon.

7.2.5 Paired-pulse stimulation

In a final set of simulations, we tested whether our model can also explain findings from so-called paired-pulse stimulation protocols.

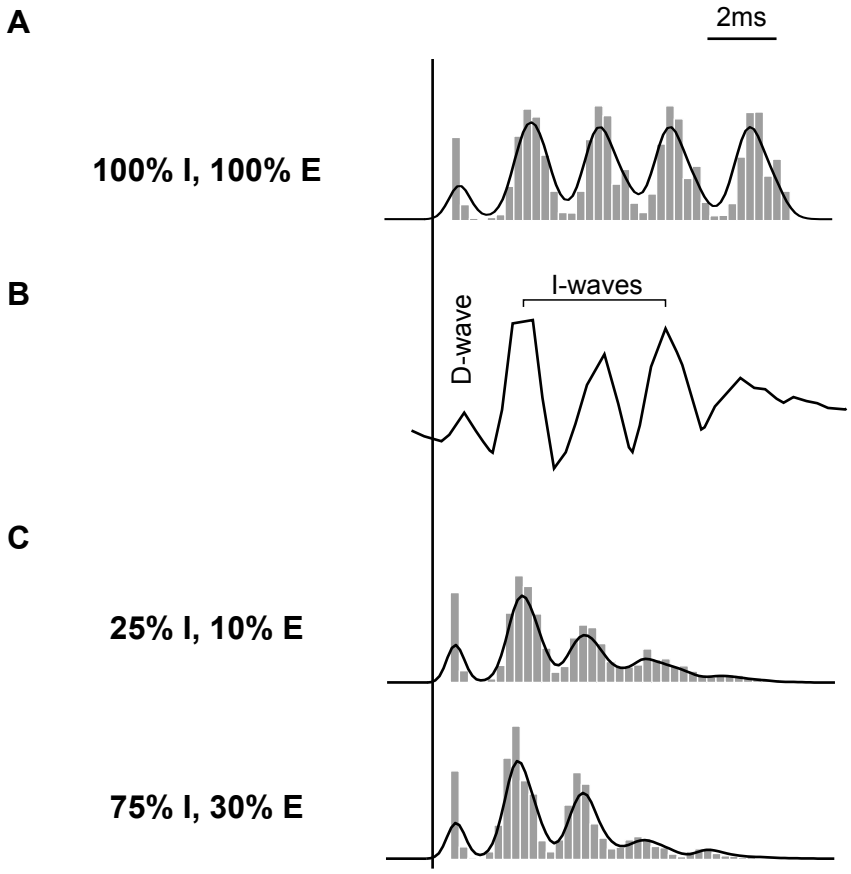

Figure 7.4: Epidural response to a TMS pulse. (A) Simulated epidural recording consisting of a small D-wave followed by 4 larger I-waves. For clarity the distributions of spike timings in the L5 cells corresponding to each epidural response are also shown. (B) Epidural recording of the response generated by the human motor cortex (adapted from Di Lazzaro et al. (1998b)). (C) Simulated responses for different TMS strengths show that in the presence of increased excitation the second I-wave increased in size. The vertical bar indicates the timing of the TMS pulse.

7.3 Conclusion

Our goal was to develop a simple computational model capable of explaining the phenomena of D- and I-waves. There is still a debate on the exact mechanisms which facilitate the production of I-waves in the motor cortex with several theoretical models proposed over the years (Ziemann and Rothwell, 2000). Although each of these models elegantly captures different features of I-waves, none is fully satisfactory. Of the ones proposed, the model in which TMS induces a large activation of L2/3 cells which fire according to their intrinsic membrane properties and act as a resonating circuit directly activating the L5 neurons (Phillips, 1987) is especially relevant for our discussion. Such a model has been argued to require a second generation of spikes in L2/3 cells triggered indirectly by TMS. By using a detailed compartmental model with a rich dendritic tree we showed that such second spikes are not necessary. Our model is fully feed-forward featuring no lateral connections

or loops. Depending on synaptic receptor kinetics and synapse position a spike generated in L2/3 arrives at the soma of the L5 cell with a certain delay. These dendritic inputs interact with the intrinsic membrane properties to cause the generation of periodic I-waves. Another relevant model proposes that TMS activates a chain of inhibitory and excitatory neurons providing waves of activation and inhibition to L5 (Patton and Amassian, 1960; Amassian et al., 1987). Again, our model shows that such chains are not required to explain the basic phenomena of D- and I-waves.

Our model was also shown to reproduce the effects of pharmacological interventions on the size and number of I-waves. Such effects were modeled as changes in $GABA_A$ and AMPA channel conductance. Furthermore, paired-pulse stimulation protocols were modeled. We assumed a lower activation threshold for inhibitory neurons compared to excitatory neurons and short-term synaptic depression of synaptic connections from L2/3 to L5. We showed that under these conditions a paired-pulse stimulation protocol consisting of a sub-threshold TMS pulse followed by a suprathreshold activation caused depression of epidural responses at short ISIs and facilitation at large ISIs, in agreement with what was reported in the literature (Hanajima et al., 1998).

Our models shares many similarities with a previous model presented by Esser *et al.* (Esser et al., 2005). Impressively, this model simulated a complex multi-layerd cortex but it only used point neuron models, neglecting the complex dendritic tree structure of L5 cells. It also required the assumption of an ad hoc refractory mechanism to explain the frequency of I-waves. In our model this is a result of the intrinsic membrane properties of L5 cells and their complex anatomical structure. In a complementary line of research, theoretical studies (Kamitani et al., 2001; Silva et al., 2008; Pashut et al., 2011) have also focused on modeling the effects of TMS on pyramidal cells with arbitrary morphology but without specifically modeling epidural responses.

Our current model has made a large number of simplifications. First, we have very much simplified cortical anatomy in our model. For example, we only modeled one type of fast spiking interneuron despite the large diversity observed in cortex. We have also placed excitatory and inhibitory synapses randomly on the L5 cell dendritic tree. A more careful placement based on anatomical data would be desirable. Furthermore, we have completely neglected recurrent connectivity — either within L2/3 or between L2/3 and L5. Adding such features to our model would certainly make it more realisitc. Beyond these issues, our future work will focus on modeling the induction of long-term plasticity with TMS protocols. Our hope is that this will pave the way for optimizing such protocols for specific clinical applications.

7.4 Materials and methods

7.4.1 The model

We present the model of the L5 cell used.

7.4.2 Simulating TMS pulses

We present a simple model of TMS pulses.

7.4.3 Simulating epidural recordings

We present a simple model of epidural recordings.

7.4.4 L5 cell active conductances

7.4.5 L5 Cell active ionic currents

CHAPTER 8

CONCLUSIONS

We were interested in studying how artificial spiking neural networks process temporal information. In this context, our goal was on one side to study the functional role of synaptic plasticity and excitatory and inhibitory neurons and on the other to develop analytical tools which enable the analysis of the neural code in order to decipher what information is encoded in the sequences of action potentials exchanged by neurons. Such studies are especially relevant in the context of robotic control where the choice of neural code is crucial to agent performance and one usually needs to benchmark different neural models and architectures. Overall, our findings presented in this Thesis might lead to a better design of artificial intelligent systems controllers and can be used to measure neuron synchrony, neural response variability and reliability and finally in the study of how information might be represented in the brain.

We have succinctly reviewed the major characteristics of the two basic computational units of neural microcircuits –neurons and synapses– and how they are reflected in state of the art computational models. Most of our understanding of neural dynamics stems from large simulation of artificial networks (Brunel, 2000; Wielaard et al., 2001; Shelley et al., 2002; Delorme and Thorpe, 2003; Mehring et al., 2003; Hill and Tononi, 2005). We have presented several simple spiking neuron models which capture basic properties of complex biologic models at a low computational cost. Such models are argued to offer a balance between computational efficiency and biological resemblance and are thus suitable for simulations of large networks and analytical treatment. A short comparison of artificial neural networks from the viewpoint of their biological relevance to realistic microcortical circuits was also presented.

We have presented the Liquid State Machine (LSM) (Maass et al., 2002b) and Echo State Network (ESN) (Jaeger, 2001b) computational paradigms for universal, online computing on continuous input streams. They have a number of interesting computational properties such as short term memory, parallel processing and high resistance to noise and were successfully applied to path planning, object tracking, motion prediction and control of autonomous robotic agents. In this context, we have developed “Robby”, a framework for distributed robotic control by spiking neural networks. The framework offers support for different robotic devices and several types of spiking neurons and plasticity rules. We have also presented Particle Swarm Optimization (PSO), an evolutionary optimization technique, which was successfully used to solve a number of optimization problems. We

have shown how PSO could be used to solve such problems and also proposed how it could be applied to increase LSM performance by either optimizing its computational properties such as the separation property (Huang et al., 2009) or its topology in order to maximize the mutual information passed to the readout layer (Chapter 3; Rusu (2011); Rusu and Ahn (2011)).

The distance between two spike trains reflects their similarity. Spike train measures were successfully used in classification of neural recordings in response to different stimuli in an attempt to predict the presented stimulus, to measure the variability of neural responses to same stimulus across multiple recordings and to quantify the degree of synchrony between neurons. We have introduced a new class of spike train measures inspired by the Pompeiu-Hausdorff distance. The measures are sensitive to the precise timing of differences across spike trains with each individual spike considered to be carrying information. Depending on a kernel $\mathscr{H}$ the measures exhibit different behaviors. On one side the kernel served the purpose of providing a local perspective, around each point in a spike sequence while on the other it may be regarded as a magnifying glass to be used to focus on one specific area of the spike trains. From a strictly mathematical point of view $\mathscr{H}$ can be just about any function because the generated metrics are commensurable. Some, however, will have a lesser physiological interpretation than others. Because the metrics generate the same topologies regardless of the choice of kernel, topological properties are the same across all spike train spaces. In simple simulations, a single spike either inserted or shifted in a spike train, the introduced measures were shown to depend on the precise timing of the spike in contrast to popular metrics such as Victor & Purpura and van Rossum (Chapter 4; Rusu and Florian (2010)).

Rao and Sejnowski (2001) have shown that Temporal Difference (TD) learning in the case of a complex biophysical model of a cortical neuron leads to a Hebbian-like Spike-timing Dependent Plasticity (STDP) rule. For simple neural models, such as the integrate-and-fire (IAF) and Izhikevich neurons we have shown the precise conditions under which TD learning leads to a Hebbian STDP-like window of plasticity. Essentially, we have found that TD learning does not lead to Hebbian STDP. In a simple setup, using IAF neurons we have also verified that such a spike-timining based plasticity mechanism implemented as a form of TD learning can be used to predict input sequences moments prior to their expected arrival. Because STDP provides a mixture of homeostatic equilibrium and synaptic competition through its LTP and LTD components we obtained a similar result by using a plasticity rule that reproduces a plasticity window with only a causal LTP part used together with a homeostatic plasticity mechanism in place of the anti-causal LTD part. In this context, we have also shown that the synaptic modifications are achieved in an optimal way when the plasticity functions are correlated to the EPSPs caused by input spikes. Essentially, this result suggests that the prediction capabilities of neurons as expressed here and in Rao and Sejnowski (2001) are not the specific consequence of a TD based plasticity rule but rather any rule that provides and STDP-like causality and some intrinsic, homeostatic regulatory mechanism (Chapter 5; Rusu (2008, 2009); Florian and Rusu (2009); Rusu and Florian (2009)).

When equipped with STDP a single spiking neuron receiving input on several afferents was shown to be capable of detecting spiking activity patterns (Masquelier et al., 2008, 2009). The spatio-temporal patterns of activity were embedded into equally dense distractor parts of the input thus making them invisible in terms of firing rate. We have analytically derived, under weak assumptions, a set of weights which enables such a detection of spiking patterns. In addition, it was shown that when a single neuron was presented with more than one embedded input pattern it randomly became selective to one (Masquelier et al., 2008). To this end, we have shown that in the presence of a homeostatic plasticity mechanism a single neuron was able to learn to respond to two patterns.

The case of multiple neurons and input patterns was studied in Masquelier et al. (2009) where it was shown that lateral connections between neurons ensured a competitive learning mechanism which forced neurons become selective to different patterns. We were however interested in the prediction capabilities of a single neuron. In contrast to similar results which hold in the context of supervised learning (Guetig and Sompolinsky, 2006; Florian, 2010a), such a STDP-based, unsupervised detection mechanism is extremely desirable since it is not only simple, computationally cheap and biologically plausible but also allows a fast, online implementation (Chapter 6).

Transcranial Magnetic Stimulation (TMS) allows to manipulate neural activity non-invasively and has been hypothesized to improve learning, facilitate stroke rehabilitation, treat depression, schizophrenia, chronic pain, or addictions such as alcoholism. Despite recent success in clinical treatments little is known about the cellular mechanisms underlying such stimulation techniques or the nature of the high-frequency repetitive responses (I-waves) they induce along descending motor pathways. Moreover, assessing the nature of I-waves or establishing the biophysical basis underlying magnetic stimulation in purely experimental settings remains difficult given the scarce recording opportunities and high variability of results across healthy subjects. There is still a debate on the exact mechanisms which facilitate the production of I-waves in the motor cortex with several theoretical models proposed over the years (Ziemann and Rothwell, 2000). Although each of these models elegantly captures different features of I-waves, none is fully satisfactory. We have introduced a model consisting of a complex compartmental layer 5 (L5) pyramidal cell stimulated by a pool of inhibitory and excitatory layer 2/3 (L2/3) fast- and regular-spiking cells. In our model, I-waves essentially appeared as a result of action potentials from the L2/3 cells impinging onto the dendritic tree of L5 cells. Our model reproduced I-waves similar to those observed in epidural responses during in vivo experiments on conscious humans and explained their formation, frequency, and timing. Furthermore, our model reproduced findings from a range of experiments with different stimulation protocols and pharmacological interventions (Chapter 7; Rusu et al. (2011a,b)).

SELECTED REFERENCES

V. E. Amassian, M. Stewart, G. J. Quirk, and J. L. Rosenthal. Physiological basis of motor effects of a transient stimulus to cerebral cortex. *Neurosurgery*, 20:74–93, 1987. 49

C. Beaulieu and M. Colonnier. A laminar analysis of the number of round-asymmetrical and flat-symmetrical synapses of spines,l dendritic trunks and cell bodies in area 17 of the cat. *Journal of Computational Neurology*, 231:180–189, 1985. 44

W. Bialek and F. Rieke. Reliability and information transmision in spiking neurons. *Trends in neuroscience*, 15:428–434, 1992. 1

M. H. Brickhard. Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:285–333, 1993. 2

N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8:183–208, 2000. 51

H. Burgsteiner. Training networks of biological realistic spiking neurons for real-time robot control. In *Proceedings of the 9th international conference on engineering applications of neural networks*, 2005. 12

H. Burgsteiner, M. Kroell, A. Leopold, and G. Steinbauer. Movement prediction from real-world images using a liquid state machine. *Applied Intelligence*, 26(2):99–109, 2007. 12

B. Cessac. A discrete time neural network model with neurons. rigorous results on the spontaneous dynamics. *Journal of Mathematical Biology*, 56:311–345, 2008. 7, 39

B. Cessac, H. Rostro, J. C. Vasquez, and T. Vieville. How gibbs distributions may naturally arise from synaptic adaptation mechanisms. *Journal of Statistical Physics*, 136(3):565–602, 2009. 1

J. Cheesman and J. Daniels. *UML components: a simple process for specifying component-based software*. Addison-Wesley Publishing, 2000. 10

H. J. Chiel and R. D. Beer. The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20:553–557, 1997. 2

G. Daoudal and D. Debanne. Long-term plasticity of intrinsic excitability: Learning rules and mechanisms. *Learning and Memory*, 10:456–465, 2003. 37

P. Del Giudice, S. Fusi, and M. Mattia. Modeling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. *Journal of Physiology*, 97:659–681, 2003. 1

A. Delorme and S. J. Thorpe. Spikenet: An event-driven simulation package for modeling large networks of spiking neurons. *Network: Computational Neural Systems*, 14:613–627, 2003. 51

N. S. Desai, L. C. Rutherford, and G. G. Turrigiano. Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nature Neuroscience*, 2(6):515–520, 1999. 37

M. M. Deza and E. Deza. *Encyclopedia of Distances*. Springer, 2009. 14

V. Di Lazzaro, A. Oliviero, P. Profice, E. Saturno, F. Pilato, A. Insola, P. Mazzone, P. Tonali, and J. C. Rothwell. Comparison of descending volleys evoked by transcranial magnetic and electric stimulation in conscious humans. *Electroencephalogr Clin Neurophysiol*, 109:397–401, 1998a. 43

V. Di Lazzaro, D. Restuccia, A. Oliviero, P. Profice, L. Ferrara, A. Insola, P. Mazzone, P. Tonali, and J. C. Rothwell. Effects of voluntary contraction on descending volleys evoked by transcranial stimulation in conscious humans. *The Journal of Physiology*, 508:625–633, 1998b. 43, 46, 48

V. Di Lazzaro, A. Oliviero, M. Meglioni, B. Cioni, G. Tamburrini, P. Tonali, and J. C. Rothwell. Direct demonstration of the effect of lorazepam on the excitability of the human motor cortex. *Clinical Neurophysiology*, 111:794–799, 2000. 43

V. Di Lazzaro, A. Oliviero, E. Saturno, F. Pilato, A. Insola, P. Mazzone, P. Profice, P. Tonali, and J. C. Rothwell. The effect on corticospinal volleys of reversing the direction of current induced in the motor cortex by transcranial magnetic stimulation. *Exp Brain Res*, 138:268–273, 2001. 43

S. K. Esser, S. L. Hill, and G. Tononi. Modeling the effects of transcranial magnetic stimulation on cortical circuits. *Journal of Neurophysiology*, 94:622–639, 2005. 44, 49

J. M. Fellous, P. H. Tiesinga, P. J. Thomas, and T. J. Sejnowski. Discovering spike patterns in neuronal responses. *Journal of Neuroscience*, 24:2989–3001, 2004. 37

R. Florian. The chronotron: a neuron that learns to fire temporally-precise spike patterns. *Available from Nature Precedings*, 2010a. 3, 42, 53

R. V. Florian. Challenges for interactivist-constructivist robotics. *New Ideas in Psychology*, 28(3): 350–353, 2010b. 2

R. V. Florian and C. V. Rusu. Temporal difference learning does not always lead to STDP. *Front. Syst. Neurosci. Conference Abstract: Computational and systems neuroscience 2009*, 2009. doi: 10.3389/conf.neuro.06.2009.03.205. 25, 52

R. D. Frostig, R. C. Frysinger, and R. M. Harper. Recurring discharge patterns in multiple spike trains. ii. application in forebrain areas related to cardiac and respiratory control during different sleep-waking states. *Biological Cybernetics*, 62:495–502, 1990. 37

W. Gerstner and W. Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge University Press, 2002. 3, 6, 7, 10, 26

E. Goodman and D. Ventura. Time invariance and liquid state machines. In *Proceedings of the Eighth Joint Conference on Information Sciences*, 2005a. 12

E. Goodman and D. Ventura. Effectively using recurrently-connected spiking neural networks. In *Proceedings of IJCNN*, 2005b. 12

E. Goodman and D. Ventura. Spatiotemporal pattern recognition via liquid state machines. In *Proceedings of IJCNN*, 2006. 12

R. Guetig and H. Sompolinsky. The tempotron : a neuron that learns spike timing-based decisions. *Nature Neuroscience*, 9(3):420–428, 2006. 3, 42, 53

R. Hanajima, Y. Ugawa, Y. Terao, K. Sakai, T. Furubayashi, K. Machii, and I. Kanazawa. Paired-pulse magnetic stimulation of the human motor cortex: differences among i waves. *J Physiol*, 509:607–618, 1998. 49

R. Haralick and L. Shapiro. *Computer and Robot Vision*. Addison-Wesley Publishing, 1992. 10

F. Hausdorff. *Grundzuege der Mengenlehre*. Veit, Leipzig, 1914. 13

D. Hebb. *The organization of behavior*. Wiley, New York, 1949. 2

S. Hill and G. Tononi. Modeling sleep and wakefulness in the thalamocortical system. *Journal of Neurophysiology*, 93:1671–1698, 2005. 51

J. Huang, Y. Wang, and J. Huang. The separation property enhancement of liquid state machine by particle swarm optimization. In W. Yu, H. He, and N. Zhang, editors, *Part III, LNCS 5553*, pages 67–76. Springer-Verlag, 2009. 11, 52

E. M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14:1569–1572, 2003. 6, 25, 30, 31, 34

H. Jaeger. Short term memory in echo state networks. Technical report, German National Research Center for Information Technology, 2001a. 12

H. Jaeger. The echo state approach to analysing and training recurrent neural networks. Technical report, Technical report, German National Research Center for Information Technology, 2001b. 1, 9, 12, 51

H. Jaeger. Adaptive nonlinear system identiffication with echo state networks. In *Proceeding of NIPS*, 2003. 12

H. Jaeger and H. Haas. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*, 304:78–80, 2004. 12

P. Joshi and W. Maass. Movement generation with circuits of spiking neurons. *Neural Computation*, 17:1715–1738, 2005. 12

Y. Kamitani, V. M. Bhalodia, and Y. Kubota S. Shimojo. A model of magnetic stimulation of neocortical neurons. *Neurocomputing*, 38–40:697–703, 2001. 49

R. Kempter, W. Gerstner, and J. L. van Hemmen. Hebbian learning and spiking neurons. *Physical*

Review E, 59:4498–4514, 1999. 1, 7

J. Kennedy and R. C. Eberhart. Particle swarm optimization. In *Proc. IEEE International Conference on Neural Networks*, pages 1942–1948, 1995. 11

A. Lazar, G. Pipa, and J. Triesch. Fading memory and time series prediction in recursive networks with different forms of plasticty. *Neural Networks*, 20(3):312–322, 2007. 38

J. Liepert, H. Bauder, H. R. Wolfgang, W. H. Miltner, E. Taub, and C. Weiller. Treatment-induced cortical reorganization after stroke in humans. *Stroke*, 31(6):1210–1216, Jun 2000. 43

C. K. Loo and P. B. Mitchell. A review of the efficacy of transcranial magnetic stimulation (TMS) treatment for depression, and current and future strategies to optimize efficacy. *J Affect Disord*, 88(3):255–267, Nov 2005. doi: 10.1016/j.jad.2005.08.001. URL `http://dx.doi.org/10.1016/j.jad.2005.08.001`. 43

W. Maass. On the computational power of noisy spiking neurons. In D. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems*. MIT Press, 1996. 1

W. Maass. Networks of spiking neurons: The third generation of neural networks models. *Neural Networks*, 10(9):1659–1671, 1997. 1

W. Maass, R. A. Legenstein, and H. Markram. A new approach towards vision suggested by biologically realistic neural microcircuit models. *Lecture notes in computer science*, 2525:282–293, 2002a. 12

W. Maass, T. Natschlaeger, and H. Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14:2531–2560, 2002b. 1, 9, 51

W. Maass, T. Natschlaeger, and H. Markram. Fading memory and kernel properties of generic cortical microcircuit models. *Journal of Physiology*, 98:315–330, 2004. 12

H. Markram, M. Toledo-Rodriguez, Y. Wang, A. Gupta, G. Silberberg, and C. Wu. Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci*, 5(10):793–807, Oct 2004. doi: 10.1038/nrn1519. URL `http://dx.doi.org/10.1038/nrn1519`. 46

T. Masquelier, R. Guyonneau, and S. J. Thorpe. Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS ONE*, 3(1):e1377, 2008. 3, 37, 38, 52

T. Masquelier, R. Guyonneau, and S. J. Thorpe. Competitive stdp-based spike pattern learning. *Neural Computation*, 21(5):1259–1276, 2009. 3, 37, 38, 52, 53

O. Mazor and G. Laurent. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48(4):661–673, 2005. 1

C. Mehring, U. Hehl, M. Kubo, M. Diesmann, and A. Aertsen. Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biological Cybernetics*, 88:395–408, 2003. 51

N. Oka, K. Morikawa, T. Komatsu, K. Suzuki, K. Hiraki, K. Ueda, and T. Omori. *Embodiment without a physical body, in R. Pfeifer, G. Westermann, C. Breazeal, Y. Demiris, M. Lungarella, R. Nunez and L. Smith, eds.* Proceedings of the Workshop on Developmental Embodied Cognition, Edinburgh, 2001. 2

M. O'Searcoid. *Metric Spaces*. Springer Undergraduate Mathematics Series. Springer, 2007. 15

T. Pashut, S. Wolfus, A. Friedman, M. Lavidor, I. Bar-Gad, Y. Yeshurun, and A. Korngreen. Mechanisms of magnetic stimulation of central nervous system neurons. *PLoS Comput Biol*, 7(3): e1002022, Mar 2011. doi: 10.1371/journal.pcbi.1002022. URL `http://dx.doi.org/10.1371/journal.pcbi.1002022`. 49

H. D. Patton and V. E. Amassian. The pyramidal tract: its excitation and functions. In *The handbook of physiology*, pages 837–861. American Physiology Society, 1960. 49

W. Paulus, J. Classen, L. G. Cohen, C. H. Large, V. Di Lazzaro, M. Nitsche, A. Pascual-Leone, F. Rosenow, J. C. Rothwell, and U. Ziemann. State of the art: Pharmacologic effects on cortical excitability measures tested by transcranial magnetic stimulation. *Brain Stimul*, 1(3):151–163, Jul 2008. doi: 10.1016/j.brs.2008.06.002. URL `http://dx.doi.org/10.1016/j.brs.2008.06.002`. 46

C. G. Phillips. Epicortical electrical mapping of motor areas in primates. In G. Bock, M. O'Connor, and J. Marsh, editors, *Motor Areas of Cerebral Cortex*, pages 5–20. John Wiley, 1987. 48

D. Pompeiu. Sur la continuité des fonctions de variables complexes. *Annales de la Faculté des Sciences de Toulouse*, 2/7(3):265–315, 1905. 13

F. Ponulak. *ReSuMe-new supervised learning method for Spiking Neural Networks*. International Conference on Machine Learning, ICML, 2005. 14

M. Pospischil, M. Toledo-Rodriguez, C. Monier, Z. Piwkowska, T. Bal, Y. Fregnac, H. Markram, and A. Destexhe. Minimal Hodgkin-Huxley type models for different classes of cortical and thalamic neurons. *Biological Cybernetics*, 99:427–441, 2008. 46

M. H. Protter. *Basic elements of real analysis*. Springer-Verlag, New York, USA, 1998. 14

Y. Prut, E. Vaadia, H. Bergman, I. Haalman, H. Slovin, and M. Abeles. Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *Journal of Neurophysiology*, 79:2857–2874, 1998. 37

R. P. N. Rao and T. J. Sejnowski. *Predictive sequence learning in recurrent neocortical circuits*. in S. A. Solla, T. K. Leen and K. R. Müller, eds, Advances in Neural Information Processing Systems 12: Proceedings of the 1999 Conference Vol. 17, MIT Press, Cambridge, MA, pp. 164–170, 2000. 25, 26

R. P. N. Rao and T. J. Sejnowski. Spike-timing-dependent hebbian plasticity as temporal difference learning. *Neural Computation*, 13(10):2221–2237, 2001. 3, 25, 26, 35, 52

R. P. N. Rao and T. J. Sejnowski. Self-organizing neural systems based on predictive learning. *Philosophical Transactions of the Royal Society London A*, 361, 2003. 25, 26

F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the neural code*. MIT Press, Cambridge, MA, 1997. 2, 13

R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Springer, 2009. ISBN 9783540627722. 14

J. Rosenthal, H. J. Waller, and V. E. Amassian. An analysis of the activation of motor cortical neurons by surface stimulation. *Journal of Neurophysiology*, 30:844–858, 1967. 46

C. V. Rusu. Relevance of STDP for TD Learning in spiking neurons. In *Conferinta Nationala ZAC2008 (Zilele Academice Clujene)*, 2008. 25, 52

C. V. Rusu. Modeling stdp as td learning. Technical report, Coneural Technical Reports, 2009. URL `http://www.rusu.coneural.org/phd/raport_STDP_TD_2009.pdf`. 25, 52

C. V. Rusu. Robby: a lightweight neurobotics framework. submitted IJCCC, 2011. 9, 52

C. V. Rusu and H. S. Ahn. Optimal Network Localization by Particle Swarm Optimization. In *Proceedings of IEEE Multi-Conference on Systems and Control*, pages 620–625, 2011. doi: 10.1109/ISIC.2011.6045394. 9, 52

C. V. Rusu and R. V. Florian. Exploring the link between temporal difference learning and spike-timing-dependent plasticity. *BMC Neuroscience*, 10(suppl. 1):p. 201, 2009. doi: 10.1186/1471-2202-10-S1-P201. 25, 52

C. V. Rusu and R. V. Florian. A new spike train metric. *BMC Neuroscience*, 11(suppl. 1):p. 161, 2010. doi: 10.1186/1471-2202-11-S1-P169. 13, 52

C. V. Rusu, U. Ziemann, and J. Triesch. A model of TMS-induced I-waves in Motor Cortex. submitted PLoS Computational Biology, 2011a. 43, 53

C. V. Rusu, U. Ziemann, and J. Triesch. A model of I-wave generation during Transcranial Magnetic Stimulation (TMS). In *Cosyne Abstracts 2012, Salt Lake City, USA*, 2011b. 43, 53

B. Schrauwen, M. D'Haene, D. Verstraeten, and J. Van Campenhout. Compact hardware for real-time speech recognition using a liquid state machine. In *Proceedings of IJCNN*, 2007. 12

M. Shelley, D. McLaughlin, R. Shapley, and D. J. Wielaard. States of high conductance in a large-scale model of the visual cortex. *Journal of Computational Neuroscience*, 13:93–109, 2002. 51

S. Silva, P. J. Basser, and P. C. Miranda. Elucidating the mechanisms and loci of neuronal excitation by transcranial magnetic stimulation using a finite element model of a cortical sulcus. *Clin Neurophysiol*, 119(10):2405–2413, Oct 2008. doi: 10.1016/j.clinph.2008.07.248. URL `http://dx.doi.org/10.1016/j.clinph.2008.07.248`. 46, 49

H. Soula, G. Beslon, and O. Mazet. Spontaneous dynamics of asymmetric random recurrent spiking neural networks. *Neural Computation*, 18(1), 2006. 7

L. Steels and R. Brooks. *The artificial life route to artificial intelligence: Building embodied, situated agents*. Lawrence Erlbaum Associates, Hillsdale, 1995. 2

J. Steil. Memory in backpropagation-decorrelation o(n) efficient online recurrent learning. In *Proceedings of ICANN*, 2005. 12

J. Steil. Online stability of backpropagation-decorrelation recurrent learning. *Neurocomputing*, 69: 642–650, 2006. 12

R. S. Sutton and A. G. Barto. *Reinforcement learning*. MIT Press, 1998. 2

G. G. Turrigiano and S. B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews. Neuroscience*, 5(2):97–107, 2004. 38

M. C. W. van Rossum. A novel spike distance. *Neural Computation*, 13:751–763, 2001. 20

D. Verstraeten, B. Schrauwen, D. Stroobandt, and J. Van Campenhout. Isolated word recognition

with the liquid state machine: a case study. *Information Processing Letters*, 95(6):521–528, 2005. 12

J. D. Victor and K. P. Purpura. Metric-space analysis of spike trains: Theory,algorithms and application. *Network: Computation in Neural Systems*, 8:127–164, 1997. 1

T. P. Vogels, K. Rajan, and L. F. Abbott. Neural network dynamics. *Annual Review of Neuroscience*, 28:357–376, 2005. 1

D. J. Wielaard, M. Shelley, D. McLaughlin, and R. Shapley. How simple cells are made in a nonlinear network model of the visual cortex. *Journal of Neuroscience*, 21:5203–5211, 2001. 51

G. Wolberg and S. Zokai. *Robust image registration using log-polar transform.* IEEE International Conference on Image Processing, 2000. 10

Y. Zhang and K. Wang. The application of liquid state machines in robot path planning. *Journal of Computers*, 4(11):1182–1186, 2009. 12

U. Ziemann. TMS and drugs. *Clin Neurophysiol*, 115(8):1717–1729, Aug 2004. doi: 10.1016/j.clinph.2004.03.006. URL `http://dx.doi.org/10.1016/j.clinph.2004.03.006`. 46

U. Ziemann and J. C. Rothwell. I-waves in motor cortex. *Journal of Clinical Neurophysiology*, 17: 397–405, 2000. 44, 48, 53