

Memory formation and recall in recurrent spiking neural networks

THÈSE N° 6260 (2014)

PRÉSENTÉE LE 14 NOVEMBRE 2014

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE DE CALCUL NEUROMIMÉTIQUE (IC/SV)
PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Friedemann ZENKE

acceptée sur proposition du jury:

Prof. B. Falsafi, président du jury
Prof. W. Gerstner, directeur de thèse
Prof. S. Fusi, rapporteur
Prof. R. Schneggenburger, rapporteur
Prof. W. Senn, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2014

Acknowledgments

I would first like to thank my supervisor, Wulfram Gerstner, for giving me the opportunity to change fields and to work in his lab, which I will always remember as buzzing and inspiring place. Thank you for the support, for helping me develop a fascinating project, and for providing such a rich intellectual environment. The opportunities to visit numerous international conferences, to utilize high end computers, and to learn from all of the excellent and frequently invited speakers have all left a lasting impression that helped to shape the scientist I have become.

Many thanks go to the experts on my jury – Stefano Fusi, Walter Senn, Ralf Schneggenburger and Babak Falsafi – for dedicating their time and energy into supporting my thesis work.

Thank you Tim Vogels for being a mentor and friend. It is hard to overestimate the role you played in this work.

I also would like to thank my colleagues in the LCN, Alex, Carlos, Chantal, Christian, Claudia, Colin, Daan, Dane, Danilo, David, Eilif, Eszter, Everton, Felipe, Guillaume, Henning, Hesam, Julien, Kerstin, Lorric, Marco, Mohammad, Moritz, Nicolas M, Nicolas F, Poz, Richard, Samuele, Skander, Tilo and Tim for giving me a warm welcome into what used to be a completely new field for me. Thank you for making the lab what it is. I am happy I found my way into it and ultimately into computational neuroscience. A lot of the aspects concerned with high performance computation would not have been possible without the support of Stéphane Ecuier. Thank you.

Special thanks are due to my parents, Hiltrud and Martin Zenke. Your unfailing moral support and advise in both scientific and general matters was invaluable to my success in graduate school. And finally, I would like to thank Koshika Yadava for whom the same and more can be said.

Abstract

Our brain has the capacity to analyze a visual scene in a split second, to learn how to play an instrument, and to remember events, faces and concepts. Neurons underlie all of these diverse functions. Neurons, cells within the brain that generate and transmit electrical activity, communicate with each other through chemical synapses. These synaptic connections dynamically change with experience, a process referred to as synaptic plasticity. These synaptic changes are thought to be at the core of the brain's ability to learn and process the world in sophisticated ways.

Our understanding of the rules of synaptic plasticity remain quite limited. To enable efficient computations among neurons or to serve as a trace of memory, synapses must create stable connectivity patterns between neurons. However there remains an insufficient theoretical explanation as to how stable connectivity patterns can exist in the presence of synaptic plasticity. What complicates and limits our understanding is that the dynamics of recurrently connected neurons depend upon their connections, which themselves change in response to the network dynamics. The recursive nature of the problem necessitates that the network connectivity and the synaptic plasticity be treated as a single compound system. Due to the nonlinear nature of the problem this quickly becomes analytically challenging. Utilizing network simulations that model the interplay between the network connectivity and synaptic plasticity can provide insight into this problem. However, most existing network models that implement biologically relevant forms of plasticity become unstable, developing seizure like activity. This suggests that these models do not accurately describe the biological networks, which have no difficulty functioning without succumbing to exploding network activity.

The instability in these network simulations could originate from the fact that theoretical studies have, almost exclusively, focused on Hebbian plasticity at excitatory synapses. Hebbian plasticity causes connected neurons that are active together to increase the connection strength between them. Biological networks, however, display a large variety of different forms of synaptic plasticity and homeostatic mechanisms, beyond Hebbian plasticity. Furthermore, inhibitory cells can undergo synaptic plasticity as well. These diverse forms of plasticity are active at the same time, and our un-

derstanding of the computational role of most of these synaptic dynamics remains elusive. This raises the important question as to whether forms of plasticity that have not been previously considered could – in combination with Hebbian plasticity – lead to stable network dynamics.

To better understand the stability of neural circuits in the presence of synaptic plasticity, we explore how different forms of plasticity and homeostasis interact with the dynamics of biologically inspired spiking networks. Specifically we assess which compensatory or homeostatic mechanisms are required and on what timescale they have to act to be able to stabilize network dynamics in the presence of plausible forms of synaptic plasticity. In addition we investigate the effect of plasticity at inhibitory synapses and its role in creating and maintaining stable network dynamics. Finally we illustrate that by combining multiple forms of plasticity with distinct roles, a recurrently connected spiking network model self-organizes to distinguish and extract multiple overlapping external stimuli. Moreover we show that the network structures remain stable over hours while plasticity is active. This long-term stability allows the network to function as an associative memory, since it classifies distorted or partially cued stimuli according to the previously learned stimuli. During intervals in which no stimulus is shown the network dynamically remembers the last stimulus as selective delay activity.

Taken together this work suggest that multiple forms of plasticity and homeostasis on different timescales have to work together to create stable connectivity patterns in neuronal networks which enable them to perform relevant computation.

Key words: spiking neural networks, synaptic plasticity, homeostasis, Hebbian learning, cell assembly, inhibitory plasticity, heterosynaptic plasticity, STDP, secreted factors, selective delay activity

Zusammenfassung

Das menschliche Gehirn benötigt lediglich den Bruchteil einer Sekunde um ein Bild zu analysieren. Andererseits erlaubt es uns ein Instrument spielen zu lernen oder uns an Ereignisse, Gesichter oder nützliche Konzepte zu erinnern. Einer der Hauptbestandteile unseres Nervensystems sind Neuronen die elektrische Signale verstärken und über chemische Synapsen miteinander austauschen. Diese synaptischen Verbindungen sind veränderbar, ein Prozess der synaptische Plastizität genannt wird. Es wird vermutet, dass Plastizität von zentraler Bedeutung für das Langzeitgedächtnis und das Erlernen neuer Fähigkeiten ist.

Bis zum heutigen Tag ist unser Verständnis der zugrundeliegenden Mechanismen von synaptischer Plastizität nur sehr rudimentär. Damit Netzwerke von Neuronen Informationen verarbeiten können oder um als Langzeitgedächtnis dienen zu können müssen stabile Strukturen in den synaptischen Verbindungen geschaffen werden. Für das Zustandekommen solcher Verbindungen in ständiger Gegenwart von plastisch veränderbaren Synapsen gibt es derzeit keine grundlegende Theorie. Erschwert wird die analytische Behandlung des Problems aufgrund seiner rekursiven Natur, die erfordert, dass neuronale Netzwerke und Plastizität als ein einheitliches dynamisches System aufgefasst werden müssen. Da dieses System hochgradig nichtlinear ist, ist die analytische Betrachtung anspruchsvoll oder oft unmöglich. Hier helfen Netzwerksimulationen teilweise die entstehende Dynamik zu analysieren und zu verstehen. Jedoch sind die meisten solcher Simulationen instabil und entwickeln epileptische Aktivität, wie sie in gesunden biologischen Netzwerken nicht beobachtet wird. Daraus lässt sich ableiten, dass die existierenden Modelle das zugrundeliegende biologische System nicht akkurat genug beschreiben.

Die Instabilität in Netzwerksimulationen könnte darin begründet sein, dass sich bisher existierende Studien fast ausschliesslich mit Hebb'scher Plastizität und exzitatorischen Synapsen beschäftigt haben. Hebb'sche Plastizität bewirkt, dass die synaptische Verbindung zwischen zwei Neuronen verstärkt wird, wenn diese zu wiederholtem Male gleichzeitig aktiv sind.

Dementgegen stehen biologische Netzwerke, in denen eine Vielzahl unterschiedlicher Formen von Plastizität beobachtet werden. Beispielsweise sind inhibitorische Synapsen auch plastisch. Daraus ergibt sich die interessan-

te Frage, ob diese bisher wenig betrachteten Formen der Plastizität, wie beispielsweise heterosynaptische Plastizität, eine wichtige Rolle beim Erlernen von stabilen Netzwerkstrukturen zufällt. Hier betrachten wir die Stabilität von biologisch inspirierten neuronalen Netzen mit plastischen Synapsen und unterschiedlichen homeostatischen Mechanismen. Insbesondere gehen wir der Frage aufgrund welcher Formen von Homeostasis in der Lage sind Hebb'sche Plastizität unter Kontrolle zu halten und auf welcher Zeitskala die Homeostase dafür aktiv sein muss. Darüberhinaus betrachten wir die Netzwerkeffekte, die sich aus plastischen inhibitorischen Synapsen ergeben und welche Rolle solchen Formen der synaptischen Plastizität im Sinne der Netzwerkstabilität zukommen könnten.

Schliesslich zeigen wir in einer abschliessenden Simulationstudie, dass die Kombination mehrerer unterschiedlicher Formen von Plastizität und Homeostasis ein Netzwerk stabilisieren kann, und es ihm gleichzeitig erlaubt selbständig zu erlernen und zwischen mehreren überlappenden externen Stimuli zu unterscheiden. Die dafür notwendigen Strukturen bleiben über Stunden stabil im Netzwerk erhalten und erlauben es auch fehlerhafte Stimuli oder Bruchstücke derer richtig dem ursprünglichen Stimulus zuzuordnen. Dafür entstehen im Netzwerk verschiedene dynamische Zustände für jeden bekannten Stimulus, in denen das Netzwerk verweilt solange es nicht erneut stimuliert wird. Das Netz arbeitet daher als assoziativer Kurzzeitspeicher.

Zusammenfassend lässt sich sagen, dass die Zusammenarbeit verschiedener Formen von Plastizität und Homeostasis entscheidend darüber ist ob neuronale Netzwerke stabile Verbindungsstrukturen erlernen mit denen sie nützliche Funktionen ausführen können.

Stichwörter: Pulsgekoppelte neuronale Netze, synaptische Plastizität, Homeostasis, Hebb'sches Lernen, inhibitorische Plastizität, heterosynaptische Plastizität, STDP, sekretierte Wachstumsfaktoren, selektive anhaltende Aktivität

Contents

Foreword	1
1 Introduction	3
1.1 Existing plasticity models in large recurrent network simulations	9
1.1.1 Balanced networks	9
1.1.2 Background activity	10
1.2 Plasticity models in high noise networks	11
1.3 Plasticity in large networks	12
1.3.1 Multiplicative STDP	14
1.3.2 Additive STDP	15
1.3.3 Inhibitory Plasticity	16
1.4 Summary and Conclusion	16
2 Plasticity needs rapid homeostasis	19
2.1 Introduction	21
2.2 Results	22
2.2.1 Simulation results	23
2.2.2 Mean field model	26
2.2.3 Model comparison	29
2.2.4 Weight decay	31
2.2.5 Synaptic scaling	32
2.3 Discussion	34
2.3.1 Homeostasis and plasticity	35
2.3.2 Influence of the model design	37
2.3.3 Experimental evidence	38
2.3.4 Limitations	39
2.3.5 Conclusion	40
2.4 Methods	42
2.4.1 Neuron model	42
2.4.2 Network model	43
2.4.3 Plasticity model	44
2.4.4 The time constant of plasticity	46

Contents

2.4.5	Numerical simulations	46
2.4.6	Derivation of the stability condition in the mean field model	46
3	Inhibitory plasticity in recurrent networks	51
3.1	Introduction	51
3.2	Results	52
3.2.1	Attractor dynamics and inhibitory synaptic plasticity	56
3.3	Discussion	58
3.4	Methods	60
3.4.1	Neuron Model	60
3.4.2	Inhibitory STDP model	61
3.4.3	Large network model	61
3.4.4	Data analysis and measuring spike-spike correlations .	62
4	Memory formation and recall	63
4.1	Introduction	63
4.2	Results	65
4.2.1	Effect of impaired individual forms of plasticity	73
4.2.2	Role of consolidation dynamics	73
4.3	Discussion	75
4.4	Methods	78
4.4.1	Neuron model	78
4.4.2	Synaptic plasticity	79
4.4.2.1	Short term plasticity	79
4.4.2.2	Plasticity of excitatory synapses	80
4.4.2.3	Consolidation dynamics	80
4.4.2.4	Homeostatic regulation of LTD	81
4.4.2.5	STDP model of inhibitory synapses	81
4.4.3	Simulations including plasticity	82
4.4.3.1	Simulation of postsynaptic tetanization protocols	82
4.4.3.2	Stimulation paradigm	82
4.4.3.3	Details of feed-forward network simulations .	82
4.4.3.4	Balanced network model	83
4.4.3.5	Simulation details	83
4.4.4	Determining readout populations	83
4.4.5	Analysis of learning rule	84
5	Speed-up limits in network simulations	87
5.1	Introduction	88
5.2	Materials and Methods	90
5.2.1	Simulation code and hardware	90
5.2.2	Network models	91

5.2.2.1	Vogels-Abbott benchmark	91
5.2.2.2	Other network models	91
5.2.2.3	Brunel network	92
5.2.3	Implementation of spike-timing-dependent plasticity	92
5.3	Results	94
5.3.1	Poisson spike trains	95
5.3.2	Vogels-Abbott network benchmark	96
5.3.3	A minimal run time model	99
5.3.4	Computational cost of STDP	102
5.4	Discussion	105
 Conclusion		 107
 A Supplements to Chapter 2		 109
A.1	Supplementary Figure S1	109
A.2	Supplementary Text S1: Rate fluctuations	111
 B Supplements to Chapter 3		 113
B.1	Annotated Protocol for Figure 3.2	113
B.2	Robustness to parameter changes	114
B.3	Supplementary Figures	115
B.4	Simulation parameters in tabular format	121
 C Supplements to Chapter 4		 125
C.1	Supplementary figures	125
C.2	Tabular network description	133
C.3	Derivation of the moments of a synaptic trace for Poisson firing statistics	138
 D Supplements to Chapter 5		 141
D.1	Network details for Vogels-Abbott Benchmark	141
D.2	Network details for Brunel network	144
D.3	Network details for Plastic 25,000 cell network at 3Hz	147
 Bibliography		 164

Contents

List of Figures

1.1	Hebbian plasticity and STDP.	5
1.2	Formation of Hebbian assemblies in neural networks.	7
2.1	The balanced network model.	23
2.2	Network stability during ongoing synaptic plasticity depends crucially on the homeostatic time constant.	25
2.3	Mean field theory predicts the stability of background activity.	27
2.4	The mean field predictions agree with results from direct sim- ulation of the spiking network.	29
2.5	Slow synaptic weight decay renders weight distribution uni- modal, but hardly affects global stability.	31
2.6	Triplet STDP with synaptic scaling requires a fast rate detector.	33
2.7	Postsynaptic priming affects STDP protocols.	41
3.1	Schematic of the inhibitory STDP.	52
3.2	Emergent global balance in networks and cell assemblies through inhibitory STDP.	53
3.3	Schematic of embedding a single cell assembly into a random network.	54
3.4	Cells in the assembly respond approximately linearly to ex- ternal stimulation.	56
3.5	Attractor dynamics of graded activity patterns in a large re- current spiking network model.	57
3.6	Schematic of the input-output curve with bistable behavior in a cell assembly with inhibitory plasticity enabled.	58
4.1	Classic synaptic learning rules fail to maintain stable cell as- semblies in recurrent neural networks.	66
4.2	Heterosynaptic plasticity can create additional fixed points for the learning rule.	68
4.3	Formation of cell assemblies and stable working memory states.	70
4.4	Formation of cell assemblies and stable working memory states (continued).	71
4.5	Network states are associative and serve as working memory.	72

List of Figures

4.6	Block of consolidation causes cell assemblies to decouple from external input.	74
4.7	The effect of heterosynaptic plasticity is small in classic pairing protocols.	76
4.8	Weight fixed point w^{FP} for different pre (x) and postsynaptic (y) firing rates.	85
4.9	Fixed points of the synaptic weight (Eq. (4.27)) as a function of $\xi \equiv \beta(w - \tilde{w})$	85
5.1	Efficient generation of Poisson spike trains from a population of input units.	96
5.2	Vogels-Abbott benchmark.	97
5.3	Scaling behavior in Aurnyn.	100
5.4	Scaling of a plastic network with STDP.	102
5.5	Run times of a balanced network with plasticity in Aurnyn and NEST.	104
A.1	Evolution of the population rate for metaplastic triplet STDP model.	110
B.1	Inhibitory plasticity self-tunes balanced state in recurrent neural networks.	114
B.2	Emergence of asynchronous irregular activity is robust to changes in conductance variables.	116
B.3	Varying excitatory synapses onto excitatory (\bar{g}^{EE}) and inhibitory neurons (\bar{g}^{EI}) independently does not disrupt emergence of AI activity for a wide range of parameters.	117
B.4	Robustness to changes in the learning rate η	118
B.5	Asynchronous network state.	118
B.6	Cell assemblies can be recalled through an external cue.	119
B.7	Supplementary statistics to accompany Fig. 3.2.	120
C.1	Firing rates are bistable under reset triplet and depend on the number of presynaptic contacts.	126
C.2	Assembly structure does not reveal itself immediately from the heterogeneous excitatory-to-excitatory weight matrix.	127
C.3	Cell assemblies form also when individual neurons only have partial access to the input.	128
C.4	Network states are associative and serve as working memory in the network with homeostatic metaplasticity.	129
C.5	Failure modes arising from blocking individual plasticity or homeostatic mechanisms.	130
C.6	Failure modes arising from blocking single plasticity mechanisms.	131

List of Figures

C.7 A slow adaptation current causes spontaneous switching between attractors.	132
--	-----

List of Figures

List of Tables

2.1	Neuron model and synaptic parameters	42
2.2	Network model parameters	44
2.3	Plasticity model parameters	44
B.1	Tabular description of network model	121
B.2	Tabular parameter summary	123
C.1	Tabular description of network model	133
C.2	Tabular parameter summary	137
C.3	The first five moments of a synaptic trace for Poisson firing. .	139
D.1	Tabular description of network model	141
D.2	Tabular parameter summary	143
D.3	Tabular description of network model	144
D.4	Tabular parameter summary	146
D.5	Tabular description of network model	147
D.6	Tabular parameter summary	149

List of Tables

Foreword

In my work I am trying understand how neural circuits form, maintain and recall memories by using learning rules which can accurately describe experimental data. A driving concept for my work has been the Hebbian cell assembly (Hebb, 1949), which has become the stereotype of an associative memory on a neural substrate for many experimentalists and theorists alike. Because cell assemblies depend on acquired recurrent connectivity, this thesis is centered around plausible plasticity models in large recurrent spiking networks.

Chapter 1 reviews the state of the art in this field at the time when I started to work on the subject. At the time only few studies had commenced upon studying realistic plasticity rules in large recurrent neural networks. Most likely because one quickly faces the problem of network instability when introducing plasticity into recurrent neural networks. The question of stability has therefore become a central topic in this thesis.

Chapter 2 deals with analyzing the effect of triplet STDP (Pfister and Gerstner, 2006) – one of the state of the art plasticity rules – on the activity in a large and strongly recurrent balanced network. The chapter lays out a simple mean-field theory that allows to understand the dominant instabilities arising in recurrent spiking networks with plastic synapses. The central result of this analysis is that physiological forms of spike-timing-dependent plasticity (STDP) in recurrent networks have to be accompanied by compensatory mechanisms which react much faster (seconds to minutes) than what is the generally accepted time domain of homeostatic mechanisms (hours to days).

This seemingly paradox result served as my initial motivation to become interested in plasticity of inhibitory synapses (Chapter 3) and lately in heterosynaptic plasticity (Chapter 4).

Chapter 3 summarizes my main results concerning inhibitory synaptic plasticity (ISP) in recurrent network models. The chapter is special in the sense that most of the results shown here are published in Vogels et al. (2011) in a broader context. Although my initial motivation to study inhibitory synaptic plasticity was driven by the knowledge that a fast compensatory mechanism is needed for stable plasticity in networks (Chapter 2), I realized that despite the fact that ISP acts as a powerful stabilizing mechanism

Foreword

for the network activity it does not automatically lead to stable weight dynamics.

Chapter 4 proposes a possible way out of this dilemma by combining standard triplet STDP with a form of heterosynaptic plasticity. The work in this chapter combines many of the lessons learned in the previous chapters and suggests that stable learning of Hebbian cell assemblies is only possible when different plasticity and homeostatic mechanisms synergetically work together.

Finally, Chapter 5, describes the technical aspects of my research. To be able to perform large scale simulations, which have to run over extended periods of time to capture the different timescales of plasticity and homeostasis, I developed the simulation environment Auryn¹. The final chapter presents several performance benchmarks and illustrates that there exists an intrinsic limit which prevents simulations from running significantly faster than real-time.

¹Auryn is freely available as Open Source at <https://github.com/fzenke/auryn>.

Chapter 1

Introduction

Our nervous system serves the sole purpose of controlling our actions and behavior in the complex and changing environment we live in. To interact with the external environment our nervous system receives input from sensory organs on the one side and generates motor output on the other side. The information processing which lies between input and output can range from fundamental reflexes, such as the eye blink reflex that we are born with, to highly complex acquired skills such as speaking a language or riding a bicycle. In either case the nervous system receives external input, processes it and outputs behavior and in most situations, it does so surprisingly quickly.

At the cellular level the nervous system consists of nerve cells or neurons which follow similar functional principles. Neurons, exist in a large diversity, however, they are all capable of receiving input, to process it and to readily communicate information fast and reliably over large distances. In particular neurons allow information flow in one direction. To that end each neuron has designated input structures called dendrites and an axon with synaptic terminals which serves as an output. Neurons are both electrically and chemically excitable. Specifically, they possess special proteins in their cell membranes, called ion pumps, which establish and maintain an electrical potential between the intracellular and the extracellular medium. By actively influencing the conductance of the cell membrane, other proteins – ion channels and receptors – control the current flow through the membrane and can thus affect this potential. Neurons use this property for signal transmission over long distances. The intricate interplay of opening and closing dynamics of distinct voltage-gated ion channels allows the generation and the fast propagation of short stereotypical electrical pulses called action potentials or spikes. These action potentials travel at high speed, up to several meters through the axon until they arrive either at a neuromuscular junction, where they cause muscle movement, or at a synapse from where they are communicated to other neurons.

In most cases a synapse consists of three major components. First, the terminal from the presynaptic axon. Second, a postsynaptic target cell

Chapter 1. Introduction

and third, a zone of the apposition. Depending on the specific makeup of the latter, synapses can be classified into two categories. Gap junctions and chemical synapses. In the case of gap junctions ionic current can flow directly through the synapse into the postsynaptic neuron and thus achieve information transmission. At the more abundant chemical synapses, the pre and postsynaptic neuron are separated by the synaptic cleft. The arrival of a presynaptic action potential triggers the secretion of a chemical transmitter compound – a neurotransmitter – into the synaptic cleft. This chemical diffuses to the postsynaptic cell and binds to specific receptor molecules which control the opening and closing of ion channels in the membrane of the postsynaptic cell. Depending on the specifics of the neurotransmitter and the ion channels, this can either excite or inhibit the firing of action potentials in the postsynaptic cell.

In the mammalian nervous system most excitatory synapses use glutamate as neurotransmitter, while the chief inhibitory neurotransmitter is γ -Aminobutyric acid (GABA). There are several different types of ion channels involved in glutamatergic transmission which includes AMPA (α -Amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid) and NMDA (N-methyl-D-aspartic acid) gated ion channels. Their ion channel opening can cause rapid excitatory postsynaptic potentials (EPSPs) with rise times well below 1ms. However, their temporal effect is short lived and depending on the specifics of the channel most EPSPs mediated by AMPA receptors decay on a timescale of 1-20ms. The fast excitatory transmission mediated through AMPA receptors is complemented by the slower mode of transmission of NMDA receptors (NMDAR). The slower opening and closing kinetics of NMDAR have typical EPSP rise times in the order of ~ 10 ms or longer. They generate long-lasting postsynaptic excitatory effects which may persist up to several hundred milliseconds. Moreover the conductance of most NMDA receptors is strongly dependent on the postsynaptic voltage. When the postsynaptic cell not been excited previously by synaptic input, NMDA channels are generally blocked by magnesium ions. Only when the cell is excited – or more specifically, depolarized – this block is gradually removed and allows substantial currents to pass through the channel, making the receptor an ideal molecular coincidence detector of pre and postsynaptic activation.

The computations that a single neuron can perform are highly limited. These limitations are presumably overcome in the brain by combining many neurons in intricately connected networks. An average neuron in the mammalian brain forms between 1000 to 10000 connections to other neurons which amounts to a total of $10^{14} - 10^{15}$ synapses in the human brain (Kandel et al., 2000). Since the information content of the human genome is insufficient to provide a detailed connectivity plan for all neurons, organisms have to provide the algorithmic means to change synaptic connections as needed to form functional neural circuits. This phenomenon is called *synaptic plasticity* and it describes the ability of synapses to change and

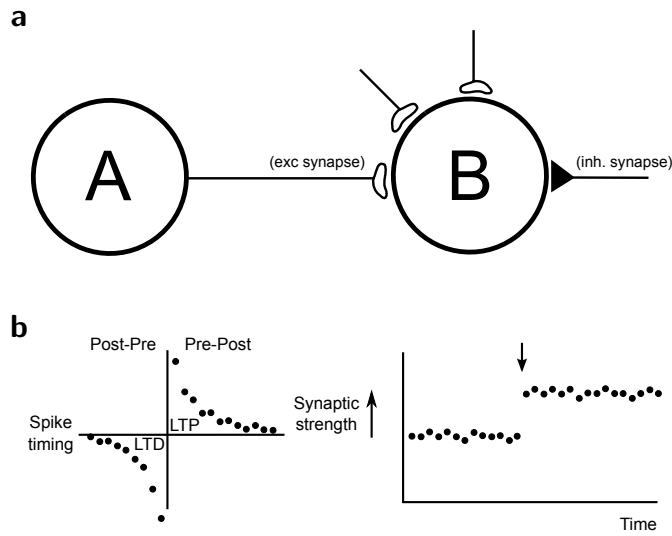


Figure 1.1 – (a) Schematic of an excitatory synaptic connection between neuron A and neuron B. Neuron B receives excitatory synaptic as well as inhibitory input from other neurons. (b) Schematic reproduction of findings on spike-timing-dependent plasticity (STDP) by Zhang et al. (1998). A fixed number of pairings with causal pre-before-post spike timing trigger a long-lasting increase in synaptic strength, while acausal post-before-pre timings cause a decrease in synaptic strength (left). Since the changes are long-lasting, they have been termed long-term potentiation (LTP) and long-term depression (LTD). Right: Schematic of the synaptic strength over time for a typical of an LTP induction protocol. The actual plasticity induction protocol is short (indicated by the downward pointing arrow) compared to the overall time of the experiment. Synaptic changes persist for hours or longer.

consequently modify the neuronal circuit they are part of. Importantly, synaptic plasticity is maintained throughout adulthood, which is the reason why it is believed to underlie learning and memory formation in the brain. To be able to work in concert neurons have to be intricately connected. This requires the synaptic changes that determine how neural circuits are wired together to follow some common organizational principles or *plasticity rules*. In uncovering these rules lies the key to understanding neural computation and memory in the brain.

The current understanding of the different plasticity rules is quite limited. To form neural circuits that compute or serve as a memory, synapses must create novel connectivity patterns between neurons. This task, however, is complicated by the fact that single synapses have only limited access to information about their surrounding world. More specifically, the information they can access is confined in space and time. These spatial and temporal limits are bound to the scales of the chemicals involved in the plasticity process. Based on these constraints Donald Hebb (Hebb, 1949) postulated more than 60 years ago a learning rule that strengthens the con-

Chapter 1. Introduction

nections between co-active neurons. In particular if neuron A repeatedly takes part in firing neuron B, the synapse between both neurons is strengthened (Figure 1.1 (a)). Today the so called forms of Hebbian plasticity are believed to constitute a central role in learning. Moreover Hebbian plasticity has since been found almost ubiquitously in the brain (Bliss and Lømo, 1973; Artola et al., 1990; Markram et al., 1997; Martin et al., 2000; Abbott and Nelson, 2000; Bi and Poo, 2001; Bliss et al., 2003; Caporale and Dan, 2008; Lisman, 2003; Sjöström et al., 2008). Despite a constantly growing body of experimental evidence our understanding of how neural circuits form and compute remains limited at best.

Most advances have been made in understanding the selective response of neurons in the early sensory systems to stereotypical stimuli. For instance, single units in early visual cortex respond preferentially to oriented-bar stimuli shown in a particular area of the visual field (Hubel and Wiesel, 1962). These so called receptive fields are formed during development and it is widely believed that Hebbian learning underlies their formation. Moreover, models of Hebbian plasticity have successfully demonstrated that receptive fields emerge under certain conditions in feed-forward network models (Bienenstock et al., 1982; Oja, 1982; Clopath et al., 2010). The brain, however, is a highly recurrent network. In fact, most of the synaptic input to a cortical neuron originates from other cortical regions and not from sensory input (DeFelipe and Fariñas, 1992). As one moves away from early sensory brain areas, the notion of receptive field becomes increasingly difficult to define. In particular one finds cells which respond to very abstract, high level stimuli which in some cases do not even depend on external stimulation (Quiroga et al., 2005; Gelbard-Sagiv et al., 2008; Li and DiCarlo, 2008). It is tempting to speculate that such spontaneous activations are the neural correlates of thought and other types of high level processing in the brain. However, it remains entirely unclear which connectivity patterns underlie these brain states and how they emerge through synaptic plasticity.

One hypothesis is that subgroups of neurons in an otherwise randomly connected network could form groups of strongly connected neurons, or *cell assemblies* (Hebb, 1949). Importantly the required assembly structure could emerge through Hebbian plasticity triggered by the repeated co-activation of a subset of neurons in an initially naïve network (Figure 1.2 (a)). Such assemblies are associative because the reactivation of only a handful of neurons could be sufficient to spread neuronal activity throughout the entire assembly by dint of the strong synaptic connections inside of the assembly. These *Hebbian* assemblies have since been thought of as the physiological equivalent of associative memories. The concept has multiple appealing features which seem in agreement with biology. First, the required learning rule could be entirely local in space in time. Second, a given circuit can support many many different assemblies to be stored in the same network (Hopfield, 1982). Third, since memories are stored in a distributed fashion,




Figure 1.2 – Formation of Hebbian assemblies in random neural networks. (a) Schematic of the formation of a cell assembly in an initially random network. The initially random connectivity in the naïve network (left) is reshaped by experience through Hebbian plasticity to form a cell assembly (right). (b) Schematic population view of a network with an embedded Hebbian assembly. The population of all excitatory cells is represented by the black circle. Similarly excitatory synapses are summarized by black (and red) arrows. Inhibitory cells and synapses are depicted in blue. A subset of excitatory cells receives correlated input (red), which causes them to form strong recurrent excitatory connections through Hebbian plasticity. (c) Spike raster of a simulation of a spiking network showing run-away potentiation. The cells in the assembly (100 units at top) become highly active due to strengthened weights between them. Eventually the rest of the network destabilizes in run-away potentiation.

Chapter 1. Introduction

the resulting memory network is fault tolerant to the death of individual neurons. Finally, neurons partaking in Hebbian assemblies retain their ability to be highly specialized for the subset of assemblies they are part of, but at the same time the overall connectivity would appear mostly random to an external observer who does not know the “code book”, i.e. the identity of the patterns stored.

Yet there is still little conclusive experimental evidence for the notion of Hebbian assemblies (Quiroga et al., 2005; Gelbard-Sagiv et al., 2008; Bathellier et al., 2012; Fuster and Jervey, 1982; Goldman-Rakic, 1995). The main hindrance is, that it would require simultaneous recordings from prohibitively large numbers of neurons in-vivo. At the same time theoretical approaches to reproducing the formation of cell assemblies in realistic models of spiking neural networks have had limited success (Amit and Mongillo, 2003b; Del Giudice et al., 2003).

To test Hebb’s hypothesis of the cell assembly (Hebb, 1949) most theoretical studies have focused on creating models with learning rules in the context of a plausible network model. Here, generally random balanced network models, which are networks consisting of excitatory and inhibitory neurons, were used. These models give a natural explanation for the high trial-by-trial variability observed in cortical circuits and mimic activity conditions similar to the ones observed in biological networks (for details see Section 1.1.1; van Vreeswijk and Sompolinsky (1996); Brunel (2000); Vogels et al. (2005)).

Biologically plausible learning rules are generally found in models of spike-timing-dependent plasticity (STDP; Figure 1.1 (b)) which is the form of synaptic plasticity ubiquitously found in the brain. STDP has been reported in a variety of different types at excitatory (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001; Caporale and Dan, 2008; Sjöström et al., 2008; Markram et al., 2011) and inhibitory synapses (Woodin et al., 2003; Haas et al., 2006; Vogels et al., 2013). While today’s state-of-the art plasticity models can very accurately describe STDP at the cell-to-cell level (Pfister and Gerstner (2006); Clopath et al. (2010)), the compound systems of learning rules and recurrent network models result in highly unstable models (Chapter 2; Zenke et al. (2013)). Although highly constrained models can reproduce to some extent the formation of cell assemblies (Figure 1.2 (b); Amit and Mongillo (2003b); Del Giudice et al. (2003)), less constrained and therefore more realistic models mostly succumb to irreversible run-away potentiation (Figure 1.2 (c)).

Hebbian plasticity is known to be unstable and therefore prone to developing run-away behavior. The reason for this is that it acts as correlation detector which further amplifies correlations. Hence, a neural population which is highly active, gets even more activated through Hebbian plasticity. Eventually this leads to an explosive activity increase (Fig. 1.2 (c)) with catastrophic effects on the information previously stored in synaptic changes

1.1. Existing plasticity models in large recurrent network simulations

(Fusi, 2002). To avoid this behavior Hebbian plasticity has to be constrained (Miller and MacKay, 1994). In a model, this can be enforced algorithmically (von der Malsburg, 1973). In biological networks, however, it is less clear which mechanisms constrain plasticity. Although a large variety of different compensatory or homeostatic mechanisms is present in real neural networks (Turrigiano, 1999, 2011; Davis, 2013, 2006; Marder and Goaillard, 2006) their individual contribution to network stability remains elusive at best. Importantly, not every plausible form of homeostasis guarantees stable network dynamics. Since explosive run-away potentiation can occur very rapidly in a plastic network (cf. Figure 1.2 (c)), Hebbian plasticity needs to be compensated by an equally rapid mechanism (Chapter 2; Zenke et al. (2013)). However, most known forms of homeostasis are slow and act on timescales of hours or days. Understanding this problem analytically and in simulations and ultimately suggesting possible ways out is one of the quintessential questions addressed in this thesis.

1.1 Existing plasticity models in large recurrent network simulations

The remainder of this chapter is dedicated to describing the state-of-the-art in the field at the beginning of this thesis project. In particular this puts the focus on plasticity models in large recurrent network simulations. Since most of these models build on top of models of balanced networks, their basic terminology will be introduced first (Section 1.1.1). This will be followed by a brief discussion of spontaneous activity and the neural background state in balanced network models (Section 1.1.2) which will entail a more detailed discussion of common sources of noise and stochasticity. Importantly we will distinguish sources of noise that arise naturally in balanced network models and noise that is added algorithmically to the simulation as external input. Finally we will turn to reviewing work on plasticity in large recurrent spiking networks which was published at the time the work presented the remaining chapters of this thesis was started (Section 1.2).

1.1.1 Balanced networks

Most computational studies have focused on balanced networks because they closely resemble two important facts from living systems. First, the circuits of the brain are constantly active and even in the absence of a stimulus, cortical neurons spike in response to recurrent input they receive from other cortical neurons (DeFelipe and Fariñas, 1992). Second, this so called background activity is highly asynchronous and irregular (AI; Section 1.1.2; Burns and Webb (1976)). Furthermore the observed high variability is not simply an effect of neuronal noise, because neuronal responses can be evoked and pre-

Chapter 1. Introduction

dicted reliably (Mainen and Sejnowski, 1995; Badel et al., 2008; Mensi et al., 2012).

In the framework of balanced networks the above observations are a result of an intricate balance between excitatory and inhibitory synaptic currents, which similar to experimental observations (Wehr and Zador, 2003; Okun and Lampl, 2008), cancel each other almost completely. Under these conditions network activity exhibits chaos (van Vreeswijk and Sompolinsky, 1996) which provides an elegant explanation for the asynchronous and irregular activity observed in biological networks. Apart from asynchronous irregular (AI) activity, recurrent networks of spiking neurons also exhibit, depending on the choice of parameters, several other mostly pathological dynamical states (Brunel, 2000; Vogels et al., 2005).

In many model studies the underlying synaptic connectivity of the balanced network is chosen randomly to make minimal model assumptions. By adding structure to random networks, by manually changing synaptic weights (Amit and Brunel, 1997b) or through the addition or removal of connections (Litwin-Kumar and Doiron, 2012), balanced networks have been demonstrated to perform memory recall from Hebbian assemblies (Amit and Brunel, 1997b; Brunel and Wang, 2001; Renart et al., 2007; Mongillo et al., 2008; Hansel and Mato, 2013; Vogels et al., 2011), information transmission (Vogels and Abbott, 2005; Kumar et al., 2008a) and gating (Vogels, 2007) as well as simple logic computations (Vogels and Abbott, 2005).

Since balanced networks are sensitive to synaptic weight changes (Brunel, 2000), one fundamental requirement to plasticity models arises directly: Any synaptic plasticity rule deployed to a balanced network needs to be compatible with the ongoing background activity. That is, the learning rule is required to at least maintain the present dynamical state of the network, to be considered as a credible and plausible plasticity model.

1.1.2 Background activity

About the role of background activity one can only speculate. Sources of stochasticity and neural variability are manifold (Hubbard et al., 1967; Calvin and Stevens, 1968; Rolls and Deco, 2010). Neurons and synapses in the brain are intrinsically noisy and to a certain degree unreliable. It is common practice to model this stochasticity, by providing noise as an external input to the otherwise deterministic network (Morrison et al., 2007; Kunkel et al., 2011; Vogels and Abbott, 2005; Vogels et al., 2011; Vogels and Abbott, 2009; Litwin-Kumar and Doiron, 2012; Amit and Brunel, 1997a; Kumar et al., 2008b; El Boustani et al., 2012). While real neural circuits exhibit activity even in the absence of external input, most balanced network model do not. Rather a network baseline activity, which resembles neuronal background activity, is ensured by providing noisy external input.

However, it is also known that completely deterministic balanced net-

1.2. Plasticity models in high noise networks

works can self-sustain AI activity even in the absence of any external input (Vogels and Abbott, 2005; Kumar et al., 2008a). The initial stochasticity in such networks is frozen in the random and sparse connectivity matrix like a random seed. The resulting activity resembles stochastic firing and hence can be seen as an intrinsic source of noise (van Vreeswijk and Sompolinsky, 1996; Fusi, 2002). This means that balanced networks have the inherent capacity to generate stochasticity. Although the emerging background activity from both approaches might look very similar, treating them the same would probably be an oversimplification.

For clarity we will distinguish between simulations with “high noise”, meaning networks that receive substantial stochastic external input or have some form of stochasticity at the neuronal level. We will use the term “low noise” to describe networks with a high fidelity at the neuronal level and little external input. Since for the latter case the networks need to sustain healthy background activity, these are essentially large random balanced networks. When considering plasticity in recurrent networks this distinction is important. In the first case the background state is practically given. In the second case (low noise) one takes a more global viewpoint. Here plasticity has to maintain and ideally also achieve background activity.

For both cases it therefore interesting to ask how plasticity and spontaneous activity interact. Spontaneous activity in balanced networks alone is an ongoing field of study (Marre et al., 2009). Despite the fact that correlations in asynchronous irregular activity are low, they still prevail (Renart et al., 2010), and thus affect synaptic plasticity. The potential implications of this are unclear, because it could cause structure to arise spontaneously.

There are several options to distinguish between high and low noise cases. Most network simulations include deterministic neuron models, synapses and learning rules. Stochasticity enters only via external input and thus it is often sufficient to weight the amount of recurrent (and therefore deterministic) drive against the amount of external input (noise). In a scenario where postsynaptic potentials are all of roughly the same amplitude, a good criterion is the ratio between the number of spikes having their origin within the network and the number of “external spikes” – if the number of internal spikes is larger, the network can be categorized as strongly recurrent which translates into “low noise” in our terminology. Vice versa, for a smaller number of internal spikes the network is mostly driven externally and therefore categorized as “high noise”.

1.2 Plasticity models in high noise networks

High noise networks mostly take the role of a sub-network that is thought to be part of a larger network, modeled mostly as stochastic process, which provides the necessary input to maintain background activity. This makes

Chapter 1. Introduction

them useful to investigate plastic changes to the connectivity matrix without taking too much care about how this new connectivity affects the network dynamics in return. Networks of this kind are generally not balanced networks in the traditional sense as they do not even require inhibitory cells to function. Furthermore they do not necessarily have to be large either since findings from small networks can, assuming stability, easily be generalized to large networks. They can therefore be simulated efficiently over extended periods of time.

Plasticity studies for spiking networks focus primarily on spike timing dependent plasticity (STDP; Markram et al. (1997); Bi and Poo (1998); Sjöström et al. (2001); Caporale and Dan (2008); Sjöström et al. (2008)). STDP in recurrent networks was first studied extensively for connected pools of stochastic linear Poisson neurons (Gilson et al., 2009b,c,a, 2010). Other studies were performed on networks with more deterministic neuron models, driven by a continuous stream of external Poisson spike trains (Bush et al., 2010; Clopath et al., 2010). All these studies have in common that they make statements about the emergent weight structure, but not about the effect of the modified weights on the network dynamics.

Whenever dynamical effects of the plastic changes are to be studied, a strong form of recurrence is required. For small networks this is usually achieved by using postsynaptic potentials that are substantially larger than those observed in-vivo. Such an approach was used to point out a novel potential role of STDP in networks (Lubenov and Siapas, 2008). Here the authors show that Hebbian STDP in conjunction with significant connection delays can stably drive a recurrent network to the border between synchronous irregular and asynchronous irregular activity. In contrast to that, anti-Hebbian STDP always causes correlated and highly synchronous activity.

There has been another study with strong postsynaptic currents (Szatmáry and Izhikevich, 2010), where the authors present a working memory model based on so called polychronous groups (Izhikevich, 2006). The 1,000 cell model is able to learn and recall such groups on-line. However, the network model relies on STDP, a novel form of spike timing dependent short term plasticity and a switch-like hysteresis of the NMDA controlled postsynaptic current, whereas background activity is largely achieved through noise.

1.3 Plasticity in large networks

An average neuron in the brain receives between 1000 to 10000 synaptic contacts from other neurons. To approximate these numbers a realistic network simulation would need at least as many neurons. Moreover, since real neurons are not connected in an all-to-all manner, the neuron number would even have to be larger (in the order of hundreds of thousands

1.3. Plasticity in large networks

neurons). Simulating such large networks becomes technically difficult and computationally expensive. However, for plastic networks, a high computational cost is orthogonal to another important requirement, which is that plastic networks generally need to be simulated on the extended timescales of plasticity (up to days) to yield useful insights (Chapter 5).

Most simulation studies of plastic networks therefore compromise on the size of network models. They are generally large enough to gain the ability to exhibit chaotic AI activity (despite being very deterministic: low noise), but small enough to be computationally tractable. Although the boundary is quite fuzzy one can state that large networks begin at some 1,000 cells, while being only upper-bounded by computational constraints. The current state-of-the-art lies at around a few hundred thousand cells (Hoang et al., 2013; Ananthanarayanan et al., 2009; Helias et al., 2012) in modular networks which consist of modules of random networks. Each random network consists of simplified integrate-and-fire neurons which represent a crude approximation to real neurons with complicated morphologies. While other modeling attempts which use more realistic neuron models as well as experimentally observed synaptic connection probabilities are currently underway (Markram, 2006) it is still unclear in how far these models will incorporate synaptic plasticity.

According to above definition of large plastic network simulations, one of the first studies of plasticity in a recurrent network model was performed by Amit and Mongillo (2003b). In this study the authors demonstrate how a Hebbian form of spike-triggered plasticity (Fusi et al., 2000) leads to restructuring of synaptic weights which gives rise to multi-stability in the network dynamics. In such networks multiple Hebbian assemblies compete to activate at elevated firing rates. This mechanism is commonly thought to be the neural correlate of working memory. Amit and Brunel (1997a) had previously demonstrated this behavior in network models with static synapses (as opposed to plastic synapses), and shown that it relies on strong recurrent feedback.

Although Amit and Mongillo (2003b) successfully demonstrate that learning of such working memory modules is in principle possible, they also describe a range of difficulties that arise when bringing Hebbian plasticity into a recurrent network model (see Del Giudice et al. (2003) for a review).

The main problem of Hebbian forms of plasticity in models is that they are unstable, a fact already appreciated in simple feed-forward networks (Oja, 1982; Bienenstock et al., 1982; Miller and MacKay, 1994). In recurrent networks the effect is largely amplified because their global activity state typically has a strong dependence on a well tuned balance between excitation and inhibition (Brunel, 2000; Vogels et al., 2005). Even small changes to the overall excitation can disrupt this balance and dramatically change the global network activity. The resulting increase in spiking correlations and overall high activity typically lead to catastrophic memory loss (Fusi, 2002).

Chapter 1. Introduction

While Amit and Mongillo (2003b) used a spike-based plasticity rule which qualitatively captures rate based induction protocols and in some cases certain aspects of STDP (Fusi et al., 2000) the focus of later studies shifted towards models designed to quantitatively describe spike-timing-dependent plasticity (STDP; Song et al. (2000); Senn et al. (2001); Pfister and Gerstner (2006); Clopath et al. (2010); van Rossum et al. (2000)). STDP is an experimentally observed form of long-term plasticity which has first been observed in slice in-vitro preparations (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001; Caporale and Dan, 2008; Sjöström et al., 2008) and later also in-vivo experiments (Gambino and Holtmaat, 2012; Pawlak et al., 2013). The existing models of STDP can be coarsely classified in multiplicative models in multiplicative STDP, in which the weight update has an explicit multiplicative dependence on the current synaptic weight, and additive STDP models in which this is not the case.

1.3.1 Multiplicative STDP

Probably the first thorough study of multiplicative STDP in a large balanced network of $\sim 10^5$ neurons was undertaken by Morrison et al. (2007). In this particular network only approximately every fourth spike evoked synaptic potential experienced by a neuron in the network was caused by an external spike. We can therefore think of it as a strongly recurrent network with low noise. Here the authors approach the question whether spontaneous background activity alone can create structure in the synaptic weights and whether or not this structure is stable. The learning rule used in the study incorporates a multiplicative weight dependence, that fits experimental data. The authors found that the network shows stable asynchronous irregular activity on long timescales (hundreds of seconds) with a stable unimodal weight distribution. However, despite the fact that the weight distribution is found to be stable, single synapses show the opposite behavior. The authors study the lifespan of strong synapses (i.e. the time a strong synapse remains a strong synapse during ongoing activity) in their simulation and find that it is rather short (~ 200 s). Thus, synaptic weights wander around rather freely in the unimodal weight distribution and the network forgets quickly. It has been shown recently that the occurrence of unimodal weight distributions and the short memory retention time can be linked directly to multiplicative learning rules (Billings and van Rossum, 2009).

In the same publication Morrison et al. (2007) study the network's response to synchronous external stimulation of a sub-population within the network. They illustrate that although the network initially reacts very moderately to external stimulation, at a given point in time this behavior changes drastically. The whole network suddenly destabilizes to an epileptic state of activity. This study was later generalized by using a similar network and the same learning rule by Kunkel et al. (2011). Here the authors

1.3. Plasticity in large networks

study the effect of synchronous stimulation of sub-populations and whether it supports the development of syn-fire chains (Gewaltig et al., 2001b; Kumar et al., 2008a), embedded feed-forward structures which propagate a volley of synchronized spiking activity through a recurrent network. The authors show analytically and numerically that feed-forward structures cannot stably propagate under the STDP rule they use. Either the stimulated sub-population is too small and the synchronous stimulus does not propagate at all – or it is too large and synchronization grows and eventually recruits the whole network. This illustrates that although the network exhibits stable spontaneous activity, external stimulation can still irreversibly compromise stability.

1.3.2 Additive STDP

Additive learning rules do not share properties of short memory retention times and a unimodal weight distribution (Billings and van Rossum, 2009). However, thorough studies on additive STDP in large recurrent networks are rare. Being unaware of the current weight value, additive STDP is free to respond to whatever correlation structure it is exposed to. Therefore additive rules in particular are prone to develop weight drift (Gerstner and Kistler, 2002b) and run-away behavior.

One of the few studies of additive STDP in a large network was undertaken to explore synaptic pruning (Iglesias et al., 2005). The authors report a bimodal steady-state weight distribution with maxima at zero (pruned) and the maximum allowed synaptic value. This is the stereotypical signature of purely additive STDP (Billings and van Rossum, 2009). However, since the authors do not give details on the network dynamics it is hard to judge how the pruning actually proceeds.

One novel approach to additive STDP is taken by El Boustani et al. (2012). Here the authors present a new phenomenological meta-plastic STDP rule, capable of explaining a large body of experimental results. Similar to earlier plasticity models their model relies on rectifying plasticity thresholds (Clopath et al., 2008). However in this study these thresholds can move. In dependence of the synaptic history the threshold terms adapt and thus limit the amount of long-term potentiation (LTP) to the benefit of long-term depression (LTD) or vice versa. Eventually this can lead to a stable equilibrium where synapses are still plastic, but run-away potentiation is ruled out. The authors illustrate this stabilizing effect of their learning rule in a network of about 5000 cells. Finally they report long memory retention times, a desirable feature for memory and hallmark of additive STDP.

1.3.3 Inhibitory Plasticity

All models mentioned so far exclusively cover plasticity at the excitatory synapses. However, as shown in this thesis (Chapter 3; Vogels et al. (2011)), plasticity at inhibitory synapses provides a powerful mechanism in establishing excitatory and inhibitory current balance at the single neuron level in feed-forward circuits and recurrent networks. Here the authors use a symmetric Hebbian STDP rule, in which co-occurring pre and postsynaptic spikes cause LTP at inhibitory synapses, irrespectively of their temporal order. In addition there is a non-Hebbian component that depresses synapses if there is presynaptic spiking only. The authors show that the learning rule they propose establishes and maintains the balance of excitation and inhibition for a broad range of scenarios in networks consisting of up to 250,000 neurons. In particular, the learning rule is capable of self-tuning a recurrent network to the AI activity regime. Furthermore it is able to restore the balance when disrupted. This is demonstrated by a sudden substantial change of the excitatory weight matrix that creates two Hebbian cell assemblies. Hereupon the inhibitory plasticity rule quickly restores the balance and silences the memories, which can later be recalled by dint of a partial external stimulus.

1.4 Summary and Conclusion

While many biologically plausible plasticity models in feed-forward networks have been proven capable of explaining experimentally observed features in early sensory systems (Bienenstock et al., 1982; Clopath et al., 2010; Pfister and Gerstner, 2006), the deployment of plasticity rules to recurrent and in particular balanced networks has had limited success. Memory formation and the role of plasticity in more profound domains of cortical processing remains in large parts an enigma.

To stably maintain a memory trace in a network in which it has little overall effect on the network dynamics already poses a challenging problem because old memories are quickly overwritten by new ones (Fusi, 2002; Fusi et al., 2005). However, as soon as a single memory trace in a network becomes strong enough to be actively recalled through a change in the overall network dynamics Hebbian plasticity amplifies this memory further. This run-away effect generally overwrites other memories in the network. This problem can be seen as yet another plasticity-stability dilemma (Fusi, 2002), in which stability now refers to network-stability as opposed to the stability of the synaptic memory trace.

Some models – multiplicative STDP in particular – have had limited success regarding stability in very large balanced networks. However, they perform rather poorly at retaining synaptic structures once they are formed. On the other hand the question of memory retention could not even be

1.4. Summary and Conclusion

addressed in large networks with additive STDP since stability issues are generally devastating.

To understand how stable connectivity patterns can emerge in stable networks through synaptic plasticity is left open as one of the central questions which will be addressed in this thesis.

Chapter 1. Introduction

Chapter 2

Synaptic plasticity in spiking neural networks needs homeostasis with a fast rate detector

Authors: Friedemann Zenke, Guillaume Hennequin and Wulfram Gerstner

A version of this chapter has been published as Zenke et al. (2013). **Author contributions:** Conceived and designed the experiments: FZ GH WG. Performed the experiments: FZ. Analyzed the data: FZ. Contributed reagents/materials/analysis tools: FZ. Wrote the paper: FZ GH WG.

Abstract

Hebbian changes of excitatory synapses are driven by and further enhance correlations between pre- and postsynaptic activities. Hence, Hebbian plasticity forms a positive feedback loop that can lead to instability in simulated neural networks. To keep activity at healthy, low levels, plasticity must therefore incorporate homeostatic control mechanisms. We find in numerical simulations of recurrent networks with a realistic triplet-based spike-timing-dependent plasticity rule (triplet STDP) that homeostasis has to detect rate changes on a timescale of seconds to minutes to keep the activity stable. We confirm this result in a generic mean-field formulation of network activity and homeostatic plasticity. Our results strongly suggest the existence of a homeostatic regulatory mechanism that reacts to firing rate changes on the order of seconds to minutes.

Author summary

Learning and memory in the brain are thought to be mediated through Hebbian plasticity. When a group of neurons is repetitively active together, their connections get strengthened. This can cause co-activation even in the absence of the stimulus that triggered the change. To avoid run-away behavior it is important to prevent neurons from forming excessively strong connections. This is achieved by regulatory homeostatic mechanisms that constrain the overall activity. Here we study the stability of background activity in a recurrent network model with a plausible Hebbian learning rule and homeostasis. We find that the activity in our model is unstable unless homeostasis reacts to rate changes on a timescale of minutes or faster. Since this timescale is incompatible with most known forms of homeostasis, this implies the existence of a previously unknown, rapid homeostatic regulatory mechanism capable of either gating the rate of plasticity, or affecting synaptic efficacies otherwise on a short timescale.

2.1 Introduction

The awake cortex is constantly active, even in the absence of external inputs. This baseline activity, commonly referred to as the “background state”, is characterized by low synchrony at the population level and highly irregular firing of single neurons. While the direct implications of the background state are presently unknown, several neurological disorders such as Parkinson’s disease, epilepsy or schizophrenia have been linked to various disruptions thereof (Filion and Tremblay, 1991; Zhang and Kaltenbach, 1998; McCormick and Contreras, 2001; Spencer et al., 2003; Uhlhaas and Singer, 2006). Theoretically, the background state is currently understood as the asynchronous and irregular (AI) firing regime resulting from a dynamic balance of excitation and inhibition in recurrent neural networks (van Vreeswijk and Sompolinsky, 1996; Brunel, 2000; Vogels et al., 2005; Renart et al., 2010). Balanced networks exhibit low activity and small mean pairwise correlations (Brunel, 2000; Renart et al., 2010). However, even small changes in the amount of excitation can disrupt the background state (Brunel, 2000; Kumar et al., 2008b). Changes in excitation can arise from Hebbian plasticity of excitatory synapses: Subsets of jointly active neurons form strong connections with each other which is thought to be the neural substrate of memory (Hebb, 1949). However, Hebbian plasticity has the unwanted side effect of further increasing the excitatory synaptic drive into cells that are already active. The emergent positive feedback loop renders this form of plasticity unstable and makes it hard to reconcile with the stability of the background state (Morrison et al., 2007).

To stabilize neuronal activity, homeostatic control mechanisms have been proposed theoretically (von der Malsburg, 1973; Oja, 1982; Bienenstock et al., 1982; Miller and MacKay, 1994; Del Giudice et al., 2003; Lazar et al., 2009; Clopath et al., 2010) and various forms have indeed been found experimentally (Abraham and Bear, 1996; Turrigiano et al., 1998; Abraham, 2008). The term homeostasis comprises any compensatory mechanism that stabilizes neural firing rates in the face of plasticity induced changes. This includes compensatory changes in the overall synaptic drive (e.g. synaptic scaling (Turrigiano et al., 1998)), the neuronal excitability (intrinsic plasticity (Desai, 2003)) or changes to the plasticity rules themselves (i.e. metaplasticity (Abraham and Bear, 1996)). Common to all experimentally found homeostatic mechanisms is their relatively slow response compared to plasticity. While synaptic weights can change on the timescale of seconds to minutes (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001), noticeable changes caused by homeostasis generally take hours or even days (Turrigiano et al., 1998; Turrigiano, 1999; Turrigiano and Nelson, 2004; Watt and Desai, 2010). This is thought to be crucial since it allows neurons to detect their average firing rate by integrating over long times. While fluctuations on short timescales cause Hebbian learning and alter synapses in

Chapter 2. Plasticity needs rapid homeostasis

a specific way to store information, at longer timescales homeostasis causes non-specific changes to maintain stability (Desai, 2003). The required homeostatic rate detector acts as a low-pass filter and therefore induces a time lag between the rate estimate and the true value of neuronal activity. As a result, homeostatic responses based on this detector become inert to sudden changes. The longer the filter time constant is, the more sluggish the homeostatic response becomes.

Here we formalize the link between stability of network activity and the timescales involved in homeostasis in the presence of Hebbian plasticity. We first study the stability of the background state during long episodes of ongoing plasticity in direct numerical simulations of large balanced networks with a metaplastic triplet STDP rule (Pfister and Gerstner, 2006) in which the timescale of homeostasis is equal to the one of the rate detector. This allows us to determine the critical timescale beyond which stability is lost. In a second step we reduce the system to a generic two-dimensional mean-field model amenable to analytical considerations. Both the numerical and the analytical approach show that homeostasis has to react to rate changes on a timescale of seconds to minutes. We then show analytically and in simulations that these stability requirements are not specific to metaplastic triplet STDP, but generalize to the case of triplet STDP in conjunction with synaptic scaling.

In summary we show that the stability of the background state requires the ratio between the timescales of homeostasis and plasticity to be smaller than a critical value τ^{crit} which is determined by the network properties. For realistic network and plasticity parameters this requires the homeostatic timescale to be short, meaning that homeostasis has to react quickly to changes in the neuronal firing rate (on the order of seconds to minutes). Our results suggest that plasticity must either be gated rapidly by a third factor, or be accompanied by a yet unknown homeostatic control mechanism that reacts on a short timescale.

2.2 Results

In the following we first discuss our results obtained from simulating spiking neural networks in the balanced state with a Hebbian learning rule subject to a plausible learning rate. In the beginning we focus on a metaplastic mechanism that regulates the amount of synaptic long term depression (LTD) homeostatically. By systematically varying the time constant of the homeostatic rate detector, we find that stability of the background state requires homeostasis to act on a timescale of minutes. We then strive to understand the underlying mechanism of the instability from a generic mean field model, which we use to analytically confirm the critical time constant found in the spiking network simulations. Finally, to explore the generality of this mean

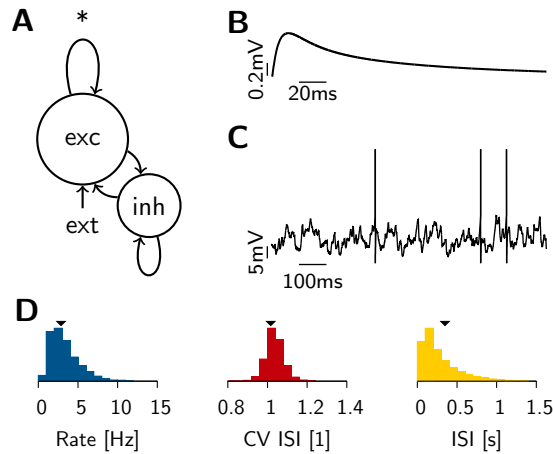


Figure 2.1 – The balanced network model. (A) Schematic of the network model. Recurrent synapses in the population of excitatory neurons (*) are subject to the homeostatic triplet STDP rule. (B) Typical magnitude and time course of a single excitatory postsynaptic potential from rest. (C) Membrane potential trace of a cell during background activity. (D) Histogram of single neuron firing rates (blue) and coefficient of variation (CV ISI, red) across neurons as well as the ISI distribution of all neurons (yellow) of the network during background activity. Arrowheads indicate mean values.

field approach, we apply the analysis to two variations of the triplet learning rule. First, we add a slow weight decay to metaplastic triplet STDP and second we switch from homeostatic metaplasticity to synaptic scaling in combination with triplet STDP. In both cases we confirm analytically and in simulations that a fast rate detector is required to assure stability.

2.2.1 Simulation results

To study the stability of the background state in balanced networks with plastic excitatory-to-excitatory (EE) synapses we simulate networks of 25000 randomly connected integrate-and-fire neurons (Fig. 2.1 (A)). Prior to any synaptic modification by plasticity, we set the network to the balanced state in which membrane potentials exhibit large sub-threshold fluctuations (Fig. 2.1 (C)), giving rise to irregular activity at low rates (≈ 3 Hz) and asynchronous firing at the population level (Fig. 2.1 (D)). In our model more than 90% of the input to each neuron comes from within the network, thus closely resembling conditions found in cortex (DeFelipe and Fariñas, 1992).

Plasticity of all recurrent EE synapses is modeled as an additive triplet STDP rule (see Pfister and Gerstner (2006) and Methods) which accurately describes experimental data from visual cortex (Sjöström et al., 2001; Pfister and Gerstner, 2006). In this metaplastic triplet STDP rule the amount of LTD is chosen such that LTP and LTD cancel on average, when the pre- and

Chapter 2. Plasticity needs rapid homeostasis

postsynaptic neurons fire with Poisson statistics at rate $\kappa = 3$ Hz. Therefore, under the assumption of low spike-spike correlations and irregular firing, κ becomes a fixed point of the network dynamics (see (Hennequin et al., 2010) and Methods). We begin with a fixed learning rate $\eta = 6.25$, which is chosen as a compromise between biological plausibility and computational feasibility (Methods). To go towards the fixed point, all neurons constantly estimate their firing rate as the moving average $\bar{\nu}$ with exponential decay constant τ , given by

$$\bar{\nu}_i(t) = \frac{1}{\tau} \sum_{k|t_i^k < t} \exp\left(-\frac{t - t_i^k}{\tau}\right) \quad (2.1)$$

where t_i^k corresponds to the k -th firing time of neuron i (see also Methods, Eq. (2.19)). If the rate estimate $\bar{\nu}_i$ of the postsynaptic neuron i lies above (below) κ , homeostasis increases (decreases) the LTD amplitude. The homeostatic time constant τ is the only free parameter of our model.

We then explore systematically for how a particular choice of τ affects the stability of the background state in the network. To allow the moving averages to settle, we run the network for an initial period of duration 3τ , during which synaptic updates are not carried out. After that, plasticity is switched on. To check whether the network dynamics remain stable, simulations are run for 24 h of biological time during which we constantly monitor the evolution of the population firing rate (Fig. 2.2 (A)). The network is considered unstable if the mean population firing rate either drops to zero or increases above 60 Hz which happens when run-away potentiation occurs (Fig. 2.2 (B)). By systematically varying the time constant τ in 1 s steps, we find that for the background state to remain stable (Fig. 2.2 (C)), τ must be shorter than some critical value $\tau^{\text{crit}} \approx 25$ s. Moreover, we find a sharp transition to instability when τ is increased beyond τ^{crit} . For $\tau < 3$ s the network has a tendency to fall silent (Fig. 2.2 (A), black line).

During stable simulation runs ($3 \text{ s} < \tau < 25 \text{ s}$), some synapses grow from their initial value w_0 up to the maximum allowed value w^{max} , while the rest of the synapses decay to zero. The resulting bimodal distribution of synaptic efficacies (Fig. 2.2 (F)) remains stable until the end of the run. This is a known phenomenon for purely additive learning rules (Toyoizumi et al., 2007; Billings and van Rossum, 2009) and we will see later that unimodal weight distributions arise by the inclusion of a weight decay or by choosing synaptic scaling as the homeostatic mechanism (van Rossum et al., 2000).

Despite the qualitative change in the weight distribution, the inter-spike-interval (ISI) distribution remains largely unaffected, while the coefficient of variation of the ISI distribution (CV ISI) is shifted to slightly higher values (Fig. 2.2 (D)). However, we noted that the single-neuron average firing rates, which are widely spread out initially, are at the end clustered slightly above the homeostatic target rate of ($\kappa = 3$ Hz) with a weak dependence




Figure 2.2 – Network stability during ongoing synaptic plasticity depends crucially on the homeostatic time constant. (A) Temporal evolution of the average firing rate in the excitatory population for different homeostatic time constants τ . Explosion of firing rate indicated by dashed lines. Curves for $\tau = 3$ s (dark blue), $\tau = 10$ s (light blue), and $\tau = 24$ s (turquoise) overlap on the interval from 2 h to 24 h indicating stability. With $\tau = 2$ s (black) we show one of the cases with very short τ where the activity spontaneously dies. (B) Spike raster of 200 randomly selected excitatory neurons. The last two seconds are shown before the network activity destabilizes ($\tau = 50$ s). (C) For $\tau = 20$ s, the activity stays asynchronous and irregular even after 24 h of simulated time. (D) Firing statistics in a stable network ($\tau = 15$ s) measured after 24 h of simulated time. Histogram of single neuron firing rates (blue) and coefficient of variation (CV ISI, red) across neurons and the ISI distribution of all neurons (yellow). Arrowheads indicate mean values. Black lines represent the corresponding statistics prior to any synaptic modifications (copied from Fig. 2.1). (E) Population firing rate for stable simulation runs at $t = 24$ h as a function of the homeostatic time constant. The dashed line indicates the target firing rate κ . (F) Evolution of the synaptic weight distribution during the first 8 hours of synaptic plasticity ($\tau = 15$ s).

Chapter 2. Plasticity needs rapid homeostasis

on the actual value of τ (Fig. 2.2 (E)). This behavior is characteristic for homeostatic firing rate control in single cells.

We conclude that metaplastic triplet STDP with a homeostatic mechanism as presented here can lead to stable dynamics in models of balanced networks exhibiting asynchronous irregular background activity. However, the timescale τ of the homeostatic mechanism critically determines stability. It has to be on the order of seconds to minutes and therefore comparable to the timescale of plasticity itself (here $\frac{\tau w}{\eta} = 476$ s). This finding is in contrast to most known homeostatic mechanisms that have experimentally been found to act on effective timescales of hours or days (Abraham and Bear, 1996; Turrigiano and Nelson, 2000; Turrigiano, 2008; Watt and Desai, 2010).

2.2.2 Mean field model

To understand why the critical time constant τ^{crit} above which homeostasis cannot control plasticity is so short, we here analyze the stability of the background state in a mean field model. In line with the spiking network model we consider a single population of neurons that fires with the mean population firing rate ν (Fig. 2.3 (A)). To find an analytic expression that characterizes the response of the background activity to changes in the recurrent weights w around the initial value w_0 , we begin with a linear neuron model

$$\nu = \Theta + \gamma x \quad (2.2)$$

with the offset Θ and the slope parameter γ . Since we are interested in weight changes around the initial value w_0 , the natural choice for x would be $\frac{w}{w_0}$. However, here we set $x = \frac{w}{w_0}\nu$ to take into account the recurrent feed-back. This choice makes γ dimensionless while Θ is measured in units of Hz. Because weights evolve slowly, while population dynamics are fast we can solve for ν and obtain the self-consistent solution

$$\nu = \frac{\Theta}{1 - \frac{\gamma w}{w_0}} \quad (2.3)$$

As we will show later, a better qualitative fit to the spiking model can be achieved with this heuristic, which will facilitate choosing the right parameters Θ and γ .

To introduce plasticity into the mean field model, we use the corresponding rate-based plasticity rule

$$\begin{aligned} \tau_w \frac{dw}{dt} &= \frac{\eta w_0}{\kappa^3} \nu_{\text{pre}} \nu_{\text{post}} (\nu_{\text{post}} - g_\kappa(\bar{\nu}_{\text{post}})) \\ &= \frac{\eta w_0}{\kappa^3} \nu^2 (\nu - g_\kappa(\bar{\nu})) \end{aligned} \quad (2.4)$$

which can be directly derived from the triplet STDP rule (Pfister and Gerstner, 2006) and also can be interpreted as a BCM model (Bienenstock et al.,




Figure 2.3 – Mean field theory predicts the stability of background activity. (A) Schematic of the mean field model. Plastic synapses are indicated by *. (B) Eigenvalues of the Jacobian evaluated at the non-trivial fixed point $\nu = \bar{\nu} = \kappa$. (C) Phase portrait for $\tau = 0.1\tau^{\text{crit}}$, a choice where background activity is stable. Nullclines are drawn in black. Arrows indicate the direction of the flow. Two prototypical trajectories starting close to Δ are shown. Blue line: Typical example of a solution that returns to the stable fixed point. Solutions starting in the shaded area, such as the red line, diverge to infinity. (D) The separatrix for four different values of τ . (E) Population firing rate of the spiking network model (simulations: red dots) for different values of weight w for connections from excitatory to excitatory neurons. Black line: Least-square fit of Eq. (2.3) on the interval $[0.98w_0, 1.02w_0]$ as indicated by the black bar. Extracted parameters are $\Theta = (0.163 \pm 0.002)$ Hz and $\gamma = (0.9476 \pm 0.0004)$ (cf. Eq. (2.3)).

Chapter 2. Plasticity needs rapid homeostasis

1982; Pfister and Gerstner, 2006; Gjorgjieva et al., 2011). Here, η is the relative learning rate and $\frac{w_0}{\kappa^3}$ sets the scale of the system. The second equality in Eq. (2.4) follows because in the recurrent model pre- and postsynaptic rates are the same ($\nu = \nu_{\text{pre}} = \nu_{\text{post}}$ and $\bar{\nu}_{\text{post}} = \bar{\nu}$). The function $g_\kappa(\bar{\nu}) = \frac{\bar{\nu}^2}{\kappa}$ scales the strength of LTD relative to LTP just as in the spiking case (cf. Methods, Eq. (2.18)). In the mean field model, the rate detector $\bar{\nu}$ (Eq. (2.1)) becomes the low pass filtered version of the population firing rate

$$\tau \frac{d\bar{\nu}}{dt} = \nu - \bar{\nu}. \quad (2.5)$$

To link the network dynamics with synaptic plasticity we take the derivative of Eq. (2.3), $\frac{d\nu}{dt} = (\nu^2 \frac{\gamma}{\Theta}) \frac{dw}{dt}$ and combine it with Eq. (2.4) to arrive at

$$\tau_w \frac{d\nu}{dt} = \frac{\eta}{\kappa^3} \frac{\gamma}{\Theta} \nu^4 (\nu - g_\kappa(\bar{\nu})) \quad (2.6)$$

which describes the temporal evolution of the mean firing rate as governed by synaptic plasticity. Taken together, equations (2.5) and (2.6) define a two-dimensional dynamical system with two fixed points. One lies at $\nu = \bar{\nu} = 0$ and represents the quiescent network. The remaining non-trivial fixed point is $\nu = \bar{\nu} = \kappa$, which we interpret as the network in its background state.

Given these choices, we now ask whether this fixed point can be linearly stable (Methods) and find that the stability of the background state requires

$$\tau < \tau^{\text{crit}} \equiv \frac{\Theta \tau_w}{\eta \gamma \kappa}. \quad (2.7)$$

For $\tau > \tau^{\text{crit}}$ infinitesimal excursions from the fixed point diverge, which corresponds to run-away potentiation in this model. We note that τ^{crit} crucially depends on the parameters $\Theta, \gamma, \tau_w, \eta$ and the target rate κ . However, we can rescale the system to natural units, by expressing firing rates in units of κ and time in units of τ^{crit} , and plot the eigenvalues as a function of τ (Fig. 2.3 (B)). The fact that the fixed point of background activity loses stability for too large values of τ is in good qualitative agreement with what we observe in the spiking model. One should further note that Eq. (2.7) is independent of the power of $\bar{\nu}$ appearing in $g_\kappa(\bar{\nu})$, as long as the fixed point of background activity exists ($\bar{\nu} > 1$) and under the condition that at criticality the imaginary parts of the eigenvalues are always non-vanishing (see Methods). This indicates the presence of oscillations which are indeed observed in the spiking network (cf. Fig. 2.2 (A), $\tau = 26$ s). The fact that the network falls silent for very small values of τ (e.g. $\tau = 2$ s in Fig. 2.2 (A)) is not captured by the mean field model.

We can make further use of the mean field model to qualitatively understand the behavior of the system far from equilibrium. Figure 2.3 (C) shows the phase plane of a network with a stable fixed point ($\tau = 0.1 \tau^{\text{crit}}$). When




Figure 2.4 – The mean field predictions agree with results from direct simulation of the spiking network. (A) Solid line: $\tau_{\text{crit}}(\eta)$ as a function of the learning rate η (cf. Eq. (2.7)), with simulation data (red points) for $\kappa = 3$ Hz. The arrow indicates the value used throughout the rest of this figure (the dotted line corresponds to the learning rate $\eta = 1$ as used in Figure A.1). (B) Same as before but as a function of κ for $\eta = \frac{1}{w_0} = 6.25$ fixed. (C) Lifetime values for the spiking network (red points) with a scaled step function as predicted by mean field theory ($\eta = \frac{1}{w_0} = 6.25$ and $\kappa = 3$ Hz). All error bars are smaller than the data points.

the system is driven away from it, and perturbations are small, the dynamics converge back towards the fixed point. However, when excursions become too large, the network activity diverges (compare Fig. 2.3 (C), dotted solution) since the fixed point of background activity is only locally stable. A numerical analysis shows that the basin of attraction is small when τ approaches τ^{crit} from below (Fig. 2.3 (D)). Hence the system is very sensitive to perturbations which easily lead to run-away potentiation. Although we expect the basin of attraction of the mean-field model and the spiking model only to be comparably where Eq. (2.3) describes the firing rates of the spiking network accurately we can assume that for robust stability $\tau \ll \tau^{\text{crit}}$ has to be satisfied.

2.2.3 Model comparison

To be able to make more quantitative predictions for the spiking network we have to choose values for the parameters on the right hand side of Eq. (2.7). These are the effective timescale of plasticity τ_w on the one hand, and Θ and γ , which characterize the network dynamics, on the other hand. We will now show that the latter can be determined from the static network model, which is independent of plasticity. Note that the parameters κ and η in our mean field model are shared with the spiking model which we will use to quantitatively compare the two.

First, we relate the variables Θ and γ to the response of the spiking network when all its EE synapses are modified. Since this is not feasible analytically, we extract the response numerically by systematically varying the EE weights around the initial state with $w_0 = 0.16$. While doing so, plasticity is disabled and we record the steady state population rate of the

Chapter 2. Plasticity needs rapid homeostasis

network (Fig. 2.3 (E)). We then minimize the mean square error for Eq. (2.3) over a small interval $[0.98w_0, 1.02w_0]$ and determine the following values: $\Theta = (0.163 \pm 0.002)$ Hz and $\gamma = (0.9476 \pm 0.0004)$. For the stability analysis only the derivative of Eq. (2.3) at w_0 matters. However, it is worth noting that the response of the balanced network is well captured by Eq. (2.3) over a much wider range than the one used for the fit. This behavior is an expected consequence of the balanced state, which is known to linearize network responses (van Vreeswijk and Sompolinsky, 1996; Van Vreeswijk and Sompolinsky, 1998). Our approximation by a linear rate model breaks down for higher rates since it does not incorporate refractory effects.

Second, under the assumption of independent and irregular firing in the background state, the plasticity time constant τ_w is fully determined by the target rate κ and known parameters of the triplet STDP model (see Methods and (Pfister and Gerstner, 2006)). For $\kappa = 3$ Hz we find $\tau_w = 2975$ s.

Using these results together with Eq. (2.7) we predict the critical timescale of homeostasis for different values of η and κ and compare it to the results that we obtain as before from direct simulations of the spiking network. Figure 2.4 (A) shows that the dependence of τ^{crit} on the learning rate η is remarkably well captured by the mean field model. The fourth power dependence on the background firing rate κ is described well for $3 \text{ Hz} < \kappa < 5 \text{ Hz}$ (Fig. 2.4 (B)), but the theory fails for smaller values, where we start to observe synchronous events in the population activity, which introduce correlations that are not taken into account in the mean field approach. In Fig. 2.4 (C) we plot the typical lifetimes (i.e. the time when the spiking simulations are stopped, because they either show run-away potentiation or the maximum simulated time $t = 24$ h is reached) as a function of τ . The figure illustrates nicely that the critical time constant τ^{crit} coincides with the sharp transition in lifetimes observed in the spiking network.

When running additional simulations with smaller learning rates ($\eta = 1$ as opposed to $\eta = \frac{1}{w_0} = 6.25$) we observe that the network destabilizes occasionally for values of τ smaller than τ^{crit} , but only after 22h of activity (see Figure A.1). We find, however, that this “late” instability can be avoided by either initializing the EE weights with a weight matrix obtained from a stable run ($\eta = 6.25$ at $t = 24$ h) or by reducing the maximally allowed synaptic weight ($w^{\text{max}} = 0.5$). Since these changes do not affect the “early” instability ($\tau > \tau^{\text{crit}}$), the “late” instability seems to have a different origin and might be linked to the spontaneous emergence of structure in the network.

Here we focus on the “early” instability which is seen in all simulations that do not respect the analytical criterion $\tau < \tau^{\text{crit}}$, after less than one hour of biological time, and therefore puts a severe stability constraint on τ . Moreover the theory is able to quantitatively confirm the timescale τ^{crit} emerging from the spiking network simulations and allows us to see the




Figure 2.5 – Slow synaptic weight decay renders weight distribution unimodal, but hardly affects global stability. (A) Evolution of the synaptic weight distribution over 8 h of background activity. (B) Synaptic weight distribution at $t = 8$ h. (C) Predictions for $\tau_{\text{decay}}^{\text{crit}}$ of mean field theory (solid line) and values obtained from direct simulation (points). (D) Final population firing rate as a function of η for values of τ where the background state is a stable fixed point (dashed line: target rate κ ; error bars: standard deviation over 100 bins of 1 s).

detailed parameter dependence. In particular for a background rate of 3 Hz and the learning rate $\eta = 1$ we find a critical timescale of $\tau^{\text{crit}} \approx 3$ min (simulations: (166.5 ± 0.5) s, mean field model: 170.6 s).

In summary, our mean field model discussed here makes accurate quantitative predictions about the stability of a large spiking network model with plastic synapses for a given timescale of homeostasis. Furthermore it gives useful insights into parameter dependencies which are computationally costly to obtain from parameter sweeps in simulations of spiking networks. Our theory confirms that metaplastic triplet STDP with biological learning rates has to be matched by a homeostatic mechanism that acts on a timescale of seconds to minutes. In the next sections we will show that the mean field framework described here can be readily extended to other forms of homeostasis.

2.2.4 Weight decay

The induction of synaptic plasticity is only a first step towards the formation of long-term memory. In the absence of neuromodulators necessary to consolidate early LTP into late LTP, these modifications have been found to decay away with a time constant of $\tau_d \approx 1$ h (Frey and Morris, 1997). To study the effect of a slow synaptic decay on the stability of the background state we focus on the early phase of plasticity. In particular we neglect consolidation in the model and introduce a slow decay term

$$\frac{dw(t)}{dt} = \underbrace{\frac{1}{\tau_w} \frac{\eta w_0}{\kappa^3} \nu^2 \left(\nu - \frac{\bar{v}^2}{\kappa} \right)}_{\text{Homeostatic triplet}} + \underbrace{\frac{\eta}{\tau_d} (w_0 - w(t))}_{\text{Decay term}} \quad (2.8)$$

where we already replaced the STDP rule by its equivalent rate based rule (see Pfister and Gerstner (2006) and Methods, Eq. (2.17)), while the effect of the decay term can be written identically in the rate based model and the STDP model. Note that for $\tau_d \rightarrow \infty$ we retrieve the model studied in Figs. 2.1-2.4. Again we determine the critical timescale of homeostasis in numerical simulations of the spiking network by systematically varying τ for different values of η . We further find that the slow weight decay causes the synaptic weights to stabilize in a unimodal distribution (Fig. 2.5 (A) and (B)) which is fundamentally different to what we observed for the decay-free case. However, the critical time constant of homeostasis τ_d^{crit} is only marginally larger than in the decay-free case (Fig. 2.5 (C)).

To assess the impact of the decay on the critical timescale, the mean field approach, as it was derived above, can be adapted to take into account the constant synaptic decay (Methods). Provided the decay time constant is sufficiently long, we find the critical time constant to be

$$\tau_d^{\text{crit}} = \left(\frac{1}{\tau^{\text{crit}}} - \frac{1}{\tau_d} \right)^{-1} \quad (2.9)$$

which is in good agreement with the results from direct simulations (Fig. 2.5 (C)). From Eq. (2.9) we can further confirm that the decay term only causes a small positive shift in the critical time constant as it was also observed in the spiking network. Furthermore, we see that the population firing rate settles to values closer to the actual target rate κ (Fig. 2.5 (D)) than this was the case in the decay-free scenario.

In summary, adding a slow synaptic weight decay to the plasticity model is sufficient to cause substantial change to the steady state weight distribution in the network. Nevertheless this slow process does not affect the need for a rapid homeostatic mechanism.

2.2.5 Synaptic scaling

To test whether the previous findings are limited to our particular choice of metaplastic homeostatic mechanism, or whether they are also meaningful in the case of synaptic scaling (Turrigiano et al., 1998) we now adapt the model by van Rossum et al. (2000) and combine it with triplet STDP

$$\frac{dw(t)}{dt} = \underbrace{\frac{1}{\tau_w} \frac{\eta w_0}{\kappa^3} \nu^2 (\nu - \kappa)}_{\text{Triplet term}} + \underbrace{\frac{1}{\tau_s} \frac{\eta}{\kappa} \left(\kappa - \left(\frac{\bar{v}^m}{\kappa^{m-1}} \right) \right)}_{\text{Scaling term}} w \quad (2.10)$$

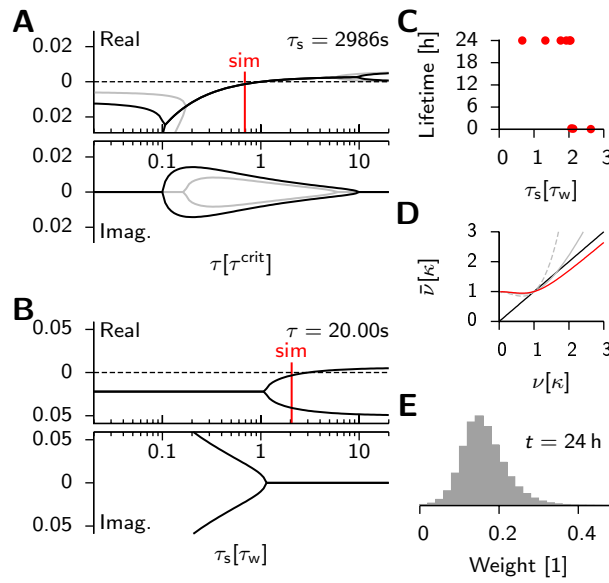


Figure 2.6 – Triplet STDP with synaptic scaling requires a fast rate detector. (A) Black line: Eigenvalues of the Jacobian ($\nu = \bar{\nu} = \kappa$) for different values of τ ($\tau_s = 2986$ s). Gray curve: Values from Fig. 2.3 (B) for reference. The red line (“sim”) indicates the critical value as obtained from simulating the full spiking network. (B) As before, but for different values of τ_s ($\tau = 20$ s). (C) Lifetimes of the background state in simulated networks of spiking neurons for different values of τ_s ($\tau = 20$ s). (D) Phase plane with nullclines. $\bar{\nu}$ -nullcline in black; ν -nullclines: dashed ($m = 1$), gray ($m = 2$) and red ($m = 3$). The latter was used in the rest of the figure. (E) Synaptic weight distribution after $t = 24$ h of simulation.

Chapter 2. Plasticity needs rapid homeostasis

where the rate of LTD is fixed in the triplet term (cf. Eq. (2.17)) and synaptic scaling is the only form of homeostasis. One important difference to the previous metaplastic STDP model is the addition of the scaling time constant τ_s which controls the timescale of synaptic scaling. In the metaplastic model we analyzed above this time constant is implicit, since it is the same as the one of plasticity (τ_w). In contrast to the original model of synaptic scaling ($m = 1$ (van Rossum et al., 2000)) here we choose $m = 3$ to avoid additional unstable fixed points in the phase plane (Fig. 5.3 (D)).

Bearing this in mind we move on to linearizing the system around the fixed point of background activity (Methods). We find that for $\tau_s \approx \tau_w$ the eigenvalues of the linearized system qualitatively have the same shape as for the plasticity rule with homeostatically modulated LTD (Fig. 5.3 (A)). In fact for sensible values of τ_s , the stability condition is exactly the same: $\tau < \tau^{\text{crit}}$ (cf. Eq. (2.7)). However, in the case of synaptic scaling Eq. (2.7) represents a necessary, but not a sufficient condition for stability. For too large values of τ_s stability is lost also in the case of $\tau < \tau^{\text{crit}}$ (Fig. 5.3 (B)). On the other hand decreasing τ_s indefinitely leads to oscillations without any further effect on stability (see Methods and (van Rossum et al., 2000)).

To compare these findings with the equivalent STDP rule we perform numerical simulations with the full spiking network in which we set $\eta = 1$ and choose τ_s on the order of τ_w ($\tau_s = 2986$ s). By changing τ systematically (Fig. 5.3 (C)) we determine the critical value to be smaller than predicted ($\approx 0.7 \tau^{\text{crit}}$), but within the same order of magnitude (Fig. 5.3 (A,C)). Conversely when we start with $\tau = 20$ s held fixed, we determine the critical value of τ_s to be on the same order as τ_w (Fig. 5.3 (B)). At the end of a stable simulation run ($t = 24$ h) we find that synaptic weights have formed a unimodal distribution (Fig. 5.3 (E)), an expected behavior of synaptic scaling (van Rossum et al., 2000).

In summary we have shown here that a fast rate detector is necessary to produce fast homeostatic responses to guarantee stable network dynamics also for the case of synaptic scaling. Although the quantitative agreement between the mean field model and the full spiking simulation is less accurate than in the case of for the metaplastic model above, both models confirm that the rate detector has to act on a timescale of seconds to minutes. Furthermore the time constant of the scaling term τ_s has to be comparable to the time scale of plasticity ($\tau_w = 2975$ s) or stability is compromised, when τ_s is chosen too large (and oscillations occur, when chosen too small).

2.3 Discussion

In this paper we have shown that a realistic additive triplet STDP rule (Pfister and Gerstner, 2006) can sustain a stable background state in balanced networks provided there is a homeostatic mechanism with a fast rate detector

that acts on a timescale of seconds to minutes. We confirmed this result in a generic two dimensional mean field model in which the stability of the background state is interpreted as the linear stability of a non-zero fixed point of the system for which the timescale of the homeostatic rate detector τ plays the role of a bifurcation parameter. These results are generic, i.e. independent of model details. In particular, we showed that similar results are obtained for triplet STDP with a form of metaplastic homeostasis, where homeostasis was implemented as a modulation of the LTD rate, or alternatively in combination with synaptic scaling. The mean field formalism produces accurate quantitative predictions for metaplastic triplet STDP. Although, in the case of triplet STDP in combination with synaptic scaling, the match of mean field model and direct simulations was less accurate, both support the notion that a fast rate detector is required for stability. For the case of synaptic scaling we found additionally that the homeostatic changes have to be implemented on a timescale comparable to the one of plasticity itself ($\tau_s \approx \tau_w \approx 1$ h), which is fast compared to most homeostatic mechanisms reported in the experimental literature, but consistent with earlier simulation studies that used fast homeostasis (von der Malsburg, 1973; Miller and MacKay, 1994; van Rossum et al., 2000; Del Giudice et al., 2003; Lazar et al., 2009; Clopath et al., 2010).

2.3.1 Homeostasis and plasticity

The fact that Hebbian learning has to be opposed by some kind of compensatory mechanism has long been known (von der Malsburg, 1973; Bienenstock et al., 1982; Oja, 1982; Miller and MacKay, 1994) and such mechanisms indeed have been found (Huang et al., 1992; Abraham and Bear, 1996; Turrigiano and Nelson, 2000). In the following we will briefly review the different types of homeostasis affecting synaptic weights and how they relate to what was used in the present study.

Homeostasis can be classified in two main categories. We call models “weight homeostasis” if they try to keep all afferent weights into a cell normalized (von der Malsburg, 1973). Such models have been criticized because they are non-local (Bienenstock et al., 1982), i.e. they require cell wide spatial averaging over synapses, which can only be achieved in a plausible way if all synaptic weights decay at a global rate modulated by the total afferent synaptic strength (Miller and MacKay, 1994). To avoid this, “rate homeostasis” models have been proposed (Bienenstock et al., 1982) which strive to maintain a certain postsynaptic firing rate. This approach, which we chose in the present study, has more experimental support (Turrigiano and Nelson, 2004; Watt and Desai, 2010). In contrast to the spatial filtering as described above, this mechanism requires temporal filtering of the postsynaptic rate over a given time window (represented by τ in this study). We can further distinguish between two principal types of homeostasis. A homeostatic

Chapter 2. Plasticity needs rapid homeostasis

mechanism can either act on the synaptic weights directly (e.g. synaptic scaling), or indirectly through metaplasticity (Abraham and Bear, 1996), by changing parameters of the plasticity model over time. The former, direct form of homeostasis allows for synaptic changes even in the absence of activity as it is seen in synaptic scaling experiments (Turrigiano et al., 1998) on a timescale of days. This is in contrast to theoretical models that apply scaling by algorithmically enforcing weight normalization (von der Malsburg, 1973; Lazar et al., 2009) on the timescale of one or a few simulation time-steps.

In our study we looked at both approaches. In the metaplastic triplet STDP model homeostasis manifest itself as a shift in the plasticity threshold between LTD and LTP (Pfister and Gerstner, 2006; Clopath et al., 2008, 2010; El Boustani et al., 2012). This is achieved by modulating the rate of LTD induction using the temporal average of the postsynaptic firing rates over a given time window (τ). As we have shown, this average has to follow the neuronal spiking activity very rapidly, meaning that plasticity parameters change on a short timescale, which is comparable to the duration of many standard STDP protocols (Sjöström et al., 2001). We therefore predict that if biological circuits rely on such a metaplastic homeostatic mechanism, weight changes are different for cells that are silent prior to a plasticity induction than for cells that have been primed by postsynaptic firing (over an extended period before the induction protocol). In Fig. 2.7 (A) we demonstrate this idea in the model of metaplastic triplet STDP ($\tau = 60$ s) for a typical LTD induction protocol (75 pairs at 5 Hz with -10 ms spike offset). Figure 2.7 (B) shows the relative differences between primed and unprimed experiments in dependence of the length of the priming duration or the priming frequency respectively. Since this plasticity rule implements homeostasis as an activity dependent change of the LTD learning rate, the amount of LTD changes dramatically while LTP is unaffected by priming. However, we expect that the main results of our mean field analysis also hold for cases in which LTP is affected, as long as the net synaptic weight change decreases with the intensity of priming. In either case the functional form of the dependence allows us to draw conclusions on the order of magnitude of τ and the exponent of \bar{v} appearing in $A_i^-(t)$ (cf. Eq. 2.18). Conversely, if homeostasis was exclusively mediated by synaptic scaling, we would expect that it manifests as a heterosynaptic effect. Its impact, however, would likely be smaller than in the case of metaplastic triplet STDP, because synaptic scaling does not have an explicit dependence on the presynaptic firing rate.

Since stability requires τ to be relatively short, it is also worth considering the extreme case where it is on the timescale of a few hundred milliseconds. In that case the learning rule can be interpreted as a quadruplet STDP rule combining a triplet term for LTP (e.g. post-pre-post) with a quadruplet term for LTD (e.g. post-post-post-pre). While such a choice of τ would make sense from a stability point of view, this behavior is not seen in experiments (Sjöström et al., 2001).

2.3.2 Influence of the model design

The timescales of synaptic plasticity and the time constants behind most homeostatic mechanisms reported in experiments are far apart. While plasticity can cause substantial synaptic changes in less than one minute (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001), homeostatic responses typically differ on the order of several magnitudes (hours or days) (Turrigiano, 2008; Watt and Desai, 2010). In this paper we have shown that even if homeostatic changes manifest relatively slowly they have to be controlled by a fast rate detector, else triplet STDP is incompatible with the low background activity observed in cortical circuits. We argue that this statement is likely not to be limited to our particular model, but rather applies to an entire family of existing plasticity models.

The basic building blocks of our study were a network model and a homeostatic plasticity rule. We used a generic balanced network model (Brunel, 2000; Compte et al., 2000; Brunel and Wang, 2001; Vogels and Abbott, 2005; Kumar et al., 2008b) to mimic brain-like spiking activity in a recurrent neural network. It is clear that the particular choice of network model does affect our results in a quantitative way and absolute predictions would require a more accurate and detailed network model. Nevertheless, we expect homeostasis to have similar timescale requirements in more detailed models as well. Indeed, as long as a strengthening of the excitatory synapses yields increased firing rates without a major change in the correlations, the qualitative predictions of the mean field model hold. However, our simulations were limited to roughly 1000 recurrent inputs per neuron, which is presumably less than what real cortical neurons receive (DeFelipe and Fariñas, 1992), so that excitatory run-away could build up even more rapidly in real networks than in our simulations.

The second building block of our model was the plasticity rule. Here we chose triplet STDP (Pfister and Gerstner, 2006) as a plasticity model that quantitatively captures a large body of experiments (Sjöström et al., 2001; Markram et al., 2012). One key feature of this model, which is seen across a range of in-vitro plasticity studies, is the fact that it yields LTP for high postsynaptic firing rates. The emergence of a critical timescale for homeostasis is mainly rooted in this fact and it is largely relaxed for pair-based STDP, be it additive or multiplicative (Morrison et al., 2007). However, such models do not capture experimental data as well as triplet STDP.

With the models we analyzed, namely the metaplastic triplet STDP and triplet STDP with synaptic scaling, we combined a realistic STDP learning rule with two quite different, but commonly used synaptic homeostatic mechanisms (Bienenstock et al., 1982; van Rossum et al., 2000; Pfister and Gerstner, 2006; Clopath et al., 2008; Lazar et al., 2009; Clopath et al., 2010; Tetzlaff et al., 2011; Gjorgjieva et al., 2011; El Boustani et al., 2012; Tetzlaff

Chapter 2. Plasticity needs rapid homeostasis

et al., 2012). The fact that we were able to show in both cases, either using a generic mean field model or numerical simulations of large balanced networks, that a fast rate detector is needed for stability, suggests that these results are quite general. The argument is further strengthened by the fact that existing computational models demonstrating stable background activity in plastic recurrent network models either use a form of multiplicative STDP which can be intrinsically stable (Morrison et al., 2007), but has poor memory retention (Morrison et al., 2007; Billings and van Rossum, 2009), or rely on a fast homeostatic mechanism (Lazar et al., 2009; El Boustani et al., 2012). In fact one of the first studies that illustrates stable learning in large recurrent networks combined with long memory retention times (El Boustani et al., 2012) is a model of metaplasticity built on top of the triplet model (Pfister and Gerstner, 2006). To describe effects observed in priming experiments (Huang et al., 1992; Christie and Abraham, 1992; Mockett et al., 2002), the authors introduce two floating plasticity thresholds that modulate the rate of LTP and LTD depending on the low-pass filtered neuronal activity. El Boustani et al. (2012) obtain the time constants behind these filters by fitting their model to experimental data. It is striking, and in agreement with what we report here, that the timescales they find are on the order of 1 s (El Boustani et al., 2012).

We conclude that current plasticity models that capture experimental data well require homeostasis to be able to react fast in order to maintain a stable background state. Likewise, if there is no rapid homeostatic control, most current plasticity models are probably missing a key ingredient to what makes cortical circuits stable.

2.3.3 Experimental evidence

The metaplastic triplet STDP rule we used makes use of an homeostatically modulated rate of LTD and can be mapped to a BCM-like learning rule (Pfister and Gerstner, 2006; Gjorgjieva et al., 2011). The BCM theory relies on a plasticity rule with a neuron wide sliding threshold (Bienenstock et al., 1982; Cooper et al., 2004). There seems to be some experimental ground for this idea (Wang and Wagner, 1999; Hulme et al., 2012) and it is intriguing, that the effects reported there are on the order of 30 min or less which points towards a relatively fast mechanism. We should further point out, that the arguments that led us to the critical timescale of homeostasis are not limited to a neuron wide sliding threshold. In fact the mean field equations for a global or local synaptic sliding threshold, or even one based on local dendritic compartments, are identical. Therefore the arguments we put forward also hold for the latter cases, which have experimental support through priming experiments (Huang et al., 1992; Christie and Abraham, 1992; Mockett et al., 2002). Priming experiments highlight changes in the induction of plasticity which depends on the synaptic activity over some

30 min.

With synaptic scaling we studied another possibility of introducing homeostasis into the triplet STDP model. Homeostatic scaling of synapses has good experimental support (Turrigiano et al., 1998; Turrigiano, 2008; Watt and Desai, 2010). Although it is generally associated with long timescales (order of days), also more rapid forms of scaling are known (Sutton and Schuman, 2006; Riegle and Meyer, 2007; Ibata et al., 2008) of which some indeed act on the order of minutes (Frank et al., 2006). Further modeling is required to test the ability of these rapid forms of homeostasis to guarantee stability in recurrent networks.

Finally one should note that the critical time scale of the rate detector strongly depends on the firing rates of the background state ($\tau^{\text{crit}} \sim \kappa^{-4}$, cf. Eq. (2.7) and Methods). The low firing rates reported experimentally (Burns and Webb, 1976; Koch and Fuster, 1989; Barth and Poulet, 2012) are therefore potentially necessary to guarantee the stability of the network. Conversely, cells or sub-networks with higher mean firing rates should have lower learning rates in order to be stable.

2.3.4 Limitations

Despite the mean field formalism being a drastic simplification of the original spiking model, the results we were able to derive from it were surprisingly accurate in the case of metaplastic triplet STDP and off by a factor of two in the case of triplet STDP with synaptic scaling. In all cases our mean field predictions overestimate the critical timescale obtained from simulations. This discrepancy has multiple potential reasons. First, in the mean field model we completely omit the existence of noise, fluctuations, and correlations. That these factors do play a role follows from the observation that the spiking network does not stabilize at the target rate κ , but at higher values (cf. Fig. 2.2 e). Although correlations in the AI state are small, they are on average positive (Renart et al., 2010). When we estimated τ_w we explicitly ignored correlations and required that LTD and LTP cancel at a firing rate κ . Adding correlations causes this cancellation to take place at slightly higher rates, which reduces the effective critical time constant. In the rate formulation of the STDP rule we make the simplifying assumption that the synaptic traces are perfect estimates of the postsynaptic firing rates. Indeed it can be shown that fluctuations that are present in the rates, bias the learning rule towards LTP (see Section A.2). Finally, any deviation of the population activity from its target value, initial or spontaneous, can be thought of as perturbations around the fixed point of background activity in the mean field model. This can compromise stability when the basin of attraction is small, as is the case when τ is close to criticality (Fig. 2.3 (D)). Again, such perturbations bias the critical value for the spiking network towards lower values. All the above points concern the simplifications made

Chapter 2. Plasticity needs rapid homeostasis

when going from the spiking model to the mean field model.

More importantly, the spiking model itself already represents a drastic simplification of the biological reality. For instance, we did not include neuronal firing rate adaptation or synaptic short-term plasticity (STP) in the present model. The timescales involved in firing rate adaptation are typically short (on the order of 100 ms) and their effect therefore negligible at the low firing rates of background activity (Benda and Herz, 2003; Brette and Gerstner, 2009). While the time constants behind STP can be longer than that, their stabilizing effect is somewhat less clear since they can be facilitating and depressing (Markram et al., 1998). Although we do not expect STP to have a strong impact on our main results, it would be an interesting avenue to verify this in future studies.

All our present studies were limited to spontaneous background activity. In a more realistic scenario we would expect the network to receive external input with spatio-temporal correlations. Such input will generally cause synaptic weights to change, which in the mean field model corresponds to a perturbation of the dynamical network state around the stable fixed point. If the perturbation leaves the system in the basin of attraction of background activity, equilibrium will be restored over time. If, however, the perturbation is strong, or perturbations are in rapid succession and start to pile up, the system loses stability once its dynamical state reaches the separatrix (cf. Fig. 2.3 (C,D)).

Another possibility worth mentioning is homeostatic regulation through inhibitory synaptic plasticity (ISP) (Lamsa et al., 2010; Woodin et al., 2003; Woodin and Maffei, 2010; Castillo et al., 2011; Kullmann et al., 2012; Vogels et al., 2013). Recent theoretical studies (Vogels et al., 2011; Luz and Shamir, 2012; Srinivasa and Jiang, 2013) suggest that ISP could produce an intrinsically stable feed-back system. Although we cannot exclude ISP as an important factor in network homeostasis, we have excluded it in the current study. It is likely that to stabilize Hebbian plasticity at excitatory synapses, ISP has to act on a comparable timescale (Sprekeler et al., 2012) and it will be interesting to integrate future experimental findings into a similar framework as presented here.

2.3.5 Conclusion

In summary, homeostatic mechanisms are necessary to stabilize the background activity in network models subject to Hebbian plasticity. Homeostasis needs to react faster than what is experimentally observed. This raises the important question of how the background activity in the brain can be stable. Our results suggest that the existence of a rapid homeostatic mechanism could be one possible answer. That, however, would require this mechanism to act on the same timescale as most STDP induction protocols. This then raises the question, why it has not been observed so far.




Figure 2.7 – Postsynaptic priming affects STDP protocols. Simulation of the metaplastic triplet STDP rule (Pfister and Gerstner, 2006). **(A)** Top: Typical protocol for the induction of LTD (75 pairs (post-pre) at 5 Hz with -10 ms spike offset) in the triplet STDP model ($\tau = 60$ s) with a postsynaptic cell which is quiescent prior to the LTD protocol (black) compared to induction after postsynaptic priming (blue). Top: Pre- and postsynaptic spikes for priming and pairing. Top, right: LTD induction. Middle: postsynaptic rate estimate $\bar{\nu}$ of the postsynaptic cell. Bottom: Weight change Δw over time. Postsynaptic priming period (duration 100 s): regular firing at $\kappa = 3$ Hz terminated by one second of silence (\star) to avoid triplet effects. **(B)** Relative differences in final weight change between quiet (Δw_q) and primed protocol (Δw_p) at the end of a LTD (gray) plasticity protocol. LTP protocol for reference (hollow, same pairing protocol, with reversed timing, +10 ms spike offset). Left: For different durations of the priming period and fixed priming frequency of 3 Hz. Right: Different priming frequencies with fixed priming duration of 60 s. The black line is a RMS fit to LTD data points of: (left) an exponential function; (right) of a quadratic function.

Chapter 2. Plasticity needs rapid homeostasis

Suitable plasticity protocols to detect such a mechanism should be similar to priming experiments (Huang et al., 1992; Christie and Abraham, 1992), but on the timescale of 1 min (Fig. 2.7). Another possibility would be, that the plasticity rate $\frac{\eta}{\tau_w}$ is not a constant after all, but subject to some neuromodulatory change (Pawlak et al., 2010). This could be possible, since it cannot be excluded that conditions in slice preparations, like the ones used to obtain the parameters of triplet STDP (Sjöström et al., 2001), are different from in-vivo conditions. Finally, also fast forms of ISP could play a role in network stability.

No matter whether through ISP or additional, hitherto unseen excitatory homeostatic effects, a variation of current models of homeostasis and plasticity seem inevitable, to achieve stability in plastic network models whilst making them biologically plausible.

2.4 Methods

To study stability in plastic spiking recurrent networks we simulated networks of 25000 integrate-and-fire neurons with conductance-based synapses (Fig. 2.1 (A)). The size of the network was chosen large enough to allow for an asynchronous irregular (AI) background state with low spiking correlations, but still small enough to enable simulations over long periods of biological time.

2.4.1 Neuron model

Table 2.1 – Neuron model and synaptic parameters

	Membrane	Threshold	Synapse
U^{exc}	0 mV	τ^{thr} 5 ms	τ^{ampa} 5 ms
U^{rest}	-70 mV	ϑ^{rest} -50 mV	τ^{gaba} 10 ms
U^{inh}	-80 mV	ϑ^{spike} 100 mV	τ^{nmda} 100 ms
τ^{m}	20 ms (10 ms*)		α 0.5

*) only inhibitory neurons

The networks we study consist of leaky integrate-and-fire neurons with a relative refractory mechanism connected by conductance-based synapses (Vogels and Abbott, 2005). The membrane voltage U_i of neuron i evolves according to

$$\begin{aligned}
 \tau^{\text{m}} \frac{dU_i}{dt} = & (U^{\text{rest}} - U_i) \\
 & + g_i^{\text{exc}}(t)(U^{\text{exc}} - U_i) \\
 & + g_i^{\text{inh}}(t)(U^{\text{inh}} - U_i)
 \end{aligned} \tag{2.11}$$

A spike is triggered when U_i crosses the spiking threshold ϑ_i . After a spike U_i is reset to U_i^{rest} and the threshold ϑ_i is increased $\vartheta_i \rightarrow \vartheta_i^{\text{spike}}$ to implement refractoriness. In the absence of spikes the threshold relaxes back to its resting value ϑ^{rest} according to

$$\tau^{\text{thr}} \frac{d\vartheta_i}{dt} = \vartheta^{\text{rest}} - \vartheta_i \quad (2.12)$$

with $\tau^{\text{thr}} = 5$ ms similar to Clopath et al. (2008). Inhibitory neurons were modeled identically except for a shorter membrane time constant τ^{m} . All relevant parameters are summarized in Table 2.1.

The spike train $S_j(t)$ of neuron j is defined as $S_j(t) = \sum_k \delta(t - t_j^k)$, where the sum runs over all k corresponding firing times t_j^k of neuron j . It affects the synaptic conductances of downstream neurons as

$$\frac{dg_i^{\text{inh}}}{dt} = -\frac{g_i^{\text{inh}}}{\tau^{\text{gaba}}} + \sum_{j \in \text{inh}} w_{ij} S_j(t) \quad (2.13)$$

if the index j corresponds to an inhibitory neuron or

$$\frac{dg_i^{\text{ampa}}}{dt} = -\frac{g_i^{\text{ampa}}}{\tau^{\text{ampa}}} + \sum_{j \in \text{exc}} w_{ij} S_j(t) \quad (2.14)$$

$$\frac{dg_i^{\text{nmda}}}{dt} = -\frac{g_i^{\text{nmda}}}{\tau^{\text{nmda}}} + g_i^{\text{ampa}} \quad (2.15)$$

in the case of an excitatory cell. Here w_{ij} is the weight of the synapse connecting neuron j with i ($w_{ij} = 0$ if the connection does not exist). Excitatory synapses contain a fast rising AMPA component with exponential decay and a slowly rising NMDA component with its respective exponential decay with time constant 100 ms. For simplicity we implemented the NMDA component as a low pass filtered version of the AMPA conductance (Eq. (2.15)). The complete excitatory postsynaptic potential (EPSP) is then given by a weighted sum of the AMPA and NMDA conductances

$$g_i^{\text{exc}}(t) = \alpha g_i^{\text{ampa}}(t) + (1 - \alpha) g_i^{\text{nmda}}(t) \quad (2.16)$$

With the chosen parameters (cf. Table 2.1), a typical EPSP has an amplitude of about 0.7 mV, as shown in Fig. 2.1 (B). For computational efficiency the voltage dependence of NMDA channels was omitted.

2.4.2 Network model

All units (20000 excitatory and 5000 inhibitory units, see Table 2.2 for details) are connected randomly with a sparse connectivity of 5%. Additionally each excitatory cell receives external input from a pool of 2500 independent Poisson processes firing at 2 Hz that are connected with 5% probability.

Chapter 2. Plasticity needs rapid homeostasis

Table 2.2 – Network model parameters

Neuron groups and connectivity		Synaptic weight structure	
Neural population	Size	Connection	Weight
Excitatory (E)	20000	E → E	$w^{\text{EE}} = w_0 = 0.16$
Inhibitory (I)	5000	E → I	$w^{\text{EI}} = w_0$
External Poisson (ext)	2500 at 2 Hz	I → E	$w^{\text{IE}} = 1.00$
Network connectivity	5%	I → I	$w^{\text{II}} = 1.00$
Connectivity from ext	5%	ext Poisson → E	$w^{\text{PE}} = w_0$

Table 2.3 – Plasticity model parameters

Plasticity window	A^+	6.5×10^{-3}
	τ^+	16.8 ms
	τ^-	33.7 ms
	τ^{slow}	114 ms
Initial weight	w_0	0.16
Weight limits	w^{min}	0
	w^{max}	1
Target firing rate	κ	3 Hz
Rel. learning rate	η	$\frac{1}{w_0} = 6.25^*$
	η	1 (Fig. A.1)

*) As used in Figs. 2.2 and 2.4 (B,C).

The relevant synaptic weight values are summarized in Table 2.2. Due to the high recurrence (on average 1000 out of 1125 connections are from within the network) the mean firing rate and network activity are sensitive to small changes in the recurrent synaptic strength. By appropriate choice of the excitatory weights ($w_0 = 0.16$) the network is initially tuned to the balanced state with AI activity at a mean population activity of approximately 3 Hz.

2.4.3 Plasticity model

We model synaptic plasticity after the triplet STDP model of Pfister and Gerstner (2006), using the minimal parameter set corresponding to in-vitro visual cortex data Sjöström et al. (2001). Plasticity only affects the EE recurrent connections. Weight updates Δw_{ij} act additively on the matrix elements w_{ij} and are given by

$$\begin{aligned} \frac{dw_{ij}}{dt} = & \eta w_0 A^+ z_j^+(t) z_i^{\text{slow}}(t - \epsilon) S_i(t) \\ & - \eta w_0 A_i^-(t) z_i^-(t) S_j(t) \end{aligned} \quad (2.17)$$

where ϵ is a small positive number and $z_n^+(t)$, $z_n^-(t)$ and $z_n^{\text{slow}}(t)$ are synaptic traces of neuron n defined as $\frac{dz_n^x}{dt} = -\frac{z_n^x}{\tau^x} + S_n(t)$ with associated time con-

starts τ^+ , τ^- and τ^{slow} respectively (see Table 2.3 and Pfister and Gerstner (2006)). Since the original triplet model describes relative synaptic changes, weight updates in Eq. (2.17) are scaled by the factor ηw_0 , where w_0 is the initial synaptic weight and η is an additional parameter that can be interpreted as a learning rate, or a conversion factor between the weight scales of the model and the true biological scale. In the model we approximate the biological scale by choosing plausible values for w_0 (cf. Fig. 2.1 (B)) and therefore expect η to be of the order of one. For a synapse with an initial weight of w_0 , a value of $\eta = 1$ corresponds to the learning rate that best fits visual cortex data (Pfister and Gerstner, 2006). However, since small values of η are computationally expensive we used $\eta = \frac{1}{w_0} = 6.25$ in Fig. 2.2 to ensure that a stable weight distribution can be observed within a day of simulated biological time (~ 4 d of computation time). Note that for $\eta = 1$ we would expect a comparable degree of convergence after 6.25 days of simulated time (roughly four weeks of computation). During ongoing plasticity the allowed weight values are limited to the interval $0 < w_{ij} < w^{\text{max}}$. Note that to avoid the creation of new synapses, connections that have zero weight initially, remain absent ($w_{ij} = 0$) throughout the entire simulation.

In simulations with metaplastic triplet STDP the amount of long term synaptic depression (LTD) $A_i^-(t)$ is varied homeostatically as a function of the moving average \bar{v}_i of the postsynaptic firing rate (Bienenstock et al., 1982; Pfister and Gerstner, 2006; Clopath et al., 2010; Gjorgjieva et al., 2011) with

$$A_i^-(t) = \frac{A^+ \tau^+ \tau^{\text{slow}}}{\tau^- \kappa} \bar{v}_i(t)^2 \quad (2.18)$$

This choice of $A_i^-(t)$ ensures that for uncorrelated Poisson firing at the rate κ LTP and LTD cancel on average. The moving average \bar{v}_i of the firing rate of neuron i is implemented as a low pass filtered version of its spike train

$$\tau \frac{d\bar{v}_i}{dt} = -\bar{v}_i + S_i(t) \quad \Leftrightarrow \quad \bar{v}_i = \frac{1}{\tau} \sum_{k|t_i^k < t} \exp\left(-\frac{t-t_i^k}{\tau}\right) \quad (2.19)$$

where τ is the timescale which controls of the temporal evolution of $A_i^-(t)$ (cf. Eq. (2.18)).

In simulations that require an additional slow weight decay of the weights we approximate this exponential decay, to avoid the costly operation of updating all weights after each time step, by periodically (period $\Delta \approx 10$ s) multiplying all weights by the factor $\exp(-\Delta/\tau_d) \approx 0.997$. Finally, simulations of synaptic scaling are performed using a fixed value $A_i^- = \frac{A^+ \tau^+ \tau^{\text{slow}}}{\tau^-} \kappa$. The scaling of the weights is approximated with the same approach as for weight decay. In such cases Δ is adapted appropriately according to the occurring scaling time constant τ_s .

2.4.4 The time constant of plasticity

We determine the timescale of plasticity in the mean field model by approximating τ_w from the plasticity parameters of the triplet STDP model (Pfister and Gerstner, 2006). To do so we consider the expectation value of the mean weight update averaged over many spike pairs, and we assume that pre- and postsynaptic firing is uncorrelated with stationary rates ν_j and ν_i respectively. The average relative weight change over time then reads

$$\begin{aligned}
 \left\langle \frac{dw_{ij}}{dt} \right\rangle &= \eta w_0 \left\langle A^+ z_j^+(t) z_i^{\text{slow}}(t - \epsilon) S_i(t) - A_i^-(t) z_i^-(t) S_j(t) \right\rangle \\
 &= \eta w_0 \left(A^+ \tau^+ \tau^{\text{slow}} \nu_j \nu_i^2 - \tau_- \nu_j \nu_i \langle A_i^-(t) \rangle \right) \\
 &= \frac{\eta w_0}{\kappa^3} \underbrace{A^+ \tau^+ \tau^{\text{slow}} \kappa^3}_{\equiv \frac{1}{\tau_w}} \underbrace{\nu_j \nu_i \left(\nu_i - \frac{\bar{\nu}_i(t)^2}{\kappa} \right)}_{\text{BCM like}} \quad (2.20)
 \end{aligned}$$

The resulting differential equation can be directly identified with Eq. (2.4) to obtain the effective time constant $\tau_w = \frac{1}{A^+ \tau^+ \tau^{\text{slow}} \kappa^3} \approx 2975$ s.

2.4.5 Numerical simulations

All differential equations were integrated using forward Euler integration with a 0.1 ms time step. Spiking simulations were written in C++ using Open MPI and the Boost libraries. The sources were compiled using the GNU C compiler. Simulations were run on 5 Linux workstations equipped with Intel(R) Core(TM)2 Duo E8400 CPUs and 24 GB of RAM each. It took approximately four and a half days to simulate one day of biological time.

Numerical results for the phase plane analysis, such as the position of the separatrix, were obtained by integrating the ODEs of the mean field model numerically using custom-written Python code.

2.4.6 Derivation of the stability condition in the mean field model

To analyze the stability of the fixed point of background activity ($\nu = \bar{\nu} = \kappa$) in the case of the metaplastic triplet STDP rule, we consider the Jacobian J of the two dimensional system (cf. Eqs. (2.5),(2.6)) in the general case of $g_\kappa(\bar{\nu}) = \frac{\bar{\nu}^n}{\kappa^{n-1}}$ for $n > 1$.

$$J = \begin{pmatrix} \Delta \left(5\nu^4 - 4\nu^3 \frac{\bar{\nu}^n}{\kappa^{n-1}} \right) & -n\Delta\nu^4 \left(\frac{\bar{\nu}^{n-1}}{\kappa^{n-1}} \right) \\ \frac{1}{\tau} & -\frac{1}{\tau} \end{pmatrix} \quad (2.21)$$

where we introduced the auxiliary variable $\Delta \equiv \frac{1}{\tau_w} \frac{\eta}{\kappa^3} \frac{\gamma}{\Theta}$. When evaluated at the fixed point J reduces to

$$J|_{\nu=\bar{\nu}=\kappa} = \begin{pmatrix} \Delta\kappa^4 & -n\Delta\kappa^4 \\ \frac{1}{\tau} & -\frac{1}{\tau} \end{pmatrix} \quad (2.22)$$

with characteristic polynomial

$$(\Delta\kappa^4 - \lambda) \left(-\frac{1}{\tau} - \lambda \right) + \frac{n\Delta\kappa^4}{\tau} = \lambda^2 - \lambda \left(\Delta\kappa^4 - \frac{1}{\tau} \right) + \frac{(n-1)\Delta\kappa^4}{\tau} \quad (2.23)$$

which determines the eigenvalues to be

$$\lambda_{1,2} = \frac{1}{2} \left(\Delta\kappa^4 - \frac{1}{\tau} \right) \pm \sqrt{\frac{1}{4} \left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2 - \frac{(n-1)\Delta\kappa^4}{\tau}} \quad (2.24)$$

of the linearized system at the fixed point of background activity. Stability of the fixed point requires all eigenvalues to have negative real parts (e.g. (Strogatz, 2001)). We now prove that the real part of both eigenvalues is negative if and only if $\tau < \frac{1}{\Delta\kappa^4}$. The square root in Eq. (2.24) is either purely imaginary, in which case $\tau < \frac{1}{\Delta\kappa^4}$ follows directly. For the case in which the square root is real we can express the larger of the two eigenvalues as

$$2\lambda_1 = \left(\Delta\kappa^4 - \frac{1}{\tau} \right) + \sqrt{\left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2 \left[1 - \frac{4(n-1)\Delta\kappa^4}{\tau \left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2} \right]} \quad (2.25)$$

$$= \left(\Delta\kappa^4 - \frac{1}{\tau} \right) + \left| \Delta\kappa^4 - \frac{1}{\tau} \right| \sqrt{c} \quad (2.26)$$

where we introduced the variable c for the term in the square brackets (Eq. (2.25)). If $\Delta\kappa^4 > \frac{1}{\tau}$ then $\lambda_1 = \frac{1}{2} \left(\Delta\kappa^4 - \frac{1}{\tau} \right) (1 + \sqrt{c}) > 0$ and the fixed point is unstable. If, however, $\Delta\kappa^4 < \frac{1}{\tau}$ then we know

$$\begin{aligned} 2\lambda_1 &= \left(\Delta\kappa^4 - \frac{1}{\tau} \right) + \sqrt{\left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2 - \frac{4(n-1)\Delta\kappa^4}{\tau}} \\ &< \left(\Delta\kappa^4 - \frac{1}{\tau} \right) + \sqrt{\left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2} \\ &= 0 \quad . \end{aligned} \quad (2.27)$$

Here, we used the fact that all occurring constants are positive, $n > 1$ and the argument in the square root is positive as well. Finally we can conclude the fixed point is stable if $\tau < \frac{1}{\Delta\kappa^4} = \frac{\Theta\tau_w}{\eta\gamma\kappa} \equiv \tau^{\text{crit}}$. This identifies τ^{crit} as an important limiting case for the stability of the fixed point. It is interesting to note that τ^{crit} is independent of n .

Chapter 2. Plasticity needs rapid homeostasis

Stability condition for weight decay

If we are to include an additional weight decay in the above model we replace Eq. (2.6) by

$$\tau_w \frac{dw}{dt} = \frac{\eta w_0}{\kappa^3} \nu^2 \left(\nu - \frac{\bar{\nu}^n}{\kappa^{n-1}} \right) + \frac{\eta}{\tau_d} (w_0 - w) \quad (2.28)$$

and proceed similarly as before by replacing all occurrences of w . In the decay term we can use the identities $w = \frac{w_0}{\gamma} \left(1 - \frac{\Theta}{\nu}\right)$ and since $\gamma = 1 - \frac{\Theta}{\kappa}$ (c.f. Eq. (2.3)) to rewrite

$$\frac{\eta}{\tau_d} \frac{\gamma}{\Theta} \nu^2 \left(1 - \frac{w}{w_0}\right) = \frac{\eta}{\tau_d} \frac{\nu}{\kappa} (\kappa - \nu) \quad . \quad (2.29)$$

We use this expression together with our results from Eq. (2.21) and the abbreviation $\Delta \equiv \frac{1}{\tau_w} \frac{\eta}{\kappa^3} \frac{\gamma}{\Theta}$, to arrive at

$$\frac{dv}{dt} = \Delta \nu^4 \left(\nu - \frac{\bar{\nu}^n}{\kappa^{n-1}} \right) + \frac{\eta}{\tau_d} \nu \left(1 - \frac{\nu}{\kappa}\right) \quad (2.30)$$

which leads to the following Jacobian at the fixed point

$$J|_{\nu=\bar{\nu}=\kappa} = \begin{pmatrix} \Delta \kappa^4 - \frac{\eta}{\tau_d} & -n \Delta \kappa^4 \\ \frac{1}{\tau} & -\frac{1}{\tau} \end{pmatrix} \quad . \quad (2.31)$$

The corresponding eigenvalues are given by

$$\lambda_{1/2} = \frac{1}{2} \left(\Delta \kappa^4 - \frac{\eta}{\tau_d} - \frac{1}{\tau} \right) \pm \sqrt{\frac{1}{4} \left(\Delta \kappa^4 - \frac{\eta}{\tau_d} - \frac{1}{\tau} \right)^2 - \left(\frac{\Delta (n-1) \kappa^4}{\tau} + \frac{\eta}{\tau \tau_d} \right)} \quad . \quad (2.32)$$

As we have seen earlier the stability is determined by the first term since the square root is purely imaginary around criticality. This leads us to the relaxed stability condition $\tau < \left(\Delta \kappa^4 - \frac{\eta}{\tau_d} \right)^{-1}$ and therefore with $\eta = 1$ and $\tau^{\text{crit}} = \frac{\Theta \tau_w}{\eta \gamma \kappa}$ we get $\tau < \left(\frac{1}{\tau^{\text{crit}}} - \frac{1}{\tau_d} \right)^{-1}$.

Stability condition for synaptic scaling

Here we will derive the critical time constant τ^{crit} for yet another variation of the triplet rule

$$\frac{dw}{dt} = \frac{1}{\tau_w} \frac{\eta w_0}{\kappa^3} \nu^2 (\nu - \kappa) + \frac{1}{\tau_s} \frac{\eta}{\kappa} \left(\kappa - \left(\frac{\bar{\nu}^m}{\kappa^{m-1}} \right) \right) w \quad (2.33)$$

which uses synaptic scaling to achieve the target rate κ (c.f. (van Rossum et al., 2000)). With the same transformations as before (i.e. $w = \frac{w_0}{\gamma} (1 - \frac{\Theta}{\nu})$) we can bring Eq. (2.33) to the form

$$\frac{d\nu}{dt} = \underbrace{\frac{1}{\tau_w} \frac{\eta\gamma}{\kappa^3 \Theta}}_{\equiv \Delta} \nu^4 (\nu - \kappa) + \underbrace{\frac{1}{\tau_s} \frac{\eta}{\Theta}}_{\equiv \Xi} \nu^2 \left(1 - \left(\frac{\bar{\nu}}{\kappa}\right)^m\right) \left(1 - \frac{\Theta}{\nu}\right) \quad (2.34)$$

which taken together with Eq. (2.5) yields the following Jacobian at the fixed point

$$J|_{\nu=\bar{\nu}=\kappa} = \begin{pmatrix} \Delta\kappa^4 & -\Xi(\kappa - \Theta)m \\ \frac{1}{\tau} & -\frac{1}{\tau} \end{pmatrix} \quad (2.35)$$

with associated eigenvalues

$$\lambda_{1/2} = \frac{1}{2} \left(\Delta\kappa^4 - \frac{1}{\tau} \right) \pm \sqrt{\frac{1}{4} \left(\Delta\kappa^4 - \frac{1}{\tau} \right)^2 - \frac{\Xi}{\tau} (\kappa - \Theta) m + \frac{\Delta\kappa^4}{\tau}} \quad (2.36)$$

We can appreciate directly from Eq. (2.36) that the real part of the largest eigenvalue is lower bounded $\text{Re}(\lambda_1) \geq \frac{1}{2} \left(\Delta\kappa^4 - \frac{1}{\tau} \right)$ and therefore we find that stability requires $\tau < \frac{1}{\Delta\kappa^4} = \tau^{\text{crit}}$, which is the same condition as above for the case of metaplastic triplet STDP. However, in the case of synaptic scaling this stability condition is necessary, but not sufficient. This we can see in Eq. (2.36) for given n , when Ξ becomes sufficiently small (τ_s sufficiently large) eventually we get $\text{Re}(\lambda_1) > 0$, where the background state loses stability (cf. Fig. 5.3 (B)). Hence, in addition to $\tau < \tau^{\text{crit}}$ there is also a critical value for τ_s which can be on a comparable scale like τ_w , but not arbitrarily large.

Acknowledgments

The authors would like to thank T.P. Vogels and H. Sprekeler for helpful discussions.

Chapter 2. Plasticity needs rapid homeostasis

Chapter 3

Inhibitory Plasticity Balances Excitation and Inhibition in Memory Networks


This chapter presents the author's (FZ) contributions to a manuscript published as Vogels et al. (2011) of which FZ is an author. In particular FZ contributed Figure 4 to the manuscript (Fig. 3.2 in the present thesis) and several Supplementary figures, all of which represent FZ's original work. That is, FZ performed all the relevant network simulations and data analysis. Moreover in the method section of this thesis modified parts of the text originally written by FZ in the Supplementary material of Vogels et al. (2011) were included. Tim P Vogels together with Wulfram Gerstner supervised the work. Section 3.2.1 constitutes unpublished work from a follow up study that was performed independently.

3.1 Introduction

Excitatory and inhibitory input currents received by cortical neurons are balanced and largely cancel each other (Shu et al., 2003; Wehr and Zador, 2003; Okun and Lampl, 2008; Froemke et al., 2007). This experimentally observed balance has given rise to numerous theoretical studies of balance in network models (Brunel, 2000; van Vreeswijk and Sompolinsky, 1996; Tsodyks and Sejnowski, 1995; Renart et al., 2010; Murphy and Miller, 2009; Vogels et al., 2005) which give an appealing theoretical explanation for high trial by trial variability of neural responses (van Vreeswijk and Sompolinsky, 1996).

Since their discovery balanced networks or derivatives thereof have served as a substrate for a range of theoretical models for rich dynamics (Murphy and Miller, 2009; Litwin-Kumar and Doiron, 2012), signal transmission (Vo-

Figure 3.1 – Schematic of the inhibitory STDP rule. Regions leading to long-term potentiation of the inhibitory synapse are marked in red. Net long-term depression is caused by presynaptic spikes only (blue). Adapted from Vogels et al. (2011).



gels and Abbott, 2005; Vogels, 2007; Gewaltig et al., 2001a; Kumar et al., 2010) and working memory (Roudi and Latham, 2007; Mongillo et al., 2008; Curti et al., 2004; Amit and Mongillo, 2003b; Compte et al., 2000; Yakovlev et al., 1998). Most existing models are built on top of the concept of random balanced networks, in which initially random connections are shaped manually and hand tuned to yield the desired network behavior. Only few studies have embarked on the question of how the underlying connectivity could arise through synaptic plasticity (Yakovlev et al. (1998); Amit and Mongillo (2003b); Clopath et al. (2010); see also Chapter 4).

One potential reason for this is that establishing biologically plausible activity in a balanced state in recurrent networks requires sensible parameter choices of synaptic connection weights which are easily compromised through current models of synaptic plasticity (Morrison et al., 2007; Brunel, 2000; Kunkel et al., 2011; Zenke et al., 2013). The resulting run-away potentiation leads to highly active and synchronous network states which are not observed in healthy biological systems. Past theoretical studies, however, have focused on plasticity at excitatory synapses and although experimentally observed (Woodin et al., 2003; Maffei et al., 2006; Woodin and Maffei, 2010; Vogels et al., 2013), plasticity of inhibitory synapses has not been targeted in theoretical studies before.

Two recent theoretical studies showed (Vogels et al., 2011; Luz and Shamir, 2012; Vogels et al., 2013) that Hebbian STDP rules at inhibitory synapses provides a simple explanation for a range of experimentally observed phenomena of balanced synaptic input to single cortical neurons (Froemke et al., 2007; Wehr and Zador, 2003). We were wondering if inhibitory STDP could also provide a potential answer as to how balance arises in recurrent network models. In particular we were interested if inhibitory STDP would establish a balanced state at the network level and maintain it even if excitatory weights change and cause a local disruption of this balance.

3.2 Results

To study the behavior of inhibitory STDP in recurrent network models we implemented a balanced network model consisting of 8,000 excitatory and





Figure 3.2 – Emergent global balance in networks and cell assemblies through inhibitory STDP. (**A-E**) Five snapshots of network activity at different stages in the simulation. Each snapshot shows (from top to bottom) 1) Instantaneous firing rates of all neurons laid out on a two dimensional grid. Squares mark designated sub-populations of neurons of which the spiking statistics were monitored during the simulation. 2) Spike raster over time of 50 neurons each from red and black sub-population. 3) Histograms of the coefficient of variation of the ISI distribution from red and black sub-population. 4) Histogram of pairwise spiking correlations of cells within the red and the black assembly. (**A**) Initial network state where excitatory neurons receive too weak inhibitory input. (**B**) Same as in (**A**) showing the network activity one hour later. (**C**) Network state as shown before but after introducing two overlapping Hebbian assemblies into the excitatory weight matrix (red and blue sub-populations). (**D**) Network state one hour after the two assemblies have been introduced. (**E**) Recall of one of the patterns by injecting additional external excitatory current into one fourth of red sub-population. See Section B.1 for an annotated protocol of the entire simulation.

Chapter 3. Inhibitory plasticity in recurrent networks

Figure 3.3 – Schematic of embedding a single cell assembly into a random network. A subset of existing connections between excitatory neurons (gray) is strengthened to form a Hebbian cell assembly (black).



2,000 inhibitory neurons with conductance based synapses (cf. Vogels and Abbott (2005) and Methods). Each neuron in the network received the same noise-free external input current which alone was sufficient to cause neurons to spike in the absence of other input. Inhibitory-to-excitatory synapses were endowed with a spike-timing-dependent plasticity rule in which synapses potentiate when they experience a pair of pre and a postsynaptic spikes in close temporal vicinity, but independent of their temporal order (Fig. 3.1). Presynaptic spikes alone on the other hand cause a depression of the synaptic weight. This rule can be written schematically as

$$\Delta w_{ij} = (\text{pre})_j \times ((\text{post})_i - \alpha) . \quad (3.1)$$

To test whether or not inhibitory STDP is compatible with a global background state at low activity we initially turned down all inhibition to the excitatory population. This caused excitatory neurons to fire synchronously at high rate (Fig. 3.2 (A)). Neural firing was highly regular with values of coefficient of variation of the inter-spike-interval distribution (CV ISI) substantially smaller than one (Fig. 3.2 (A)).

The network was then allowed to evolve freely while inhibitory STDP was active. After one hour of simulated time the network settled into a low activity and highly asynchronous state in which single cells fire highly irregularly (Fig. 3.2 (B); cf. Brunel (2000)) due to the highly anti-correlated excitatory and inhibitory input currents (Fig. B.5 and van Vreeswijk and Sompolinsky (1996)). The transition from synchronous regular activity to asynchronous irregular activity was robust to parameter changes and did not depend on fine tuning (see Appendix B.2) .

Excitatory synapses in biological neuronal networks are plastic (Bliss and Lømo, 1973; Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001) and change through experience. At the level of individual neurons synaptic plasticity can easily disrupt the balance between excitation and inhibition and cause regular firing as well as elevated firing rates. We were wondering if inhibitory STDP could stabilize the balanced state despite substantial changes to the excitatory connections. To simulate one possible outcome of such plastic changes we manually imprinted two overlapping Hebbian assemblies into the excitatory weight matrix by strengthening the existing connections between the neurons within the same assembly (Fig. 3.3). This change immediately caused assembly neurons to fire at elevated firing rates

with regular spike trains (CV ISI smaller than one) and with high synchrony (Fig. 3.2(C)). The rest of the network was only mildly affected by the change. The CV ISI of neurons outside one of the assemblies was elevated, while spiking correlations were comparably low with a mean around zero (Fig. 3.2(C)).

After one hour of ongoing inhibitory STDP the network settles again to a low activity asynchronous irregular background state similar to the one observed before the cell assemblies were introduced (Fig. 3.2(D)). Since excitatory connections are not plastic the assembly structure remains unchanged in the excitatory-to-excitatory weights, but their existence is not revealed from the network activity.

To serve as an associative memory it has to be possible to retrieve the information stored in a Hebbian cell assembly. We tested if it was possible to recall the memories stored in our network by externally stimulating a subset of neurons in one of the two assemblies. This partial cue of the previously stored memory lead to the activation of the other neurons within the same assembly (Fig. 3.2(E)). During recall cells not being directly stimulated externally, exhibited increased CV ISI values while correlations between cells inside the assembly were still low (Fig. 3.2(E)). After stimulation the activity inside the assembly dropped back to baseline almost immediately. Recall was not limited to one assembly, but the second assembly could be recalled similarly (Fig. B.6(F)).

We were wondering whether or not the formed assemblies could be activated at the same time. To achieve this we predominantly stimulated cells from the overlapping population of both cell assemblies. This caused either assembly to display elevated activity similar to what was observed for a single assembly (Fig. B.6(G)).

The stimulation of background neurons had a small effect on the overall network dynamics, but slightly decreased activity of cells taking part in the two cell assemblies (Fig. 3.4(A)). The stimulation of cells within either assembly, however, caused assembly wide elevated firing rates also in the unstimulated cells. The increase of firing rate within the unstimulated part of the assembly was nearly linear in the number of stimulated neurons (Fig. 3.4(B,C)). When stimulating the overlap of both assemblies the firing rate of both assemblies increased close to linearly with approximately half the slope compared to the case were non-overlapping cells were stimulated (Fig. 3.4(D)).

In all cases the brief stimulation of a subset of neurons taking part in a given cell assembly caused an activation of the remaining cells. Prolonged simulation over several seconds caused inhibitory plasticity to decrease firing rates inside of the respective assembly which resulted in a negative shadow pattern in the activity upon removal of the stimulation (not shown).

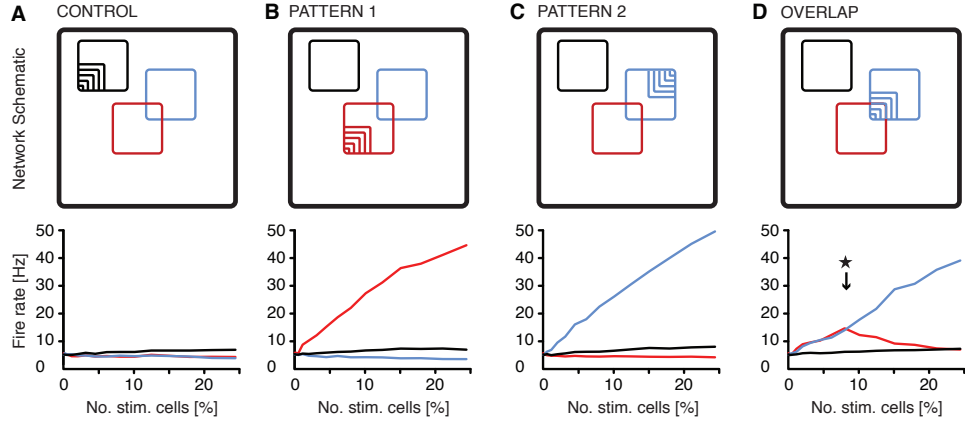


Figure 3.4 – Cells in the assembly respond approximately linearly to external stimulation. (A-D) Top: Schematic of stimulation paradigm. Bottom: Firing rate in neurons within the three sub-populations which do not receive additional external input for different numbers of neurons being stimulated. (A) Stimulation of cells inside the control population. (B) Stimulation of cells in the red assembly. (C) Stimulation of cells in the blue assembly. (D) Stimulation of cells in the overlap. At \star the stimulation has spill-over into the blue assembly which breaks the symmetry.

3.2.1 Attractor dynamics and inhibitory synaptic plasticity

In the simulations discussed so far the activity of a cell assembly dropped back immediately to baseline at the offset of an external partial stimulation of part of the assembly. It has been suggested in a series of theoretical works that cell assemblies could give rise to persistent activity (Hopfield, 1982; Amit and Brunel, 1997b,a; Gerstner and van Hemmen, 1992) and therefore form the neural substrate for working memory (Fuster and Jervey, 1982; Goldman-Rakic, 1995).

We were wondering if this behavior could be found in networks with stronger recurrent feedback. To do so we created a larger network model consisting of 250,000 integrate-and-fire neurons (200,000 of which are excitatory, see Methods). Similar to the protocol shown in Figure 3.2 we first allowed inhibitory STDP to drive the entire network to settle in a low activity background state. Again we introduced ad-hoc three different cell assemblies into the network by using the following simple modification rule

$$w_{ij}^{\text{exc}} \rightarrow w_{ij}^{\text{exc}} + \sum_{\mu} \gamma_i^{\mu} \Theta(\gamma_j^{\mu}) \quad (3.2)$$

in which Θ stands for the step function and γ_i^{μ} were the normalized pixel intensities from three flattened gray scale images of the Nobel laureates of Physiology or Medicine 1963: Hodgkin, Huxley and Eccles.

The resulting excitatory connectivity caused the network to initially inhibit synchronous activity (Fig. 3.5 (a)) of all three graded activity patterns





Figure 3.5 – Attractor dynamics of graded activity patterns in a large recurrent spiking network model with inhibitory STDP. Different snapshots of the activity of all excitatory cells laid out on a 2 dimensional grid. (a) Initial state. All three graded memory patterns previously wired into the excitatory weight matrix are active synchronously. Patterns from left to right: Huxley, Hodgkin and Eccles. (b) The same network after 5 minutes of simulated time. (c) Background state after burn-in period. (d) 4 snapshots of spontaneous recall of single patterns and background activity (bottom left). While the Hodgkin attractor is active, a faint shadow pattern of the the previously active Huxley pattern is visible.

Figure 3.6 – Schematic of the input-output curve with bistable behavior in a cell assembly with inhibitory plasticity (ISP) enabled. When activity in a cell assembly is too high – for instance due to the loss of stability of the background state – inhibitory plasticity drives the system back towards the region of bi-stability.



and the rest of the network. After a short period of ongoing inhibitory STDP all three patterns started to switch rapidly with sometimes more than one pattern being active (Fig. 3.5 (b)). This burn in period was followed by a period of background activity (Fig. 3.5 (c)) which is reminiscent of (Fig. 3.2 (D)). In contrast to our observations in the 10,000 cell network, however, patterns activated spontaneously without external stimulation (Fig. 3.5 (d)) and deactivated consequently after several seconds. Patterns previously active left behind a negative shadow in the network activity (Fig. 3.5 (d)). While this shadow persisted patterns were less likely to activate again. The transitions to activity were taking place spontaneously, but could also be induced by a transient partial external stimulation (not shown). However, activation of one of the attractors was in most cases unsuccessful immediately after it had been activated recently.

3.3 Discussion

In this chapter we illustrated how the balanced state can emerge naturally in networks in which inhibitory synapses obey a simple Hebbian STDP learning rule. Importantly balance does not only emerge initially but is also maintained in the face of perturbations as they could be caused through synaptic plasticity at excitatory weights. To that end we demonstrated that the ad-hoc addition of two overlapping Hebbian cell assemblies to the excitatory weights breaks the balanced state locally inside the assemblies and causes them to be permanently active. We have seen that inhibitory plasticity efficiently compensates for this broken balance and re-establishes the global AI state again in which the cell assemblies are hardly distinguishable from the network activity alone unless they are being recalled by dint of an external stimulation to part of the stored assembly.

While in a small network the externally activated assembly responds near linearly in stimulus strength and does not display delay activity, a larger network able to provide stronger recurrent feedback within a cell assembly

can give rise to multi-stable dynamics in which cell assemblies spontaneously activate in the absence of external stimulation. These dynamics were mostly autonomous, but could also be induced through external stimulation.

An intuitive explanation for this behavior comes from considering the input-output relation of a given cell assembly within the network in a mean-field framework (Amit and Brunel, 1997b). In this framework assembly dynamics can be seen as solving a set of self-consistent equations $x = F(\frac{x}{w} + \text{rest of the network})$ with the input-output non-linearity of the cell assembly F , the recurrent weight strength w and the effective output of the assembly neurons x and additional excitatory and plastic inhibitory input from the rest of the network. To yield bi-stable behavior the input output relation has to have multiple fixed points, i.e. the straight line has to have multiple intersection points with the non-linear function (Fig. 3.6). Satisfying this relationship requires fine-tuning (Amit and Brunel, 1997b). If the tuning is off only by a few percent and the recurrent feedback is too strong, the dynamics only have a single fixed point at high firing rates which causes the assembly to be permanently active.

In our simulations we chose the target rate ρ_0 of the inhibitory learning rule such that it can only be satisfied by the cell assembly in the bi-stable regime (Fig. 3.6). This ensures that inhibitory STDP fine-tunes cell assemblies into the vicinity of the bi-stability. If, however, the activity in the assembly is too high on average because 'attractor-switching' is not triggered frequently enough through external stimuli, inhibition into this pattern slowly increases until the upper stable fixed point is lost (Fig. 3.6). At this point the cells in the assembly drop to their low activity state (possibly leaving behind a shadow pattern; cf. 3.5 (d)) which now lies below the target rate ρ_0 of inhibitory plasticity. Consequently inhibition to the recently deactivated assembly is henceforth decreased until either the assembly is activated externally or eventually spontaneously when the lower fixed point is lost. Since LTD in our inhibitory STDP learning rule is independent of the postsynaptic rate (cf. Eq. (3.6)) the up-regulation is slower than the down-regulation (LTP is linear in the postsynaptic firing rate).

In summary the overall dynamics of each assembly seem to evolve on a slow limit cycle around the region of bi-stability. The mechanism gives an intuitive explanation why it is more difficult to reactivate recently deactivated assemblies and why assemblies that were inactive for long spontaneously activate.

While inhibitory STDP offers a step forward in terms of the amount of tuning required to create bistability in a balanced network, the mechanism does not offer a satisfying solution to memory formation and recall. From our daily life experience we know that some memories are not being recalled for long while others at times may be used frequently and in quick succession. However, if each memory assembly was to live on a slow limit cycle it would require each neuron to be in the active state of a cell assembly on average

Chapter 3. Inhibitory plasticity in recurrent networks

the same amount of time.

Another open question is how inhibitory plasticity interacts with ongoing excitatory plasticity. While inhibitory plasticity can act on a similar time scale as excitatory plasticity one could think of it as a rapid compensatory mechanism able to stabilize network dynamics (Zenke et al. (2013); Chapter 2). There is a catch, however. Inhibitory plasticity does not reverse LTP of excitatory synapses. This leads to a situations in which excitatory and inhibitory weights grow in equal proportion (and therefore maintaining an overall balance) until one of them hits an upper limiting bound or worse: a point at which the resulting correlations in sub-threshold input are strong enough to cause synchronization in the entire network at which point the asynchronous irregular balanced state is lost. There is currently a range of studies underway which are looking at this potential role of inhibitory plasticity.

Chapter 4 offers a different solution to the problem which involves heterosynaptic plasticity and the modulation of inhibitory plasticity by global secreted factors.

3.4 Methods

The network model we used is based on a network by Vogels and Abbott (2005). The model consists of 8000 excitatory and 2000 inhibitory integrate-and-fire neurons (Section 3.4.1) which are connected randomly with a connection probability of 2%. Following Vogels and Abbott (2005) the initial excitatory weights were chosen as $\bar{g} = 3\text{nS}$. Inhibitory-to-inhibitory weights were by a factor 10 stronger. To create assembly structure within the excitatory weight matrix we potentiated existing synaptic weights by multiplying their initial synaptic strength by $\chi = 5$ (unless mentioned otherwise). Inhibitory-to-excitatory synapses obeyed a symmetric STDP learning rule (Section 3.4.2). A complete summary of the model description and parameters can be found in tabular form in Appendix B.4.

3.4.1 Neuron Model

For all simulations we used a leaky integrate-and-fire model with conductance based synapses (Vogels and Abbott, 2005). The voltage V_i of neuron i evolves according to the following differential equation

$$\tau \frac{dV_i}{dt} = (V^{\text{rest}} - V_i) + (g_i^{\text{E}}(t)(V^{\text{E}} - V_i) + g_i^{\text{I}}(t)(V^{\text{I}} - V_i) + I_{\text{b}}) \times \frac{1}{g_{\text{leak}}}. \quad (3.3)$$

with the resting potential $V^{\text{rest}} = -60\text{mV}$, the excitatory reversal potential $V^{\text{E}} = 0$ and the inhibitory reversal potential $V^{\text{I}} = -80\text{mV}$. A spike is triggered when the voltage reaches a threshold $\vartheta = -50\text{mV}$, after which the

membrane potential remains clamped at V^{rest} during a refractory period of $\tau = 5\text{ms}$. We further fixed $g^{\text{leak}} = 10\text{nS}$ and assumed a membrane time constant of $\tau = 20\text{ms}$ (see Appendix B for a detailed tabular listing of all parameters).

When a neuron receives an excitatory presynaptic spike the conductance variable $g_i^{\text{E}}(t)$ makes a state transition $g_i^{\text{E}}(t) \rightarrow g_i^{\text{E}}(t) + \Delta g_{ij}^{\text{E}}$ or $g_i^{\text{I}}(t) \rightarrow g_i^{\text{I}}(t) + \Delta g_{ij}^{\text{I}}$ for an inhibitory spike respectively. In the absence of spikes both quantities follow an exponential decay

$$\tau_{\text{E}} \frac{dg_i^{\text{E}}}{dt} = -g_i^{\text{E}} \quad \text{and} \quad \tau_{\text{I}} \frac{dg_i^{\text{I}}}{dt} = -g_i^{\text{I}} \quad (3.4)$$

with synaptic time constants $\tau_{\text{E}} = 5\text{ms}$ and $\tau_{\text{I}} = 10\text{ms}$. The conductances Δg_{ij} are defined as $\Delta g_{ij} = \bar{g}W_{ij}$ where \bar{g} is the base weight which is fixed while W_{ij} is plastic for inhibitory-to-excitatory (see Section 3.4.2) or static for other connections. In addition to synaptic input all neurons received constant current input $I_{\text{b}} = 200\text{pA}$ which is sufficient to drive an isolated neuron without additional synaptic input to spike. All simulations were written in C++ and differential equations were integrated using the forward Euler method with a 0.1ms time step.

3.4.2 Inhibitory STDP model

We implemented a symmetric STDP rule in which co-occurring pre and postsynaptic spikes result in synaptic potentiation and presynaptic spikes alone cause depression (cf. Fig. 3.1). This learning rule was implemented using synaptic traces which are defined as

$$\frac{dz_i}{dt} = -\frac{z_i}{\tau_{\text{STDP}}} + S_i(t) \quad (3.5)$$

in which τ_{STDP} is the time constant which determines the timescale of STDP and $S_i(t)$ is the associated spike train $S_i(t) = \sum_{t'} \delta(t - t'_i)$ where the sum runs over all spike times t'_i of neuron i . The STDP learning rule acting on the variables W_{ij} can be written in the compact form

$$\frac{d}{dt}W_{ij} = +\eta(z_i - \alpha)S_j(t) + \eta z_j S_i(t) \quad (3.6)$$

where η is a learning rate and the constant $\alpha = 2\rho_0\tau_{\text{STDP}}$ is called the depression factor in which ρ_0 takes the role of a target firing rate. Allowed values for inhibitory weights W_{ij} were limited to the range $[W_{\text{min}}, W_{\text{max}}]$ (for parameter values see Table D.6).

3.4.3 Large network model

The large network model as it was used to generate Figure 3.5 was an up-scaled version of the network used before. In particular this network con-

Chapter 3. Inhibitory plasticity in recurrent networks

sisted of 200,000 excitatory and 50,000 inhibitory neurons. Most parameters and also the connection probability were unchanged at 2% which resulted in approximately 4,000 excitatory recurrent inputs to each network neuron. To compensate for the increased excitatory drive each neuron receives, we reduced the overall synaptic strength to $\bar{g} = 0.035\text{nS}$ (was $\bar{g} = 3\text{nS}$ in the 10,000 cell network). Moreover we used a lower learning rate of $\eta = 1 \times 10^{-5}$.

3.4.4 Data analysis and measuring spike-spike correlations

To characterize the global state we monitored correlations between individual spike trains as well as the population firing rate (the average of firing rates across the network), and the population rate's standard deviation σ_{Rate} , as well as average membrane potentials, and inter-spike-intervals (ISIs). Irregular asynchronous network activity which is often thought to be representative for cortical dynamics has a roughly constant population firing rate and often exhibits low spiking correlation values (Renart et al., 2010) while coefficients of variation of the inter-spike-intervals (CV ISI) are close to one.

To calculate the distributions of the pairwise spiking correlations (Renart et al., 2010) and the CV ISI from neurons within designated sub-populations we collected spiking data from all 392 neurons within each memory pattern and from the control group. Following Renart et al. (2010) we computed the spiking correlation coefficient X_{ij} between spike trains $S_i(t)$ and $S_j(t)$. We first constructed filtered spike trains F_i defined as

$$F_i(t) = S_i(t) * K(t), \quad (3.7)$$

in which the spike train $S_i = \sum_f \delta(t - t_i^f)$ is convolved with a symmetric bi-exponential kernel $K(t)$ (with $\int_{-\infty}^{\infty} K(t)dt = 0$) defined as

$$K(t) = \frac{1}{\tau_1} \exp\left(-\frac{|t|}{\tau_1}\right) - \frac{1}{\tau_2} \exp\left(-\frac{|t|}{\tau_2}\right) \quad (3.8)$$

with $\tau_1 = 50\text{ms}$ and $\tau_2 = 4 \times \tau_1$. The unnormalized covariance $V_{ij} = \sum_t F_i(t)F_j(t)$ over all discrete times t then leads to correlation coefficients

$$X_{ij} = \frac{V_{ij}}{\sqrt{V_{ii}V_{jj}}}. \quad (3.9)$$

We calculate $\sim 38,000$ pairwise correlation coefficients between the filtered spike trains of a given group (for all recall experiments we only stimulated a set of neurons that was disjunct to the neurons used for computing the spike train statistics.). All computation was done in discrete time with a resolution of $dt = 1\text{ms}$.

Chapter 4

Formation and recall of cell assemblies in recurrent networks of spiking neurons: multiple roles of plasticity

Authors: Friedemann Zenke, Everton Agnes and Wulfram Gerstner

At the time of writing of this thesis the material presented in this chapter is being prepared for submission to a peer-reviewed journal. Friedemann Zenke was the lead investigator in this study and responsible for all main areas. In particular he performed all simulations and analysis and created the figures. Wulfram Gerstner is supervising the project and wrote most of the introduction and parts of the results section. Everton Agnes participated initially in discussions on the learning rule.

4.1 Introduction

The concepts of cell assembly and Hebbian learning (Hebb, 1949), have inspired generations of experimental (Bliss and Lømo, 1973; Artola et al., 1990; Markram et al., 1997; Martin et al., 2000; Abbott and Nelson, 2000; Bi and Poo, 2001; Bliss et al., 2003; Caporale and Dan, 2008; Lisman, 2003; Sjöström et al., 2008) and theoretical work (Grossberg, 1969; von der Malsburg, 1973; Hopfield, 1982; Bienenstock et al., 1982; Song et al., 2000; Morrison et al., 2008; Gerstner et al., 1996; MacKay and Miller, 1990) (for a review see (Markram et al., 2012)). A cell assembly, loosely formulated as a group of neurons with strong connections amongst each other, can be interpreted as a functional circuit of brain activity. Delay activity of neurons during working memory tasks (Fuster and Jervey, 1982; Goldman-Rakic, 1995) or

Chapter 4. Memory formation and recall

recognition of abstract items (Quiroga et al., 2005; Bathellier et al., 2012) can be interpreted as the activation of such a cell assembly during memory recall.

While models of cell assemblies for fixed, preset, connectivity can be readily constructed (Hopfield, 1982; Gerstner and van Hemmen, 1992; Amit and Brunel, 1997a; Amit et al., 1985; Mongillo et al., 2008), the question of whether Hebbian learning rules can be used to *form* and *maintain* such assemblies in a stable fashion has, so far, not found a satisfying answer (Rochester et al., 1956; Amit and Fusi, 1994; Fusi, 2002; Mongillo et al., 2003; Fusi and Abbott, 2007; Amit and Mongillo, 2003b). The reasons for this failure of models to create functional memory assemblies in networks of spiking neurons with a single biologically plausible synaptic plasticity rule are manifold. First, neurons in the brain come in different varieties, and experimental forms of plasticity depend on the type of neuron and connection (Abbott and Nelson, 2000). Second, plasticity manifests itself in multiple forms including rate-dependent, (Bliss and Lømo, 1973), voltage-dependent (Artola et al., 1990), and spike-timing dependent (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001; Caporale and Dan, 2008; Sjöström et al., 2008) homosynaptic as well as heterosynaptic (Chistiakova et al., 2014; Lynch et al., 1977) plasticity. Third, induction of synaptic plasticity needs to be distinguished from processes of synaptic consolidation and maintenance (Lisman, 1985; Frey and Morris, 1998). Finally, nonstandard forms of plasticity such as structural plasticity (Stepanyants et al., 2002; Trachtenberg et al., 2002), short-term plasticity (Markram and Tsodyks, 1996; Abbott et al., 1997) or homeostatic synaptic changes (Turrigiano et al., 1998; Turrigiano and Nelson, 2000) complicate the picture.

Here we show that a combination of homosynaptic Hebbian plasticity with heterosynaptic and homeostatic processes gives rise to the stable formation of cell assemblies, and that these cell assemblies do not degrade or inflate during memory recall. In order to distinguish different forms of plasticity in our model, we use the following terms and criteria. First, we call contributions to synaptic plasticity that depend only on variables of the postsynaptic neuron, but not on those of the presynaptic neurons, 'heterosynaptic' whereas manifestations of synaptic plasticity that depend jointly on pre- and postsynaptic activity are called 'homosynaptic'; similarly, a change of the synapse which only depends on transmitter release, but not on the state of the postsynaptic neuron will be called 'transmitter-triggered'. By definition, heterosynaptic or transmitter-triggered plasticity is non-Hebbian. Second, our terminology takes the time scale into account on which synaptic changes manifest themselves. We refer to slow processes that show up on a time scale of tens of minutes or hours as 'homeostatic' in order to contrast them with 'induced' plasticity caused by typical protocols (lasting a few seconds to tens of seconds) for induction of long-term potentiation or depression. Third, a mathematical rule of synaptic plasticity is considered

as 'local' if it only depends on the activity of the presynaptic and the state of the postsynaptic neuron, but not on any other neuron. A local rule can be under the influence of a global factor (Crow, 1968; Izhikevich, 2007; Pawlak et al., 2010) such as the neuromodulator dopamine (Schultz et al., 1997) or other secreted factors (Turrigiano, 2012). In this terminology, our model of plasticity at excitatory synapses is local whereas that at inhibitory synapses is under the influence of a global factor. Note that a 'local' rule (e.g., (Oja, 1982; Chistiakova et al., 2014)) can give rise to 'heterosynaptic' effects.

4.2 Results

We simulated a network of 4096 excitatory and 1024 inhibitory randomly connected integrate-and-fire neurons containing a cell assembly of 400 excitatory neurons. The assembly is defined by intra-assembly synapses that are initialized at stronger values than those of the rest of the network (Fig. 4.1 (a,b)). In the absence of plasticity, the network functions as a working memory, exhibiting delay activity (Fig. 4.1 (c,d) consistent with earlier findings (Amit and Brunel, 1997a; Amit and Tsodyks, 1991; Gerstner and van Hemmen, 1992; Treves, 1993; Mongillo et al., 2008), but when we switch on a standard homosynaptic model of spike-time plasticity (Bienenstock et al., 1982; Pfister and Gerstner, 2006), the activity of neurons within the assembly, characterized by their firing rates, increases dramatically, followed by a slower increase of neuronal activity outside the assembly (Fig. 4.1 (e)).

The biologically unrealistic increase of firing rates in this, and similar (Song et al., 2000; Senn et al., 2001; Shouval et al., 2002), homosynaptic Hebbian models results from an interaction of the network dynamics with synaptic plasticity. The change of synaptic weight from a presynaptic neuron j to a postsynaptic neuron i in standard homosynaptic plasticity models such as the classic Bienenstock-Cooper-Munro rule (Bienenstock et al., 1982) or modern NNMDA-dependent (Shouval et al., 2002), spike-timing dependent (Senn et al., 2001; Pfister and Gerstner, 2006) or voltage-dependent (Clopath et al., 2010) variants requires that the activity $(\text{pre})_j$ of the presynaptic neuron is multiplied with the activation of some postsynaptic variables $(\text{post})_i$ and can be summarized as $\Delta w_{ij} \propto (\text{pre})_j \times (\text{post})_i \times F((\text{post})_i - \theta_i)$, with a function F that vanishes if $(\text{post})_i = \theta_i$, e.g. $F(x) = x$. It is a homosynaptic rule, because a synapse that is not presynaptically activated, $(\text{pre})_j = 0$ does not change. If presynaptic activity occurs $(\text{pre})_j > 0$, then it depends on the present state of the postsynaptic neuron whether the weight changes upward or downward. Even in the presence of presynaptic activity, the weight change is zero if the postsynaptic variable is zero $(\text{post})_i = 0$ or equal to the threshold θ_i . Activity values where the weight does not change are called fixed points of the synaptic dynamics.




Figure 4.1 – Classic synaptic learning rules fail to maintain stable cell assemblies in recurrent neural networks. (a) Schematic of a recurrent network with random sparse connectivity and prior to any experience (naïve; left) and with embedded cell assembly, potentially formed through experience dependent plasticity. (b) 2D histogram (40×40 binning) of the excitatory recurrent synaptic weights. Synaptic weights within the 400 cell assembly are preset to be stronger than other weights in the network (zoom). (c) Spike raster from a recurrent network model with a single embedded cell assembly. At * a 300ms external stimulation is applied to the cells within the assembly which triggers persistent activity. (d) Population rate of cells outside of the cell assembly (black) and inside the assembly (gray). (e) Same as in (d), but for a plastic network with triplet STDP. Activity inside the cell assembly (red) explodes rapidly and the activity of the other cells (black) follows more slowly. (f) 2D histogram of the synaptic weight matrix at the end of 100s of simulated time. Many presynaptic connections into the assembly have been potentiated. (g) Solutions of a nonlinear system of equations describing a stationary network state can be found graphically as the intersection between the diagonal (black) and the effective input-output-curve of neurons within the assembly (red). For the naïve network there exists only one stable solution at low firing rate (light red line). Experience dependent synaptic change can alter the network response such that three solutions exist (red line). Stable solutions are marked by dotted vertical lines. (h) Schematic of rate dependence of triplet STDP or BCM (blue). At low (high) postsynaptic activity synaptic long-term depression (potentiation) dominates. The plasticity rule has a fixed point at 10 Hz which is not congruent with the fixed points of the network dynamics (cf. g).

The synaptic dynamics interact with the neuronal dynamics of cell assemblies in a memory network. During memory recall, the assembly is strongly active while background neurons (i.e., those not participating in the assembly) show weak spontaneous activity (Fig. 4.1 (c)). In the presence of a homosynaptic plasticity rule with the above structure, the spontaneous activity of background neurons and the higher activity of the assembly neurons leads to an increase of all those synapses projecting into an assembly neuron i (Fig. 4.1 (f)). This finding can be understood in the framework of graphical network analysis (e.g. (Amit and Brunel, 1997b)) of working memory models. During memory recall, assembly neurons receive input from neurons of the same assembly. Input firing rates are transformed into output firing rates by an effective gain function, closely related, but not identical to, the neuronal f-I curve. Stable memory recall requires that the rates of input neurons (i.e. cells in the assembly) and output neurons (other cells in the same assembly) match (Fig. 4.1 (g)). Since memory recall should not change the contents of the memory, the synaptic dynamics should be at one of its fixed points discussed above. In general, there is a mismatch between the network dynamics and synaptic dynamics (Fig. 4.1 (h)). Matching the fixed points by a shift of the threshold θ (Bienenstock et al., 1982) works only if the shift is faster than the dynamics of induced synaptic plasticity (Zenke et al., 2013) and is therefore inconsistent with the notion of homeostasis as a slow adaptation towards a physiologically desired state.

Since homeostatic mechanisms are too slow, we added a simplified description of heterosynaptic and transmitter-triggered plasticity to our homosynaptic plasticity rule (see Methods). Similar to earlier models of heterosynaptic plasticity (Chen et al., 2013; Chistiakova et al., 2014), all synapses j onto neuron i are subject to change whenever the postsynaptic variable $(\text{post})_i$ reaches a high value. The direction of change depends on the present value w_{ij} of the synaptic weight which is consistent with experiments of tetanic burst induction ((Chen et al., 2013); cf. Fig. 4.2 (a)). Our model of transmitter-triggered plasticity is proportional to presynaptic activity $(\text{pre})_j$ and helps to stabilize neuronal firing rates (Kempster et al., 2001). Combination of all three terms leads to plasticity induction according to

$$\Delta w_{ij} \propto (\text{pre})_j \times (\text{post})_i \times F((\text{post})_i - \theta_i) - \beta (w_{ij} - \tilde{w}_{ij}) \times (\text{post})_i^4 + \delta (\text{pre})_j. \quad (4.1)$$

Both the reference weight \tilde{w}_{ij} and, in simulations with an additional homeostatic mechanism, also the threshold θ_i change on the time scale of slow homeostatic processes, while the other model parameters, in particular β and δ , are fixed (see Methods). The power of four in the heterosynaptic term is a convenient way to implement a non-linear threshold-like sensitivity to the postsynaptic activity. Analysis of the plasticity model shows that the stable points of the synaptic dynamics self-adjust to match that of the




Figure 4.2 – Heterosynaptic plasticity can create additional fixed points for the learning rule. (a) Bidirectional non-Hebbian plasticity in our model. Plasticity is induced by postsynaptic tetanization as described in Chen et al. (2013). At time $t = 0$ the postsynaptic neuron is forced to spike (3 burst trains at 1/60 Hz, 10 bursts per train at 1 Hz, 5 spikes per burst at 100 Hz). Tetanization increases some of the weights (i), while other weights decrease (ii) or remain unchanged (iii). (b) Schematic of input-output relationship of a cell assembly (top) and postsynaptic rate dependence of triplet STDP with heterosynaptic plasticity for three different values of the synaptic weight (bottom). Stable (unstable) fixed points are marked with dotted (dashed) vertical lines. (c) Mean firing rate over time of a single postsynaptic integrate-and-fire neuron with plastic synapses (triplet STDP and heterosynaptic plasticity). The postsynaptic neuron receives simultaneous Poisson input via an active pathway (100 synapses at 10 Hz) and a control pathway (100 synapses at 1 Hz, initial weight $w_0^{\text{ctl}} = 0.1$). Different colors signify different initial weights in the active pathway. (d) Temporal evolution of the average synaptic weight. Same color code as in (c). Active pathway (solid lines). Control pathway (dotted lines). (e) Mean firing rate of the same neuron as in (d) after 600s of simulated time for different values of the plasticity parameter β . Black (red) data points are from simulations with initially lower (higher) synaptic afferent weights. (f) Schematic of simulation of a single integrate-and-fire neuron receiving spatiotemporally correlated input. (g) Simulation results set-up described in (f). Top panel: Evolution of the synaptic weights. Middle panel: Spike raster of the input spike trains. Bottom panel: Mean firing rate of the postsynaptic neuron over time. After 100s the input spike trains follow a Gaussian profile with random mean that is shifted randomly at random intervals.

network dynamics (Fig. 4.2 (b) and Section 4.4.5). To confirm stability in simulations we first stimulated a single postsynaptic neuron with Poisson spike input from 100 presynaptic neurons at 10 Hz, while 100 other presynaptic neurons were firing at 1 Hz. Depending on the initial postsynaptic activity, the firing rates stabilize at two different values (Fig. 4.2 (c)). Importantly, synaptic weights do not saturate, but stabilize at different levels depending on their pathway and the postsynaptic firing rate. The bistability of postsynaptic activities can be found over a broad range of model parameters β (Fig. 4.2 (e)) or δ (data not shown). When stimulated with a localized stimulus which jumps on average once per second, the postsynaptic neuron develops a localized receptive field (Fig. 4.2 (f,g)), similar to earlier models (Bienenstock et al., 1982; Clopath et al., 2010).

The plasticity model of excitatory synapses was then implemented in the random network of excitatory and inhibitory neurons (Fig. 4.5 (a)) where it was combined with a model of inhibitory plasticity which was subject to a global modulating factor (Fig. 4.5 (b); see Methods). Each excitatory neuron received input from excitatory and inhibitory neurons in the network, but also from a small patch of sensory neurons which defines the spatial location of its receptive field (Fig. 4.5 (c)). All excitatory and inhibitory synapses were initiated close to the same reference value w_0 , but evolved thereafter freely according to the plasticity rules described above. The network was stimulated by applying repeatedly and stochastically one of four possible full-field input patterns (Fig. 4.5 (d,h)). Plasticity induced the development of spatially structured feature detectors within the receptive fields (Fig. 4.4 (b-d,f)) and, in parallel, to the development of strongly connected assemblies in the lateral excitatory connections (Fig. 4.4 (e)) which is reminiscent of recent experimental findings by Ko et al. (2013). However, in the present case recurrent connections grow strong enough, that two hours after the start of the stimulation paradigm, the network exhibits selective delay activity after a brief stimulation of one of the patterns (Fig. 4.4 (h)). Neurons that participate in the assembly exhibit a broad range of firing rates during delay activity (Fig. 4.4 (i)) while background neurons have a rate around 1 Hz or less. There are some neurons which respond to none of the patterns (Fig. 4.4 (b)), suggesting that there is a 'reserve' pool of neurons that could become sensitive to a novel pattern not included in the stimulation paradigm.

To check whether recall is associative, we stimulated the network with partial input by occluding to three quarters of the input field. In all cases, we found activation of the appropriate assembly corresponding to the partial information, indicating memory recall from partial cues (Fig. 4.5 (a,b)). Completely novel stimuli, unrelated to those previously encountered or a known pattern combined with partially wrong information could initiate memory recall of any pattern they shared overlap with (Fig. 4.5 (c)).

Once triggered the delay activity was stable over extended periods of




Figure 4.3 – Formation of cell assemblies and stable working memory states. (a) Schematic of the spiking network model. (b) Schematic of the input paradigm. (c) Schematic of the inhibitory STDP learning rule. (d) Top panel: Spike raster of the network activity at the beginning of a simulation. The raster shows 1024 excitatory neurons from the network. Colored bars on top indicate stimulation times and stimulus identity. Bottom panel: Population firing rate over time of neurons in assemblies coding for the respective stimuli as determined at the end of the simulation. (e) Histograms of the firing rate and the coefficient of variation of the inter-spike-intervals (CV ISI) of neurons in the network. (f) Same as (d) after about 5min of simulated time. (g) Same as (d,f) after about 30 min of simulated time. (h) Same as (d,f,g) after about 2h. Note the larger time scale and the increased mean stimulation interval $T_{\text{Off}} = 20\text{s}$. For clarity only every fourth spike is shown. Black range indicates data range used for spike statistics in (i). (i) Histograms of firing rates and CV ISI in the network during the interval marked in (h).

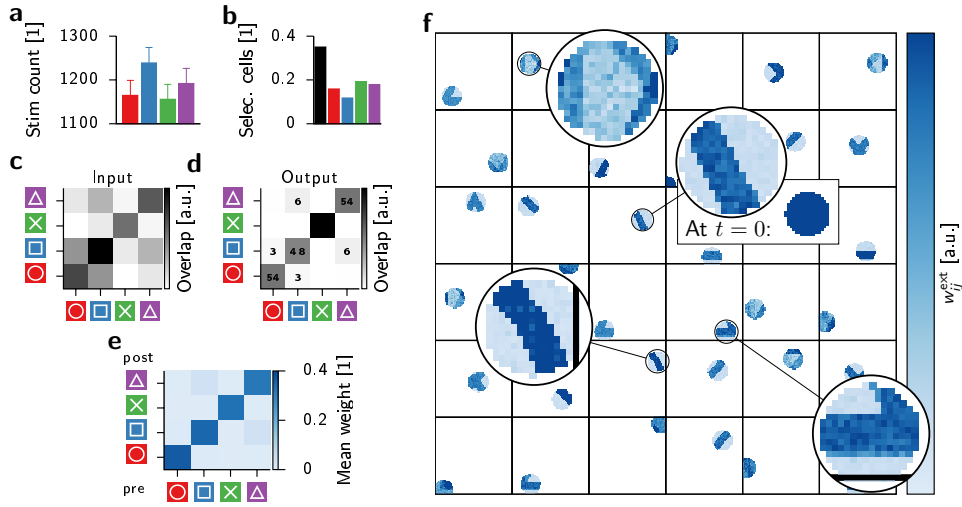


Figure 4.4 – Formation of cell assemblies and stable working memory states (continued). (a) Histogram of overall stimulation counts for the four different stimuli. Error bars: Square root of respective count corresponding to expected count variations for randomly chosen stimuli. (b) Bar plot illustrating the relative fraction of cells selective for a given stimulus. Black: No preference. (c) Matrix representation of the overlap between pairs of input patterns. (d) Same as (c), but for observed stimulus evoked network activity after learning. Numbers in selected fields are given in percent of the maximum value. (e) Mean weight strength of the recurrent weights between neurons with different stimulus preference ($t = 1h$). (f) Receptive fields with respect to external input of 36 randomly chosen network neurons after learning. Points represent existing connections and their position in the 2d input space. Color encodes the connection strength. Four zoom-ins on representative examples. Initial state for one neuron plotted for reference (box).




Figure 4.5 – Network states are associative and serve as working memory. (a) Population averaged peristimulus time histograms (PSTH) of the four relevant readout populations for all different distorted stimuli presented (keys on the right). (b) Top icons: Distorted cues fed into the network (black bars below code for time and duration of the stimulation). Middle panel: Spike raster of network activity (only every fourth spike of every fourth neuron shown for clarity). Bottom panel: Population rate of the four acquired network states. (c) Same as (b) but about one hour later in the simulation. The “F” is a new stimulus the network did not see during training.

time, although, spontaneous state transitions could occur occasionally. We were wondering if the addition of a slow adaption current could slowly destabilize the recall states over time. To that end we added a slow adaption mechanism on a timescale of 20s, similar to what has recently been reported for cortical neurons by Pozzorini et al. (2013), which caused rates inside an active cell assembly to change slowly over time and spontaneous state transitions to occur (Fig. C.7).

In the current framework the formation and recall of cell assemblies was dependent on the a sensible initial choice of afferent weights. If the weights were chosen too weak initially, not enough neurons reached stimulus evoked activity above the LTP plasticity threshold (Fig. C.5 (a)). Consequently no assembly structure was formed. In classical models of plasticity this problem is avoided by introducing a mechanisms of homeostatic metaplasticity (Abraham and Bear, 1996) which lowers the plasticity threshold when neuronal activity is low over an extended period of time (Bienenstock et al., 1982; Pfister and Gerstner, 2006). By adding a moving threshold, our model gained the ability to form cell assemblies even when the initial weights were weak (Fig. C.3) which showed similar recall behavior (Fig. C.4) to the case discussed above (cf. Fig. 4.5).

4.2.1 Effect of impaired individual forms of plasticity

We have argued that an existing assembly structure strong enough to support delay activity cannot be stable under triplet STDP alone (cf. Fig. 4.1 (c)) and proposed that heterosynaptic plasticity manifested as a postsynaptic burst detector can restore stability. However, this was so far only shown in a network with a preexisting assembly. We confirmed that indeed no stable assembly structure could be learned by deactivating heterosynaptic plasticity ($\beta = 0$) in our model (Fig. C.5 (b)).

To closer investigate the role of the other plasticity mechanisms, we specifically deactivated inhibitory synaptic plasticity in the model and repeated our training protocol (cf. Fig. 4.3). Doing so does not prevent the model to develop delay activity. However, the overall activity inside the assembly shifted to notably higher rates (Fig. C.6(a)).

We performed similar simulations in which transmitter triggered plasticity was deactivated ($\delta = 0$) and found that a large fraction of cells in the network falls silent and does not fire a single spike during one hour of simulated network activity (Fig. C.6(b)).

4.2.2 Role of consolidation dynamics

To study the role of consolidation dynamics (Section 4.4.2.3) we devised a simulation protocol which allows to conveniently follow the temporal evolution of external and internal connectivity in the network as it was used




Figure 4.6 – Block of consolidation causes cell assemblies to decouple from external input. (a,c,e) Control network with consolidation. (b,d,f) Network in which all consolidation dynamics are blocked. (a) Initial state of control network. Histogram of a 2000×2000 neuron section of input weight matrix (bin size 40 neurons) with preset block patterns (top left), same representation of recurrent weight matrix without initial structure (top middle). Letters a,b,c mark regions used to compute average weight (bar plot, bottom left). Spike raster of 2000 network neurons (top right), in which the colored bars on top indicate the stimulus identity. Activity in 3×400 neuron block patterns over time (bottom right). (b) Same as (a), but with consolidation dynamics blocked. Otherwise the network simulation is unchanged. (c,d) Same as before at later time. After 800s the recurrent connectivity in both networks has formed three cell assemblies and the network has developed working states which correlate with the last stimulus. (e,f) After 20min the stimulation interval is switched from $T_{Off} = 2s \rightarrow 20s$. (e) In the network with consolidation the input connections remain strong enough to trigger reliable state switching. (f) As external input weights fade away the network is “stuck” in a single recall state.

before. In parallel we considered an identical preparation in which the consolidation dynamics were blocked. In both cases we prepared the network in an initial state in which three blocks of 400 neurons each were connected through strong connections to equally sized blocks of neurons in the network (Fig. 4.6 (a,b)). The recurrent connectivity was initially unstructured. We then began to externally stimulate the three populations sequentially with elevated Poisson firing rates (Fig. 4.6 (a,b)). After 10 minutes both networks had formed cell assemblies and thus copied the input structure which allowed them to maintain delay activity even in the absence of an active stimulus (Fig. 4.6 (c,d)). There was no apparent difference in the dynamics of either network. After 20min of simulated time we increased the mean inter-stimulation interval ($T_{\text{Off}} = 2\text{s} \rightarrow 20\text{s}$) which causes both networks to spend a large fraction of time in the delay activity state. After 30 minutes we could appreciate distinct differences in dynamics of the two networks. While state switches are evoked reliably through external stimulus presentation in the control network (Fig. 4.6 (e)), the network in which consolidation is blocked does not respond to most external stimuli (Fig. 4.6 (f)) and delay activity is mostly decoupled from external input.

The effect is accompanied with a notable decrease of the input connection weights projecting into the respective assemblies ((Fig. 4.6 (d,f))), whereas this decrease does not occur in the control network (Fig. 4.6 (c,e)). Furthermore internal assembly weights in the network without consolidation are weaker than compared to the network with consolidation. However, this weakening did not impair the cell assemblies ability to display delay activity at around 20Hz (Fig. 4.6 (f)).

4.3 Discussion

Our model has multiple components of homosynaptic, heterosynaptic and transmitter-triggered plasticity as well as consolidation and, where this is mentioned explicitly, a slow form of homeostatic metaplasticity. Removing any one of these mechanisms destroyed the proper function as a memory module. Our results indicate that a variety of plasticity mechanisms encountered in the brain need to work together to allow proper memory function.

Integral part of our model is heterosynaptic plasticity to complement triplet STDP. Because heterosynaptic plasticity and triplet STDP act on the same timescale this can yield multi-stability in firing rates in networks with embedded cell assemblies, which cannot be achieved through slow homeostatic mechanisms alone (Zenke et al., 2013). To allow multiple stable solutions to co-exist (cf. Fig. 4.2 (b)) it is crucial that heterosynaptic plasticity acts as a burst detector. Here we chose the burst detector to appear with a fourth order postsynaptic dependence (Eq. (4.14)). While this is the lowest possible power to enable a second stable solution of the weight dynamics in a




Figure 4.7 – The effect of heterosynaptic plasticity is small in classic pairing protocols. **(a)** Relative weight change caused by a standard pairing protocol with causal (+10ms) and acausal (-10ms) spike timing at different pairing frequencies for triplet STDP with (red) and without (black) heterosynaptic plasticity. Protocol: 75 simulated pairings at +10ms and -10ms spike-timing. **(b)** Relative weight change under a plasticity protocol in which the pre and postsynaptic cell spike randomly with Poisson statistics with fixed mean firing rates. Classic triplet model (black) or with heterosynaptic plasticity (red). Error bars indicate the standard deviation from $n = 20$ randomly generated 10s Poisson spike trains for pre and postsynaptic cell.

cell assembly (cf. Fig. 4.2 (b)) we expect higher powers or a sharp threshold (Chen et al., 2013) to lead to similar behavior. Note that the effect of heterosynaptic plasticity is small in classical pairing experiments (Fig. 4.7 (a)), whereas it manifests itself clearly in experiments with postsynaptic high-frequency stimulation only (Chen et al., 2013; Chistiakova et al., 2014). Moreover our model indicates that heterosynaptic plasticity should become important during paired recordings with random spike-timing (Fig. 4.7 (b)).

In our model, similar to Clopath et al. (2008), the reference weight \tilde{w} follows w on a slower timescale than the induction of plasticity. During a certain time window (~ 20 -60min) induced changes can be reversed by uncorrelated postsynaptic activity. This is in qualitative agreement with experimental observations on the reversal of LTD and LTP (Zhou et al., 2003; Zhou and Poo, 2004; Abraham, 2003). During external stimulation network neurons become first selective to a certain stimulus through plastic changes of the external afferent connections. Consequently recurrent connections follow and form cell assemblies between neurons with the same preferred stimulus. Blocking the consolidation dynamics in the model causes the network to depotentiate the previously potentiated afferent connections once the network starts to exhibit delay activity (Fig. 4.6). This causes the assemblies to eventually decouple from external input. Thus, a network without synaptic consolidation can no longer recall memories from partial cues. Our findings suggest a potentially important computational role of consolidation.

We have defined homeostatic processes as mechanisms that act on timescales

of several minutes or longer to drive a system back into its normal physiological state. In our model three processes serve this role. First, transmitter-triggered plasticity, second, regulation of LTD amplitude, and third inhibitory synaptic plasticity. Functionally, transmitter-triggered plasticity ensures a baseline level of activity for individual cells and it opposes, in a push-pull fashion, low activity LTD induced through random spike pairings (Fig. C.6 (a)). Other homeostatic push-pull mechanisms such as synaptic scaling (Turrigiano et al., 1998) could serve the same function.

In our model inhibitory plasticity regulates the global network activity and thus constrains the number of active cells per cell assembly (Fig. C.6 (b)). While it is widely accepted that synaptic inhibition in neural networks controls the overall network activity, it is neither clear how inhibitory connections are set up nor how the strength of inhibition adapts in the presence of excitatory plasticity. Here we used a simple inhibitory STDP learning rule, which is reminiscent of Woodin et al. (2003), as a proxy for what could in fact be achieved by a range of different learning rules (Vogels et al., 2013). The global modulation in our model is justified by experiments showing that inhibitory plasticity is controlled via secreted factors and homeostatically targets global network activity rather than the activity of single neurons (Turrigiano, 2012). After stable delay activity is established in our network model, inhibitory plasticity can be turned off completely without notable impact on network function (not shown). In that sense inhibitory plasticity plays its role in establishing a stable working point.

Homeostatic regulation of LTD plays a similar role. When the initial connectivity is hand-tuned to ensure that a substantial fraction of network neurons are activated by a stimulus it is not necessary to include it into the model (cf. Fig. 4.4). However, if network neurons have no initial stimulus preference they fail to develop selectivity in the absence of some form of homeostatic regulation (Fig. C.5 (a)).

The emergence of cell assemblies through Hebbian synaptic plasticity has been studied in the past (Clopath et al., 2010; Amit and Mongillo, 2003a). Limited success has been made, as how to such assemblies can be learned and recalled without causing run-away behavior and erase the stored information in a setting where plasticity is always active. As an example, to achieve stable learning of cell assemblies Amit and Mongillo (2003a) had to manually adjust the strength of the external input to the stimulated cells to compensate for growing recurrent connections. In our model this problem does not arise for multiple reasons. First, input connections are plastic and can automatically adjust to compensate for increasing recurrent drive. Second, the combination of triplet STDP with heterosynaptic plasticity as it is presented here cannot lead to run-away dynamics as long as the number of active cells per assembly is limited by global inhibition. Third, to ensure that global inhibition and other parameters lie in the right parameter range, plasticity is complemented with a small number of homeostatic mechanisms

acting on longer timescales that drive the network to a physiologically desired operational regime.

4.4 Methods

To study the formation and recall of cell assemblies we simulated spiking neural network models with random sparse connectivity and multiple different forms of synaptic plasticity. The networks we studied consisted of 5120 integrate-and-fire neurons (4096 excitatory and 1024 inhibitory). In the following we describe the different elements of the model. For clarity we only quote the most relevant parameters in the text. A complete list of all parameter values is supplied in the Supplementary material (Section C.2).

4.4.1 Neuron model

Throughout our study we use leaky integrate-and-fire neurons with spike frequency adaptation (SFA) which receive conductance based synaptic input. The temporal evolution of the membrane voltage U_i of neuron i is described by

$$\begin{aligned} \tau^m \frac{dU_i}{dt} = & (U^{\text{rest}} - U_i) \\ & + g_i^{\text{exc}}(t)(U^{\text{exc}} - U_i) \\ & + g_i^{\text{inh}}(t)(U^{\text{inh}} - U_i) \end{aligned} \quad (4.2)$$

where the inhibitory input g_i^{inh} is defined as the dimensionless quantity

$$g_i^{\text{inh}}(t) = g_i^{\text{gaba}}(t) + g_i^{\text{a}}(t) \quad (4.3)$$

which is the sum of inhibitory synaptic input $g_i^{\text{gaba}}(t)$ and a contribution from spike triggered adaptation $g_i^{\text{a}}(t)$. Both quantities evolve according to their own dynamics

$$\frac{dg_i^{\text{gaba}}}{dt} = -\frac{g_i^{\text{gaba}}}{\tau^{\text{gaba}}} + \sum_{j \in \text{inh}} w_{ij} S_j(t) \quad (4.4)$$

$$\frac{dg_i^{\text{a}}}{dt} = -\frac{g_i^{\text{a}}}{\tau^{\text{a}}} + \Delta^{\text{a}} S_i(t) \quad (4.5)$$

and their states jump at the arrival of presynaptic spikes from upstream inhibitory neurons $S_j(t) = \sum_k \delta(t - t_j^k)$ or the occurrence of postsynaptic action potentials $S_i(t)$.

Where this is mentioned explicitly we add a second adaptation conductance to g_i^{inh} which is described by the same temporal evolution as shown in Eq. (4.5) but using different values for Δg^{adapt} and τ^{adapt} . The resulting

long lasting adaptation effect allows us to mimic cellular behavior recently characterized by Pozzorini et al. (2013).

Depolarizing current in Eq. (4.2) results from excitatory synaptic input

$$g_i^{\text{exc}}(t) = \alpha g_i^{\text{ampa}}(t) + (1 - \alpha) g_i^{\text{nmda}}(t) \quad (4.6)$$

with a fast AMPA-like component $g_i^{\text{ampa}}(t)$ and a slowly rising and decaying NMDA-like component $g_i^{\text{nmda}}(t)$. Their temporal evolution is described by the following set of equations

$$\frac{dg_i^{\text{ampa}}}{dt} = -\frac{g_i^{\text{ampa}}}{\tau^{\text{ampa}}} + \sum_{j \in \text{exc}} w_{ij} u_j(t) x_j(t) S_j(t) \quad (4.7)$$

$$\tau^{\text{nmda}} \frac{dg_i^{\text{nmda}}}{dt} = -g_i^{\text{nmda}} + g_i^{\text{ampa}} \quad (4.8)$$

where the quantities $u_j(t)$ and $x_j(t)$ appearing with the synaptic weight w_{ij} are describe the evolution of short term plasticity (Section 4.4.2.1).

An action potential is triggered when the membrane voltage of neuron i rises above the threshold value ϑ_i . Following a spike the voltage U_i is reset to U_i^{rest} . At the same time the threshold ϑ_i is transiently increased $\vartheta_i \rightarrow \vartheta_i^{\text{spike}}$ to implement an absolute and relative refractory period. In the absence of spikes the dynamic threshold ϑ_i relaxes quickly to its resting state $\vartheta_i^{\text{rest}}$ as described by

$$\tau^{\text{thr}} \frac{d\vartheta_i}{dt} = \vartheta_i^{\text{rest}} - \vartheta_i \quad (4.9)$$

4.4.2 Synaptic plasticity

Our model combines different forms of plasticity. Excitatory synapses exhibit short term plasticity (STP), spike-timing-dependent plasticity (STDP) and heterosynaptic plasticity. Inhibitory to excitatory synapses are plastic and obey a STDP rule which is globally modulated by a secreted factor.

4.4.2.1 Short term plasticity

All excitatory connections in our model exhibit short term plasticity (STP) as described previously (Markram et al., 1998; Mongillo et al., 2008). The temporal evolution of the STP state variables $u_i(t)$ and $x_i(t)$ is described by

$$\frac{d}{dt} x_j(t) = \frac{1 - x_j(t)}{\tau_d} - u_j(t) x_j(t) S_j(t) \quad (4.10)$$

$$\frac{d}{dt} u_j(t) = \frac{U - u_j(t)}{\tau_f} + U (1 - u_j(t)) S_j(t) \quad (4.11)$$

with a fixed parameter set for all synapses ($\tau_d = 200 \text{ ms}, \tau_f = 600 \text{ ms}$ and $U = 0.2$) which leads to a saturating output non-linearity.

4.4.2.2 Plasticity of excitatory synapses

Plastic excitatory connections are subject to three different mechanisms of long term plasticity: Triplet STDP (Pfister and Gerstner, 2006) as well as transmitter-triggered (Kempster et al., 2001) and heterosynaptic plasticity (Chistiakova et al., 2014). All three forms of plasticity effect the synaptic weights w_{ij} directly as follows

$$\frac{d}{dt}w_{ij}(t) = A z_j^+(t) z_i^{\text{slow}}(t - \epsilon) S_i(t) \quad (4.12)$$

$$-B_i(t) z_i^-(t) S_j(t) \quad (4.13)$$

$$-\beta (w_{ij} - \tilde{w}_{ij}(t)) (z_i^-(t - \epsilon))^3 S_i(t) \quad (4.14)$$

$$+\delta S_j(t) \quad (4.15)$$

The first two expressions (Exp. (4.12) and (4.13)) model the Triplet STDP part. Expression (4.14) represents heterosynaptic plasticity in which the high power of the postsynaptic firing rate acts as a burst detector. The ϵ offset in the time argument ensures that the current action potential is not counted in the trace. Finally Expression (4.15) represents the term responsible for transmitter-triggered plasticity. All occurrences of the state variable $z_{j/i}(t)$ denote synaptic traces which can occur as pre or postsynaptic quantities. Each of them evolves independently according to the following differential equation

$$\frac{dz_i^x}{dt} = -\frac{z_i^x}{\tau^x} + S_i(t) \quad (4.16)$$

with individual time constants τ^x . The parameters A , β and δ are fixed. Moreover $B_i(t) = A$ (unless mentioned otherwise; see Section (4.4.2.4) for details) and the reference weight $\tilde{w}_{ij}(t)$ evolves according to its own consolidation dynamics.

4.4.2.3 Consolidation dynamics

Similar to existing work \tilde{w} follows the negative gradient of a double well potential (Clopath et al., 2008; Lisman, 1985; Gerstner and Kistler, 2002b) and is biased by the difference between current weight w_{ij} and the reference weight \tilde{w}_{ij} described by the following expression

$$\begin{aligned} \tau^{\text{cons}} \frac{d}{dt} \tilde{w}_{ij}(t) = & w_{ij}(t) - \tilde{w}_{ij}(t) \\ & -P \tilde{w}_{ij}(t) \left(\frac{w^{\text{P}}}{2} - \tilde{w}_{ij}(t) \right) (w^{\text{P}} - \tilde{w}_{ij}(t)) \end{aligned} \quad (4.17)$$

in which the parameter P controls the strength of the potential, w^{P} defines the upper stable fixed point for the unbiased case (i.e. $w_{ij}(t) = \tilde{w}_{ij}(t)$); in

this case the lower stable fixed point is at $\tilde{w}_{ij}(t) = 0$). And finally $\tau^{\text{cons}} = 20$ min characterizes the rate of convergence towards the stable equilibrium solutions.

4.4.2.4 Homeostatic regulation of LTD

In most simulations we keep the rate of LTD $B_i(t)$ fixed ($B_i(t) = A$) and choose initial synaptic strength such that a subset of neurons respond with rates above the LTP threshold to external stimulation. Here we implicitly assume that the network has been prepared in this state by a homeostatic mechanism which acts on a much longer timescale than it is captured in our simulations. To test this hypothesis we ran a subset of simulations in which $B_i(t)$ was explicitly time dependent (Fig. C.3 and C.4; Bienenstock et al. (1982); Pfister and Gerstner (2006)). Where mentioned explicitly in the text we set

$$B_i(t) = \begin{cases} AC_i(t) & \text{for } C_i(t) \leq 1 \\ A & \text{otherwise} \end{cases} \quad (4.18)$$

$$\frac{d}{dt}C_i(t) = -\frac{C_i(t)}{\tau^{\text{hom}}} + \left(z_i^{\text{ht}}(t)\right)^2 \quad (4.19)$$

with $\tau^{\text{hom}} = 20$ min, in which the dynamics of the synaptic trace $z_i^{\text{ht}}(t)$ are given as before by Eq. (4.16) with $\tau^{\text{ht}} = 100$ ms.

4.4.2.5 STDP model of inhibitory synapses

In our model inhibitory-to-excitatory synapses obey a simple STDP learning rule whose intensity and sign are controlled globally through a secreted factor (Turrigiano, 2012).

The STDP rule (cf. Fig. 4.4) is given by

$$\frac{d}{dt}w_{ij}(t) = \eta G(t) ((z_i(t) + 1) S_j(t) + z_j(t) S_i(t)) \quad (4.20)$$

where η are constants, the $z_{j/i}^x$ denote pre/postsynaptic traces (cf. Eq. (4.16)), $S_{j/i}(t)$ are the pre/postsynaptic spike trains, and $G(t)$ represents the global secreted factor which we model as $G(t) = H(t) - \gamma$ where $H(t)$ is the low pass filtered version of all spikes in the excitatory population given by

$$\frac{d}{dt}H(t) = -\frac{H(t)}{\tau^H} + \sum_{i \in \text{exc}} S_i(t) \quad (4.21)$$

with characteristic time constant $\tau^H = 10$ s. We interpret $G(t)$ as a chemical signal which neurons secrete when they are active and which diffuses in the intracellular space where it can be sensed by other neurons or synapses as

Chapter 4. Memory formation and recall

a measure of the global network activity. When this activity drops below its target value γ , $G(t)$ is smaller than zero and the STDP learning rule (Eq. (4.20)) becomes reminiscent of Woodin et al. (2003). By modulating inhibitory STDP this way the learning rule becomes bidirectional and serves as a controller for the overall network activity.

4.4.3 Simulations including plasticity

4.4.3.1 Simulation of postsynaptic tetanization protocols

To mimic the plasticity protocol described in Chen et al. (2013) (Fig 4.2 (a)) we connected a single postsynaptic neuron with 1000 presynaptic connections with our plasticity rule. We used random initial values for the synaptic weights w_{ij} and their reference values \tilde{w}_{ij} which were drawn independently from a normal distribution with mean 0.3 and variance 0.09. To ensure positive values after the assignment, all weight values were then truncated at zero.

We simulated the ongoing measurement of the EPSP size from two different pathways via designated synapses (Figs. 4.2 (a) i and ii) which were stimulated at alternating intervals with two spikes (50ms time difference) every 7.5s. This stimulation was maintained during the entire protocol except during tetanization (10min $< t <$ 13min) where the postsynaptic cell was forced to spike through simulated current injection (3 trains with one minute offset consisting of 10 burst at 1Hz with 5 spikes at 100Hz each; compare Chen et al. (2013)).

4.4.3.2 Stimulation paradigm

For simulations requiring external stimulation we used the following paradigm. In the absence of a stimulus all input neurons were firing with Poisson statistics at a fixed rate of 10Hz. A designated set of stimuli was fixed at the beginning of the simulation as a graded (Fig. 4.2) or a binary activation pattern of input neurons (e.g. Fig. 4.4). During stimulation input cells which were active in a given pattern fired with a by $s = \gamma 35\text{Hz}$ increased rate (unless mentioned otherwise), in which γ was one for binary patterns or from the interval $[0,1]$ for graded activity patterns.

Stimulus order was chosen randomly with inter-stimulus-intervals and stimulus durations drawn from Exponential distributions with respective mean values T_{Off} and T_{On} . All simulations were allowed for an initial burn-in period of at least 50s during which no stimulation was triggered.

4.4.3.3 Details of feed-forward network simulations

To characterize the effect of our excitatory plasticity rule on a single postsynaptic neuron we simulated two simple feed-forward networks without

inhibition (Fig. 4.2 (c-e) and Fig. C.1). In particular we simulated a single postsynaptic neuron which received plastic excitatory input from two populations of Poisson neurons. The neurons in one population fired at 10Hz and the initial weight w_0 took different values in the interval $[0.2, 0.35]$ as stated in Figure 4.2 (c-e). Neurons in the second Poisson population fired at 1Hz and the initial weight was initialized at a value of 0.1.

In Figure 4.2 (g) we used a similar setup with 1000 presynaptic Poisson inputs (initial weight $w_0 = 0.05$) all firing at a constant background rate of 10Hz. The stimulus set consisted of ten Gaussian profiles in the presynaptic index with fixed standard deviation $\sigma = 50$. Stimulation onset was at $t = 100$ s mean stimulation interval $T_{\text{On}} = 20$ s ($T_{\text{Off}} = 100$ ms).

4.4.3.4 Balanced network model

In all network simulations we used a balanced network model consisting of 4096 (64×64) excitatory and 1024 inhibitory integrate and fire neurons. The connectivity within the network between all neurons was random sparse with an overall connection probability of 10%. Neurons in the excitatory population received additional input from an external population of equal size which provided noisy background input. These input connections were either pre-structured so that cells from within a circular area (radius $R = 8$) in the 64×64 input space projected to individual network neurons (cf. Fig. C.3). In particular the center of the circle was chosen randomly within the input space for each postsynaptic neuron. Where mentioned explicitly input connections initialized with random sparse connectivity with 5% connection probability. All excitatory afferent connections relayed stimuli from the external Poisson input population to the network (see Section 4.4.3.1).

In simulations involving plasticity all afferent connections to excitatory cells in the network were plastic (cf. Section 4.4.2 and Fig. 4.4 (a)). Moreover plasticity was always active, also during periods when the network was cued with partial stimuli (cf. Fig. 4.5).

4.4.3.5 Simulation details

All simulation code was written in C++ and based on our in-house network simulation framework Auryn. Neuronal state variables were updated using the forward Euler method with 0.1ms temporal resolution. The only exception from that was the slow evolution of the reference weights \tilde{w}_{ij} which were updated with a time step of 1.2s for efficiency reasons.

4.4.4 Determining readout populations

To determine which cells respond to a given stimulus we compute the total spike count per neuron per stimulus during all stimulus-on periods in the time interval $3000\text{s} < t < 3500\text{s}$. We then compute the population average

Chapter 4. Memory formation and recall

per stimulus and count neurons as active for a given stimulus if their average spike response is larger than the mean. We also computed the population firing rate of the corresponding sub-population (activity) with a temporal resolution of 100ms. These data are either plotted directly along with spike raster plots or used to compute peristimulus time histograms.

4.4.5 Analysis of learning rule

If the pre and postsynaptic spike trains are Poisson triplet STDP can be interpreted as a rate model (Pfister and Gerstner, 2006). The derivation for the present plasticity rule proceeds mostly along the same lines. Only the heterosynaptic term needs a somewhat special treatment to take into account the higher order correlations that arise.

We start from the formulation of the full plasticity rule

$$\frac{d}{dt}w_{ij}(t) = A z_j^+(t) z_i^{\text{slow}}(t - \epsilon) S_i(t) \quad (4.22)$$

$$- B_i(t) z_i^-(t) S_j(t) \quad (4.23)$$

$$- \beta (w_{ij} - \tilde{w}_{ij}(t)) z_i^{\text{het}}(t - \epsilon)^3 S_i(t) \quad (4.24)$$

$$+ \delta S_j(t) \quad . \quad (4.25)$$

Since consolidation dynamics are slower than the weight dynamics in w we now assume fixed $B_i(t) = B$, $\tilde{w}_{ij} = \text{const}$. Taking the expectation value on both sides yields

$$\left\langle \frac{d}{dt}w_{ij}(t) \right\rangle = A \tau^+ \tau^{\text{slow}} xy^2 - B \tau^- xy - \beta (w_{ij} - \tilde{w}_{ij}) \left\langle z_i^{\text{het}}(t)^3 \right\rangle y + \delta x$$

where we replaced all presynaptic traces by their mean firing rates x (and postsynaptic traces by y respectively). This leaves us with the third order moment $\langle z_i^{\text{het}}(t)^3 \rangle$ of the burst detector which for a Poisson spike train can be computed conveniently by taking the Laplace transform of the stationary distribution of z when interpreting it as a random variable (see Section C.3 for details).

This allows us to arrive at

$$\begin{aligned} \left\langle \frac{d}{dt}w_{ij}(t) \right\rangle &= A \tau^+ \tau^{\text{slow}} xy^2 - B \tau^- xy + \delta x \\ &\quad - \beta (w_{ij} - \tilde{w}_{ij}) \left(\frac{y\tau}{6} (2 + 9y\tau + 6y^2\tau^2) \right) y \quad . \quad (4.26) \end{aligned}$$

For given pre and postsynaptic firing rates this equation has a single stable fixed point w_{ij}^{FP} at




Figure 4.8 – Weight fixed point w^{FP} for different pre (x) and postsynaptic (y) firing rates (left panel; $\beta = 0.004$ and $\tilde{w} = 0.25$). Three representative cross sections for $x = 1$ (dashed line), $x = 30$ (dotted line) and for $x = y$ (solid line) are plotted in the right panel.

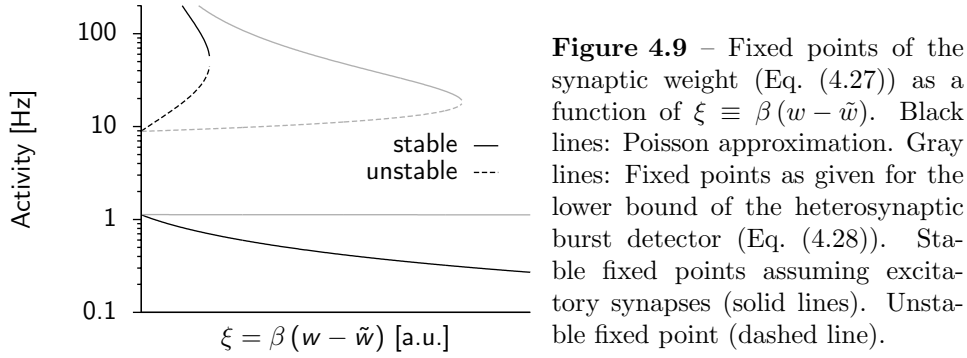


Figure 4.9 – Fixed points of the synaptic weight (Eq. (4.27)) as a function of $\xi \equiv \beta(w - \tilde{w})$. Black lines: Poisson approximation. Gray lines: Fixed points as given for the lower bound of the heterosynaptic burst detector (Eq. (4.28)). Stable fixed points assuming excitatory synapses (solid lines). Unstable fixed point (dashed line).

$$w_{ij}^{\text{FP}} = \tilde{w}_{ij} + \frac{A\tau^+\tau^{\text{slow}}xy^2 - B\tau^-xy + \delta x}{\beta\left(\frac{1}{6}y\tau(2 + 9y\tau + 6y^2\tau^2)\right)y} \quad (4.27)$$

We now consider connections inside a cell assembly ($x = y$) and the afferent connections into the assembly ($x \ll y$). From Eq. (4.27) we see that both weights are finite and the individual fixed points are well separated for the two connection types (Fig. 4.8). In particular connections between neurons with high activity potentiate to a fixed level at which they slowly saturate even if rates increase further.

We can appreciate the multiple stationary solutions of the learning rule by looking at the fixed points of Eq. (4.26) as a function of $\xi \equiv \beta(w - \tilde{w})$ (Fig. 4.9), which illustrates that for sensible choices of the parameter β weight dynamics can always be made bistable.

So far we have assumed that spiking is Poisson. However, particularly at elevated firing rates the effects of refractory period and adaptation are non-negligible any more and spiking becomes generally less irregular. The nontrivial analytical treatment of this issue will remain open at this point. Nevertheless Jensen's inequality always guarantees that the burst detector

Chapter 4. Memory formation and recall

obeys the lower bound

$$\langle z_i^{\text{het}}(t)^3 \rangle \geq \langle z_i^{\text{het}}(t) \rangle^3 \quad (4.28)$$

at which the stable fixed point moves to lower activities (Fig. 4.9). The true fixed points are likely to lie in between this lower bound and the Poisson approximation, given that the neuron does not burst. Bursts in the output activity lead to an increase of the relative strength of heterosynaptic plasticity beyond the Poisson approximation and can thus cause the upper fixed point to disappear.

Chapter 5

Inter-process communication, not plasticity, limits the speed-up in simulations of spiking neural networks

Authors: Friedemann Zenke and Wulfram Gerstner

At the time of writing this thesis, a version of this chapter was in revision for *Frontiers in Neuroinformatics*. The entire chapter represents the original work of Friedemann Zenke. In particular he is responsible for all main aspects. Wulfram Gerstner suggested minor modifications during the manuscript preparation.

Abstract

To understand how the central nervous system performs computations using recurrent neuronal circuitry, simulations have become an indispensable tool for theoretical neuroscience. To study neuronal circuits and their ability to self-organize, increasing attention has been directed towards synaptic plasticity. In particular spike-timing-dependent plasticity (STDP) creates specific demands for simulations of spiking neural networks. On the one hand a high temporal resolution is required to capture the millisecond timescale of typical STDP windows. On the other hand network simulations have to evolve over minutes, hours or even days, to capture the timescale of long-term plasticity. We have addressed this demand and created the highly specialized simulation environment “Auryn” for the simulation of synaptic plasticity in medium-size recurrent network models. Here we compare the performance of our code to other simulators like Brian, NEST and Neuron. We find that Auryn is faster than Brian and Neuron while performance is similar to NEST. In parallel simulations NEST and Auryn exhibit comparable scaling behavior. In particular the speed-up saturates for a low numbers of active cores. This saturation is then analyzed using a minimal run time

model, which shows that latencies in the inter-process communication constitute the bottleneck. Finally we show that current models of STDP do not impose an additional constraint on the scaling behavior for medium size networks. Our results show that real-time simulation of plastic networks is possible. However, the speed-up margin of parallel code is limited due to inter-process communications. This suggests that researchers addressing problems of plasticity in recurrent neural networks should take care in the selection of their simulation hardware.

5.1 Introduction

Neurons communicate with each other by short electrical pulses, called action potentials or spikes, that can be considered as unitary events. In simple neuron models of integrate-and-fire type, such events are generated by a threshold crossing process. The dynamics of a single neuron, which forms one unit of a large brain network, are therefore relatively simple.

Nevertheless, the simulation of activity in large neural networks, which has been receiving increasing interest over the past years (Markram, 2006; Ananthanarayanan et al., 2009; Lang et al., 2011; Koch and Reid, 2012; Waldrop, 2012; Kandel et al., 2013), poses multiple computational challenges. First, brain networks consist of billions of neurons (Kandel et al., 2000). Even if each neuron is described as a relatively simple dynamic processing unit (e.g., an adaptive integrate-and-fire neuron with two or three update equations per neuron (Izhikevich, 2003; Brette and Gerstner, 2009; Gerstner et al., 2014)), the sheer number of units suggests that real-time simulation of these equations will be hard to achieve on a single core. Therefore parallelization of computation is desirable. Second, each unit sends and receives signals from thousands of others (Kandel et al., 2000; DeFelipe and Fariñas, 1992), such that connectivity between units is relatively high compared to classical models in the physical sciences where interactions are mainly between nearest neighbors in physical space (Anderson, 1995). Therefore, the communication overhead in a parallel implementation could potentially be high. Third, the synaptic contact points between two connected units are not fixed but may change (Bliss and Lømo, 1973; Markram et al., 1997; Bi and Poo, 1998; Zhang et al., 1998; Bi and Poo, 2001; Markram et al., 2012). Therefore connections cannot be described with fixed parameters, but need further dynamic variables. Moreover, the evolution of these synaptic variables depends on activity of both the sending and the receiving neuron so that their treatment requires additional care and readily available parallelization approaches cannot be used. The changes in the dynamic values associated with the synaptic contact points are referred to as synaptic plasticity.

The question therefore arises whether the scaling of parallel implementa-

tions of simulated neural networks is dominated mainly by the inter-process communication or by the dynamics of the connections. This question cannot be answered in a straightforward manner, because it depends on multiple factors. First, the communication between neurons only takes place at the moment when a spike happens, leading to event-based updating schemes (Morrison et al., 2007, 2008). Therefore the number of events per unit of time plays a role for the communication overload. Second, changes of synaptic parameters, while induced by spike events, are relatively small so that they evolve on a slower time scale. Roughly speaking, a biological neuron sends out spikes that last each about one millisecond. The rate at which these spike events are generated is a few per second. The slowest dynamics are that of synaptic plasticity which typically needs several spike events to induce a measurable change. Moreover, once changes are induced they often persist for many hours. In the field of neuroscience, the behavioral phenomenon of learning and memory formation is thought to be intimately linked to the biological rules of synaptic plasticity (Bliss and Lømo, 1973; Markram et al., 1997; Bi and Poo, 1998; Zhang et al., 1998; Bi and Poo, 2001; Markram et al., 2012). To check in experiments whether a stable memory has been formed it is not uncommon to follow a biological substrate for 24 hours or more. If we want to simulate learning and memory formation, the simulator has therefore to cover time scales from milliseconds to days.

While simulation packages for networks with fixed connectivity are readily available (Gewaltig and Diesmann, 2007; Eliasmith et al., 2012; Hoang et al., 2013), simulations of plastic brain circuits have received much less attention (Gewaltig and Diesmann, 2007; Izhikevich and Edelman, 2008; Ananthanarayanan et al., 2009). For example, the NEST simulation environment has been in the released initially for fixed network connections and models of synaptic plasticity have been added later on (Gewaltig and Diesmann, 2007; Morrison et al., 2007).

Here we focus on networks of several tens of thousands of neurons. These medium-sized networks are of particular practical importance because they are used in many theory and modeling labs world wide. To achieve high simulation speeds in medium size recurrent networks we have developed the lightweight parallel simulation environment Auryn. Simulations based on this framework were successfully used in a range of studies that either involve plasticity (Vogels et al., 2011; Zenke et al., 2013) or in which long recordings of simulated network data were a requirement (Lütcke et al., 2013). Auryn is written from scratch in C++ and adopts powerful concepts from existing open source simulators such as Brian (Goodman and Brette, 2008) or NEST (Gewaltig and Diesmann, 2007). Its source has recently been released under the general public license (GPL). Here we describe the key features of our simulation environment and compare its performance to three popular simulators in the field (namely Neuron, Brian and NEST).

We show that medium size networks of point neurons with plastic synapses can be simulated in real-time or faster. In our tests on a single computer or a small cluster, simulations using Auryn generally ran faster than this was the case for other simulators in their standard configuration. We achieved comparable results with NEST when running it with the same forward Euler solver that was used in Auryn. Under these conditions both simulators showed not only comparable run times, but also similar scaling behavior when run in parallel on a cluster. In particular we show that small to medium size networks (up to 25,000 IF fire neurons with plasticity) can be run in real-time or faster. However, we observed for both simulators that run times saturate already at a relatively low number of parallel processes. This makes it difficult to speed up the simulation beyond about one tenth of real-time. We show that this saturation is due to communication latencies in the inter-process communications. In a last step we illustrate that the implementation details of spike based plasticity rules, such as STDP, do not have a large impact on run times. Latencies in the inter-process communication still constitute the main bottleneck. Consequently it is likely that the speed – instead of size – requirement can be fulfilled by single powerful compute nodes or super-computers with extremely low latencies rather than large clusters.

5.2 Materials and Methods

In this manuscript we compare results from a range of different neuron, network and plasticity models. However, there are some underlying similarities. All networks are built from integrate-and-fire neurons with either current based or conductance based synaptic input. We have summarized the detailed model description for the neuron models, plasticity rules and network models in tabular form according to (Nordlie et al., 2009) (Supplementary Material). In the following we only give a short overview of the simulation code, hardware and network models we used. In Section 5.2.3 we comment on general implementation details of STDP in simulations.

5.2.1 Simulation code and hardware

For all simulations using Auryn (version 0.3) which is publicly available on the Internet¹ we used forward Euler integration with a 0.1ms integration time step and a 0.8ms synaptic delay. Simulations are compiled against Boost (version 1.41.0), GSL (version 1.13) and MPICH2 (version 1.2.1) using the GNU C++ compiler (version 4.4.7). The code was executed on either a single node or a cluster consisting of 4 nodes. The individual nodes were technically identical and running Red Hat Enterprise Linux (version 6) on

¹<http://www.fzenke.net/auryn>

a dual CPU (Intel Xeon CPU E5-2670 0 @ 2.60GHz) board with 64GB of RAM. Nodes communicated using Ethernet link aggregation over four 1Gb connections each via a switch comprised of two Cisco Nexus Fabric Extenders (N2K-C2248TP) and two Cisco Nexus (N5K-C5548-UP).

5.2.2 Network models

5.2.2.1 Vogels-Abbott benchmark

For comparison against other simulators we adapted the `VAbenchmark2.py` benchmark code from the PyNN (Davison et al., 2009) package to run the same simulation in Brian, NEST and Neuron. To compare simulation times alone we modified the benchmark code such that it did not record any spikes or membrane potentials. Furthermore we introduced a 0.8ms synaptic delay between all synapses to reduce inter-process communication for the cases where parallel execution was possible. For Figure 5.2 we used a simulation time of 20s. For all benchmarks only the run time of the actual simulation was measured. The time to set up the weight matrices or to write data to disk was ignored.

We implemented the same benchmark network in Auryn. It is comprised as an example in the current release of the code (`sim_coba_benchmark.cpp`). The only difference between the two implementations was that Auryn initially loads the network state from a file. Hence no priming with 50ms of external Poisson input was needed.

Unless mentioned otherwise we used Neuron (version 7.3), NEST (version 2.2.2) and Brian (Goodman and Brette, 2008) (version 1.2.1) in their respective default configurations and compiled against MPI libraries where possible. Unless mentioned otherwise we used MPICH2 (version 1.2.1). However we did not encounter a notable difference in performance with OpenMPI (version 1.4.3).

5.2.2.2 Other network models

Apart from the Vogels-Abbott benchmark described in the last section (Figure 5.2 and Figure 5.3 A) we used two distinct network models that are adapted from published work. The detailed model description is available in tabular form (Nordlie et al., 2009) in the Supplementary Material. The code to all simulations used in this paper is freely available with the Auryn simulator package.

1. The 25,000 cell network used in Figure 5.3 B-D is described in Lütcke et al. (2013) to mimic a cortical circuit at low firing rates receiving temporally varying external input. The code is available as `sim_dense.cpp` with the Auryn simulator.

2. The 25,000 cell network used for Figures 5.3 C-D and Figure 5.4 B-D is described in Zenke et al. (2013). Its example code is available as `sim_background.cpp` with the Aurnyn simulator.

5.2.2.3 Brunel network

This is a 10,000 cell balanced network model based on Brunel (2000). We adapted the simulation included in the pyNEST examples in the current NEST (Gewaltig and Diesmann, 2007) release (version 2.2.2) under the `LeNovere_2012` directory (Gewaltig et al., 2012), which comes as a non-plastic network and the same network with weight dependent STDP (Supplementary Material). To make the networks comparable with the other simulations synaptic delays were set to 0.8ms and fixed random connectivity was used.

5.2.3 Implementation of spike-timing-dependent plasticity

A broad family of spike-timing-dependent plasticity (STDP) models can be written in the following form Gerstner and Kistler (2002b)

$$\begin{aligned} \frac{dw_{ij}}{dt} &= a_1^{\text{pre}} S_j(t) + a_1^{\text{post}} S_i(t) \\ &\quad + S_j(t) \int_0^\infty W(s) S_i(t-s) ds \\ &\quad + S_i(t) \int_0^\infty W(-s) S_j(t-s) ds \end{aligned} \quad (5.1)$$

where a_1^{pre} and a_1^{post} are constants, $W(t)$ is a real valued function with finite support and $S_j(t)$ is the presynaptic ($S_i(t)$ the postsynaptic) spike train given as a sum of delta functions $S_x(t) = \sum_k \delta(t-t_{xk})$ where t_{xk} runs over all spike times k of neuron x . The parameters a_1^{pre} , a_1^{post} as well as the window $W(s)$ may depend on the momentary value w_{ij} of the synaptic weights (van Rossum and Turrigiano, 2001; Gütig et al., 2003). Expression (5.1) describes a piecewise constant function of time with jumps whenever pre- or postsynaptic spikes occur. Note that STDP can also contain higher order terms (Pfister and Gerstner, 2006) which does not influence the key points of our argument. In many situations the window function $W(t)$ can be well approximated by one or multiple exponential functions. As an example

$$W(t) = \begin{cases} A^+ \exp\left(-\frac{t}{\tau_A}\right) & t > 0 \\ B^- \exp\left(+\frac{t}{\tau_B}\right) & t \leq 0 \end{cases}$$

where A , B , τ_A and τ_B are constants, yields a plausible STDP curve (Song et al., 2000; Gerstner and Kistler, 2002a; Zhang et al., 1998). Whenever the

5.2. Materials and Methods

window function can be broken down to exponential shapes, this allows for an efficient online implementation by using synaptic traces (Gerstner and Kistler, 2002b; Morrison et al., 2008). A synaptic trace $z_i(t)$ is a low pass filtered version of the spike train $S_i(t)$ of neuron i . It is described by the linear differential equation

$$\frac{dz_i^x}{dt} = -\frac{z_i^x}{\tau_x} + S_i(t) \quad (5.2)$$

with associated respective timescale τ_x . In the absence of spikes the solution is a simple exponential decay. Eq. (5.2) can either be integrated time-continuously by multiplication with the constant $\exp(-\frac{\Delta t}{\tau_x})$ in every simulation time step Δt or by using the fact that the analytical solution is known for arbitrary time intervals (event-based update).

By combining Eq. (5.1) and (5.2) synaptic weight updates can be written as follows

$$\frac{dw_{ij}}{dt} \propto A^+ z_j^+(t) S_i(t) - A^- z_i^-(t) S_j(t) + a_1^{\text{pre}} S_j(t) + a_1^{\text{post}} S_i(t) \quad (5.3)$$

which is ideally suited for event-based integration because weight changes only occur at pre- or postsynaptic spike times. To add plasticity to a network simulation one therefore simply adds the required number of traces (cf. Eq. (5.2)) and the event-based weight update. The simplest implementation of STDP now proceeds as follows: at time t_j of a presynaptic spike of neuron j the trace $z_i^-(t)$ is read out and the necessary weight update is applied to the weight w_{ij} . Since the postsynaptic trace of neuron i can be integrated alongside with the neuronal state no particular care has to be taken for parallel processing.

This changes in two ways in the case of a postsynaptic spike. First, in case of parallel processing the neuron from which the spike originated might not be integrated on the same physical computer. Hence there is no simple way of providing the value of its synaptic trace. Second, the simulator might not offer efficient means of finding all presynaptic partners of a postsynaptic neuron. Doing this efficiently generally costs memory, because each neuron needs to keep a list of all its presynaptic partners. This approach is therefore prohibitive for simulators like NEST which are aiming at very large scale simulations, where memory usage is a limiting factor (Helias et al., 2012).

Per default Auryn takes the simplest and most memory greedy approach where at each postsynaptic firing time all associated weight updates are carried out immediately. To be able to provide the value of the trace of any presynaptic neuron at the time of the update Auryn computes presynaptic traces on all nodes in a time continuous way. That means that presynaptic traces are evolved in every time step irrespectively if the value is needed or

not. Since every process needs to keep track of all presynaptic traces, it also means that some redundant work is being done.

Auryn alternatively supports an event-based approach. This approach exploits the fact that in the absence of spikes, the solution to Eq. (5.2) is an exponential decay. Since the event-based trace update cannot be vectorized efficiently this only provides an advantage at low firing rates. However, the increased overhead due to the use of the exponential function generally seems to outweigh the advantages of this approach, which is why Auryn chooses by default time-continuous updates of presynaptic traces.

NEST uses an event-based approach which is more elaborate. Synaptic weight updates are only carried out at the times of presynaptic spikes (Morrison et al., 2007). To do this, each neuron stores its past firing history in a small buffer. Whenever a presynaptic spike occurs, all post-pre updates are applied retrospectively in a batch. Since for each update all quantities appearing in Eq. (5.3) have to be known, the retrospective update requires to keep track of these values or to compute them when needed.

5.3 Results

Auryn was written with the UNIX philosophy in mind: Do one thing and do it well (Raymond, 2003). To run simulations fast Auryn simulates networks of spiking neurons and writes relevant output to human readable text files. It does not perform any analysis and the output files have to be processed and analyzed independently. At the heart Auryn is a collection of C++ classes that are combined into a compiled program to form the simulation. This allows the compiler to optimize each simulation code specifically for the hardware it runs on.

Like other simulators Auryn takes a hybrid approach between event-based and continuous integration (Morrison et al., 2005). Neuron models in Auryn are integrated continuously, while weight updates for many standard synaptic plasticity rules are implemented in an efficient event-based way.

Quantities that require time continuous integration are typically neuronal state variables describing synaptic conductance and membrane voltage. In many neuronal networks large sets of identical or similar neurons need to be integrated. The required computation can be vectorized efficiently. The advantages of vectorization are the reduction of function calls, the efficient use of multi-layer cache architectures deployed in modern CPUs, and giving the compiler the opportunity to use hardware for single instruction multiple data (SIMD) such as SSE or AVX. Vectorization is therefore widely used in existing simulators such as NEST and Brian (Brette and Goodman, 2011) as well as in our code.

The logical extension to vectorization is parallelization. To run parallel code Auryn uses the message passing interface (MPI) as a general and

versatile back-end to allow parallel simulations on a single machine with multiple cores or distributed over multiple physical machines in a cluster. Auryn currently does not support multithreading, a design decision which was mainly rooted in trying to keep things as simple as possible.

To place Auryn amongst existing software we performed a series of tests and compared its performance to existing simulators such as NEST, Brian and Neuron. We first describe the simple matter of efficiently generating Poisson spike trains and then show run time results from simulating several different balanced networks with or without plasticity. We use these networks to systematically study the scaling behavior of our code when run in parallel and discuss the emerging saturation properties with the help of a simple run time model. Finally we discuss the cost of plasticity in network simulations.

5.3.1 Poisson spike trains

Spike trains from Poisson neurons are widely used to provide noisy background input to networks of spiking neurons to achieve some baseline activity. In NEST such input is implemented as the combination of a Poisson generator with so called parrot neurons. The Brian simulator has its own dedicated object for the purpose, a `PoissonGroup`. Since we often use Poisson input in our simulations, Auryn comes with an optimized algorithm for the particular scenario where a single pool of N Poisson neurons fires with identical firing rates.

To efficiently generate Poisson spike trains from such a configuration we consider a grid spanned by N rows corresponding to the Poisson neurons and discrete-time with bins of size Δt on the x -axis (Figure 5.1 A). To create Poisson spikes we could now fill each row of the grid with spikes by drawing exponentially distributed inter-spike-intervals. This can be done on-line, but requires a certain degree of book-keeping because the algorithm has to remember the N last spike times of all Poisson units. It is more efficient to fill each column at the very time step when the spikes are needed. This can be done efficiently since the distribution of inter-spike-intervals (ISIs) is the same in x and y -direction. Therefore all spikes can be generated online when they are needed during the simulation (Figure 5.1 A). When a jump leads beyond N it is simply continued in the next time-step. This way every random number yields a spike.

To test the performance we wrote equivalent simulation code in Brian, NEST (by using pyNEST) and Auryn which simply implements 1000 Poisson neurons firing at 5Hz firing rate. The run times for simulating 100s spike trains are shown in Figure 5.1 B. This illustrates that in this particular scenario the Auryn approach is faster than the other two examples. This allows Auryn to generate Poisson spike trains at a minimal overhead, which is important when aiming at network simulations in, or faster than, real-

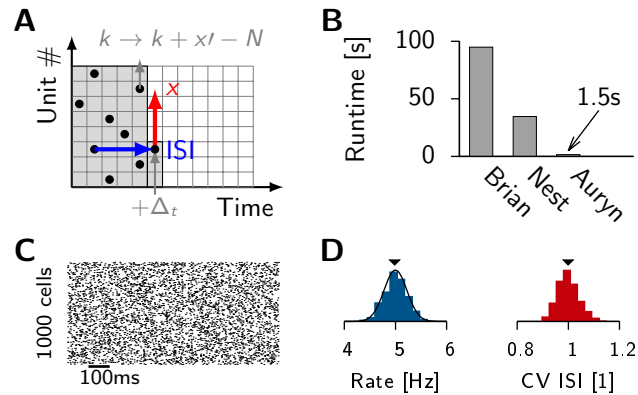


Figure 5.1 – Efficient generation of Poisson spike trains from a population of input units. (A) Illustration of the algorithm. One typical ISI interval for single Poisson neuron (blue). Typical example of a random stride of magnitude x (red arrow) or a stride x' that would lead beyond N and is consequently continued in the next time step (gray). (B) Run times of different simulators to generate 100s long spike trains from 1000 Poisson neurons spiking at 5Hz. (C) Spike raster from simulated Poisson spike trains from A. (D) From the same simulation: Distribution of firing rates (left) with the theoretical expectation from the Poisson distribution (solid line). Right: Distribution of coefficient of variation of the ISI (CV ISI). The mean values of the distributions are indicated by arrow heads.

time. The same algorithm is also used to efficiently set up random sparse connectivity matrices in our code.

5.3.2 Vogels-Abbott network benchmark

We study recurrent networks of leaky-integrate-and-fire neurons. To achieve low run times Auryn integrates the neuronal state variables (e.g. the membrane potential) using the forward Euler method. More sophisticated integration algorithms or neuron models can be readily implemented. Similar to NEST, neuron models in Auryn are included as independent C++ files, which can be extended or modified. To further improve performance in Auryn, most variables involved in the numerical integration of the neuronal state variables use single precisions floating point values. This generally allows to perform twice as many floating point operations when using SIMD as with double precision, thus yielding a theoretical speed-up by a factor of two. At the same time, the choice of single precision resolution doubles the number of variables that fit into a cache block compared to double precision and for this reason reduces costly cache misses.

To test the performance of our code and compare it with other simulators we used a down-scaled version of the conductance-based Vogels-Abbott model (Vogels and Abbott, 2005), a balanced network model which has been used as a benchmark in the past (Brette et al., 2007). The source for this




Figure 5.2 – Vogels-Abbott benchmark. (A) Spike raster of the excitatory cells in the network as simulated with Auryn. (B) Distribution of the firing rates (left) and the the coefficient of variation (right). Solid color bars are the results from Auryn. The black lines is the output from NEST. Mean values of the distributions are indicated by arrow heads. (C) Run time of different simulators when simulating 20s of the Vogels-Abbott benchmark network. The dashed line indicates a real-time simulation. (D) Zoom on the comparison of Auryn against NEST*, where a minimal forward Euler solver was used to integrate the neuronal state variables. (E) Run times of Nest and Auryn for 20s of benchmark simulation. Simulations ran in parallel on a single machine. (F) Simulations ran distributed on four nodes connected via Ethernet.

Chapter 5. Speed-up limits in network simulations

benchmark is available with pyNN (Davison et al., 2009) which allowed us to simulate the same network in Brian, NEST and Neuron.

The very same network was implemented using Auryn. The observed network activity was comparable to the other simulators (Figure 5.2 A, B). However, we also observed small differences (on the order of 0.5Hz) between the mean population firing rate in the NEST simulation and obtained with Brian or Auryn. To check if such differences could be due to the different underlying integration methods, we re-implemented the neuron model in NEST using the forward Euler method (NEST*: we use the asterisk to disambiguate the different integration methods). This change indeed resulted in the same mean population firing rates in the simulations performed with Brian, NEST* and Auryn. We further confirmed this result by comparing the rate distribution, and the distribution of the CV ISI obtained with the different simulators (Figure 5.2 B), using a two sample Kolmogorov-Smirnov test (Auryn-NEST* (rate) $D=0.0124$, $p=0.9678$. (CVISI) $D=0.0252$, $p=0.2678$; Brian-NEST* (rate) $D=0.0119$, $p=0.9775$, (CVISI) $D=0.0117$, $p=0.9804$). This suggests that the observed differences were indeed due to the different integration schemes and that the use of single precision values in Auryn has negligible effects.

In a next step we measured the execution times for the different simulators when running the network during 20s on a single core (Methods; Figure 5.2C). We found that the Auryn code runs significantly faster than the other simulators. Importantly Auryn runs roughly four times as fast as real-time, whereas the other simulators are slower than real-time (with the exception of NEST*). From the comparison of NEST and NEST* one can also appreciate why comparing the simulators in their standard configuration is not necessarily a fair comparison. In the present example NEST uses the computationally more expensive fourth order Runge-Kutta method, which puts it far behind Auryn in terms of run times. NEST* on the other hand achieves comparable run times, with Auryn being only a factor of ≈ 2.5 faster. We speculate that this difference is likely to come from the use of single precision versus double precision variables for neuronal state variables.

NEST and Neuron as well as our own code can be run in parallel by using MPI. Since neuron was much slower in the single-core comparison, we limited our study to a comparison of NEST* with our own code. In order to compare run times when running parallel code with MPI, we evaluate the run times of the same simulation on one of our servers (with a total of 16 physical cores) while varying the numbers of cores used for the computation (Figure 5.2E). Since we are interested in running simulations in real-time or faster, we introduce the relative run time T (the total run time divided by the simulated time) as a performance criterion. T equal to one corresponds to real-time simulation, while a value smaller than one indicates simulation speeds which are faster than real-time. When plotting the

run time measurements for different numbers of active cores n , we observe that both simulators exhibit comparable scaling behavior of $\propto \frac{1}{n}$. To see whether the simulation speed could be further increased we now repeat the same measurement when the code is run distributed over four nodes identical to the machine we used before (Methods). In this case the simulations scale poorly and run times start to increase already for a low number of cores (Figure 5.2F).

The fact that NEST and our own code show similar scaling and saturation behavior suggests that there is a shared underlying principle, which is presumably due to overhead in the inter-process communications.

5.3.3 A minimal run time model

To test this hypothesis and to understand the origin of the saturation we study the scaling behavior in a simple model. We assume that the run time T for simulating the network is the sum of the time T_{sim} spent on actual simulation plus the time spent on inter-process communications T_{sync}

$$T = T_{\text{sim}} + T_{\text{sync}}$$

In an ideal parallel implementation T_{sim} scales as $\propto \frac{1}{n}$ where n is the number of cores that share the work. For small message sizes most MPI implementations use either the Recursive Doubling or the Bruck algorithm for the necessary All Gather operation that communicates the spikes from one node to the other (Thakur et al., 2005). Both algorithms scale as $\propto \beta \log(n) + \gamma n^{-1/n}$. Taken together we expect the run time T to scale as

$$T = \frac{\alpha}{n} + \beta \lceil \log_2 n \rceil + \gamma \frac{n-1}{n} \quad (5.4)$$

with only positive parameters α , β and γ . We can see immediately from Expression (5.4) that it has a lower bound $T_{\text{LB}} > 0$. This is the manifestation of the plain fact that inter-process communication takes time. Simulations cannot run faster than the time they spend on communication.

To see how well this simple model captures run time data we fit Expression (5.4) to the relative run times observed for the Vogels-Abbott benchmark when run on a single node (Figure 5.3 A). Most standard algorithms for all-to-all collective communication with small message sizes (Bruck algorithm and recursive doubling) perform best for process numbers that are a power of two. We therefore limit our measurements to values of n which are powers of two or their respective midpoints. We furthermore exclude the data point at 16 cores (64 cores for the distributed runs) from the fit because these cases did not leave a single core for the system, which could influence the outcome of the time measurement.

When fitted to the single node data, Expression (5.4) yields the following parameters $\alpha = 0.279 \pm 0.002$, $\beta = 0.0 \pm 0.0023$ and $\gamma = 0.016 \pm 0.007$, which




Figure 5.3 – Scaling behavior in Auryn. (A) Relative run times T for Vogels-Abbott Benchmark distributed over multiple cores on a single machine (solid circles) or distributed over a cluster of four nodes (triangles). Lines: RMS fit of Eq. (5.4) to local (red) and distributed data (blue). (B) same as in A, but in a network of 25,000 neurons with average firing rates of 0.3Hz. Circles and red line: run locally on a single node. Triangles and blue line: same network run distributed over four nodes. (C) Bar plot of parameter α from fitting Expression. (5.4) to the run time data in A (red), B (blue) and a third 25,000 network (dark gray) with a mean firing rate of 3Hz and plastic synapses. (D) same as C but for parameters β and γ .

5.3. Results

are all in units of relative run time. We see that for 12 cores the time spent on communication, characterized by β and γ is still significantly lower than the time spent on the actual simulation code represented by α .

When the same code is run distributed over multiple machines, the $\propto \frac{1}{n}$ scaling behavior is broken (Figure 5.3 A). We again extract the following parameters from the root mean square fit of Expression (5.4): $\alpha = 0.27 \pm 0.01$, $\beta = 0.019 \pm 0.003$ and $\gamma = 0.07 \pm 0.02$. One can observe that α remains unchanged, while β is now significantly different from zero, which causes the function (Eq. (5.4)) to rise again for increasing numbers of cores. The increase in β and γ is due to the increased communication delays. Messages between processes cannot be exchanged anymore via shared memory, but have to travel over Ethernet to another physical machine. This takes time and imposes a limit onto the achievable run time. The minimum run time can be extracted from the plot or analytically from the fitted parameters. The latter approaches yields $T_{\min} = 0.15 \pm 0.02$ achieved in the range of 6-8 cores used.

To be able to interpret this value we measured the mean run time for an All Gather for a message size comparable to that of the Vogels-Abbott benchmark. We found $t_{\text{AllGather}} \approx 0.1\text{ms}$. We may interpret this as the minimal time necessary for synchronization of all processes. Since the network simulation has a minimum delay of 0.8ms after which all processes have to synchronize, the speed-up of a simulation cannot be above a factor of 8 compared to real-time which imposes a lower bound on the relative run time at $T_{\text{lb}} = \frac{t_{\text{AllGather}}}{0.8\text{ms}} = 0.125$. This is only slightly smaller than T_{\min} . In conclusion our findings suggest that starting from 6-8 cores communication delays constitute the major share of the total run time ($T_{\text{sync}} \gg T_{\text{sim}}$).

To see whether these results generalize to other networks, we performed further run time measurements for a larger network (Lütcke et al., 2013). This network consists of 25,000 cells receiving fluctuating external input (Figure 5.3 B, see Methods) with a mean population firing rate of 0.3Hz. It thus generates less spikes per unit time than the Vogels-Abbott benchmark and therefore also exchanges smaller messages via MPI. The functional shape of the run times can be well captured by Expression (5.4). When run distributed over multiple nodes the run times of the 25,000 cell network do not show a large difference between the local and the distributed case. The best possible run time lies at $T \approx 0.25$ which again is comparable to the result from the Vogels-Abbott benchmark network.

We simulated a third network with 25,000 cells and plastic excitatory-to-excitatory synapses (Zenke et al., 2013) (see next section) and fit Expression (5.4) to the run time values. All three networks have different computational load which is reflected in the differences in the extracted values for α (Figure 5.3 C). However, values for β and γ which characterize the communication latencies are comparable in all three cases.

Although the average number of spikes and therefore the average mes-

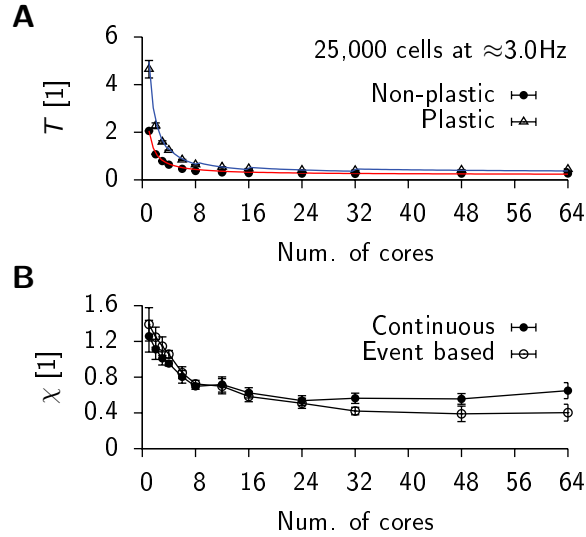


Figure 5.4 – Scaling of a plastic network with STDP. (A) Relative run times of a 25,000 cell network model (filled circles). The triangular points indicate run times for the same network with plastic excitatory-to-excitatory synapses that are subject to triplet STDP. The red (blue) solid line represents the outcome of an RMS fit of Expression (5.4) to the data for the non-plastic (plastic) case. (B) Relative difference of run times of the network from A with a naive continuous integration of the presynaptic traces and the event-based implementation as a function of the number of cores.

sage size in the two 25,000 networks differ by a factor of 10, the values for β and γ do not show a significant difference. Since β represents the communication cost of the Bruck algorithm (Thakur et al., 2005) this could mean that the lower bound T_{LB} is approaching values which are governed by the communication delays intrinsic to the hardware layer.

5.3.4 Computational cost of STDP

Auryn was developed to provide an efficient environment for simulating plastic synapses in recurrent neural networks. To that end we are particularly interested in simulating spike-timing-dependent plasticity (STDP) as it is a form of plasticity commonly found in the brain (Markram et al., 1997; Zhang et al., 1998; Bi and Poo, 1998, 2001; Markram et al., 2011). STDP can be implemented efficiently in an event-based way where synaptic weights only change when pre- or postsynaptic spikes occur (Gerstner and Kistler, 2002b; Morrison et al., 2008). We were wondering how performance and scaling are affected by introducing STDP in a network simulation.

We measured the run times of a medium size network simulation (25,000 cells, Methods) without plasticity and with triplet STDP (Pfister and Gerstner, 2006) at the excitatory-to-excitatory synapses. When run on a sin-

gle core, the non-plastic (NP) network runs at about half real-time, while the plastic (PL) network shows a run time increase by a factor of ≈ 2.3 (Figure 5.4 A). When run on multiple cores in parallel both networks show comparable scaling and saturation behavior. On 64 cores the run time of the PL network is only about $T_{\text{PL}} \approx 1.6T_{\text{NP}}$. To investigate this in more detail we now plot the relative run time difference χ defined as

$$\chi = \frac{T_{\text{PL}} - T_{\text{NP}}}{T_{\text{NP}}} \quad (5.5)$$

as a function of the number of cores (Figure 5.4 B). This illustrates that there is an initial drop in χ for $n < 24$ cores after which the fractional time spent on plasticity starts to rise again slowly. The initial drop is presumably linked to the larger amount of cache becoming available with increasing the number of cores. The rise for higher numbers of cores is expected because Auryn integrates copies of all presynaptic traces on all processes continuously (Methods). This introduces redundant computation, which destroys the $\sim \frac{1}{n}$ scaling. Although the effect is small initially it becomes significant for large numbers of parallel processes. Although in the present case this only has a minor effect on the total run time, in some cases it can become advantageous to reduce this redundant computation. To that end Auryn supports to switch to event-based integration of the presynaptic traces (Methods), which causes χ to converge towards a constant value of roughly $\chi \approx 0.4$. Overall, both strategies allow the simulation of the plastic network at about $T \approx 0.4$.

We were wondering if the simulation of STDP can be sped up further and what impact the implementation details of STDP have on the run time. Auryn implements STDP in a very simple and straight forward way. Weight updates are carried out at every pre- or postsynaptic spike times in a way which only uses atomic operations. Other simulators, such as NEST use a different strategy, in which weight updates are carried out at the arrival times of presynaptic spikes only. Both approaches have their advantages and their shortcomings (Methods).

To compare the performance of both simulations without implementing new plasticity rules in NEST, we limited our study to code that already existed Gewaltig and Diesmann (2007); Gewaltig et al. (2012). In particular we used two example simulations that come with the current NEST release. Both simulations implement the same balanced random network model based on work by Nicolas Brunel Brunel (2000). In one case all connections are static whereas in the other case excitatory-to-excitatory connections evolve according to a weight dependent STDP rule (Methods).

We created the same network model in Auryn. For a better comparison we wrote an Auryn class with the same functionality as the NEST `poisson_generator` and implemented the same neuron model as used in NEST. Network simulations were run for 20s simulated time with low learn-




Figure 5.5 – Run times of a balanced network with plasticity in Auryn and NEST. (A) Scaling behavior of the relative run time T for a classical balanced network model Brunel (2000) when simulated in NEST (black) and Auryn (red). The measurement was done for the non-plastic (NP) network and with multiplicative STDP (PL). (B) Bar plot of the end point values in A. Single core (left), 32 cores (middle) and all 64 cores (right). Dashed line: real-time simulation. (C) Relative difference between run time of the plastic and non-plastic network for Auryn (red circles) and NEST (blue triangles).

ing rates to avoid that plasticity influences the firing statistics over the time course of the simulation. The resulting network activity of both implementations was comparable at $\approx 43\text{Hz}$ population firing rate. We ran both implementations distributed over four machines for varying numbers of cores. The respective run times of the PL and NP configuration were recorded (Figure 5.5 A), in which Auryn was integrating presynaptic traces using the time-continuous approach. On a single core the addition of STDP caused an increase of the run times by a factor of ≈ 2.5 for Auryn (≈ 4 for NEST, Figure 5.5 B, left). Both implementations show good initial scaling behavior, but saturate quickly (NP as well as PL configuration). When going from 32 to 64 cores run times only decreased marginally (Figure 5.5 B). Auryn achieved real-time simulation speed for the non-plastic (NP) configuration and took about twice as long when simulated with STDP (Figure 5.5 B). The plastic simulation in NEST was still significantly slower ($T \approx 5$) than PL simulations in Auryn. We repeated this measurement for Auryn using the event-based strategy for presynaptic trace update and did not find significant differences in run time (data not shown).

Overall this analysis shows that for medium size networks, Auryn’s simple approach to STDP is faster than the method NEST uses. At present it is not clear if or how STDP implementations can be made more efficient to speed up simulation speeds even further.

5.4 Discussion

In this paper we have shown that small or medium size recurrent networks with STDP can be simulated in, or close to, real-time if performance-optimized code is used. We compared the results achieved with our in-house code to NEST and found that the simulators show comparable scaling behavior when run in parallel on multiple cores and the same underlying integration scheme is used. In the case of parallel execution inter-process communications constitute the limiting factor that prevents simulations from running substantially faster than real-time. Importantly, when simulating plastic networks with different forms of STDP these results still hold. In particular we have shown that the overhead that is added to a network simulation by forms of STDP scales well and is nearly negligible when simulations are run in parallel.

These results suggest that plasticity studies in medium size networks of point neurons are limited intrinsically by communication delays of the inter-process communication, when run on off-the-shelf hardware. In the present study the best performance was achieved on a single powerful computer where communication proceeded via shared memory. However, it is not clear at what point the cost of an increased number of cores, which partly share cache and memory resources, will outweigh its advantages. Large clusters,

Chapter 5. Speed-up limits in network simulations

on the other hand, are only advantageous if they have low communication latencies which is likely to be only achievable with specialized hardware. A continuation of this study on such machines would be insightful.

When we interpret our results in the light of GPU computing (Richert et al., 2011; Hoang et al., 2013), we may conclude that GPUs can only provide an advantage if the communication between CPU and GPU can be held at a minimum. This requires that not only the neurons, but also the entire connectivity and plasticity to be implemented on the GPU. At present it is not clear yet if this can be done efficiently for a wide variety of plasticity rules.

In summary we have shown that real time simulations of plastic networks of point neurons are achievable with appropriate and highly optimized software. However, at the same time increasing simulation speed beyond $10\times$ real-time is challenging due to limitations in the inter-process communications.

Acknowledgments

The authors would like to thank A. Seeholzer and Moritz Deger for their helpful input on the manuscript and furthermore A. Seeholzer for the implementation of the forward Euler integrate-and-fire model for NEST.

Conclusion

In this work we have shown that different forms of synaptic plasticity and homeostasis serve different purposes and they can be coarsely classified according to the timescale they act on. Certain forms of Hebbian plasticity in the brain such as STDP can be rapidly induced. To quantify *rapid*, we used analytical methods to extract the effective timescales of plasticity from triplet STDP (Pfister and Gerstner, 2006) which quantitatively describes experimental data (Sjöström et al., 2001). By using methods from dynamical systems theory we concluded that the required timescale of a necessary compensatory mechanism needs to be faster than experimentally observed homeostatic mechanisms which are widely believed to be responsible for stabilizing Hebbian plasticity dynamics (Turrigiano et al., 1998; Abraham and Bear, 1996; Bienenstock et al., 1982).

Heterosynaptic plasticity seems a promising candidate compensatory mechanism (Chistiakova et al., 2014) and we have illustrated that spiking recurrent networks endowed with triplet STDP and heterosynaptic plasticity can successfully form and recall associative memories in form of cell assemblies in the face of ongoing plasticity.

Other forms of plasticity, such as Hebbian forms of inhibitory plasticity or other types of homeostatic plasticity, are intrinsically stable. They mostly have to act on a much longer timescale to avoid canceling synaptic changes which are desired. These processes seem to perform what is classically achieved in network models through manual parameter tuning. As we have shown in simulations, inhibitory plasticity for instance, finely tunes network dynamics to a point which would be difficult to achieve through manual tuning. It seems only plausible that biological systems use similar ways of efficiently controlling and refining their activity states in an on-line way, because it allows the systems to adapt to internal or external changes, which is not possible by simply relying on an initially frozen and finely tuned parameter set.

It is clear that in this work we could only scratch the surface of the wide range of experimentally observed (or unobserved) mechanisms involving plasticity and homeostasis. It will be an interesting avenue for future studies to in a first step form a more complete theoretical understanding of what the most important ingredients of plasticity mechanisms are to achieve

Conclusion

a given learning task. The second step should be to see how these theoretical requirements overlap with the rich underlying complexity of biological synapses, their internal dynamics and their ways of interacting with the network beyond the direct electrical signaling they relay.

First steps in this direction have been made from theoretical side (Clopath et al., 2008; Fusi et al., 2005) and it will be useful to interpret these results further within a network context which might necessitate the inclusion of more global signals and factors Turrigiano (2012).

Appendix A

Supplements to Chapter 2

A.1 Supplementary Figure S1




Figure A.1 – Evolution of the population rate for metaplastic triplet STDP model. (A) Temporal evolution of mean population rate for different values of τ ($\eta = 1$). While the change in stability in the vicinity of $\tau^{\text{crit}} = 166$ s can be understood from the mean field theory, which also predicts the observed oscillations at criticality, the late destabilization of the curve $\tau = 150$ s is not captured by the theory. (B) Evolution of mean population rate for $\eta = 1$. Black: $\tau = 150$ s and weights are initialized with the weights from a stable run ($\eta = 6.25$, $\tau = 10$ s) at $t = 24$ h. Cyan: Same, but with $\tau = 500$ s. Dark blue: $\tau = 150$ s, weight initialization as in (A), but maximally allowed weights limited to $w^{\max} = 0.5$. Light blue: $\tau = 500$ s, network falls silent at $t \approx 10$ h. Purple: $\tau = 500$ s, with $w_0 = 0.158$ (the learning rate was unchanged), which reduces the initial excursion to low rates.

A.2 Supplementary Text S1: Rate fluctuations

The approximation that κ is the fixed point of our plasticity rule is only true in a rate model where all rates take their respective mean value. Here we show that in the presence of rate fluctuations the actual fixed point is shifted to higher values. Since a neuron receives many inputs at any moment in time, we ignore fluctuations in the presynaptic rate, and focus on the postsynaptic side. First we analyze the effect on the synaptic weights of a single unit with fluctuating rate ν_i ($\langle \nu_i^2 \rangle = \nu^2 + \sigma^2$). We will furthermore neglect any fluctuations in $\bar{\nu}_i$. For the purely additive learning rule weights are stable on average if $\tau_w \langle \frac{dw}{dt} \rangle = 0$. Applying the temporal average to the rate based learning rule

$$\tau_w \frac{dw}{dt} \propto \nu \nu_i \left(\nu_i - \frac{\bar{\nu}_i^2}{\kappa} \right) \quad (\text{A.1})$$

and requiring it to be equal to zero gives

$$\nu \langle \nu_i^2 \rangle - \frac{\nu^3}{\kappa} \langle \nu_i \rangle = 0 \quad (\text{A.2})$$

$$\nu^2 + \sigma^2 - \frac{\nu^3}{\kappa} = 0 \quad (\text{A.3})$$

$$\nu = \kappa + \frac{\sigma^2}{\nu^2} \quad (\text{A.4})$$

where we assumed $\nu > 0$. Therefore the stable rate ν is always higher than κ . This intuition breaks down if τ is chosen too short. For the case $\tau\nu < 1$ fluctuations in $\bar{\nu}_i$ become substantial.

By ignoring both correlations and fluctuations in the mean field model systematic errors are introduced that underestimate the actual resulting target rate.

Appendix A. Supplements to Chapter 2

Appendix B

Supplements to Chapter 3

B.1 Annotated Protocol for Figure 3.2

The simulation protocol for Fig. 3.2 proceeded as follows:

Start: $t = -1$ min: The AI network dynamics of the original network Vogels and Abbott (2005) without inhibitory plasticity. This phase serves as a reference and is not shown.

4, A: $t = 0$ min: Inhibitory to excitatory synapses are turned to 0 efficacy. The network is forced out of the AI regime and begins to fire at high rates. Simultaneously, inhibitory plasticity is turned on.

4, B: $t = 60$ min: Inhibitory plasticity has restored AI activity.

4, C: $t = 60$ min, 5 s: The excitatory non-zero weights of the 2 designated memory patterns are increased ad-hoc by a factor of 5. The neurons of the subset begin to exhibit elevated and more synchronized activity.

4, D: $t = 120$ min: Inhibitory plasticity has successfully suppressed any elevated activity from the pattern and restored the global background state.

4, E: $t = 120$ min, 5 s: By delivering an additional, 1 s long stimulus as described above to 25% of the cells within one memory pattern, the whole pattern is activated. Activity inside the pattern stays asynchronous and irregular, and the rest of the network, including the other pattern, remains nearly unaffected.

Fig. S4 continues the protocol:

S4, F: $t = 120$ min, 11 s: The other (blue) memory pattern is activated with a stimulus analogous to the one used in Fig. B.6.




Figure B.1 – Inhibitory plasticity self-tunes balanced state in recurrent neural networks. **(a-b)** Each panel shows the instantaneous firing rate of the 10,000 network neurons laid out on a two dimensional grid. The 2,000 inhibitory neurons are arranged on the right hand side of the square. **(a)** Network in its initial state where half of the excitatory neurons receive strong inhibition (top) and the other half receives no inhibition (bottom). **(b)** Same as (a), but after only 10s of plasticity. **(c)** Same as (a,b) but after one hour of simulated time.

S4, G: $t = 120$ min, 17s: Both memory patterns respond with elevated AI activity to a stimulus to 25% of the cells, including the cells shared between both patterns.

B.2 Robustness to parameter changes

To estimate the sensitivity of our results to parameter changes we ran a range of additional simulations in which the initial state of the entire network simulation was altered. To that end we first repeated the simulation in Figure 3.2 (A), but this time only for half of the excitatory neurons inhibition was turned down while the other half received strong inhibition (Fig. B.1 (a)). While it took little time to decrease the firing rates of the neurons which were firing at initially high firing rates the turning down of inhibition took substantially longer (Fig. B.1 (b)). The end result after one hour of simulated time, however, remained unchanged. After one hour of simulated time all neurons fired at low average firing rate and the initial separation of the excitatory neurons was not evident from the activity any more (Fig. B.1 (c)).

To characterize how robust this effect was with respect to changes in the synaptic parameters we ran extensive parameter sweeps for which we numerically characterized the final network activity state after one hour of simulated time. To do so we acquired at the end of all our simulations the CV ISI and the standard deviation of the population rate σ_{Rate} , which we used as a computationally cheap proxy for synchrony in the network.

While the CV ISI is an efficient measure of the regularity of individual

B.3. Supplementary Figures

spike trains, the magnitude of the fluctuations of the population rate give insight about the synchrony of the local network activity. Taken together they allowed us to efficiently decide whether a given network or sub-population thereof can be characterized as in the asynchronous irregular (AI) regime (Brunel, 2000). We characterized a network as to be in the AI regime when the following condition was true:

$$(\text{CV ISI} > 1) \wedge (\sigma_{\text{Rate}} < 5\text{Hz}). \quad (\text{B.1})$$

The results of applying this criterion on various different parameters of synaptic conductances are shown in Figures B.2 and B.3.

We were wondering how strong the emergence of asynchronous irregular (AI) activity was dependent on the learning rate η used in the plasticity model. To that end we repeated the complete range of simulations shown in Figure 3.2 (A-E) for three different orders of magnitude of η . While this scaled the timescale of the balancing effect of inhibitory STDP to take place (Fig. B.4), no effect on the final network state was observed (data not shown).

B.3 Supplementary Figures

Appendix B. Supplements to Chapter 3




Figure B.2 – Emergence of asynchronous irregular activity is robust to changes in conductance variables. **(A)** Network statistics at the end of the learning protocol as shown in Fig. 3.2, as a function of the excitatory synaptic conductance g^E . All values are given in relative units with respect to the standard value $\bar{g} = 3nS$ and excitatory conductances onto excitatory and inhibitory cells are the same. Top row: Mean firing rate of excitatory sub-populations (firing rates are obtained from filtered population spike train; exponential filter with 1s time constant). Middle row: Mean CV ISI. Bottom row: Standard deviation of the population rate (5ms filter time constant). Columns signify the three representative sub-populations that activity was recorded from. No Pattern: Random network without memory pattern as in Fig. 3.2B. Pattern (OFF): Network after the introduction and balancing of a memory as in Fig. 3.2. Recorded neurons belong to the (black) control group in which connections have not been strengthened. Pattern (ON): As in Fig. 3.2D, recorded neurons belong to the red cell assembly. Simulated parameter combinations are indicated by black dots and are interpolated using natural neighbor interpolation. **(B)** Parameter region with AI activity. Right panels: Raster plots of 200 neurons each for three representative parameter combinations as indicated on the graph on the left.

B.3. Supplementary Figures

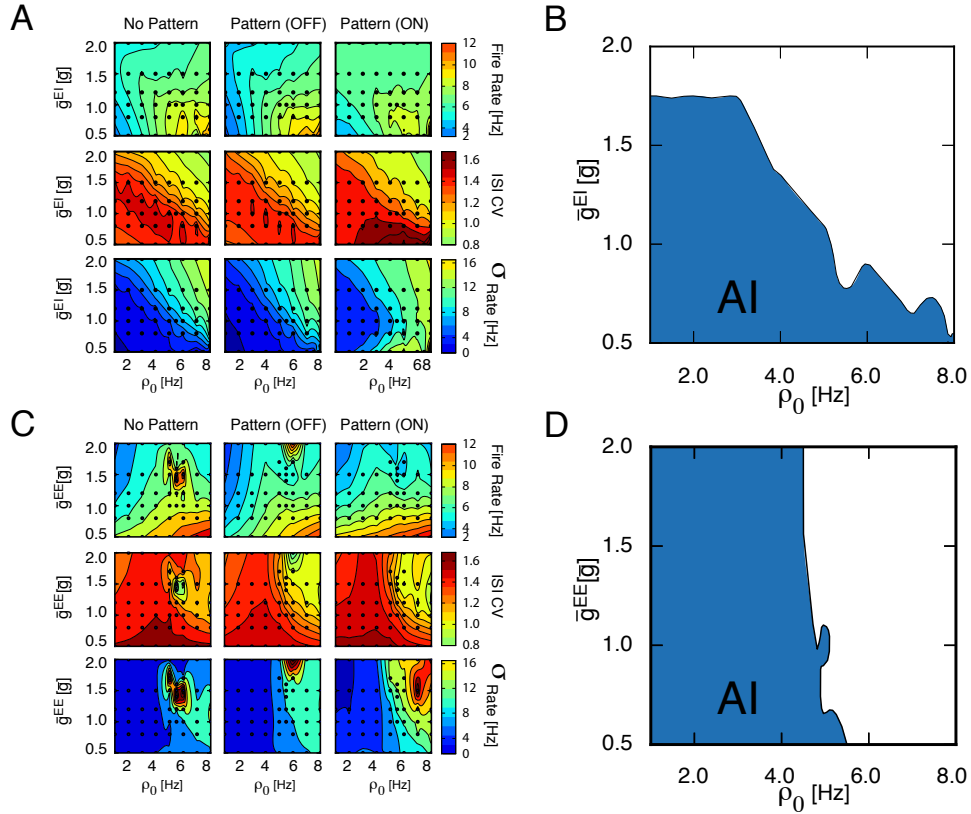


Figure B.3 – Varying excitatory synapses onto excitatory (\bar{g}^{EE}) and inhibitory neurons (\bar{g}^{EI}) independently does not disrupt emergence of AI activity for a wide range of parameters. (A) Same as Fig. B.2 but for \bar{g}^{EE} fixed. (B) AI state criterion (Eq. B.1) for (A). (C) Same as Fig. B.2 but with \bar{g}^{EI} held fixed. (D) AI state criterion (Eq. B.1) for (C).

Appendix B. Supplements to Chapter 3




Figure B.4 – Robustness to changes in the learning rate η . (A, B, C) Average population firing rate (1s bins) recorded from a set of background neurons (black) and from the red cell assembly (red) (as shown in Fig. 3.2). Top to bottom: Different learning rates η spanning three orders of magnitude. The black arrow indicates time of strengthening of synapses within the cell assemblies. Note rescaling of time axes.




Figure B.5 – Asynchronous network state. (a) Spike raster of excitatory cells over time. (b) Currents received by a single neuron in the network. Depolarizing current (red) and hyper polarizing current (blue). Net current including leak current (black).

B.3. Supplementary Figures




Figure B.6 – Cell assemblies can be recalled through an external cue. Continuation of Figure 3.2. (**F,G**) Top: Instantaneous cell activity (1s filter) of all network cells laid out on a 100×100 grid. Below: Spike raster of 50 cells in one of the cell assemblies (red, cf. Fig. 3.2 (A)) and of 50 background cells (black). Below: Histograms for spiking and other neuronal statistics (from top to bottom): 1) CV ISI. 2) Spiking correlations. 3) Synaptic conductance. 4) Mean firing rate during recall. (**F**) Recall of the blue assembly pattern (cf. Fig. 3.2 (A)). (**G**) Recall of both assemblies through stimulation of their overlap.




Figure B.7 – Supplementary statistics to accompany Fig. 3.2. Top row: Synaptic conductance. Bottom row: Firing rates. Red assembly neurons (red, top) and neurons from black control population (black, bottom).

B.4 Simulation parameters in tabular format

Table B.1 – Tabular description of network model (Fig. 3.2) following Nordlie et al. (2009).

A Model Summary	
Populations	Two: excitatory, inhibitory
Topology	—
Connectivity	Random all-to-all connections
Neuron model	Leaky integrate-and-fire, fixed voltage threshold, fixed absolute refractory time
Synapse model	conductance based inputs (exponentially decaying PSC)
Plasticity	Inhibitory plasticity
Input	Fixed input current to all units
Input (Recall)	Poisson spike trains to sub-population
Measurements	Spike activity

B Populations		
Name	Elements	Size
E	Iaf neuron	$N_E = 4N_I$
I	Iaf neuron	N_I

C Connectivity			
Name	Source	Target	Pattern
EE	E	E	Random with sparseness ϵ , weight \bar{g}^E ($\chi\bar{g}^E$ for connections in a pattern)
EI	E	I	Random with sparseness ϵ , weight \bar{g}^E
IE	I	E	Random with sparseness ϵ , weight plastic \bar{g}_{ij}^{IE}
II	I	I	Random with sparseness ϵ , weight \bar{g}^{II}

D1 Neuron and Synapse Model	
Name	Iaf neuron
Type	Leaky integrate-and-fire, exponential conductance based input
Subthreshold dynamics	$\tau \frac{dV_i}{dt} = (V^{\text{rest}} - V_i) + (g_i^E (V^E - V_i) + g_i^I (V^I - V_i) + I_b) \times \frac{1}{g_{\text{leak}}}$ $V(t) = V_{\text{rest}} \quad \text{otherwise}$
Synaptic dynamics	$\frac{dg_i^E(t)}{dt} = -\frac{g_i^E(t)}{\tau_E} + \bar{g}^E \delta(t - t_i^*)$ <p>and</p> $\frac{dg_i^I(t)}{dt} = -\frac{g_i^I(t)}{\tau_I} + \bar{g}_{ij}^{Ix} \delta(t - t_i^*)$
Spiking	If $V(t - dt) < \theta \wedge V(t) \geq \theta$ <ol style="list-style-type: none"> 1. set $t^* = t$ 2. emit spike with time-stamp t^*

Appendix B. Supplements to Chapter 3

D2 Plasticity Model	
Name	Inhibitory Spike Timing Dependent Plasticity (iSTDP)
Type	Symmetric iSTDP with a constant offset for presynaptic spikes
Acts on	IE
Synaptic traces	$\frac{dx_i}{dt} = -\frac{x_i}{\tau_{\text{STDP}}} + \delta(t - t_i^*)$
Online rule	$W_{ij} = W_{ij} + \eta(x_i - \alpha)$ <p style="text-align: center;">for presynaptic spikes of neuron j at time t_f^j</p> <p>and $W_{ij} = W_{ij} + \eta x_j$ <p style="text-align: center;">for postsynaptic spikes at time t_f^i</p> </p>

E Input	
Type	Description
Current input	Fixed current I to all neurons

Table B.2 – Simulation parameter summary for the network model (Fig. 3.2) following Kunkel et al. (2011).

Populations		
Name	Value	Description
N_E	8000	Size of excitatory population E
N_I	2000	Size of inhibitory population I

Connectivity		
Name	Value	Description
ϵ	0.02	Probability of any connection (EE,EI,IE,II)
\bar{g}	3 nS	Basic weight unit
\bar{g}^E	\bar{g}	Weight of all excitatory synapses (= 3 nS)
γ	10	Scaling factor for inhibitory weights
\bar{g}^{II}	$\gamma\bar{g}$	Weight of inhibitory to inhibitory synapses (= 30 nS)
\bar{g}_{ij}^{IE}	$W_{ij}\bar{g}^{II}$	Weight of inhibitory to excitatory synapses (= $W_{ij} \times 30$ nS)
χ	5	Potential factor of excitatory weights belonging to one or more patterns

Neuron Model		
Name	Value	Description
τ	20 ms	Membrane time constant
Θ	-50 mV	Spiking threshold
V^{rest}	-60 mV	Resting potential
V^E	0 mV	Excitatory reversal potential
V^I	-80 mV	Inhibitory reversal potential
g^{leak}	10 nS	Leak conductance
I_b	200 pA	Background current to each cell (unless stated otherwise)
τ_{ref}	5 ms	Absolute refractory period

Synapse Model		
Name	Value	Description
τ_E	5 ms	Decay constant of AMPA-type conductance
τ_I	10 ms	Decay constant of GABA-type conductance

Plasticity Model		
Name	Value	Description
τ_{STDP}	20 ms	Decay constant of (pre and post) synaptic trace
η	1×10^{-4}	Learning rate
α	0.12	Presynaptic offset
W_{min}	0	Minimum inhibitory synaptic weight
W_{max}	$10\bar{g}^{II}$	Maximum inhibitory synaptic weight

Appendix B. Supplements to Chapter 3

Appendix C

Supplements to Chapter 4

C.1 Supplementary figures




Figure C.1 – Firing rates are bistable under reset triplet and depend on the number of presynaptic contacts. (a) Mean firing rate over time of a 100 postsynaptic integrate-and-fire neurons receiving 100 random spike trains from a population of 1000 Poisson processes (10Hz) via synapses with triplet STDP and heterosynaptic plasticity. Different colors signify three different initial conditions of the afferent synaptic weights. (b) End point firing rates at $t = 600s$ vs initial firing rates ($t = 10s$) from (a). (c) Same as (a) but for a single initial weight value and varying numbers of connections (fixed connection probability between presynaptic Poisson neurons and postsynaptic cell of 10%). (d) End point firing rates at $t = 600s$ vs initial firing rates ($t = 10s$) from (a). (e) End point firing rates at $t = 600s$ vs number of afferent connections for the cluster at low rates in (d). Pearson’s correlation coefficient is denoted by letter r . (f) End point firing rates at $t = 600s$ vs number of afferent connections for the cluster at high rates in (d).




Figure C.2 – Assembly structure does not reveal itself immediately from the heterogeneous excitatory-to-excitatory weight matrix. (a) Small section of the excitatory-to-excitatory weight matrix for 200 neurons from simulation in Figure (C.3) at $t = 2h$. Point size indicates relative strength of connections (arbitrary units). (b) Histogram of weight strength w in the excitatory weight matrix. (c) Histogram of reference weight strength \tilde{w} . (d) Scatter plot of reference weight values \tilde{w} vs synaptic weight w .

Appendix C. Supplements to Chapter 4




Figure C.3 – Unstructured input weights and homeostatic regulation of LTD rate yields formation of cell assemblies and working memory. **(a)** Top panel: Spike raster of network activity after 2h of learning. Stimulation times and stimulus identity are indicated by colored bars at the top of the plot. Bottom panel: Population firing rate of cell assemblies corresponding to the four stimuli. **(b)** Histograms of network spike statistics acquired during initial activity shown in (a). Top: Firing rate, bottom: CV ISI. **(c)** Like (a), with external stimulation enabled. **(d)** Same as (a,c) after about one hour of simulation. **(e)** Same as (d) but for interval marked in (d). **(f)** Same as (d) but with increased inter-stimulus-interval ($T_{\text{off}} = 20\text{s}$). **(g)** Histogram showing the number of stimulus presentations at $t = 1\text{h}$. **(h)** Network neurons laid out on 2D grid with stimulus preference indicated in color. **(i)** Bar plot displaying the fraction of neurons coding for a specific stimulus (black: no stimulus preference). **(j)** Overlap between all pairs of stimulus activity. **(k)** Overlap of stimulus evoked network activity states. **(l)** Receptive fields of 10 randomly selected neurons within the network. **(m)** Mean weight strength between neurons according to their preferred stimulus.

C.1. Supplementary figures




Figure C.4 – Network states are associative and serve as working memory in the network with homeostatic metaplasticity. (a) Top icons: Distorted cues fed into the network (black bars below that code for time and duration of the stimulation). Middle panel: Spike raster of network activity (only every fourth spike of every fourth neuron shown for clarity). Bottom panel: Population rate of the four acquired network states. (b) Population averaged peristimulus time histograms of the four relevant readout populations for all different distorted stimuli presented (keys on the right).




Figure C.5 – Failure modes arising from blocking single plasticity mechanisms. (a,b) Top left: spike raster (all excitatory neurons) with stimulation interval and stimulus identity on top (colored bars). Bottom left: Activity in putative memory patterns as determined from Figure 4.4. Histograms (from top left to bottom right): 1) Neural firing rates determined over the interval shown in the raster plot. 2) Relative proportion of active (at least one spike during a one hour interval) and silent neurons. 3) CV ISI over entire interval shown in raster plot on the left. 4) Excitatory synaptic weight distribution as determined at the end of the simulation (after 300s for (a), 2h otherwise). (a) Blocking of homeostatic regulation of LTD rate in a network without initially structured afferent connections (cf. C.3). The network fails to develop selectivity. (b) Blocking heterosynaptic plasticity causes rapidly increasing firing rates without learning or delay activity (simulation stopped after 300s).




Figure C.6 – Failure modes arising from blocking single plasticity mechanisms. Same panels as in Fig. C.5. **(a)** Blocking inhibitory plasticity still allows the network to develop delay activity. Assembly size and steady state population firing rate depends on the fixed inhibitory weight (not shown). **(b)** Blocking of transmitter-triggered plasticity causes many neurons to fall and remain silent.




Figure C.7 – A slow adaptation current causes spontaneous switching between attractors. (a) Spike raster of 1024 excitatory neurons with a slow spike triggered adaptation current with time a time constant of 20s (cf. Pozzorini et al. (2013) and methods). At around 139min the currently active assembly (green) is switched by a brief external stimulus (red). All other state changes are spontaneous. (b) Histogram of firing rates averaged over the interval marked with a black bar in (a). (c) Histogram of coefficient of variation of the inter-spike-interval distribution (CV ISI) averaged over the interval marked in (a).

C.2 Tabular network description

Table C.1 – Tabular description of network model (Nordlie et al., 2009).

A Model Summary	
Populations	Three: excitatory, inhibitory, Poisson input
Topology	—
Connectivity	Random sparse connectivity (fixed probability)
Neuron model	Leaky integrate-and-fire, relative and absolute refractory period, spike triggered adaptation
Synapse model	Conductance based, exponentially decaying (AMPA, GABA), and double exponential (NMDA)
Plasticity	—
Input	Poisson input with additional rate coded input patterns
Measurements	Spike activity, synaptic efficacy (continuously of a subset of synapses and of all synapses at the end of a simulation)

B Populations		
Name	Size	Elements
E	4096	Adaptive IF neuron
I	1024	IF neuron
Stim	4096	Poisson stimulus population

C Connectivity			
Name	Source	Target	Pattern
EE	E	E	Fixed prob. $\epsilon = 0.1$, $w = 0.1$, excitatory, plastic
EI	E	I	Fixed prob. $\epsilon = 0.1$, $w = 0.2$, excitatory
IE	I	E	Fixed prob. $\epsilon = 0.1$, $w = 0.1$, inhibitory, plastic
II	I	I	Fixed prob. $\epsilon = 0.1$, $w = 0.2$, inhibitory
StimE	Stim	E	Fixed prob. $\epsilon = 0.05$, $w = 0.2$, excitatory, plastic
			For all above: Fixed delay $D = 0.8\text{ms}$

Appendix C. Supplements to Chapter 4

D1 Neuron and Synapse Model	
Name	AIF neuron
Type	Adaptive leaky integrate-and-fire, absolute and relative refractoriness, conductance input, spike triggered adaptation, AMPA and GABA-like conductances exponentially decaying, double exponential NMDA-like conductance
Subthr. dynamics	$\tau^m \frac{d}{dt} U_i = (U^{\text{rest}} - U_i) + g_i^{\text{exc}}(t)(U^{\text{exc}} - U_i) + g_i^{\text{inh}}(t)(U^{\text{inh}} - U_i)$
Cond.	$g_i^{\text{exc}}(t) = \alpha g_i^{\text{ampa}}(t) + (1 - \alpha) g_i^{\text{nmda}}(t)$ $g_i^{\text{inh}}(t) = g_i^{\text{gaba}}(t) + \underbrace{g_i^{\text{a}}(t) + g_i^{\text{b}}(t)}_{\text{Adaptation}}$
Excitation	$\frac{d}{dt} g_i^{\text{ampa}} = -\frac{g_i^{\text{ampa}}}{\tau^{\text{ampa}}} + \sum_{j \in \text{exc}} x_j(t) u_j(t) w_{ij} S_j(t)$ for the definition of short term plasticity variables x_j and u_j see D2 below. $\tau^{\text{nmda}} \frac{d}{dt} g_i^{\text{nmda}} = -g_i^{\text{nmda}} + g_i^{\text{ampa}}$
Inhibition	$\frac{d}{dt} g_i^{\text{inh}} = -\frac{g_i^{\text{inh}}}{\tau^{\text{gaba}}} + \sum_{j \in \text{inh}} w_{ij} S_j(t)$
Adaptation	$\frac{d}{dt} g_i^{\text{a}}(t) = -\frac{g_i^{\text{a}}}{\tau^{\text{a}}} + \Delta^{\text{a}} S_i(t)$ $g_i^{\text{b}}(t) = 0$ except in simulations with adaptation on multiple timescales: $\frac{d}{dt} g_i^{\text{b}}(t) = -\frac{g_i^{\text{b}}}{\tau^{\text{b}}} + \Delta^{\text{b}} S_i(t)$
Threshold	$\tau^{\text{thr}} \frac{d}{dt} \vartheta_i(t) = -\vartheta_i(t) + \vartheta^{\text{rest}}$
Spiking	If $U_i > \vartheta_i(t)$ then $S_i(t) \rightarrow S_i(t) + \delta(t)$ (emit a spike) and $\vartheta_i(t) \rightarrow +50\text{mV}$ and $U_i \rightarrow U^{\text{rest}}$
D2 Plasticity Model	
Name	Short term plasticity following Tsodyks and Markram (1997)
Type	Depressing and facilitating short term dynamics
Acts on	EE, EI, Stim, E
Dynamics	$\frac{d}{dt} x_j(t) = \frac{1 - x_j(t)}{\tau^{\text{d}}} - u_j(t) x_j(t) S_j(t)$ $\frac{d}{dt} u_j(t) = \frac{U - u_j(t)}{\tau^{\text{f}}} + U (1 - u_j(t)) S_j(t)$

Name	IF neuron
Type	The non-adaptive inhibitory neurons were implemented identically except of the omission of the adaptation terms $g_i^{\text{a}} = g_i^{\text{b}} = 0$.

C.2. Tabular network description

D3 Plasticity Model	
Name	Plasticity rule for excitatory synapses
Type	STDP rule with pre and post offset terms and slower consolidation dynamics
Acts on	EE and StimE
Synaptic traces	$\frac{d}{dt} z_j^{\text{pre}}(t) = -\frac{z_j^{\text{pre}}(t)}{\tau^{\text{pre}}} + S_j(t)$ with presynaptic spike train $S_j(t)$. $\frac{d}{dt} z_i^{\text{post}}(t) = -\frac{z_i^{\text{post}}(t)}{\tau^{\text{post}}} + S_i(t)$ with postsynaptic spike train $S_i(t)$. $\frac{d}{dt} z_i^{\text{post2}}(t) = -\frac{z_i^{\text{post2}}(t)}{\tau^{\text{post2}}} + S_i(t)$
Online rule	$\frac{d}{dt} w_{ij} = S_i(t) \left(A z_j^{\text{pre}} z_i^{\text{post2}}(t) - \beta (z_i^{\text{post}}(t))^3 (w_{ij} - \tilde{w}_{ij}) \right) - S_j(t) (B z^{\text{post}}(t)_i - \delta)$
Consolidation	$\tau^{\text{cons}} \frac{d}{dt} \tilde{w}_{ij} = -\tilde{w}_{ij} + w_{ij} - P \tilde{w}_{ij}(t) \left(\frac{w^{\text{p}}}{2} - \tilde{w}_{ij}(t) \right) (w^{\text{p}} - \tilde{w}_{ij}(t))$ <p>which is integrated with step size of 1.2s.</p>
Meta-plasticity	<p>Where this is mentioned explicitly we allow B to have a slow time dependence (cf. Bienenstock et al. (1982); Pfister and Gerstner (2006)):</p> $B_i(t) = \begin{cases} A C_i(t) & \text{for } C_i(t) \leq 1 \\ A & \text{otherwise} \end{cases}$ <p>with</p> $\frac{d}{dt} C_i(t) = -\frac{C_i(t)}{\tau^{\text{hom}}} + (z_i^{\text{ht}}(t))^2$

D4 Plasticity Model	
Name	Inhibitory Spike Timing Dependent Plasticity (iSTDP)
Type	Symmetric iSTDP with a constant offset for presynaptic spikes and a global modulation
Acts on	IE
Synaptic traces	$\frac{d}{dt} z_i(t) = -\frac{z_i(t)}{\tau^{\text{iSTDP}}} + S_i(t)$
Secreted factor	$G(t) = H(t) - \alpha$ with $\frac{d}{dt} H(t) = -\frac{H(t)}{\tau^{\text{H}}} + \sum_{i \in \text{exc}} S_i(t)$
Online rule	$\frac{d}{dt} w_{ij}(t) = \eta G(t) ((z_i(t) + 1) S_j(t) + z_j(t) S_i(t))$

Appendix C. Supplements to Chapter 4

E		Input
Name	Stimulus group	
Type	Rate coded Poisson input	
Size	4096 Poisson neurons	
Firing rates	$\nu_i = \nu^{\text{bg}} + \nu^{\text{scale}} \xi_i(t)$	
Input pattern	$\xi_i(t) = \begin{cases} 0 & \text{no active input pattern} \\ \xi_i^\mu & \text{if pattern } \mu \text{ is active} \end{cases}$ <p>where the ξ_i^μ are flattened 64×64 grayscale images normalized to $[0,1]$. Only a single pattern can be active at a time. A pattern stays active during a finite period drawn from an exponential distribution with mean T^{On}. Each pattern activation is followed by a period of inactivity with a duration drawn from another exponential distribution with mean T^{Off}.</p>	
F		Measurements
Type	Description	
	Spike activity for raster and activity plots and weight matrices at the end of the simulation.	

Table C.2 – Simulation parameter summary of network model.

Populations		
Name	Value	Description
N_E	4096	Size of excitatory population E
N_I	1024	Size of inhibitory population I
N_{Stim}	4096	Size of external population Stim

Connectivity		
Name	Value	Description
ϵ	0.1	Probability of connection (EE,EI,IE,II)
ϵ^{Stim}	0.05	Probability of connection (StimE)
\tilde{w}	0.0	Initial reference weight for all connections
w^{EE}	0.1	Initial excitatory weight
w^{EI}	0.2	Excitatory weight
w^{IE}	0.15	Initial inhibitory weight
w^{II}	0.2	Inhibitory weight
w^{StimE}	0.5	Initial excitatory weight

Neuron Model		
Name	Value	Description
τ	20 ms	Membrane time constant
U^{rest}	-60 mV	Resting potential
U^{exc}	0 mV	Excitatory reversal potential
U^{inh}	-80 mV	Inhibitory reversal potential
α^E	0.2	AMPA/NMDA ratio (excitatory population)
α^I	0.3	AMPA/NMDA ratio (inhibitory population)
τ^{ampa}	5ms	AMPA decay time constant
τ^{gaba}	10ms	GABA decay time constant
τ^{nmda}	100ms	NMDA decay time constant
τ^a	100ms	Adaptation time constant
Δ^a	0.1	Adaptation strength
τ^b	20s	Slow adaptation time constant
Δ^b	5×10^{-4}	Slow adaptation strength
τ^{thr}	2ms	Threshold time constant
ϑ^{rest}	-50mV	Threshold resting value

Appendix C. Supplements to Chapter 4

Short term plasticity model (excitatory synapses)		
Name	Value	Description
τ^d	200ms	Depression time constant
τ^f	600ms	Facilitation time constant
U	0.2	Initial release probability parameter

Plasticity Model (excitatory synapses)		
Name	Value	Description
τ^{pre}	20ms	Presynaptic trace for excitatory plasticity
τ^{post}	20ms	Postsynaptic trace for excitatory plasticity
$\tau^{\text{post}2}$	100ms	Slow postsynaptic trace for excitatory plasticity
A	1×10^{-3}	LTP rate
B	1×10^{-3}	LTD rate (in simulations without metaplasticity)
δ	2×10^{-5}	Transmitter triggered plasticity strength
β	0.05	Heterosynaptic plasticity strength parameter
τ^{cons}	20min	Consolidation time constant
w^P	0.5	Upper fixed point of reference weight potential
P	20	Potential strength parameter
τ^{hom}	20min	Metaplasticity time constant (only designated simulations with metaplasticity)
τ^{ht}	100ms	Activity trace time constant for metaplasticity

Plasticity Model (inhibitory synapses)		
Name	Value	Description
τ^{ISTDP}	20ms	STDP trace time constant
γ	4Hz	Target population rate
τ^H	10s	Secreted factor time constant
η	2×10^{-5}	Learning rate

Stimulus Model		
Name	Value	Description
ν^{bg}	10Hz	Background firing rate
ν^{scale}	35Hz	Maximum rate increase
T^{On}	1s	Mean stimulus-on period
T^{Off} $t < 1\text{h}$	2s	Mean stimulus-off period
T^{Off} $t \geq 1\text{h}$	20s	Mean stimulus-off period (unless stated otherwise)

C.3 Derivation of the moments of a synaptic trace for Poisson firing statistics

We can describe the probability flux of a synaptic trace z defined by the ordinary differential equation

$$\frac{dz}{dt} = -\frac{z}{\tau} + aS(t) \quad (\text{C.1})$$

with time constant τ and spike train $S(t)$ by the following partial differential equation

C.3. Derivation of the moments of a synaptic trace for Poisson firing statistics

n	$\langle z^n \rangle$
1	$\lambda\tau$
2	$\frac{1}{2}\lambda\tau(1+2\lambda\tau)$
3	$\frac{1}{6}\lambda\tau(2+9\lambda\tau+6\lambda^2\tau^2)$
4	$\frac{1}{12}\lambda\tau(3+25\lambda\tau+36\lambda^2\tau^2+12\lambda^3\tau^3)$
5	$\frac{1}{60}\lambda\tau(12+175\lambda\tau+425\lambda^2\tau^2+300\lambda^3\tau^3+60\lambda^4\tau^4)$

Table C.3 – The first five moments of a synaptic trace under the assumption of Poisson firing with rate λ where we assumed a jump size $a = 1$.

$$\partial_t p(z, t) = \frac{1}{\tau} \partial_z (z p(z, t)) - \lambda p(z, t) + \lambda p(z - a, t) \quad (\text{C.2})$$

$$= \frac{1}{\tau} p(z, t) + \frac{z}{\tau} \partial_z p(z, t) - \lambda p(z, t) + \lambda p(z - a, t) \quad (\text{C.3})$$

where the terms involving λ describe the sink and source terms that capture the jumps of size a caused by spikes at rate λ .

By now requiring stationary $\partial_t p = 0$ we have

$$z \frac{d}{dz} p(z) = -p(z) + \lambda\tau p(z) - \lambda\tau p(z - a) \quad (\text{C.4})$$

We now compute the Laplace transform of the left hand side of Eq. (C.4):

$$\int_0^\infty dz e^{-sz} z \frac{dp}{dz} = [ze^{-sz} p]_0^\infty - \int_0^\infty dz \frac{d}{dz} (e^{-sz} z) p \quad (\text{C.5})$$

$$= - \int_0^\infty dz p (e^{-sz} - sz e^{-sz}) \quad (\text{C.6})$$

$$= -\bar{p}(s) + s \int_0^\infty dz p z e^{-sz} \quad (\text{C.7})$$

$$= -\bar{p}(s) - s \frac{d}{ds} \bar{p}(s) \quad (\text{C.8})$$

which we now insert into Eq. (C.4):

$$-\bar{p} - s \frac{d\bar{p}}{ds} = -\bar{p} + \lambda\tau\bar{p} - \lambda\tau e^{-sa}\bar{p} \quad (\text{C.9})$$

$$\frac{d\bar{p}}{ds} = -\lambda\tau \frac{1 - e^{-sa}}{s} \bar{p} \quad (\text{C.10})$$

and therefore $\bar{p}(s) = \exp\left(-\lambda\tau \int_0^s \frac{1 - e^{-s'a}}{s'} ds'\right)$. The moments μ_n of z are then given by

$$\mu_n = (-1)^n \lim_{s \rightarrow 0} \frac{d^n}{ds^n} \bar{p}(s) \quad (\text{C.11})$$

Appendix C. Supplements to Chapter 4

Appendix D

Supplements to Chapter 5

D.1 Network details for Vogels-Abbott Benchmark

Table D.1 – Tabular description for Vogels-Abbott Benchmark following Nordlie et al. (2009).

A Model Summary	
Populations	Three: excitatory, inhibitory, Poisson input
Topology	—
Connectivity	Random sparse connections
Neuron model	Leaky integrate-and-fire, fixed voltage threshold, fixed absolute refractory time (voltage clamp)
Synapse model	Exponentially decaying AMPA and GABA conductances and δ -conductance pulses (discontinuous jumps)
Plasticity	—
Input	Initial poisson input (50ms) to prime/start the network
Measurements	Spike activity

B Populations		
Name	Size	Elements
E	3200	IF neuron
I	800	IF neuron
Ext	200	Poisson neurons

C Connectivity			
Name	Source	Target	Pattern
EE	E	E	Random sparse P, probability of connection $\epsilon = 0.02$ weight $w = 0.4$, excitatory
EI	E	I	Random sparse ($\epsilon = 0.02$), $w = 5.1$, inhibitory
IE	I	E	Random sparse ($\epsilon = 0.02$), $w = 0.4$, excitatory
II	I	I	Random sparse ($\epsilon = 0.02$), $w = 5.1$, inhibitory
Ext	Ext	E	Random sparse ($\epsilon = 0.02$), $w = 0.4$, excitatory
			All connections have a delay $D = 0.8$ ms

Appendix D. Supplements to Chapter 5

D1 Neuron and Synapse Model	
Name	IF neuron
Type	Leaky integrate-and-fire, exponential conductance based input
Subthreshold dynamics	<p>If not refractory ($t > t^* + \tau_{\text{ref}}$):</p> $\tau^m \frac{dU_i}{dt} = (U^{\text{rest}} - U_i) + g_i^{\text{exc}}(t)(U^{\text{exc}} - U_i) + g_i^{\text{inh}}(t)(U^{\text{inh}} - U_i)$ <p>else:</p> $U_i = U_{\text{rest}}$
Synaptic dynamics	$\frac{dg_i^{\text{exc}}}{dt} = -\frac{g_i^{\text{exc}}}{\tau^{\text{exc}}} + \sum_{j \in \text{exc}} w_{ij} S_j(t)$ <p>and</p> $\frac{dg_i^{\text{inh}}}{dt} = -\frac{g_i^{\text{inh}}}{\tau^{\text{inh}}} + \sum_{j \in \text{inh}} w_{ij} S_j(t)$
Spiking	If $U_i > \Theta \wedge t > t^* + \tau_{\text{ref}}$ set $t^* = t$

E Input	
Type	Description
Poisson neurons	Fixed homogenous rate $\nu_{\text{ext}} = 10$

F Measurements	
Type	Description
	Spike activity for raster plots. Run time of simulation (excluding time for setup and saving data).

Table D.2 – Simulation parameter summary for B.1 Network details for Vogels-Abbott Benchmark following Kunkel et al. (2011).

Populations		
Name	Value	Description
N_E	3200	Size of excitatory population E
N_I	800	Size of inhibitory population I
N_{Ext}	200	Size of Poisson population Ext

Connectivity		
Name	Value	Description
ϵ	0.02	Probability of any connection (EE,EI,IE,II)
w^{EE}	0.4	Weight excitatory→excitatory
w^{EI}	0.4	Weight excitatory→inhibitory
w^{IE}	5.1	Weight inhibitory→excitatory
w^{II}	5.1	Weight inhibitory→inhibitory
w^{Ext}	0.4	Weight excitatory→excitatory

Neuron Model		
Name	Value	Description
τ	20 ms	Membrane time constant
Θ	-50 mV	Spiking threshold
U^{rest}	-60 mV	Resting potential
U^E	0 mV	Excitatory reversal potential
U^I	-80 mV	Inhibitory reversal potential
τ_{ref}	5 ms	Absolute refractory period

Synapse Model		
Name	Value	Description
τ_E	5 ms	Decay constant of AMPA-type conductance
τ_I	10 ms	Decay constant of GABA-type conductance

D.2 Network details for Brunel network

Table D.3 – Tabular description for Brunel network following Nordlie et al. (2009).

A Model Summary	
Populations	Three: excitatory, inhibitory, Poisson input
Topology	—
Connectivity	Random convergent connections
Neuron model	Leaky integrate-and-fire, fixed voltage threshold, fixed absolute refractory time (voltage clamp)
Synapse model	δ -current pulses (discontinuous voltage jumps)
Plasticity	—
Input	Independent Poisson input
Measurements	Spike activity

B Populations		
Name	Size	Elements
E	8000	IF neuron
I	2000	IF neuron
Ext		Poisson generator

C Connectivity			
Name	Source	Target	Pattern
EE	E	E	Random convergent ($C_E \rightarrow 1$), $w = 0.1\text{mV}$
EI	E	I	Random convergent ($C_E \rightarrow 1$), $w = 0.1\text{mV}$
IE	I	E	Random convergent ($C_I \rightarrow 1$), $w = -0.5\text{mV}$
II	I	I	Random convergent ($C_I \rightarrow 1$), $w = -0.5\text{mV}$
			For all above: $C_E = 800$, $C_I = 200$ Delay $D = 0.8\text{ms}$

D.2. Network details for Brunel network

D Neuron and Synapse Model	
Name	IF neuron
Type	leaky integrate-and-fire, exponentially decaying conductance input
Subthr. dynamics	<p>If not refractory ($t > t^* + \tau_{\text{ref}}$):</p> $\frac{dU_i}{dt} = \frac{1}{\tau^m} (U^{\text{rest}} - U_i) + w_{\text{xE}} \sum_{j \in \text{E}} S_j(t) + w_{\text{xI}} \sum_{j \in \text{I}} S_j(t)$ <p>else: $U_i = U_{\text{rest}}$</p>
Spiking	If $U_i > \Theta \wedge t > t^* + \tau_{\text{ref}}$ set $t^* = t$
STDP (only simulations with plasticity)	$\frac{dw(t)}{dt} = \lambda \left(1 - \frac{w(t)}{w_{\text{max}}} \right) z_j^+(t) S_i(t) - \lambda \alpha \frac{w(t)}{w_{\text{max}}} z_i^-(t) S_j(t)$ <p>if $w_{ij} < 0$ then $w_{ij} \rightarrow 0$ if $w_{ij} > w_{\text{max}}$ then $w_{ij} \rightarrow w_{\text{max}}$</p>
Synaptic traces	$\frac{dz_n^x}{dt} = -\frac{z_n^x}{\tau^x} + S_n(t)$
E Input	
Type	Description
Poisson generator	Independent Poisson input at fixed rate $\nu_{\text{ext}} = 20\text{kHz}$
F Measurements	
Type	Description
	Spike activity for raster plots. Run time of simulation (excluding time for setup and saving data).

Table D.4 – Simulation parameter summary for B.1 Network details for Brunel network following Kunkel et al. (2011).

Populations		
Name	Value	Description
N_E	8000	Size of excitatory population E
N_I	2000	Size of inhibitory population I

Connectivity		
Name	Value	Description
ϵ	0.1	Probability of any connection (EE,EI,IE,II)
w^{EE}	0.1mV	excitatory→excitatory
w^{EI}	0.1mV	excitatory→inhibitory
w^{IE}	-0.5mV	inhibitory→excitatory
w^{II}	-0.5mV	inhibitory→inhibitory
w^{ExtE}	0.1mV	Poisson Generator→excitatory
w^{ExtI}	0.1mV	Poisson Generator→inhibitory

Neuron Model		
Name	Value	Description
τ	20 ms	Membrane time constant
Θ	-50 mV	Spiking threshold
U^{rest}	-60 mV	Resting potential
U^E	0 mV	Excitatory reversal potential
U^I	-80 mV	Inhibitory reversal potential
τ_{ref}	2 ms	Absolute refractory period

Plasticity Model		
Name	Value	Description
λ	1×10^{-9}	Learning rate
α	2.02	Relative strength of LTD
w_{min}	0.0mV	Minimum weight
w_{max}	0.3mV	Maximum weight
$\tau^+ = \tau^-$	20ms	STDP time constant

D.3 Network details for Plastic 25,000 cell network at 3Hz

For a more detailed description see Chapter 2 or Zenke et al. (2013).

Table D.5 – Tabular description for 25k cell network Nordlie et al. (2009).

A Model Summary	
Populations	Three: excitatory, inhibitory, Poisson input
Topology	—
Connectivity	Random sparse connections
Neuron model	Leaky integrate-and-fire, moving voltage threshold, absolute and relative refractory time)
Synapse model	Exponentially decaying AMPA, GABA and a slow NMDA conductances. Synaptic input produces δ -conductance pulses (discontinuous jumps) in AMPA and GABA conductances.
Plasticity	—
Input	Poisson input from a population of input units
Measurements	Spike activity

B Populations		
Name	Size	Elements
E	20000	IF neuron
I	5000	IF neuron
Ext	2500	Poisson neurons

C Connectivity			
Name	Source	Target	Pattern
EE	E	E	Fixed probability ϵ , $w = 0.16$, excitatory
EI	E	I	Fixed probability ϵ , $w = 0.16$, excitatory
IE	I	E	Fixed probability ϵ , $w = 1$, inhibitory
II	I	I	Fixed probability ϵ , $w = 1$, inhibitory
Ext	Ext	E	Fixed probability ϵ , $w = 0.16$, excitatory
			For all connections: delay $D = 0.8\text{ms}$, probability of connection $\epsilon = 0.05$

Appendix D. Supplements to Chapter 5

D Neuron and Synapse Model	
Name	IF neuron
Type	leaky integrate-and-fire, exponentially decaying conductance input from AMPA and GABA conductances plus a slowly rising and NMDA current with 100ms decay time (voltage dependence ignored)
Subthr. dynamics	$\tau^m \frac{dU_i}{dt} = (U^{\text{rest}} - U_i) + g_i^{\text{exc}}(t)(U^{\text{exc}} - U_i) + g_i^{\text{inh}}(t)(U^{\text{inh}} - U_i)$
Cond. dynamics	$g_i^{\text{exc}}(t) = \alpha g_i^{\text{ampa}}(t) + (1 - \alpha) g_i^{\text{nmda}}(t)$
Excitation	$\frac{dg_i^{\text{ampa}}}{dt} = -\frac{g_i^{\text{ampa}}}{\tau^{\text{ampa}}} + \sum_{j \in \text{exc}} w_{ij} S_j(t)$ $\tau^{\text{nmda}} \frac{dg_i^{\text{nmda}}}{dt} = -g_i^{\text{nmda}} + g_i^{\text{ampa}}$
Inhibition	$\frac{dg_i^{\text{inh}}}{dt} = -\frac{g_i^{\text{inh}}}{\tau^{\text{gaba}}} + \sum_{j \in \text{inh}} w_{ij} S_j(t)$
Thr. dyn.	$\tau^{\text{thr}} \frac{d\vartheta_i}{dt} = \vartheta^{\text{rest}} - \vartheta_i$
Spiking	If $U_i > \vartheta$ then $t^* = t$ and $\vartheta \rightarrow +50\text{mV}$ and $U_i \rightarrow U^{\text{rest}}$
Triplet STDP	$\frac{dw_{ij}}{dt} = \eta w_0 A^+ z_j^+(t) z_i^{\text{slow}}(t - \epsilon) S_i(t) - \eta w_0 A_i^-(t) z_i^-(t) S_j(t)$ <p>and $A_i^-(t) = \frac{A^+ \tau^+ \tau^{\text{slow}}}{\tau^- \kappa} \bar{v}_i(t)^2$ if $w_{ij} < 0$ then $w_{ij} \rightarrow 0$ if $w_{ij} > w_{\text{max}}$ then $w_{ij} \rightarrow w_{\text{max}}$</p>
Synaptic traces and hom. rate estimate	$\frac{dz_n^x}{dt} = -\frac{z_n^x}{\tau^x} + S_n(t) \quad \text{and} \quad \tau \frac{d\bar{v}}{dt} = -\bar{v} + S_n(t)$

E Input	
Type	Description
Poisson neurons	Fixed rate $\nu_{\text{ext}} = 2$

F Measurements	
Type	Description
	Spike activity for raster plots. Run time of simulation (excluding time for setup and saving data).

Table D.6 – Simulation parameter summary for B.1 Network details for 25k cell network following Kunkel et al. (2011).

Populations		
Name	Value	Description
N_E	20000	Size of excitatory population E
N_I	5000	Size of inhibitory population I
N_{Ext}	2500	Size of external Poisson neuron population

Connectivity		
Name	Value	Description
ϵ	0.05	Probability of any connection (EE, EI, IE, II, Ext)
w^{EE}	$w = 0.16$	excitatory \rightarrow excitatory
w^{EI}	$w = 0.16$	excitatory \rightarrow inhibitory
w^{IE}	$w = 1$	inhibitory \rightarrow excitatory
w^{II}	$w = 1$	inhibitory \rightarrow inhibitory
w^{Ext}	$w = 0.16$	Poisson neurons \rightarrow excitatory

Neuron Model		
Name	Value	Description
τ	20 ms	Membrane time constant
U^{rest}	-60 mV	Resting potential
U^E	0 mV	Excitatory reversal potential
U^I	-80 mV	Inhibitory reversal potential
ϑ^{rest}	-50 mV	Threshold value at rest
τ^{thr}	5 ms	Time constant of threshold dynamics
τ^{ampa}	5 ms	AMPA conductance decay time constant
τ^{nmda}	100 ms	NMDA conductance decay time constant
τ^{gaba}	10 ms	GABA conductance decay time constant
α	0.5	Ratio between AMPA and NMDA conductance

Plasticity Model		
Name	Value	Description
η	6.25	Relative learning rate
A^+	$6.5 \times 10^{-3} \eta$	LTP strength
τ^+	16.8 ms	Trace time constant LTP
τ^+	33.7 ms	Trace time constant LTD
τ^{slow}	114 ms	Long trace time constant LTP
w_{max}	1	Maximum weight
κ	3 Hz	Target rate

Appendix D. Supplements to Chapter 5

Bibliography

- Abbott, L. F. and Nelson, S. B. Synaptic plasticity: taming the beast. *Nat Neurosci*, pages 1178 – 1183, 2000.
- Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. Synaptic depression and cortical gain control. *Science*, 275(5297):220–224, 1997.
- Abraham, W. C. How long will long-term potentiation last? *Phil Trans R Soc Lond B*, 358(1432):735–744, 2003.
- Abraham, W. C. Metaplasticity: tuning synapses and networks for plasticity. *Nat Rev Neurosci*, 9(5):387–387, 2008.
- Abraham, W. C. and Bear, M. F. Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci*, 19(4):126–130, 1996.
- Amit, D. J. and Brunel, N. Dynamics of a recurrent network of spiking neurons before and following learning. *Network*, 8(4):373–404, 1997a.
- Amit, D. J. and Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex*, 7(3): 237–252, 1997b.
- Amit, D. J. and Fusi, S. Learning in neural networks with material synapses. *Neural Comput*, 6(5):957–982, 1994.
- Amit, D. J. and Mongillo, G. Selective delay activity in the cortex: Phenomena and interpretation. *Cereb Cortex*, 13(11):1139–1150, 2003a.
- Amit, D. J. and Mongillo, G. Spike-driven synaptic dynamics generating working memory states. *Neural Comput*, 15(3):565–96, 2003b.
- Amit, D. J. and Tsodyks, M. V. Quantitative study of attractor neural network retrieving at low spike rates: I. substrate—spikes, rates and neuronal gain. *Network*, 2(3):259–273, 1991.
- Amit, D. J., Gutfreund, H., and Sompolinsky, H. Spin-glass models of neural networks. *Phys Rev A*, 32(2):1007–1018, 1985.
- Ananthanarayanan, R., Esser, S. K., Simon, H. D., and Modha, D. S. The cat is out of the bag: Cortical simulations with 109 neurons, 1013 synapses. SC '09, page 63:1–63:12, New York, NY, USA, 2009. ACM.

Bibliography

- Anderson, J. D. *Computational fluid dynamics: the basics with applications*. McGraw-Hill, New York, 1995.
- Artola, A., Bröcher, S., and Singer, W. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature*, 347(6288):69–72, 1990.
- Badel, L., Lefort, S., Brette, R., Petersen, C. C. H., Gerstner, W., and Richardson, M. J. E. Dynamic i-v curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *J Neurophysiol*, 99(2):656–666, 2008.
- Barth, A. L. and Poulet, J. F. Experimental evidence for sparse firing in the neocortex. *Trends Neurosci*, 35(6):345–355, 2012.
- Bathellier, B., Ushakova, L., and Rumpel, S. Discrete neocortical dynamics predict behavioral categorization of sounds. *Neuron*, 76(2):435–449, 2012.
- Benda, J. and Herz, A. V. M. A universal model for spike-frequency adaptation. *Neural Comput*, 15(11):2523–2564, 2003.
- Bi, G.-Q. and Poo, M.-M. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24):10464–10472, 1998.
- Bi, G.-q. and Poo, M.-m. Synaptic modification by correlated activity: Hebb’s postulate revisited. *Annu Rev Neurosci*, 24(1):139–166, 2001.
- Bienenstock, E., Cooper, L., and Munro, P. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci*, 2(1):32–48, 1982.
- Billings, G. and van Rossum, M. C. W. Memory retention and spike-timing-dependent plasticity. *J Neurophysiol*, 101(6):2775–2788, 2009.
- Bliss, T. V. P. and Lømo, T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol*, 232(2):331–356, 1973.
- Bliss, T. V. P., Collingridge, G. L., and Morris, R. G. M. Introduction. long-term potentiation and structure of the issue. *Philos Trans R Soc Lond B Biol Sci*, 358(1432):607–611, 2003.
- Brette, R. and Gerstner, W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J Neurophysiol*, page 3637–3642, 2009.
- Brette, R. and Goodman, D. F. M. Vectorized algorithms for spiking neural network simulation. *Neural Comput*, 23(6):1503–1535, 2011.
- Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J., Diesmann, M., Morrison, A., Goodman, P., Harris, F., Zirpe, M., Natschläger, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A., El Boustani, S., and Destexhe, A. Simulation of networks of spiking neurons: A review of tools and strategies. *Front Comput Neurosci*, 23(3):349–398, 2007.

Bibliography

- Brunel, N. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8(3):183–208, 2000.
- Brunel, N. and Wang, X.-J. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J Comput Neurosci*, 11(1):63–85, 2001.
- Burns, B. D. and Webb, A. C. The spontaneous activity of neurones in the cat’s cerebral cortex. *Proc R Soc Lond B Biol Sci*, 194(1115):211–223, 1976.
- Bush, D., Philippides, A., Husbands, P., and O’Shea, M. Dual coding with STDP in a spiking recurrent neural network model of the hippocampus. *PLoS Comput Biol*, 6(7):e1000839, 2010.
- Calvin, W. H. and Stevens, C. F. Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J Neurophysiol*, 31(4):574–587, 1968.
- Caporale, N. and Dan, Y. Spike Timing–Dependent plasticity: A hebbian learning rule. *Annu Rev Neurosci*, 31(1):25–46, 2008.
- Castillo, P. E., Chiu, C. Q., and Carroll, R. C. Long-term plasticity at inhibitory synapses. *Curr Opin Neurobiol*, page 1–11, 2011.
- Chen, J.-Y., Lonjers, P., Lee, C., Chistiakova, M., Volgushev, M., and Bazhenov, M. Heterosynaptic plasticity prevents runaway synaptic dynamics. *J Neurosci*, 33(40):15915–15929, 2013.
- Chistiakova, M., Bannon, N. M., Bazhenov, M., and Volgushev, M. Heterosynaptic plasticity: Multiple mechanisms and multiple roles. *Neuroscientist*, 2014.
- Christie, B. R. and Abraham, W. C. Priming of associative long-term depression in the dentate gyrus by theta frequency synaptic activity. *Neuron*, 9(1):79–84, 1992.
- Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L., and Gerstner, W. Tag-trigger-consolidation: A model of early and late long-term-potential and depression. *PLoS Comput Biol*, 4(12):e1000248, 2008.
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci*, 13(3):344–352, 2010.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X.-J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex*, 10(9):910–923, 2000.
- Cooper, L. N., Intrator, N., Blais, B. S., and Shouval, H. Z. *Theory of Cortical Plasticity*. World Scientific, New Jersey, 2004.
- Crow, T. J. Cortical synapses and reinforcement: a hypothesis. *Nature*, 219(5155):736–737, 1968.

Bibliography

- Curti, E., Mongillo, G., Camera, G. L., and Amit, D. J. Mean field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Comput*, 16(12):2597–2637, 2004.
- Davis, G. W. Homeostatic control of neural activity: From phenomenology to molecular design. *Annu Rev Neurosci*, 29(1):307–323, 2006.
- Davis, G. W. Homeostatic signaling and the stabilization of neural function. *Neuron*, 80(3):718–728, 2013.
- Davison, A. P., Brüderle, D., Eppler, J., Kremkow, J., Müller, E., Pecevski, D., Perrinet, L., and Yger, P. PyNN: a common interface for neuronal network simulators. *Front Neuroinform*, 2:11, 2009.
- DeFelipe, J. and Fariñas, I. The pyramidal neuron of the cerebral cortex: Morphological and chemical characteristics of the synaptic inputs. *Prog Neurobiol*, 39(6):563–607, 1992.
- Del Giudice, P., Fusi, S., and Mattia, M. Modelling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. *J Physiol Paris*, 97(4-6):659–681, 2003.
- Desai, N. S. Homeostatic plasticity in the CNS: synaptic and intrinsic forms. *J Physiol Paris*, 97(4-6):391–402, 2003.
- El Boustani, S., Yger, P., Frégnac, Y., and Destexhe, A. Stable learning in stochastic network states. *J Neurosci*, 32(1):194–214, 2012.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. A large-scale model of the functioning brain. *Science*, 338(6111):1202–1205, 2012.
- Filion, M. and Tremblay, L. Abnormal spontaneous activity of globus pallidus neurons in monkeys with MPTP-induced parkinsonism. *Brain Res*, 547(1):140–144, 1991.
- Frank, C. A., Kennedy, M. J., Goold, C. P., Marek, K. W., and Davis, G. W. Mechanisms underlying the rapid induction and sustained expression of synaptic homeostasis. *Neuron*, 52(4):663–677, 2006.
- Frey, U. and Morris, R. G. M. Weak before strong: dissociating synaptic tagging and plasticity-factor accounts of late-LTP. *Neuropharmacology*, 37(4-5):545–552, 1998.
- Frey, U. and Morris, R. G. M. Synaptic tagging and long-term potentiation. *Nature*, 385(6616):533–6, 1997.
- Froemke, R. C., Merzenich, M. M., and Schreiner, C. E. A synaptic memory trace for cortical receptive field plasticity. *Nature*, 450(7168):425–429, 2007.
- Fusi, S. Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol Cybern*, 87(5):459–470, 2002.

- Fusi, S. and Abbott, L. F. Limits on the memory storage capacity of bounded synapses. *Nat Neurosci*, 10(4):485–493, 2007.
- Fusi, S., Annunziato, M., Badoni, D., Salamon, A., and Amit, D. J. Spike-driven synaptic plasticity: Theory, simulation, VLSI implementation. *Neural Comput*, 12(10):2227–2258, 2000.
- Fusi, S., Drew, P. J., and Abbott, L. F. Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611, 2005.
- Fuster, J. M. and Jervey, J. P. Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *J Neurosci*, 2(3):361–375, 1982.
- Gambino, F. and Holtmaat, A. Spike-timing-dependent potentiation of sensory surround in the somatosensory cortex is facilitated by deprivation-mediated disinhibition. *Neuron*, 75(3):490–502, 2012.
- Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. Internally generated reactivation of single neurons in human hippocampus during free recall. *Science*, 322(5898):96–101, 2008.
- Gerstner, W., Kempter, R., van Hemmen, J. L., and Wagner, H. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–81, 1996.
- Gerstner, W. and Kistler, W. M. Mathematical formulations of hebbian learning. *Biol Cybern*, 87(5-6):404–415, 2002a.
- Gerstner, W. and Kistler, W. M. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 1 edition, 2002b.
- Gerstner, W. and van Hemmen, J. Associative memory in a network of ‘spiking’ neurons. *Network*, 3(2):139–164, 1992.
- Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. *Neuronal dynamics: from single neurons to networks and models of cognition*. Cambridge University Press, Cambridge, 2014.
- Gewaltig, M.-O. and Diesmann, M. NEST (NEural simulation tool). *Scholarpedia*, 2(4):1430, 2007.
- Gewaltig, M.-O., Diesmann, M., and Aertsen, A. Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. *Neural Networks*, 14(6-7):657–673, 2001a.
- Gewaltig, M.-O., Diesmann, M., and Aertsen, A. Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. *Neural Networks*, 14(6-7):657–673, 2001b.
- Gewaltig, M.-O., Morrison, A., and Plesser, H. E. NEST by example: An introduction to the neural simulation tool NEST. In Le Novère, N., editor, *Computational Systems Neurobiology*, pages 533–558. Springer Netherlands, 2012.

Bibliography

- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. A., and Hemmen, J. L. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. i. input selectivity–strengthening correlated input pathways. *Biol Cybern*, 101(2):81–102, 2009a.
- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. a., and van Hemmen, J. L. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. II. input selectivity–symmetry breaking. *Biol Cybern*, 101(2):103–14, 2009b.
- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. A., and van Hemmen, J. L. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks III: partially connected neurons driven by spontaneous activity. *Biol Cybern*, 101(5-6):411–26, 2009c.
- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. a., and van Hemmen, J. L. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks v: self-organization schemes and weight dependence. *Biol Cybern*, page 365–386, 2010.
- Gjorgjieva, J., Clopath, C., Audet, J., and Pfister, J.-P. A triplet spike-timing–dependent plasticity model generalizes the Bienenstock–Cooper–Munro rule to higher-order spatiotemporal correlations. *Proc Natl Acad Sci U S A*, 108(48):19383–19388, 2011.
- Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron*, 14(3):477–485, 1995.
- Goodman, D. and Brette, R. Brian: a simulator for spiking neural networks in python. *Front Neuroinform*, 2:5, 2008.
- Grossberg, S. On learning of spatiotemporal patterns by networks with ordered sensory and motor components. 1. excitatory components of the cerebellum. *Stud Appl Math*, page 105–132, 1969.
- Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J Neurosci*, 23(9):3697–3714, 2003.
- Haas, J. S., Nowotny, T., and Abarbanel, H. Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J Neurophysiol*, 96(6):3305–3313, 2006.
- Hansel, D. and Mato, G. Short-term plasticity explains irregular persistent activity in working memory tasks. *J Neurosci*, 33(1):133–149, 2013.
- Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory*. Wiley & Sons New York, 1949.
- Helias, M., Kunkel, S., Masumoto, G., Igarashi, J., Eppler, J. M., Ishii, S., Fukai, T., Morrison, A., and Diesmann, M. Supercomputers ready for use as discovery machines for neuroscience. *Front Neuroinform*, 6, 2012.

- Hennequin, G., Gerstner, W., and Pfister, J.-P. STDP in adaptive neurons gives close-to-optimal information transmission. *Front Comput Neurosci*, 4, 2010.
- Hoang, R. V., Tanna, D., Bray, L. C. J., Dascalu, S. M., and Jr, F. C. H. A novel CPU/GPU simulation environment for large-scale biologically realistic neural modeling. *Front Neuroinform*, 7:19, 2013.
- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*, 79(8):2554, 1982.
- Huang, Y. Y., Colino, A., Selig, D. K., and Malenka, R. C. The influence of prior synaptic activity on the induction of long-term potentiation. *Science*, 255(5045):730–733, 1992.
- Hubbard, J. I., Stenhouse, D., and Eccles, R. M. Origin of synaptic noise. *Science*, 157(3786):330–331, 1967.
- Hubel, D. H. and Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol*, 160(1):106–154.2, 1962.
- Hulme, S. R., Jones, O. D., Ireland, D. R., and Abraham, W. C. Calcium-dependent but action potential-independent BCM-Like metaplasticity in the hippocampus. *J Neurosci*, 32(20):6785–6794, 2012.
- Ibata, K., Sun, Q., and Turrigiano, G. G. Rapid synaptic scaling induced by changes in postsynaptic firing. *Neuron*, 57(6):819–826, 2008.
- Iglesias, J., Eriksson, J., Grize, F., Tomassini, M., and Villa, A. E. P. Dynamics of pruning in simulated large-scale spiking neural networks. *BioSystems*, 79(1-3):11–20, 2005.
- Izhikevich, E. M. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, 2003.
- Izhikevich, E. M. Polychronization: Computation with spikes. *Neural Computation*, 18(2):245–282, 2006.
- Izhikevich, E. M. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb Cortex*, 17(10):2443–2452, 2007.
- Izhikevich, E. M. and Edelman, G. M. Large-scale model of mammalian thalamo-cortical systems. *Proc Natl Acad Sci U S A*, 105(9):3593–3598, 2008.
- Kandel, E. R., Schwartz, J. H., and Jessell, T. M. *Principles of neural science*. McGraw-Hill, Health Professions Division, New York, 2000.
- Kandel, E. R., Markram, H., Matthews, P. M., Yuste, R., and Koch, C. Neuroscience thinks big (and collaboratively). *Nat Rev Neurosci*, 14(9):659–664, 2013.
- Kempter, R., Gerstner, W., and van Hemmen, J. L. Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Comput*, 13(12):2709–2741, 2001.
- Ko, H., Cossell, L., Baragli, C., Antolik, J., Clopath, C., Hofer, S. B., and Mrsic-Flogel, T. D. The emergence of functional microcircuits in visual cortex. *Nature*, 496(7443):96–100, 2013.

Bibliography

- Koch, C. and Reid, R. C. Neuroscience: Observatories of the mind. *Nature*, 483 (7390):397–398, 2012.
- Koch, K. W. and Fuster, J. M. Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Exp Brain Res*, 76(2):292–306, 1989.
- Kullmann, D. M., Moreau, A. W., Bakiri, Y., and Nicholson, E. Plasticity of inhibition. *Neuron*, 75(6):951–962, 2012.
- Kumar, A., Rotter, S., and Aertsen, A. Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J Neurosci*, 28(20):5268–5280, 2008a.
- Kumar, A., Schrader, S., Aertsen, A., and Rotter, S. The high-conductance state of cortical networks. *Neural Comput*, 20(1):1–43, 2008b.
- Kumar, A., Rotter, S., and Aertsen, A. Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nat Rev Neurosci*, 11(9):615–627, 2010.
- Kunkel, S., Diesmann, M., and Morrison, A. Limits to the development of feed-forward structures in large recurrent neuronal networks. *Front Comput Neurosci*, 4:160, 2011.
- Lamsa, K. P., Kullmann, D. M., and Woodin, M. A. Spike-timing dependent plasticity in inhibitory circuits. *Front Synaptic Neurosci*, 2:8, 2010.
- Lang, S., Dercksen, V. J., Sakmann, B., and Oberlaender, M. Simulation of signal flow in 3D reconstructions of an anatomically realistic neural network in rat vibrissal cortex. *Neural Networks*, 24(9):998–1011, 2011.
- Lazar, A., Pipa, G., and Triesch, J. SORN: a self-organizing recurrent neural network. *Front Comput Neurosci*, 3:23, 2009.
- Li, N. and DiCarlo, J. J. Unsupervised natural experience rapidly alters invariant object representation in visual cortex. *Science*, 321(5895):1502–1507, 2008.
- Lisman, J. E. A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylating kinase. *Proc Natl Acad Sci U S A*, 82(9):3055–3057, 1985.
- Lisman, J. Long-term potentiation: outstanding questions and attempted synthesis. *Phil Trans R Soc Lond B*, 358(1432):829–842, 2003.
- Litwin-Kumar, A. and Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat Neurosci*, 15(11):1498–1505, 2012.
- Lubenov, E. V. and Siapas, A. G. Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron*, 58(1):118–131, 2008.
- Lütcke, H., Gerhard, F., Zenke, F., Gerstner, W., and Helmchen, F. Inference of neuronal network spike dynamics and topology from calcium imaging data. *Front Neural Circuits*, 7(201), 2013.

Bibliography

- Luz, Y. and Shamir, M. Balancing feed-forward excitation and inhibition via hebbian inhibitory synaptic plasticity. *PLoS Comput Biol*, 8(1):e1002334, 2012.
- Lynch, G. S., Dunwiddie, T., and Gribkoff, V. Heterosynaptic depression: a post-synaptic correlate of long-term potentiation. *Nature*, 266(5604):737–739, 1977.
- MacKay, D. J. C. and Miller, K. D. Analysis of linsker’s application of hebbian rules to linear networks. *Network*, 1:257–297, 1990.
- Maffei, A., Nataraj, K., Nelson, S. B., and Turrigiano, G. G. Potentiation of cortical inhibition by visual deprivation. *Nature*, 443(7107):81–84, 2006.
- Mainen, Z. F. and Sejnowski, T. J. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–6, 1995.
- Marder, E. and Goaillard, J.-M. Variability, compensation and homeostasis in neuron and network function. *Nat Rev Neurosci*, 7(7):563–574, 2006.
- Markram, H. The blue brain project. *Nat Rev Neurosci*, 7(2):153–160, 2006.
- Markram, H. and Tsodyks, M. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382(6594):807–810, 1996.
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297):213–215, 1997.
- Markram, H., Wang, Y., and Tsodyks, M. Differential signaling via the same axon of neocortical pyramidal neurons. *Proc Natl Acad Sci U S A*, 95(9):5323–5328, 1998.
- Markram, H., Gerstner, W., and Sjöström, P. J. A history of spike-timing-dependent plasticity. *Front Synaptic Neurosci*, 3:4, 2011.
- Markram, H., Gerstner, W., and Sjöström, P. J. Spike-timing-dependent plasticity: a comprehensive overview. *Front Synaptic Neurosci*, page 2, 2012.
- Marre, O., Yger, P., Davison, A. P., and Frégnac, Y. Reliable recall of spontaneous activity patterns in cortical networks. *J Neurosci*, 29(46):14596–606, 2009.
- Martin, S. J., Grimwood, P. D., and Morris, R. G. M. Synaptic plasticity and memory: An evaluation of the hypothesis. *Annu Rev Neurosci*, 23(1):649–711, 2000.
- McCormick, D. A. and Contreras, D. On the cellular and network bases of epileptic seizures. *Annu Rev Physiol*, 63(1):815–846, 2001.
- Mensi, S., Naud, R., Pozzorini, C., Avermann, M., Petersen, C. C. H., and Gerstner, W. Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms. *J Neurophysiol*, 107(6):1756–1775, 2012.
- Miller, K. D. and MacKay, D. J. The role of constraints in hebbian learning. *Neural Comput*, 6(1):100–126, 1994.

Bibliography

- Mockett, B., Coussens, C., and Abraham, W. C. NMDA receptor-mediated meta-plasticity during the induction of long-term depression by low-frequency stimulation. *Eur J Neurosci*, 15(11):1819–1826, 2002.
- Mongillo, G., Amit, D. J., and Brunel, N. Retrospective and prospective persistent activity induced by hebbian learning in a recurrent cortical network. *Eur J Neurosci*, 18(7):2011–2024, 2003.
- Mongillo, G., Barak, O., and Tsodyks, M. Synaptic theory of working memory. *Science*, 319(5869):1543–1546, 2008.
- Morrison, A., Mehring, C., Geisel, T., Aertsen, a. D., and Diesmann, M. Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput*, 17(8):1776–801, 2005.
- Morrison, A., Aertsen, A., and Diesmann, M. Spike-timing-dependent plasticity in balanced random networks. *Neural Comput*, 19(6):1437–67, 2007.
- Morrison, A., Diesmann, M., and Gerstner, W. Phenomenological models of synaptic plasticity based on spike timing. *Biol Cybern*, 98(6):459–478, 2008.
- Murphy, B. K. and Miller, K. D. Balanced amplification: A new mechanism of selective amplification of neural activity patterns. *Neuron*, 61(4):635–648, 2009.
- Nordlie, E., Gewaltig, M.-O., and Plesser, H. E. Towards reproducible descriptions of neuronal network models. *PLoS Comput Biol*, 5(8):e1000456, 2009.
- Oja, E. Simplified neuron model as a principal component analyzer. *J Math Biol*, 15(3):267–273, 1982.
- Okun, M. and Lampl, I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci*, 11(5):535–7, 2008.
- Pawlak, V., Wickens, J. R., Kirkwood, A., and Kerr, J. N. D. Timing is not everything: Neuromodulation opens the STDP gate. *Front Synaptic Neurosci*, 2:146, 2010.
- Pawlak, V., Greenberg, D. S., Sprekeler, H., Gerstner, W., and Kerr, J. N. Changing the responses of cortical neurons from sub- to suprathreshold using single spikes in vivo. *elife*, 2:e00012, 2013.
- Pfister, J.-P. and Gerstner, W. Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci*, 26(38):9673–9682, 2006.
- Pozzorini, C., Naud, R., Mensi, S., and Gerstner, W. Temporal whitening by power-law adaptation in neocortical neurons. *Nat Neurosci*, 16(7):942–948, 2013.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, 2005.
- Raymond, E. S. *The art of UNIX programming*. Addison-Wesley, Boston, 1st edition edition, 2003.

- Renart, A., Moreno-Bote, R., Wang, X.-J., and Parga, N. Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput*, 19(1):1–46, 2007.
- Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K. D. The asynchronous state in cortical circuits. *Science*, 327(5965):587–590, 2010.
- Richert, M., Nageswaran, J. M., Dutt, N., and Krichmar, J. L. An efficient simulation environment for modeling large-scale cortical processing. *Front Neuroinform*, 5, 2011.
- Riegle, K. C. and Meyer, R. L. Rapid homeostatic plasticity in the intact adult visual system. *J Neurosci*, 27(39):10556–10567, 2007.
- Rochester, N., Holland, J., Haibt, L., and Duda, W. Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IEEE Trans Inf Theory*, 2(3):80–93, 1956.
- Rolls, E. T. and Deco, G. *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford University Press, USA, 1 edition, 2010.
- Roudi, Y. and Latham, P. E. A balanced memory network. *PLoS Comput Biol*, 3(9):e141, 2007.
- Schultz, W., Dayan, P., and Montague, P. R. A neural substrate of prediction and reward. *Science*, 275(5306):1593–1599, 1997.
- Senn, W., Markram, H., and Tsodyks, M. An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput*, 13(1):35–67, 2001.
- Shouval, H. Z., Bear, M. F., and Cooper, L. N. A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci U S A*, 99(16):10831–10836, 2002.
- Shu, Y., Hasenstaub, A., and McCormick, D. A. Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–293, 2003.
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, 2001.
- Sjöström, P., Rancz, E., Roth, A., and Häusser, M. Dendritic excitability and synaptic plasticity. *Physiol Rev*, 88(2):769, 2008.
- Song, S., Miller, K. D., and Abbott, L. F. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3(9):919–926, 2000.
- Spencer, K. M., Nestor, P. G., Niznikiewicz, M. A., Salisbury, D. F., Shenton, M. E., and McCarley, R. W. Abnormal neural synchrony in schizophrenia. *J Neurosci*, 23(19):7407–7411, 2003.
- Sprekeler, H., Clopath, C., and Vogels, T. P. Interactions of excitatory and inhibitory synaptic plasticity, 2012.

Bibliography

- Srinivasa, N. and Jiang, Q. Stable learning of functional maps in self-organizing spiking neural networks with continuous synaptic plasticity. *Front Comput Neurosci*, 7, 2013.
- Stepanyants, A., Hof, P. R., and Chklovskii, D. B. Geometry and structural plasticity of synaptic connectivity. *Neuron*, 34(2):275–288, 2002.
- Strogatz, S. H. *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering*. Westview Press, 1 edition, 2001.
- Sutton, M. A. and Schuman, E. M. Dendritic protein synthesis, synaptic plasticity, and memory. *Cell*, 127(1):49–58, 2006.
- Szatmáry, B. and Izhikevich, E. M. Spike-timing theory of working memory. *PLoS Comput Biol*, 6(8):e1000879, 2010.
- Tetzlaff, C., Kolodziejski, C., Timme, M., and Wörgötter, F. Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Front Comput Neurosci*, 5:47, 2011.
- Tetzlaff, C., Kolodziejski, C., Timme, M., and Wörgötter, F. Analysis of synaptic scaling in combination with hebbian plasticity in several simple networks. *Front Comput Neurosci*, 6:36, 2012.
- Thakur, R., Rabenseifner, R., and Gropp, W. Optimization of collective communication operations in MPICH. *Int J High Perform C*, 19(1):49–66, 2005.
- Toyoizumi, T., Pfister, J.-P., Aihara, K., and Gerstner, W. Optimality model of unsupervised spike-timing-dependent plasticity: Synaptic memory and weight distribution. *Neural Comput*, 19(3):639–671, 2007.
- Trachtenberg, J. T., Chen, B. E., Knott, G. W., Feng, G., Sanes, J. R., Welker, E., and Svoboda, K. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 420(6917):788–794, 2002.
- Treves, A. Mean-field analysis of neuronal spike dynamics. *Network*, 4(3):259–284, 1993.
- Tsodyks, M. V. and Sejnowski, T. Rapid state switching in balanced cortical network models. *Methods*, 6, 1995.
- Tsodyks, M. V. and Markram, H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94(2):719–723, 1997.
- Turrigiano, G. Too many cooks? intrinsic and synaptic homeostatic mechanisms in cortical circuit refinement. *Annu Rev Neurosci*, 34(1):89–103, 2011.
- Turrigiano, G. G. Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci*, 22(5):221–227, 1999.
- Turrigiano, G. G. The self-tuning neuron: Synaptic scaling of excitatory synapses. *Cell*, 135(3):422–435, 2008.

Bibliography

- Turrigiano, G. G. Homeostatic synaptic plasticity: Local and global mechanisms for stabilizing neuronal function. *Cold Spring Harb Perspect Biol*, 4(1):a005736, 2012.
- Turrigiano, G. G. and Nelson, S. B. Hebb and homeostasis in neuronal plasticity. *Curr Opin Neurobiol*, 10(3):358–364, 2000.
- Turrigiano, G. G. and Nelson, S. B. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci*, 5(2):97–107, 2004.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896, 1998.
- Uhlhaas, P. J. and Singer, W. Neural synchrony in brain disorders: Relevance for cognitive dysfunctions and pathophysiology. *Neuron*, 52(1):155–168, 2006.
- van Rossum, M. C. W., Bi, G. Q., and Turrigiano, G. G. Stable hebbian learning from spike timing-dependent plasticity. *J Neurosci*, 20(23):8812–8821, 2000.
- van Rossum, M. and Turrigiano, G. Correlation based learning from spike timing dependent plasticity. *Neurocomputing*, 38-40:409–415, 2001.
- van Vreeswijk, C. and Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.
- Van Vreeswijk, C. and Sompolinsky, H. Chaotic balanced state in a model of cortical circuits. *Neural Comput*, 10(6):1321–1371, 1998.
- Vogels, T. P. *Signal Processing in Neural Networks*. PhD thesis, Brandeis University, 2007.
- Vogels, T. P. and Abbott, L. F. Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci*, 25(46):10786, 2005.
- Vogels, T. P. and Abbott, L. F. Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat Neurosci*, 12(4):483–491, 2009.
- Vogels, T. P., Rajan, K., and Abbott, L. F. Neural network dynamics. *Annu Rev Neurosci*, 28:357–376, 2005.
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573, 2011.
- Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., Maffei, A., Miller, P., Wierenga, C., Woodin, M. A., Zenke, F., and Sprekeler, H. Inhibitory synaptic plasticity - spike timing dependence and putative network function. *Front Neural Circuits*, 7(119), 2013.
- von der Malsburg, C. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14(2):85–100, 1973.

Bibliography

- Waldrop, M. M. Computer modelling: Brain in a box. *Nature*, 482(7386):456–458, 2012.
- Wang, H. and Wagner, J. J. Priming-induced shift in synaptic plasticity in the rat hippocampus. *J Neurophysiol*, 82(4):2024–2028, 1999.
- Watt, A. J. and Desai, N. S. Homeostatic plasticity and STDP: keeping a neuron’s cool in a fluctuating world. *Front Synaptic Neurosci*, 2:5, 2010.
- Wehr, M. and Zador, A. M. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–6, 2003.
- Woodin, M. A. and Maffei, A. *Inhibitory Synaptic Plasticity*. Springer, 1st edition, 2010.
- Woodin, M. A., Ganguly, K., and Poo, M.-m. Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in cl- transporter activity. *Neuron*, 39(5):807–820, 2003.
- Yakovlev, V., Fusi, S., Berman, E., and Zohary, E. Inter-trial neuronal activity in inferior temporal cortex: a putative vehicle to generate long-term visual associations. *Nat Neurosci*, 1(4):310–7, 1998.
- Zenke, F., Hennequin, G., and Gerstner, W. Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput Biol*, 9(11):e1003330, 2013.
- Zhang, J. S. and Kaltenbach, J. A. Increases in spontaneous activity in the dorsal cochlear nucleus of the rat following exposure to high-intensity sound. *Neurosci Lett*, 250(3):197–200, 1998.
- Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., and Poo, M.-m. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395(6697):37–44, 1998.
- Zhou, Q. and Poo, M.-m. Reversal and consolidation of activity-induced synaptic modifications. *Trends Neurosci*, 27(7):378–383, 2004.
- Zhou, Q., Tao, H. W., and Poo, M.-m. Reversal and stabilization of synaptic modifications in a developing visual system. *Science*, 300(5627):1953–1957, 2003.

Curriculum Vitae

Friedemann Zenke

Academic title Dipl. Phys.

Nationality German

Education

09/2009 – present **Ecole Polytechnique Fédérale de Lausanne (EPFL)**,
Switzerland
PhD student, Computational neuroscience
Advisor: Prof. Gerstner

04/2007 – 03/2009 **University of Bonn**, Germany
Diploma, Physics, Summa Cum Laude
Advisor: Prof. Beck
Thesis (experimental hadron physics): A new avalanche photodiode readout for the Crystal Barrel experiment

02/2006 – 03/2007 **Australian National University**, Canberra, Australia
Physics, High Distinction

10/2003 – 02/2006 **University of Bonn**, Germany
Physics, Summa Cum Laude

Scholarships & Awards

2012 Outstanding Teaching Assistant Award 2012, The School of
Computer and Communication Sciences, EPFL

09/2009 – 12/2012 Marie Curie PhD fellowship

06/2007 – 07/2009 Scholarship, Bonn-Cologne Graduate School of Physics and
Astronomy

12/2005 – 04/2009 Scholarship, Friedrich Ebert Foundation

02/2006 – 12/2006 Scholarship, DAAD (The German Academic Exchange Service)

Curriculum Vitae

Referee Activities

PLoS Computational Biology, Frontiers in Computational Neuroscience (guest editor), Neural Computation, Journal of Computational Neuroscience

Research and Work Experience

- 10/2011 – 01/2012 Visiting Researcher, Group of Prof. Maass
Technical University of Graz, Austria
- 04/2009 – 07/2009 Research assistant, Crystal Barrel experiment (hadron physics)
Helmholtz Institute for Radiation and Nuclear Physics, Bonn,
Germany
- 05/2008 – 06/2008 Summer fellow, Hampton University Graduate School (hadron
physics)
Jefferson National Laboratory, Newport News, Virginia, USA
- 08/2007 – 09/2007 Summer fellow, detector development (nuclear physics)
GSI, Institute for Heavy Ion Research, Darmstadt, Germany
- 10/2007 – 03/2008 Database developer (Oracle), part time
Friedrich Ebert Foundation, Bonn, Germany
- 08/2006 – 11/2006 Internship, Nuclear Reaction Dynamics Group (nuclear physics)
Department of Nuclear Physics, Australian National Univer-
sity, Canberra, Australia
- 08/2004 – 09/2004 Internship, compound database development (bio informat-
ics)
Fraunhofer-Institute for Algorithms and Scientific Comput-
ing, Sankt Augustin, Germany
- 09/2002 – 06/2003 Public service, Alexianer psychiatric hospital, Cologne, Ger-
many

Teaching Experience

- 2013 Spring TA, Neural Networks and Biological Modeling, Prof. Gerst-
ner, EPFL, Lausanne, Switzerland
- Cosupervisor Julien Duc, Master thesis on the role of Asynchronous Re-
lease in working memory
- Cosupervisor Eszter Vertes, Master rotation project on the effect of pow-
erlaw adaptation of firing thresholds in models of working
memory
- 2012 Fall TA, Unsupervised and Reinforcement Learning, Prof. Ger-
stner, EPFL, Lausanne, Switzerland
- 2012 Summer Tutor, Advanced Course in Computational Neuroscience (ACCN),
Bedlwo, Poland

2012 Spring	TA, Pattern Recognition and Machine Learning, Prof. Seeger, EPFL, Lausanne, Switzerland
2011 Spring	TA, Neural Networks and Biological Modeling, Prof. Gerstner, EPFL, Lausanne, Switzerland
Cosupervisor	3 Bachelor Projects (Reinforcement learning within the TagTriC model), EPFL, Lausanne, Switzerland
2010 Fall	TA, Computer Science for Life Sciences (programming project in C++), J. Sam, EPFL, Lausanne, Switzerland
2010 Spring	TA, Neural Networks and Biological Modeling, Prof. Gerstner, EPFL, Lausanne, Switzerland
Cosupervisor	3 Bachelor Projects (Extensions to the Hopfield model), EPFL, Lausanne, Switzerland
2009 Fall	Supervisor for Experimental Student Projects, Detectors and Instrumentation in Hadron Physics, Profs Beck and Thoma, University of Bonn, Germany
2008 Spring	TA, Experimental Physics II, Prof. Desch, University of Bonn, Germany
Languages	German (mother tongue), English (fluent), French (fluent)