

# **Probabilistic Models of Computation and Self-organization in Recurrent Networks of Spiking Neurons**

by  
Johannes BILL

**DISSERTATION**  
submitted for the degree of  
**Doctor Technicae**



**Institute for Theoretical Computer Science  
Graz University of Technology**

Thesis Advisor  
Assoc.Prof. Dipl.-Ing. Dr.techn. Robert LEGENSTEIN

Graz, April 2015

This document is set in Palatino, compiled with [pdfL<sup>A</sup>T<sub>E</sub>X2<sub>ε</sub>](#) and [Biber](#).

The L<sup>A</sup>T<sub>E</sub>X template from Karl Voit is based on [KOMA script](#) and can be found online: <https://github.com/novoid/LaTeX-KOMA-template>

## Statutory Declaration

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources.

Graz, \_\_\_\_\_

Date

\_\_\_\_\_  
Signature

## Eidesstattliche Erklärung<sup>1</sup>

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst, andere als die angegebenen Quellen/Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommenen Stellen als solche kenntlich gemacht habe.

Graz, am \_\_\_\_\_

Datum

\_\_\_\_\_  
Unterschrift

---

<sup>1</sup>Beschluss der Curricula-Kommission für Bachelor-, Master- und Diplomstudien vom 10.11.2008; Genehmigung des Senates am 1.12.2008



# Abstract

Variability and stochasticity constitute ubiquitous phenomena in mammalian neural processing, ranging from channel noise on the molecular level, to response variability in neural populations, and trial-to-trial variability in perception. In recent years, Bayesian probability theory has emerged as a successful description in cognitive science and neuroscience to accommodate the apparent stochastic nature of perception, reasoning and learning of humans and other mammals within a coherent mathematical picture. This finding has led to the hypothesis that mammals may perform probabilistic rather than logical inference based on a learned internal model of their environment. However, our knowledge, how probabilistic computations could be supported by the neural substrate, is rudimentary at best. Joint research in cognitive science, experimental and computational neuroscience has set out to identify how probability distributions could be represented in neural activity patterns, and how the involved probabilistic calculations could be performed by cortical microcircuits.

In this thesis, principles of Bayesian computation and self-organization in models of recurrent spiking neural networks are investigated from the perspective of probability theory and machine learning. Using the framework of Markov chain Monte Carlo sampling, it is shown that recurrent networks of stochastic spiking neurons can represent complex high-dimensional probability distributions in a distributed spike code. In computer simulations, it is demonstrated how statistically correct network-wide inference can be performed through simple neural operations and local synaptic communication. The stochastic description of neural dynamics thereby incorporates noise as a computational resource and could explain several experimentally observed phenomena, such as neural response variability, spontaneous activity and perceptual multistability, as natural features of an ongoing sampling process.

The Bayesian description further permits the investigation of experience-dependent learning from presented spiking input streams. Specifically, it is shown that neural networks can develop and refine an internal probabilistic model of their environment through purely local plasticity rules. Using machine learning theory, it is first derived that in small winner-take-all network architectures the interplay of homeostatic intrinsic and Hebbian synaptic plasticity can approximate a powerful unsupervised learning algorithm, namely generalized stochastic online Expectation Maximization. In a next step, this finding is generalized to a spatially extended neural sheet model for inference and learning with a connectivity structure that is reminiscent of cortical microcircuits. The neural sheet model combines aspects of local competitive and large-scale associative learning under a unified Bayesian account and gives rise to an important theoretical insight: stochastic adaptations, acting on only locally and instantaneously available information in every neuron

and synapse, can accomplish a well-defined network-wide optimization process for long-term experience-dependent learning. A well-adapted neural sheet performs variational sample-based inference through its inherent response dynamics by transforming high-dimensional, and potentially noisy, spiking input streams into a sparse and reliable distributed spike code. This enables a network to draw statistically correct conclusions even in case of incomplete or ambiguous input.

Beyond computational neuroscience, the findings of this thesis are furthermore relevant to the operation of biologically inspired neuromorphic hardware devices. These artificial systems implement physical models of neurons and synapses in microscale electronic circuitry with the aim to establish novel inherently parallel, fault-tolerant, self-adapting computing platforms. Here, particular attention is given to bistable memristive materials, which exhibit voltage-dependent stochastic switching of their electric conductance, as plastic neuromorphic synapses. It is shown that, when multiple stochastic memristors are operated in parallel to jointly form a plastic synapse, neuromorphic winner-take-all architectures are endowed with statistical learning capabilities of remarkable robustness to device imperfections. In summary, the network architectures for distributed probabilistic inference and self-organized learning, presented in this thesis, provide an auspicious operation paradigm for neuromorphic computing systems.

# Zusammenfassung

Variabilität und Stochastizität sind in der neuronalen Informationsverarbeitung von Säugetieren allgegenwärtig. Sie reichen von molekularen Zufallsprozessen in Ionenkanälen über Aktivitätsschwankungen in neuronalen Populationen bis zu zeitlich veränderlichen Wahrnehmungen. In den letzten Jahren entwickelte sich die Bayessche Wahrscheinlichkeitsrechnung zu einer erfolgreichen Beschreibungssprache in den Kognitions- und Neurowissenschaften, um die scheinbar stochastische Beschaffenheit von Wahrnehmungen und Schlussfolgerungen sowie des Lernens von Menschen und anderen Säugetieren in einem kohärenten mathematischen Rahmen zusammenzuführen. Diese Erkenntnis führte zu der Hypothese, dass Säugetiere möglicherweise statistisch anstatt formallogisch schlussfolgern – basierend auf einem erlernten internen Modell ihrer Lebensumwelt. Allerdings kann unser Wissen darüber, wie probabilistische Berechnungen von neuronalen Netzen als physikalische Grundlage unterstützt werden könnten, bestenfalls als rudimentär bezeichnet werden. Fachübergreifende Forschungsarbeiten aus den Kognitionswissenschaften sowie aus den experimentellen und theoretischen Neurowissenschaften haben zu erkunden begonnen, wie Wahrscheinlichkeitsverteilungen in neuronalen Aktivitätsmustern repräsentiert, und wie die dafür benötigten probabilistischen Rechenoperationen von kortikalen Netzwerken durchgeführt werden könnten.

In dieser Dissertation werden Prinzipien der Bayesschen Informationsverarbeitung und Selbstorganisation in Modellen rekurrenter Netzwerke spikender Neuronen mit den Methoden der Wahrscheinlichkeitstheorie und des Maschinellen Lernens untersucht. Auf Grundlage der Theorie von Markov-Chain-Monte-Carlo-Verfahren wird bewiesen, dass rekurrente Netze stochastisch spikender Neuronen komplexe, hochdimensionale Wahrscheinlichkeitsverteilungen in räumlich verteilten Aktivitätsmustern repräsentieren können. Mithilfe von Computersimulationen wird gezeigt, wie statistisch korrekte, netzwerkweite Inferenz durch elementare neuronale Operationen und lokalen synaptischen Informationsaustausch erzielt werden kann. Durch die stochastische Beschreibung neuronaler Dynamik werden Zufallsprozesse dabei als integraler Bestandteil der Berechnung begriffen. In dieser Betrachtung stellen sich zahlreiche experimentell beobachtbare Phänomene, wie zum Beispiel Aktivitätsschwankungen, spontane Aktivität und zeitlich schwankende Wahrnehmungen, als sichtbare Merkmale eines permanent laufenden Berechnungsprozesses dar.

Die Bayessche Beschreibung ermöglicht darüber hinaus eine Untersuchung erfahrungsbasierter Lernvorgänge, die auf der Integration zugeführter, spikender Eingangsdatenströme beruhen. Es wird gezeigt, dass neuronale Netzwerke ein internes statistisches Modell ihrer Umgebung unter Verwendung von strikt lokalen Plastizitätsmechanismen ausbilden und verfeinern können. Mithilfe der Theorie des Maschinellen Lernens wird zunächst für den Fall kleiner Winner-Take-All

Netzwerkarchitekturen bewiesen, dass das Zusammenspiel von homöostatischer intrinsischer Plastizität und Hebbscher synaptischer Plastizität einen leistungsstarken Algorithmus für unüberwachtes Lernen (generalized stochastic online Expectation Maximization) approximieren kann. Dieses Ergebnis wird im Folgenden auf ein räumlich ausgedehntes Modell neuronalen Zellgewebes generalisiert. Die Verbindungsstruktur des räumlich ausgedehnten Modells weist Ähnlichkeiten mit kortikalen Netzwerken im Gehirn auf. Außerdem besitzt es die Fähigkeit zu statistischer Inferenz sowie zu unüberwachtem Lernen. Dabei werden Aspekte kurzreichweitiger kompetitiver und langreichweitiger assoziativer Lernverfahren unter einer einheitlichen Bayesschen Betrachtung vereint. Dies führt zu einem wichtigen theoretischen Resultat: Stochastische Adaptionen in jedem Neuron und jeder Synapse, lediglich auf Grundlage lokaler und instantan verfügbarer Information, kann einen wohldefinierten, netzwerkweiten Optimierungsprozess für langfristiges erfahrungsbasiertes Lernen verwirklichen. Ein hinreichend adaptiertes Netzwerk approximiert durch seine natürliche Dynamik einen sample-basierten Inferenzprozess und transformiert dabei hochdimensionale, und möglicherweise störungsbehaftete, Eingangssignale in effiziente, zuverlässige Aktivierungsmuster. Dies befähigt das Netzwerk auch im Falle unvollständiger oder mehrdeutiger Eingangssignale, statistisch fundierte Schlussfolgerungen zu ziehen.

Die Ergebnisse dieser Dissertation sind – jenseits der theoretischen Neurowissenschaften – auch für den Einsatz biologisch inspirierter neuromorpher Computerarchitekturen von Bedeutung. Diese künstlichen Systeme enthalten physische Modelle von Neuronen und Synapsen in hochintegrierten elektronischen Schaltkreisen. Dies dient dem Ziel, neuartige inhärent parallel arbeitende, fehlertolerante, anpassungsfähige Berechnungsmaschinen zu entwickeln. In dieser Arbeit gilt die besondere Aufmerksamkeit der Verwendbarkeit bistabiler Memristoren, deren elektrische Leitfähigkeit in Abhängigkeit der anliegenden Spannung stochastisch zwischen den beiden stabilen Zuständen springen kann, als plastische neuromorphe Synapsen. Es stellt sich heraus, dass neuromorphe Winner-Take-All Netzwerkarchitekturen die Fähigkeit zu statistischem Lernen erlangen, wenn mehrere stochastische Memristoren parallel zu einer gemeinsamen Synapse verschaltet werden, und dabei eine bemerkenswerte Robustheit gegen Bauteilmängel aufweisen. Zusammenfassend lässt sich erwarten, dass die in dieser Dissertation vorgestellten Netzwerkarchitekturen für verteilte statistische Inferenz und selbstorganisiertes Lernen ein vielversprechendes Leitbild für den Betrieb neuromorpher Berechnungssysteme darstellen.



# Acknowledgements

I wish to express my deepest gratitude to Robert Legenstein, my supervisor, for his extraordinary support and the excellent guidance he provided me with during my PhD studies. I also thank Karlheinz Meier for taking the long way from Heidelberg and for agreeing to be the second referee of this thesis.

I wish to express my sincere thanks to Wolfgang Maass, head of the institute, for giving me the opportunity to participate in several international conferences, scientific workshops and summer courses. I am extremely thankful and indebted to all my colleagues from the Institute for Theoretical Computer Science for many hours of intense and insightful discussions, for their friendly support and, of course, for our joint leisure activities and the frolic atmosphere in the office. I am deeply grateful to my co-authors for the exciting joint research and for many years of close collaboration. This thesis would not have been possible without their brilliant contributions and insightful comments. Thank you Lars Büsing, Ilja Bytschok, Stefan Habenschuss, Robert Legenstein, Wolfgang Maass, Karlheinz Meier, Bernhard Nessler, Dejan Pecevski, Mihai Petrovici, Dimitri Probst and Johannes Schemmel.

I wish to thank Karlheinz Meier, Johannes Schemmel, Mihai Petrovici and Ilja Bytschok from Ruprecht-Karls-Universität Heidelberg for kindly hosting me during my research stays in 2011 and 2013. Similarly, I thank Themis Prodromakis, Alexander Serb and Radu Berdan from the University of Southampton and Imperial College London for introducing me to the world of memristors and for their hospitality during my research stays in 2013 and 2014. I thank Alan Gelperin, David Tank and Michael Berry from Princeton University for the organization of the BCNN summer course in 2012 which offered me an invaluable insight into experimental neuroscience.

This thesis would have been impossible without the financial support of FACETS-ITN and the PNEUMA project. I would like to thank the coordinators and lecturers of FACETS-ITN for organizing the interdisciplinary scientific training and complementary skill courses. Also, I wish to thank the Technische Universität Graz for providing the facilities for my PhD studies. Furthermore, I would like to take this opportunity to express my gratitude to the countless developers of free software. This thesis was almost entirely conducted with free software, covering programming and simulation environments, office and drawing applications, operating systems, network infrastructure and countless tools.

I would like to thank my friends, my flat mate Jakob, and the singers of the university choir for their cheerful company in joyful times and their patience and friendship in less bright days. Finally, I feel immensely thankful to my parents, Helga and Albrecht, for their loving care and unyielding support during all these years.



# Contents

<b>Preface</b>	<b>1</b>
<b>1 Introduction</b>	<b>3</b>
<b>2 Neural dynamics as sampling</b>	<b>13</b>
2.1 Introduction . . . . .	14
2.2 Recapitulation of MCMC sampling . . . . .	17
2.3 Neural sampling . . . . .	18
2.4 Neural sampling in discrete time . . . . .	21
2.5 Neural sampling with a relative refractory mechanism . . . . .	23
2.6 Neural sampling in continuous time . . . . .	27
2.7 Demonstration of perceptual multistability in sampling networks . . . . .	30
2.8 Discussion . . . . .	34
<b>3 Homeostatic and Hebbian plasticity united under a Bayesian account</b>	<b>37</b>
3.1 Introduction . . . . .	38
3.2 Homeostatic plasticity in winner-take-all networks . . . . .	40
3.3 Theory for the WTA model . . . . .	42
3.4 Spiking network dynamics with homeostatic plasticity . . . . .	45
3.5 Homeostatic plasticity in recurrent spiking networks . . . . .	46
3.6 Discussion . . . . .	48
<b>4 Statistical learning with compound memristive synapses</b>	<b>49</b>
4.1 Introduction . . . . .	50
4.2 Stochastic memristors as plastic synapses . . . . .	52
4.3 Compound memristive synapses in winner-take-all networks . . . . .	57
4.4 Memristive synapses support inference and online learning . . . . .	58
4.5 Demonstration of unsupervised learning . . . . .	63
4.6 Influence of synaptic resolution . . . . .	64
4.7 Robustness to device variations . . . . .	66
4.8 Discussion . . . . .	69
<b>5 Distributed Bayesian computation and learning in spiking neural sheets</b>	<b>75</b>
5.1 Introduction . . . . .	76
5.2 Probabilistic inference in spatially extended spiking networks . . . . .	81
5.3 Emergence of local experts through synaptic plasticity . . . . .	89
5.4 Plastic recurrent synapses integrate structural knowledge . . . . .	95
5.5 Emergence of excitatory subnetworks in neural sheets . . . . .	102
5.6 Discussion . . . . .	104

<b>Appendices</b>	<b>114</b>
<b>A List of publications</b>	<b>117</b>
<b>B Appendix to chapter 2: Neural dynamics as sampling</b>	<b>119</b>
B.1 Derivation of the neural sampling theory . . . . .	119
B.2 Details to the computer simulations . . . . .	130
B.3 Firing statistics of neural sampling networks . . . . .	135
B.4 Approximation quality of different neuron and synapse models . . . . .	136
<b>C Appendix to chapter 4: Statistical learning with compound memristive synapses</b>	<b>141</b>
C.1 Probabilistic model definition . . . . .	141
C.2 Inference . . . . .	142
C.3 Learning via generalized Expectation-Maximization . . . . .	143
C.4 Spiking network implementation . . . . .	144
C.5 Details to the computer simulations . . . . .	146
C.6 Implementation with leaky integrator neurons . . . . .	147
<b>D Appendix to chapter 5: Distributed Bayesian computation and learning in spiking neural sheets</b>	<b>149</b>
D.1 Generative model . . . . .	149
D.2 Inference in the generative model (Corollary 1) . . . . .	152
D.3 Model optimization via Generalized Expectation Maximization . . . . .	153
D.4 Approximate plasticity rule for recurrent synapses . . . . .	157
D.5 Interaction of time scales . . . . .	158
D.6 Details to the computer simulations . . . . .	159
<b>Bibliography</b>	<b>169</b>

# List of Figures

1.1	Illustration of the thesis according to Marr’s levels of analysis . . . . .	5
2.1	Neuron model with absolute refractory mechanism. . . . .	21
2.2	Neuron model with relative refractory mechanism. . . . .	24
2.3	Sampling from a Boltzmann distribution by spiking neurons with relative refractory mechanism. . . . .	26
2.4	Perceptual multistability as probabilistic inference with neural sampling. . . . .	31
3.1	Homeostatic plasticity in WTA circuits . . . . .	39
3.2	Theory for the WTA model . . . . .	44
3.3	Dynamical properties of homeostatically regulated spiking networks	45
4.1	Compound synapse model with stochastic memristors . . . . .	53
4.2	Spiking network for probabilistic inference and online learning . . . . .	57
4.3	Learning of hand-written digits . . . . .	63
4.4	Influence of the synaptic resolution . . . . .	65
4.5	Robustness to spatial and temporal noise . . . . .	66
4.6	Robustness to unbalanced switching probabilities . . . . .	67
5.1	Neural sheet model for distributed inference and self-organized learning . . . . .	78
5.2	Neural sheets can perform Bayesian inference on distributed spiking input . . . . .	87
5.3	Emergence of probabilistic local experts through synaptic plasticity	92
5.4	Plastic recurrent synapses integrate structural knowledge . . . . .	98
5.5	Emergence of excitatory subnetworks in neural sheets . . . . .	101
B.1	Firing statistics of neural sampling networks. . . . .	135
B.2	Comparison of neural sampling with different neuron and synapse models. . . . .	138
B.3	Sampling from a Boltzmann distribution with more realistic PSP shapes. . . . .	139
D.1	Sampling quality and heuristic learning rule . . . . .	161



# List of Tables

4.1	Translation of synapse parameters between the hardware and theory domain . . . . .	71
B.1	List of parameters of the computer simulations . . . . .	131
B.2	Approximation quality of networks with different refractory mechanisms . . . . .	137
D.1	Parameters of the computer simulations . . . . .	159
D.2	List of symbols used in Appendix D . . . . .	166





# Preface

One of the most fascinating skills of humans and other mammals is their ability to adapt, as an individual, to the characteristics of their environment. This ability extends well beyond a mere reaction to immediate sensory stimuli. Instead, it appears that humans and mammals acquire custom knowledge on their specific environment in that they develop and maintain a model of their surroundings which enables them to classify impressions, to put perceptions into context, to make predictions, and to perform goal-directed actions. Be it a human who grew up on the countryside and moves to a vibrant city as an adult, be it a dog that understands the instructions of its master, or be it a laboratory mouse that swiftly finds its path to reward in a labyrinth task: Learning from experiences greatly enhances the faculties passed across generations in a genetic code. It is safe to assume that these abilities crucially rely on the formation of complex circuits in our brains. Particularly, the mammalian cerebral cortex with its dense and highly plastic recurrent networks of millions (mouse) to billions (human) of nerve cells (ROTH and DICKE, 2005) is believed to play an integral role in the emergence of “intelligent” behavior. The primal impact of experience-shaped neural processing on our perception of the world has been epitomized by neuroscientist Henry Markram in a quote that inevitably triggers associations that reach back to Plato and Descartes.

*“The brain builds a version of the universe and projects this version of the universe like a bubble all around us.” – “The vast majority of what you see is not what comes in through the eyes. It is what you infer about that room.”*

HENRY MARKRAM

It is an intriguing – intellectually beautiful and emotionally flattering – idea that the brain maintains a faithful internal model of the outside world, that it makes sense of the manifold sensory stimuli by developing abstract, generalizing representations, that it identifies temporal, causal and statistical relations among these abstract representations, and that it makes predictions which also incorporate consequences of anticipated actions and future events. That it conceives plans, envisions alternatives, and that it will never stop acquiring even more knowledge. The brain, a mighty self-organized simulator of the world, serving a creature to attain its goals.

The particular goal behind this thesis is to contribute to this intriguing hypothesis by exploring how models of plastic recurrent spiking neural networks could develop capabilities of analytically correct inference and self-organized learning from experience within a statistical description of brain computation.



# Chapter 1

## Introduction

Understanding the organization of computations in the brain displays one of the cardinal scientific challenges in the early twenty-first century. Ambitious multidisciplinary research collaborations have been initiated to promote this endeavor, such as the US-centered BRAIN initiative, the EU-centered Human Brain Project, or the Swiss-based Blue Brain Project. For computational neuroscience, a subdiscipline that attempts to mathematically model information processing in the brain, two aspects of brain function turn out to be particularly challenging. First, there likely exists no single appropriate level of detail for neuroscientific modeling. Instead theoretical research must accommodate the complex and often subtle interactions between multiple scales, ranging from the molecular level, via the cell level, up to collective network phenomena. Second, neural activity exhibits substantial variability *in vivo*, even in response to (almost) identical stimuli (FISER et al., 2004). This trial-to-trial variability does not originate only from deterministic circuit interactions. Neural activity is also subject to numerous microscopic noise sources (FAISAL et al., 2008) and thus appears to be, at least to some extent, truly stochastic. This observation has led to the idea to perceive stochasticity as a computational resource and to incorporate variability into models of brain function (R. P. N. RAO et al., 2002; ROLLS and DECO, 2010). Many studies have found that perception, reasoning and learning of humans and other mammals appear to be consistent with a Bayesian interpretation of probability theory (KERSTEN et al., 2004; KÖRDING and WOLPERT, 2004; GRIFFITHS and TENENBAUM, 2006; T. YANG and SHADLEN, 2007; ORBAN et al., 2008; ANGELAKI et al., 2009; DENISON et al., 2009; SCHULZ, 2012). These findings have led to the hypothesis that higher animals may perform probabilistic rather than logical inference, an idea that has spawned a new branch in cognitive science and neuroscience and that is often referred to as the Bayesian brain (T. S. LEE and MUMFORD, 2003; DOYA et al., 2007; TENENBAUM et al., 2011).

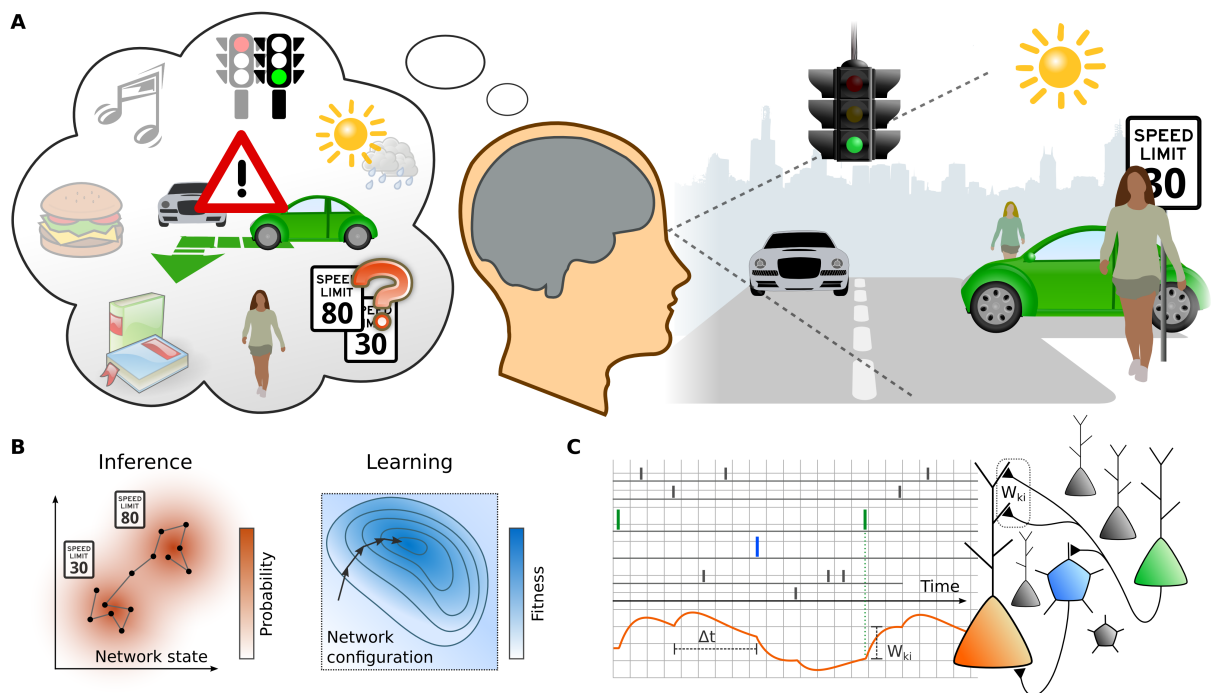
### Bayesian probability theory as a normative framework of brain computation

In the Bayesian brain, mammalian perception and reasoning are described by means of probability distributions that integrate instantaneous observations, internal goals or environmental context (a formal definition will be provided below). The success of the probabilistic description on the behavioral level naturally raises the question how probability distributions could be represented in the brain, and how the

## 1 Introduction

involved computations are supported by neural circuitry as a physical substrate. In recent years, several experimental studies identified potential neural correlates of probabilistic computations (T. YANG and SHADLEN, 2007; BERKES et al., 2011; FETSCH et al., 2012). In general, however, *in vivo* measurements are challenging due to the requirement to record from entire neural populations with high spatio-temporal resolution. Novel recording techniques, such as two photon imaging in behaving animals (HARVEY et al., 2012), could provide promising methods for future research. In order to facilitate targeted research, experimental work is being paralleled by theoretical studies that investigate possible neural implementations of probabilistic computations (R. S. ZEMEL et al., 1998; R. P. RAO, 2004; R. ZEMEL et al., 2005; DENEVE, 2005; MA et al., 2006; R. P. RAO, 2007; DENEVE, 2008b; STEIMER et al., 2009; PECEVSKI et al., 2011; SAVIN and DENEVE, 2014). Such theoretical studies focus especially on the algorithmic constraints imposed, for instance, by a spike-based communication, the spatially and temporally confined availability of information to neurons and synapses, or the integration and storage of information across multiple time scales. In this way, researchers from cognitive science, experimental and computational neuroscience cooperatively explore the qualification of a Bayesian perspective on neural processing as a unifying, normative framework of brain computation.

Another aspect of theoretical research addresses how the probabilistic relations, that underlie perception and reasoning, could emerge in the brain from experiences. Thoughts on probabilistic approaches to experience-dependent learning can be traced back to Thomas Bayes and Hermann von Helmholtz (WESTHEIMER, 2008). Interestingly, it was a young engineering discipline, namely statistical machine learning, that broached the topic again and became tremendously successful in training artificial information processing systems during the last decades (BISHOP, 2006; BARBER, 2012; SHUKOFF and AHLQUIST, 2012). Particularly in fields that deal with typical tasks for biological organisms, such as image classification (SALAKHUTDINOV and G. HINTON, 2009; SRIVASTAVA and SALAKHUTDINOV, 2012) or speech recognition (GRAVES et al., 2013), unsupervised statistical model optimization algorithms rank among the most powerful training methods. In recent years, multiple studies in computational neuroscience have set out to reintegrate probabilistic learning approaches from machine learning theory into biologically inspired models of neural networks (DENEVE, 2008a; NESSLER et al., 2009a; BREA et al., 2012; REZENDE et al., 2012; KECK et al., 2012; KAPPEL et al., 2014). Of particular interest are machine learning algorithms that can be formulated as distributed systems and which are structurally reminiscent of biological networks. For this thesis, we will focus on sample-based approaches since these were favored by theoretical work as being well-suited for neural learning (ACKLEY et al., 1985; FISER et al., 2010), were proposed to mimic the dynamics of neural networks (HOYER and HYVÄRINEN, 2003) and perception (GERSHMAN et al., 2009), were demonstrated to be applicable to a variety of tasks (G. HINTON, 2002; G. HINTON and SALAKHUTDINOV, 2006), and appear to be consistent with cortical firing statistics during learning (BERKES et al., 2011). Specifically, we will make use of sample-based approaches for maximum likelihood learning to explore self-organization principles in models of spiking neural networks.



**Fig. 1.1: Illustration of the thesis according to Marr's levels of analysis.** (A) Computational level: The Bayesian brain hypothesis proposes that humans and other mammals maintain a probabilistic model of their environment. Inference amounts to calculating a posterior distribution according to Bayes rule. Experience-dependent learning adapts the internal model to better suit the statistics of the environment. (B) Algorithmic/representational level: For inference, the network trajectory is understood as an MCMC sampling process. Self-organized learning is realized by small updates in the neuron and synapse parameters, thereby implementing an online version of a Generalized Expectation Maximization algorithm. (C) Physical level: The neural substrate for inference and learning is described on the level of individual spiking neurons which communicate via plastic synapses. To keep the analysis tractable, idealized neuron and synapses models are employed. Copyright information: The illustration uses artwork from [www.openclipart.org](http://www.openclipart.org) that had been released into the public domain.

## The rationale behind this thesis

The general ideas and research questions behind this thesis can be explained well within Marr's levels of analysis who proposed that a comprehensive understanding of information processing in the brain comprises three levels of analysis: a computational level, an algorithmic/representational level, and a physical level (MARR and VISION, 1982). The rationale of this thesis according to the three levels is illustrated in Fig. 1.1 and will be discussed in the following. On the *computational level*, sketched in Fig. 1.1A, we build upon the concept of generative models (BISHOP, 2006). The basic idea is that the brain maintains an abstract model of the environment, including both representations of real-world objects, as well as representations of less tangible entities such as velocities, predicted situations, goals, feelings or potential threats. Furthermore, the internal model of the environment captures causal and statistical relations among the entities. For the formal description, the abstract entities are enumerated by an index  $k$  and identified with

## 1 Introduction

*latent random variables*  $z_k$ . All knowledge about the expected state of entities – or formally the value of the associated random variables  $z_k$  – and about relations among them is described by a joint probability distribution  $p(\mathbf{z})$  over the entire latent state  $\mathbf{z} = (z_1, \dots, z_K)^T$ . The distribution  $p(\mathbf{z})$  is termed the *prior* as it is independent of sensory stimuli or real-world observations (also called data, input or observables). In the example of the traffic scene in Fig. 1.1A, the brain could maintain an abstract representation of traffic lights, combined with the prior knowledge that red and green signals are mutually exclusive. In order to connect the abstract representations to the data, the Bayesian brain maintains a model of the expected stimulus given a current state of the latent variables  $\mathbf{z}$ . To this end, also the observables are formally enumerated and associated with random variables  $y_i$ . The expectation, how likely an input configuration  $\mathbf{y}$  occurs for a latent state  $\mathbf{z}$ , is described by means of a *likelihood distribution*  $p(\mathbf{y} | \mathbf{z})$ . The question what entities are eventually considered to be latent or observable depends on the specific modeling attempt. To give a simplified example from vision, a variable  $y_i$  could represent the firing activity of a retinal ganglion cell, and variables  $z_k$  could be associated with face-selective cells in cortical area IT (PURVES et al., 2008). The full probabilistic model  $p(\mathbf{y}, \mathbf{z}) = p(\mathbf{z}) \cdot p(\mathbf{y} | \mathbf{z})$ , formed jointly by prior and likelihood, contains all knowledge required for (hypothetically) generating the expected distribution of inputs  $p(\mathbf{y})$  by eliminating the internal representations through marginalization,

$$p(\mathbf{y}) = \sum_{\mathbf{z}} p(\mathbf{y}, \mathbf{z}) , \quad (1.1)$$

where the summation runs over all possible latent state configurations. This property explains the term *generative model* and offers a mathematical correspondent to the more vivid “universe in our brain”. In the generative view, the latent variables  $\mathbf{z}$  play the role of hypothetical hidden causes of the input  $\mathbf{y}$ . For this reason, the terms latent variable and hidden cause are often used interchangeably for the variables  $z_k$ . A key benefit of generative models is that they provide a principled method for inferring the hidden causes  $\mathbf{z}$  behind an observation: For any given observation  $\mathbf{y}$ , the distribution  $p(\mathbf{y}, \mathbf{z})$  can be viewed as a function of  $\mathbf{z}$  that assigns beliefs to all possible states  $\mathbf{z}$  according to the laws of probability theory. This idea is formalized in Bayes rule,

$$p(\mathbf{z} | \mathbf{y}) = \frac{p(\mathbf{z}) \cdot p(\mathbf{y} | \mathbf{z})}{p(\mathbf{y})} , \quad (1.2)$$

which assigns a *posterior* belief  $p(\mathbf{z} | \mathbf{y})$  to each possible latent state  $\mathbf{z}$  for any given observation  $\mathbf{y}$ . The inference process consists of two components: The likelihood  $p(\mathbf{y} | \mathbf{z})$  aligns the internal state  $\mathbf{z}$  with the observation  $\mathbf{y}$ ; the prior  $p(\mathbf{z})$  reconciles the beliefs inferred during the first step with general knowledge on typical latent state configurations. The denominator  $p(\mathbf{y})$  only provides a normalization constant<sup>1</sup> and has little conceptual importance.

The traffic scene in Fig. 1.1A illustrates the inference process. The posterior, sketched in the thought bubble, captures the most essential aspects of the scene and neglects

---

<sup>1</sup>The denominator is fully determined by the numerator,  $p(\mathbf{y}) = \sum_{\mathbf{z}'} p(\mathbf{z} = \mathbf{z}') p(\mathbf{y} | \mathbf{z} = \mathbf{z}')$ , where the summation runs over all possible latent states  $\mathbf{z}'$ .

arguably unimportant details, such as the texture of wheel rims or the detailed shape of the traffic signal. Many generally possible abstract representations, e.g. books and music (grayed out), play no role in the inferred state. Others, such as a red traffic light or rainy weather, are actively suppressed since they are known to be irreconcilable with competing concepts promoted by the observation (green light and sunshine). Beyond these immediate conclusions, the posterior can also contain enriched information that is not directly present in the data, for instance, on the expected evolution of the scene and arising threats. The traffic scene example further illustrates a conceptual advantage of the Bayesian approach: Since the posterior assigns a belief not only to the most likely latent state configuration, but also to all possible alternatives, it naturally contains information on the (un-)certainty of the inferred result. This property turns out to be particularly useful in cases of incomplete or ambiguous observations, such as for the illustrated case of a partially occluded speed limit sign. In such ambiguous situations, the likelihood  $p(\mathbf{y} | z)$  will assign similar values to multiple possible latent states. In combination with previously obtained knowledge stored in the prior (e.g., the set of typical speed limits in the particular country), the likelihood contribution could be modulated, and thus yield a knowledgeable posterior belief on the different alternatives.

We next turn to the *algorithmic/representational level*, illustrated in Fig. 1.1B. In order to be a successful inference device, the Bayesian brain has to solve (at least) two general tasks. First, it is supposed to calculate and represent the posterior distribution in response to sensory stimuli – or, to be a little lenient, at least a close approximation  $q(z | \mathbf{y})$  to the correct posterior  $p(z | \mathbf{y})$ . Second, it must be capable of adapting its implicit generative model  $p(\mathbf{y}, z)$ , which underlies the inferences, to suit the true statistics of the environment well. For the formal description, we refer to the true statistics of the environment as  $p^*(\mathbf{y})$ . Again, we do not postulate a perfect match between the implicit model and reality. Instead, we assume that the generative model  $p(\mathbf{y}, z | \theta)$  can be shaped to some degree by a set of parameters  $\theta$ . Learning then amounts to adapting the parameters  $\theta$  such that  $p^*(\mathbf{y}) \approx p(\mathbf{y} | \theta) = \sum_z p(\mathbf{y}, z | \theta)$ .

Regarding the first task, the calculation and representation of the posterior distribution, we follow a sampling approach (see Fig. 1.1B left). The basic idea is to interpret the network trajectory as a sequence of latent state realizations  $z(t)$ . Then, the relative frequency, how often the network visits the different states  $z$  for a given stimulus  $\mathbf{y}$ , gives rise to a probability distribution  $q(z | \mathbf{y})$ . The statement that a network visits each state  $z$  proportionally to  $q(z | \mathbf{y})$  is, by definition, synonym to the statement that the network samples from the distribution  $q(z | \mathbf{y})$ . To describe the temporal evolution of the system, we will employ the mathematical framework of Markov chain Monte Carlo (MCMC) sampling that is often used in machine learning (ANDRIEU et al., 2003; BISHOP, 2006). In contrast to dynamical systems theory, which deals with the evolution of deterministic systems, MCMC techniques describe the evolution of a system by means of a stochastic transition operator. Through this fundamentally stochastic formulation, MCMC techniques are naturally suited to treat stochasticity as a computational resource (GERSHMAN et al.,

## 1 Introduction

2009). The resulting sampling process during inference is sketched in Fig. 1.1B: the network is expected to spend most of the time in regions of the state space that are associated with the different likely percepts, and to evade other regions that appear inconsistent with the data.

The second task, learning of the parameters  $\theta$ , addresses the self-organized adaptation of the generative model. We will pursue the simple objective that the learning process aims at refining the generative model to mirror the input distribution, i.e., to develop a faithful probabilistic model  $p(\mathbf{y} | \theta)$  of the outside world  $p^*(\mathbf{y})$ . For this thesis, we restrict the investigation to unsupervised, experience-dependent learning of a non-temporal generative model. In other words, we do not address learning of temporal sequences in the environment, nor do we include any form of external feedback, such as reward for actions. The rationale behind the focus on unsupervised model optimization as a first step is that a faithful internal representation of the environment can provide a foundation for complimentary learning approaches at a later stage. Such complimentary approaches, e.g., reinforcement learning algorithms (SUTTON and BARTO, 1998), could then operate on a sparse and stable representation of the input for temporal sequence learning, planning and goal-directed behavior. To measure the fitness of the full generative model  $p(\mathbf{y}, z | \theta)$  to the statistics of the environment  $p^*(\mathbf{y})$ , we employ a standard objective function from the machine learning literature, namely the *log-likelihood*  $\mathcal{L}(\theta)$  of the data under the model (BISHOP, 2006):

$$\mathcal{L}(\theta) := \langle \log p(\mathbf{y} | \theta) \rangle_{p^*(\mathbf{y})} = \sum_{\mathbf{y}} p^*(\mathbf{y}) \cdot \log \sum_z p(\mathbf{y}, z | \theta) . \quad (1.3)$$

The log-likelihood function  $\mathcal{L}(\theta)$  assigns a value to each possible configuration  $\theta$  of the generative model, and has a maximum when the internal model matches the true input distribution as closely as possible (in terms of the Kullback-Leibler divergence) within the parameter domain of  $\theta$ .

In summary, we pursue to unite two goals: (1) the maximization of  $\mathcal{L}(\theta)$  with respect of  $\theta$ , leading to a well-adapted generative model  $p(\mathbf{y}, z | \theta)$ , and (2) the ability to perform sample-based inference in this model, i.e.,  $q(z | \mathbf{y}) \approx p(z | \mathbf{y}, \theta)$ , through the inherent response dynamics of a spiking network. In order to achieve these goals, we will employ a *Generalized Expectation Maximization algorithm* (GEM), a powerful learning framework from machine learning theory (NEAL and G. HINTON, 1998; SATO, 1999). The fitness function  $\mathcal{F}(\theta, q)$  of GEM depends on the parameters  $\theta$  and on the variational posterior distribution  $q(z | \mathbf{y})$  of the network:

$$\mathcal{F}(\theta, q) = \mathcal{L}(\theta) - \langle D_{\text{KL}}(q(z | \mathbf{y}) || p(z | \mathbf{y}, \theta)) \rangle_{p^*(\mathbf{y})} . \quad (1.4)$$

The fitness  $\mathcal{F}(\theta, q)$  aims, on the one hand, to develop a faithful model by maximizing the log-likelihood and, on the other hand, to preserve the ability to exploit the model for Bayesian inference by minimizing the average Kullback-Leibler divergence between the network response and the analytically exact posterior distribution. Thus,  $\mathcal{F}(\theta, q)$  accommodates both of the above-described general tasks



of a Bayesian brain in a single objective function. For the algorithmic implementation, we will employ an online version of the GEM algorithm that operates on a continuous stream of input samples  $\mathbf{y}(t) \sim p^*(\mathbf{y})$  and increases  $\mathcal{F}(\boldsymbol{\theta}, q)$  by means of small parameter updates in  $\boldsymbol{\theta}$  and, simultaneously, small adaptations of the network distribution  $q(\mathbf{z} | \mathbf{y})$ .

Finally, on the *physical level* sketched in Fig. 1.1C, we seek for spiking network implementations which are specific in the sense that we aim to model neural activity on the level of individual neurons which evolve in continuous time and communicate via action potentials and plastic synapses. On this level of description, the particular connectivity structure of the network and the spread of information in the network via local synaptic interactions become essential aspects of the investigation. Indeed, the constraints on the local availability of information in space and time, imposed by a spike-based synaptic communication between neurons, display one of the cardinal challenges addressed in this thesis.

## Organization of the thesis

Each chapter of the thesis tackles a certain aspect of the broad question how collective network-wide function can emerge from purely local integration of information and communication. The chapters are based on peer-reviewed publications I authored together with Lars Büsing, Stefan Habenschuss, Bernhard Nessler, Wolfgang Maass, and Robert Legenstein. Detailed references and author contributions are provided at the beginning of each chapter.

In **chapter 2 “Neural dynamics as sampling”**, it is shown how a recurrent spiking neural network can sample from a well-defined joint probability distribution. To this end, neural dynamics are described in the theory of Markov chain Monte Carlo sampling. Each network neuron monitors the spiking activity of a subset of other, synaptically connected neurons and adapts its own instantaneous firing probability such that the network-wide spike pattern instantiates a Gibbs sampling-inspired sampling algorithm.

In **chapter 3 “Homeostatic and Hebbian plasticity united under a Bayesian account”**, it is explored how Hebbian synaptic plasticity can interact with forms of homeostatic plasticity. The chapter extends previous work of (NESSLER et al., 2009a; NESSLER et al., 2013) that linked STDP-type synaptic plasticity to the Expectation Maximization learning algorithm. Here it is shown that, in conjunction with local homeostatic processes, the network-wide effect of neural plasticity can be described in the framework of Generalized Expectation Maximization. The resulting network dynamics exhibit improved robustness during learning. Furthermore, the generalized theory extends the scope of application by overcoming several limiting assumptions of the previous work on the connectivity structure and the input presentation.

## 1 Introduction

In **chapter 4 “Statistical learning with compound memristive synapses”**, the investigation, how collective function can emerge from stochastic local computations, is extended to the structure of plastic synaptic connections. When the total synaptic weight between two neurons is composed of multiple plastic constituents, as sketched in Fig. 1.1C, changes in the efficacy of each individual constituent will affect the spike response of the post-synaptic neuron. This altered spike response can be observed by the other constituents and establishes an indirect link between the otherwise independent components. In the chapter it is shown that this seemingly weak link can be utilized for cooperative learning when plastic changes in the constituents occur stochastically. While originally developed for neuromorphic applications, the theory could also be applied to biological models where pairs of neurons can maintain multiple synaptic contacts.

In the final **chapter 5 “Distributed Bayesian computation and learning in spiking neural sheets”**, the neural sampling theory of chapter 2 and the plasticity dynamics of chapter 3 are combined in a spatially structured network model for Bayesian inference and self-organized learning.<sup>2</sup> The recurrent network model integrates distributed spiking input streams and unites aspects of competitive learning and associative learning in a continuous neural sheet. On the physical level, these learning aspects are established via local lateral inhibition and sparse recurrent excitation, respectively. While precise plasticity rules can be derived for input synapses, an approximate solution is proposed for recurrent connections. Over the course of learning, the network develops a sparse distributed spike code which carries compressed information on salient input features combined with knowledge on statistical relations among them. The ability to combine instantaneous observations with previously obtained knowledge within a well-defined generative model enables the network to perform sample-based inference even on incomplete or locally ambiguous input streams.

As we will show, connecting the physical level with the algorithmic/representational and computational level allows to embed the spike response of individual neurons (acting on the time scale of milliseconds) within a network-wide learning theory of experience-dependent plasticity (spanning time scales of minutes to hours). In order to keep this connection analytically tractable, simplifying modeling assumptions have to be made. The most important assumptions made throughout this thesis are as follows. For interpreting the neuronal spike response on the physical level as network states  $z(t)$  on the representational level, we will employ a straightforward one-to-one mapping: each latent variable  $z_k$  is directly associated with the spiking activity of one network neuron. For the neurons, we will employ a simple stochastic spike response point neuron model that emits action potentials based on an abstract membrane potential. For synaptic signaling, neurons communicate via non-additive rectangular post-synaptic potentials. In most chapters, we will only model excitatory neurons explicitly. The effect of interneurons, in contrast, will typically be described by abstract forms of inhibition. Plasticity enters the network models in

---

<sup>2</sup>The findings presented in this chapter have been submitted for publication and are under review by the time of writing this thesis.

two ways: STDP-type plasticity changes the weight of input synapses and recurrent synapses, and thereby adapts the parameters  $\theta$  of the generative model  $p(\mathbf{y}, \mathbf{z} | \theta)$ ; homeostatic plasticity controls the neurons' predisposition to spike in the absence of input, and thereby contributes to keeping the network response distribution  $q(\mathbf{z} | \mathbf{y})$  within a stable activity regime. Other forms of plasticity, such as short-term synaptic plasticity (MARKRAM et al., 1998), dendritic plasticity (TAVOSANIS, 2012), or third-factor rules (VASILAKI et al., 2009), are not addressed in the thesis. Further assumptions of the different models are described in the respective chapters.

## Other publications not included in the thesis

All chapters of this thesis employ roughly the same level of abstraction, using simple stochastic neuron models and highly idealized synaptic communication. On the one hand, the high level of abstraction permits a rigorous analytical treatment of inference and learning in the network. On the other hand, however, the abstract description disregards many well-established dynamical properties of neurons and synapses. This raises the question how the abstract models relate to distinct features of more elaborate neural models, and whether the identified computational phenomena are preserved in more detailed network models. These questions are addressed, at least to some extent, in two publications I contributed to in collaboration with colleagues from Heidelberg university. In PETROVICI et al. (2013), it is investigated how recurrent networks of leaky integrate-and-fire (LIF) neurons with conductance-based synapses and short-term plasticity dynamics can attain the correct firing statistics for sampling from known joint probability distributions, namely Boltzmann machines. To this study, I contributed as a functional first author, and the manuscript on our findings is under review by the time of writing this thesis. In a follow-up study (PROBST et al., 2015), networks of LIF neurons were endowed with the connectivity structure proposed by PECEVSKI et al. (2011) for sample-based inference in general Bayesian networks. Our findings were published in *Frontiers in computational neuroscience* in 2015.

## Relevance to computational neuroscience and neuromorphic engineering

The probabilistic models of computation and self-organization in recurrent networks of spiking neurons explored in this thesis could provide new insight into how computations can be organized in cortical microcircuits. Particularly, the neural sheet model investigated in chapter 5 bears some resemblance to salient connectivity motifs and response characteristics reported for superficial layers of cerebral cortex. Yet, in view of the evolutionary process our brains have undergone, it is unlikely that the brain "really" performs any mathematically pristine inference algorithm or that it "actually" pursues one particular universal learning objective. The main contribution of a comprehensive theoretical description of brain computation is found in a refinement of our conception of experimental results and in the guidance it provides for the design of future experiments.

## 1 Introduction

Beyond computational neuroscience, the network architectures explored in this thesis could provide promising operation principles for novel, neuro-inspired computing platforms. These platforms, often termed neuromorphic hardware, implement physical models of plastic neurons and synapses in nanoscale silicon architectures (MEAD and ISMAIL, 1989; SCHEMMEL et al., 2010) and complement traditional, strictly deterministic and serial, von Neumann computing designs. A key intention of neuromorphic engineering is the development of fault-tolerant, fundamentally parallel, self-adapting computing platforms. The network models of this thesis could likely contribute to this endeavor since network neurons evolve in parallel and exchange information fully asynchronously without a central clock signal via a spike event-based communication scheme. Furthermore, the simplicity of the employed idealized neuron and synapse models could facilitate neuromorphic implementations.

A major challenge for neuromorphic circuit design is the high-density integration of plastic synapse arrays (SCHEMMEL et al., 2006) which typically allocate most of the chip area. Recently, memristive materials (CHUA, 1971; STRUKOV et al., 2008; J. J. YANG et al., 2013) have gained increasing attention as potential plastic neuromorphic synapses. Memristors are novel nanoscale circuit elements which change their electrical conductance based on the history of voltage flux through the device. Due to this inherent plasticity, memristors are expected to overcome the requirement of extensive supporting plasticity circuitry in neuromorphic chip design (INDIVERI et al., 2013). In chapter 4, it is explored how a certain device class, namely memristors that exhibit stochastic switching between bistable states (GABA et al., 2013; SURI et al., 2013), could be utilized as plastic synapses for robust, self-organized statistical learning in spiking networks. On these grounds, the Bayesian spiking network models proposed in this thesis could likely contribute to the “software” of neuromorphic hardware platforms.

# Chapter 2

## Neural dynamics as sampling

### Contents

---

2.1	Introduction . . . . .	14
2.2	Recapitulation of MCMC sampling . . . . .	17
2.3	Neural sampling . . . . .	18
2.4	Neural sampling in discrete time . . . . .	21
2.5	Neural sampling with a relative refractory mechanism . . . . .	23
2.6	Neural sampling in continuous time . . . . .	27
2.7	Demonstration of perceptual multistability in sampling networks . . . . .	30
2.8	Discussion . . . . .	34

---

**Abstract.** The organization of computations in networks of spiking neurons in the brain is still largely unknown, in particular in view of the inherently stochastic features of their firing activity and the experimentally observed trial-to-trial variability of neural systems in the brain. In principle there exists a powerful computational framework for stochastic computations, probabilistic inference by sampling, which can explain a large number of macroscopic experimental data in neuroscience and cognitive science. But it has turned out to be surprisingly difficult to create a link between these abstract models for stochastic computations and more detailed models of the dynamics of networks of spiking neurons. Here we create such a link, and show that under some conditions the stochastic firing activity of networks of spiking neurons can be interpreted as probabilistic inference via Markov chain Monte Carlo (MCMC) sampling. Since common methods for MCMC sampling in distributed systems, such as Gibbs sampling, are inconsistent with the dynamics of spiking neurons, we introduce a different approach based on non-reversible Markov chains, that is able to reflect inherent temporal processes of spiking neuronal activity through a suitable choice of random variables. We propose a neural network model and show by a rigorous theoretical analysis that its neural activity implements MCMC sampling of a given distribution, both for the case of discrete and continuous time. This provides a step towards closing the gap between abstract functional models of cortical computation and more detailed models of networks of spiking neurons.

**Acknowledgments and author contributions.** This chapter is based on the publication

LARS BUESING, JOHANNES BILL, BERNHARD NESSLER, WOLFGANG MAASS (2011). “Neural Dynamics as Sampling: A Model for Stochastic Computation in Recurrent Networks of Spiking Neurons.” In: *PLoS Computational Biology*.

To this study, I contributed as second author. The neural sampling theory was developed by LB. The computer experiments were designed by LB, JB and WM; implemented and conducted by JB; and analyzed by LB and JB. All authors contributed to writing the manuscript. The authors thank Mihai Petrovici, Robert Legenstein and Samuel Gershman for helpful discussions.

### 2.1 Introduction

Attempts to understand the organization of computations in the brain from the perspective of traditional, mostly deterministic, models of computation, such as attractor neural networks or Turing machines, have run into problems: Experimental data suggests that neurons, synapses, and neural systems are inherently *stochastic* (ROLLS and DECO, 2010), especially in vivo, and therefore seem less suitable for implementing deterministic computations. This holds for ion channels of neurons (CANNON et al., 2010), synaptic release (FLIGHT, 2010), neural response to stimuli (trial-to-trial variability) (AZOUZ and GRAY, 1999; GERSTNER and KISTLER, 2002), and perception (BRASCAMP et al., 2006). In fact, several experimental studies arrive at the conclusion that external stimuli only modulate the highly stochastic spontaneous firing activity of cortical networks of neurons (FISER et al., 2004; RINGACH, 2009). Furthermore, traditional models for neural computation have been challenged by the fact that typical sensory data from the environment is often noisy and ambiguous, hence requiring neural systems to take *uncertainty* about external inputs into account. Therefore many researchers have suggested that information processing in the brain carries out probabilistic, rather than logical, inference for making decisions and choosing actions (S. GEMAN and D. GEMAN, 1984; R. P. N. RAO et al., 2002; DOYA et al., 2007; KÖRDING and WOLPERT, 2004; KERSTEN et al., 2004; GOPNIK and TENENBAUM, 2007; T. S. LEE and MUMFORD, 2003; HOYER and HYVÄRINEN, 2003; SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009; GRIFFITHS et al., 2008; T. YANG and SHADLEN, 2007; GOLD and SHADLEN, 2007; SADAGHIANI et al., 2010). Probabilistic inference has emerged in the 1960’s (PEARL, 1988), as a principled mathematical framework for reasoning in the face of uncertainty with regard to observations, knowledge, and causal relationships, which is characteristic for real-world inference tasks. This framework has become tremendously successful in real-world applications of artificial intelligence and machine learning. A typical computation that needs to be carried out for probabilistic inference on a high-dimensional joint distribution  $p(z_1, \dots, z_l, z_{l+1}, \dots, z_K)$  is the evaluation of the conditional distribution  $p(z_1, \dots, z_l | z_{l+1}, \dots, z_K)$  (or marginals

thereof) over some variables of interest, say  $z_1, \dots, z_l$ , given variables  $z_{l+1}, \dots, z_K$ . In the following, we will call the set of variables  $z_{l+1}, \dots, z_K$ , which we condition on, the *observed* variables and denote it by  $\mathbf{o}$ .<sup>1</sup>

Numerous studies in different areas of neuroscience and cognitive science have suggested that probabilistic inference could explain a variety of computational processes taking place in neural systems (see R. P. N. RAO et al. (2002); DOYA et al. (2007)). In models of perception the observed variables  $\mathbf{o}$  are interpreted as the sensory input to the central nervous system (or its early representation by the firing response of neurons, e.g., in the LGN in the case of vision), and the variables  $z_1, \dots, z_l$  model the interpretation of the sensory input, e.g., the texture and position of objects in the case of vision, which might be encoded in the response of neurons in various higher cortical areas (T. S. LEE and MUMFORD, 2003). Furthermore, in models for motor control the observed variables  $\mathbf{o}$  often consist not only of sensory and proprioceptive inputs to the brain, but also of specific goals and constraints for a planned movement (FRISTON et al., 2010; TOUSSAINT, 2009; TOUSSAINT and GOERICK, 2010), whereas inference is carried out over the variables  $z_1, \dots, z_l$  representing a motor plan or motor commands to muscles. Recent publications show that human reasoning and learning can also be cast into the form of probabilistic inference problems (TENENBAUM et al., 2006; GRIFFITHS and TENENBAUM, 2006; OAKSFORD and CHATER, 2007). In these models learning of concepts, ranging from concrete to more abstract ones, is interpreted as inference in lower and successively higher levels of hierarchical probabilistic models, giving a consistent description of inductive learning within and across domains of knowledge.

In spite of this active research on the functional level of neural processing, it turned out to be surprisingly hard to relate the computational machinery required for probabilistic inference to experimental data on neurons, synapses, and neural systems. There are mainly two different approaches for implementing the computational machinery for probabilistic inference in “neural hardware”. The first class of approaches builds on deterministic methods for evaluating exactly or approximately the desired conditional and/or marginal distributions, whereas the second class relies on sampling from the probability distributions in question. Multiple models in the class of deterministic approaches implement algorithms from machine learning called message passing or belief propagation (R. P. RAO, 2007; STEIMER et al., 2009; DENEVE, 2008b; LITVAK and ULLMAN, 2009). By clever reordering of sum and product operators occurring in the evaluation of the desired probabilities, the total number of computation steps are drastically reduced. The results of subcomputations are propagated as “messages” or “beliefs” that are sent to other parts of the computational network. Other deterministic approaches for representing distributions and performing inference are probabilistic population

---

<sup>1</sup>In this chapter, the probability distribution a network samples from is directly denoted by  $p$  (instead of  $q$  as suggested in the introduction of the thesis). The distinction between the two distributions  $p$  and  $q$  will become important only in the context of learning, as addressed in subsequent chapters. Furthermore, we denote the subset of observed variables by  $\mathbf{o}$  in order to distinguish them from inputs  $\mathbf{y}$ , which will be modeled explicitly as being external in later chapters.

## 2 Neural dynamics as sampling

code (PPC) models (SAHANI and DAYAN, 2003). Although deterministic approaches provide a theoretically sound hypothesis about how complex computations can possibly be embedded in neural networks and explain aspects of experimental data, it seems difficult (though not impossible) to conciliate them with other aspects of experimental evidence, such as stochasticity of spiking neurons, spontaneous firing, trial-to-trial variability, and perceptual multistability.

Therefore other researchers (e.g., FISER et al. (2010); HOYER and HYVÄRINEN (2003); SUNDARESWARA and SCHRATER (2008); GERSHMAN et al. (2009)) have proposed to model computations in neural systems as probabilistic inference based on a different class of algorithms, which requires stochastic, rather than deterministic, computational units. This approach, commonly referred to as sampling, focuses on drawing *samples*, i.e., concrete values for the random variables that are distributed according to the desired probability distribution. Sampling can naturally capture the effect of apparent stochasticity in neural responses and seems to be furthermore consistent with multiple experimental effects reported in cognitive science literature (SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009). On the conceptual side, it has proved to be difficult to implement learning in message passing and PPC network models. In contrast, following the lines of (ACKLEY et al., 1985), the sampling approach might be well suited to incorporate learning.

Previous network models that implement sampling in neural networks are mostly based on a special sampling algorithm called Gibbs (or general Metropolis-Hastings) sampling (S. GEMAN and D. GEMAN, 1984; SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009; G. E. HINTON et al., 2006). The dynamics that arise from this approach, the so-called Glauber dynamics, however are only superficially similar to spiking neural dynamics observed in experiments, rendering these models rather abstract. Building on and extending previous models, we propose here a family of network models, that can be shown to exactly sample from any arbitrary member of a well-defined class of probability distributions via their inherent network dynamics. These dynamics incorporate refractory effects and finite durations of postsynaptic potentials (PSPs), and are therefore more biologically realistic than existing approaches. Formally speaking, our model implements Markov chain Monte Carlo (MCMC) sampling in a spiking neural network. In contrast to prior approaches however, our model incorporates irreversible dynamics (i.e., no detailed balance) allowing for finite time PSPs and refractory mechanisms. Furthermore, we also present a continuous time version of our network model. The resulting stochastic dynamical system can be shown to sample from the correct distribution. In general, continuous time models arguably provide a higher amount of biological realism compared to discrete time models.

The paper is structured in the following way. First we provide a brief introduction to MCMC sampling. We then define the neural network model whose neural activity samples from a given class of probability distributions. The model will be first presented in discrete time together with some illustrative simulations. An extension of the model to networks of more detailed spiking neuron models which feature a relative refractory mechanism is presented. Furthermore, it is shown how the



neural network model can also be formulated in continuous time. Finally, as a concrete simulation example we present a simple network model for perceptual multistability.

## 2.2 Recapitulation of MCMC sampling

In machine learning, sampling is often considered the “gold standard” of inference methods, since, assuming that we can sample from the distribution in question, and assuming enough computational resources, any inference task can be carried out with arbitrary precision (in contrast to some deterministic approximate inference methods such as variational inference). However sampling from an arbitrary distribution can be a difficult problem in itself, as, e.g., many distributions can only be evaluated modulo a global constant (the partition function). In order to circumvent these problems, elaborate MCMC sampling techniques have been developed in machine learning and statistics (ANDRIEU et al., 2003). MCMC algorithms are based on the following idea: instead of producing an ad-hoc sample, a process that is heuristically comparable to a global search over the whole state space of the random variables, MCMC methods produce a new sample via a “local search” around a point in the state space that is already (approximately) a sample from the distribution.

More formally, a Markov chain  $M$  (in discrete time) is defined by a set  $S$  of states (we consider for discrete time only the case where  $S$  has a finite size, denoted by  $|S|$ ) together with a transition operator  $T$ . The operator  $T$  is a conditional probability distribution  $T(s|s')$  over the next state  $s$  given a preceding state  $s'$ . The Markov chain  $M$  is started in some initial state  $s(0)$ , and moves through a trajectory of states  $s(t)$  via iterated application of the stochastic transition operator  $T$ . More precisely, if  $s(t-1)$  is the state at time  $t-1$ , then the next state  $s(t)$  is drawn from the conditional probability distribution  $T(s|s(t-1))$ . An important theorem from probability theory (see, e.g., p. 232 in (GRIMMETT and STIRZAKER, 2001)) states that if  $M$  is irreducible (i.e., any state in  $S$  can be reached from any other state in  $S$  in finitely many steps with probability  $> 0$ ) and aperiodic (i.e., its state transitions cannot be trapped in deterministic cycles), then the probability  $p(s(t) = s|s(0))$  converges for  $t \rightarrow \infty$  to a probability  $p(s)$  that does not depend on the initial state  $s(0)$ . This state distribution  $p$  is called the invariant distribution of  $M$ . The irreducibility of  $M$  implies that it is the only distribution over the states  $S$  that is invariant under its transition operator  $T$ , i.e.

$$p(s) = \sum_{s' \in S} T(s|s') \cdot p(s') \quad . \quad (2.1)$$

Thus, in order to carry out probabilistic inference for a given distribution  $p$ , it suffices to construct an irreducible and aperiodic Markov chain  $M$  that leaves  $p$  invariant, i.e., satisfies equation (2.1). Then one can answer numerous probabilistic

## 2 Neural dynamics as sampling

inference questions regarding  $p$  without any numerical computations of probabilities. Rather, one plugs in the observed values for some of the random variables (RVs) and simply collects samples from the conditional distribution over the other RVs of interest when the Markov chain approaches its invariant distribution.

A convenient and popular method for the construction of an operator  $T$  for a given distribution  $p$  is looking for operators  $T$  that satisfy the following detailed balance condition,

$$T(s|s') \cdot p(s') = T(s'|s) \cdot p(s) \quad (2.2)$$

for all  $s, s' \in S$ . A Markov chain that satisfies (2.2) is said to be reversible. In particular, the Gibbs and Metropolis-Hastings algorithms employ reversible Markov chains. A very useful property of (2.2) is that it implies the invariance property (2.1), and this is in fact the standard method for proving (2.1). However, as our approach makes use of irreversible Markov chains as explained below, we will have to prove (2.1) directly.

### 2.3 Neural sampling

Let  $p(z_1, \dots, z_K)$  be some arbitrary joint distribution over  $K$  binary variables  $z_1, \dots, z_K$  that only takes on values  $> 0$ . We will show that under a certain computability assumption on  $p$  a network  $\mathcal{N}$  consisting of  $K$  spiking neurons  $\nu_1, \dots, \nu_K$  can sample from  $p$  using its inherent stochastic dynamics. More precisely, we show that the stochastic firing activity of  $\mathcal{N}$  can be viewed as a non-reversible Markov chain that samples from the given probability distribution  $p$ . If a subset  $\mathcal{o}$  of the variables are observed, modeled as the corresponding neurons being “clamped” to the observed values, the remaining network samples from the conditional distribution of the remaining variables given the observables. Hence, this approach offers a quite natural implementation of probabilistic inference. It is similar to sampling approaches which have already been applied extensively, e.g., in Boltzmann machines, however our model is more biologically realistic as it incorporates aspects of the inherent temporal dynamics and spike-based communication of a network of spiking neurons. We call this approach *neural sampling* in the remainder of the paper.

In order to enable a network  $\mathcal{N}$  of spiking neurons to sample from a distribution  $p(z_1, \dots, z_K)$  of binary variables  $z_k$ , one needs to specify how an assignment  $(z_1, \dots, z_K) \in \{0, 1\}^K$  of values to these binary variables can be represented by the spiking activity of the network  $\mathcal{N}$  and vice versa. A spike, or action potential, of a biological neuron  $\nu_k$  has a short duration of roughly 1 ms. But the effect of such spike, both on the neuron  $\nu_k$  itself (in the form of refractory processes) and on the membrane potential of other neurons (in the form of postsynaptic potentials) lasts substantially longer, on the order of 5 ms to 100 ms. In order to capture this temporally extended effect of each spike, we fix some parameter  $\tau$  that models the average duration of these temporally extended processes caused by a spike. We say

that a binary vector  $(z_1, \dots, z_K)$  is represented by the firing activity of the network  $\mathcal{N}$  at time  $t$  for  $k = 1, \dots, K$  iff:

$$z_k(t) = 1 \iff \nu_k \text{ has fired within the time interval } (t - \tau, t]. \quad (2.3)$$

In other words, any spike of neuron  $\nu_k$  sets the value of the associated binary variable  $z_k$  to 1 for a duration of length  $\tau$ .

An obvious consequence of this definition is that the binary vector  $(z_1, \dots, z_K)$  that is defined by the activity of  $\mathcal{N}$  at time  $t$  does not fully capture the internal state of this stochastic system. Rather, one needs to take into account additional non-binary variables  $(\zeta_1, \dots, \zeta_K)$ , where the value of  $\zeta_k$  at time  $t$  specifies *when* within the time interval  $(t - \tau, t]$  the neuron  $\nu_k$  has fired (if it has fired within this time interval, thereby causing  $z_k = 1$  at time  $t$ ). The neural sampling process has the Markov property only with regard to these more informative auxiliary variables  $\zeta_1, \dots, \zeta_K$ . Therefore our analysis of neural sampling will focus on the temporal evolution of these auxiliary variables. We adopt the convention that each spike of neuron  $\nu_k$  sets the value of  $\zeta_k$  to its maximal value  $\tau$ , from which it linearly decays back to 0 during the subsequent time interval of length  $\tau$ .

For the construction of the sampling network  $\mathcal{N}$ , we assume that the membrane potential  $u_k(t)$  of neuron  $\nu_k$  at time  $t$  equals the log-odds of the corresponding variable  $z_k$  to be active, and refer to this property as *neural computability condition*:

$$u_k(t) = \log \frac{p(z_k = 1 | \mathbf{z}_{\setminus k})}{p(z_k = 0 | \mathbf{z}_{\setminus k})} \quad , \quad (2.4)$$

where we write  $z_k$  for  $z_k(t)$  and  $\mathbf{z}_{\setminus k}$  for the current values  $z_i(t)$  of all other variables  $z_i$  with  $i \neq k$ . Under the assumption we make in equation (2.4), i.e., that the neural membrane potential reflects the log-odds of the corresponding variable  $z_k$ , it is required that each single neuron in the network can actually compute the right-hand side of equation (2.4), i.e., that it fulfills the neural computability condition.

A concrete class of probability distributions, that we will use as an example in the remainder, are Boltzmann distributions:

$$p(\mathbf{z}) = \frac{1}{Z} \exp \left( \sum_{ij} \frac{1}{2} W_{ij} z_i z_j + \sum_i b_i z_i \right) \quad (2.5)$$

with arbitrary real valued parameters  $b_i, W_{ij}$  which satisfy  $W_{ij} = W_{ji}$  and  $W_{ii} = 0$  (the constant  $Z$  ensures the normalization of  $p(\mathbf{z})$ ). For the Boltzmann distribution, condition (2.4) is satisfied by neurons  $\nu_k$  with the standard membrane potential


$$u_k(t) = b_k + \sum_{i=1}^K W_{ki} z_i(t) \quad , \quad (2.6)$$

## 2 Neural dynamics as sampling

where  $b_k$  is the bias of neuron  $v_k$  (which regulates its excitability),  $W_{ki}$  is the strength of the synaptic connection from neuron  $v_i$  to  $v_k$ , and  $W_{ki}z_i(t)$  approximates the time course of the postsynaptic potential in neuron  $v_k$  caused by a firing of neuron  $v_i$  with a constant signal of duration  $\tau$  (i.e., a square pulse). As we will describe below, spikes of neuron  $v_k$  are evoked stochastically depending on the current membrane potential  $u_k$  and the auxiliary variable  $\zeta_k$ .

The neural computability condition (2.4) links classes of probability distributions to neuron and synapse models in a network of spiking neurons. As shown above, Boltzmann distributions satisfy the condition if one considers point neuron models which compute a linear weighted sum of the presynaptic inputs. The class of distributions can be extended to include more complex distributions using a method proposed in (NESSLER et al., 2009b) which is based on the following idea. Neuron  $v_k$  representing the variable  $z_k$  is not directly influenced by the activities  $z_{\setminus k}$  of the presynaptic neurons, but via intermediate nonlinear preprocessing elements. This preprocessing might be implemented by dendrites or other (inter-) neurons and is assumed to compute nonlinear combinations of the presynaptic activities  $z_{\setminus k}$  (similar to a kernel). This allows the membrane potential  $u_k$ , and therefore the log-odds ratio on the right-hand side of (2.4), to represent a more complex function of the activities  $z_{\setminus k}$ , giving rise to more complex joint distributions  $p(\mathbf{z})$ . The concrete implementation of non-trivial directed and undirected graphical models with the help of preprocessing elements in the neural sampling framework is subject of current research. For the examples given in this study, we focus on the standard form of the membrane potential (2.6) of point neurons. As shown below, these spiking network models can emulate any Boltzmann machine (BM) (ACKLEY et al., 1985).

A substantial amount of preceding studies has demonstrated that BMs are very powerful, and that the application of suitable learning algorithms for setting the weights  $W_{ij}$  makes it possible to learn and represent complex sensory processing tasks by such distributions (G. E. HINTON et al., 2006; G. E. HINTON, 2010). In applications in statistics and machine learning using such Boltzmann distributions, sampling is typically implemented by Gibbs sampling or more general *reversible* MCMC methods. However, it is difficult to model some neural processes, such as an absolute refractory period or a postsynaptic potential (PSP) of fixed duration, using a reversible Markov chain, but they are more conveniently modeled using an irreversible one. As we wish to keep the computational power of BMs and at the same time to augment the sampling procedure with aspects of neural dynamics (such as PSPs with fixed durations, refractory mechanisms) to increase biological realism, we focus in the following on irreversible MCMC methods (keeping in mind that this might not be the only possible way to achieve these goals).



**Fig. 2.1: Neuron model with absolute refractory mechanism.** The figure shows a schematic of the transition operator  $T^k$  for the internal state variable  $\zeta_k$  of a spiking neuron  $v_k$  with an absolute refractory period. The neuron can fire in the resting state  $\zeta_k = 0$  and in the last refractory state  $\zeta_k = 1$ .

## 2.4 Neural sampling in discrete time

Here we describe neural dynamics in discrete time with an absolute refractory period  $\tau$ . We interpret one step of the Markov chain as a time step  $dt$  in biological real time. The dynamics of the variable  $\zeta_k$ , that describes the time course of the effect of a spike of neuron  $v_k$ , are defined in the following way.  $\zeta_k$  is set to the value  $\tau$  when neuron  $v_k$  fires, and decays by 1 at each subsequent discrete time step. The parameter  $\tau$  is chosen to be some integer, so that  $\zeta_k$  decays back to 0 in exactly  $\tau$  time steps. The neuron can only spike (with a probability that is a function of its current membrane potential  $u_k$ ) if its variable  $\zeta_k \leq 1$ . If however,  $\zeta_k > 1$ , the neuron is considered refractory and it cannot spike, but its  $\zeta_k$  is reduced by 1 per time step. To show that these simple dynamics do indeed sample from the given distribution  $p(z)$ , we proceed in the following way. We define a joint distribution  $p(\zeta, z)$  which has the desired marginal distribution  $\sum_{\zeta} p(\zeta, z) = p(z)$ . Further we formalize the dynamics informally described above as a transition operator  $T$  operating on the state vector  $(\zeta, z)$ . In Appendix B, we show that  $p(\zeta, z)$  is the unique invariant distribution of this operator  $T$ , i.e., that the dynamics described by  $T$  produce samples  $z$  from the desired distribution  $p(z)$ . We refer to sampling through networks with this stochastic spiking mechanism as *neural sampling with absolute refractory period* due to the persistent refractory process.

Given the distribution  $p(z)$  that we want to sample from, we define the following joint distribution  $p(\zeta, z)$  over the neural variables:

$$p(\zeta, z) := p(\zeta|z) \cdot p(z) \quad \text{with} \quad p(\zeta|z) := \prod_{k=1}^K p(\zeta_k|z_k)$$

$$\text{where} \quad p(\zeta_k|z_k) := \begin{cases} \tau^{-1} & \text{for} \quad z_k = 1 \wedge \zeta_k > 0 \\ 1 & \text{for} \quad z_k = 0 \wedge \zeta_k = 0 \\ 0 & \text{otherwise} \end{cases} \quad (2.7)$$

This definition of  $p(\zeta_k|z_k)$  simply expresses that if  $z_k = 1$ , then the auxiliary variable  $\zeta_k$  can assume any value in  $\{1, 2, \dots, \tau\}$  with equal probability. On the other hand  $\zeta_k$  necessarily assumes the value 0 if  $z_k = 0$  (i.e., when the neuron is in its resting state).

The state transition operator  $T$  can be defined in a transparent manner as a composition of  $K$  transition operators,  $T = T^1 \circ \dots \circ T^K$ , where  $T^k$  only updates the

## 2 Neural dynamics as sampling

variables  $\zeta_k$  and  $z_k$  of neuron  $v_k$ , i.e., the neurons are updated sequentially in the same order (this severe restriction will become obsolete in the case of continuous time discussed below). We define the composition as  $(T^k \circ T^l)(\cdot) = (T^k(T^l(\cdot)))$ , i.e.,  $T^l$  is applied prior to  $T^k$ . The new values of  $\zeta_k$  and  $z_k$  only depend on the previous value  $\zeta'_k$  and on the current membrane potential  $u_k(z_{\setminus k})$ . The interesting dynamics take place in the variable  $\zeta_k$ . They are illustrated in Fig. 2.1, where the arrows represent transition probabilities greater than 0.

If the neuron  $v_k$  is not refractory, i.e.,  $\zeta'_k \leq 1$ , it can spike (i.e., a transition from  $\zeta'_k \leq 1$  to  $\zeta_k = \tau$ ) with probability

$$T^k(\zeta_k = \tau | \zeta'_k, z_{\setminus k}) = \sigma(u_k - \log \tau) \quad , \quad (2.8)$$

where  $\sigma(x) = (1 + e^{-x})^{-1}$  is the standard sigmoidal activation function and the log denotes the natural logarithm. The term  $u_k$  is the current membrane potential, which depends on the current values of the variables  $z_i$  for  $i \neq k$ . The term  $\log \tau$  in (2.8) reflects the granularity of a chosen discrete time scale. If it is very fine (say one step equals one microsecond), then  $\tau$  is large, and the firing probability at each specific discrete time step is therefore reduced. If the neuron in a state with  $\zeta'_k \leq 1$  does not spike,  $\zeta_k$  relaxes into the resting state  $\zeta_k = 0$  corresponding to a non-refractory neuron.

If the neuron is in a refractory state, i.e.,  $\zeta'_k > 1$ , its new variable  $\zeta_k$  assumes deterministically the next lower value  $\zeta_k = \zeta'_k - 1$ , reflecting the inherent temporal process:

$$T^k(\zeta_k = \zeta'_k - 1 | \zeta'_k, z_{\setminus k}) = 1 \quad . \quad (2.9)$$

After the transition of the auxiliary variable  $\zeta_k$ , the binary variable  $z_k$  is deterministically set to a consistent state, i.e.,  $z_k = 1$  if  $\zeta_k \geq 1$  and  $z_k = 0$  if  $\zeta_k = 0$ .

It can be shown that each of these stochastic state transition operators  $T^k$  leaves the given distribution  $p$  invariant, i.e., satisfies equation (2.1). This implies that any composition or mixture of these operators  $T^k$  also leaves  $p$  invariant, see, e.g., (ANDRIEU et al., 2003). In particular, the composition  $T = T^1 \circ \dots \circ T^K$  of these operators  $T^k$  leaves  $p$  invariant, which has a quite natural interpretation as firing dynamics of the spiking neural network  $\mathcal{N}$ : At each discrete time step the variables  $\zeta_k, z_k$  are updated for all neurons  $v_k$ , where the update of  $\zeta_k, z_k$  takes preceding updates for  $\zeta_i, z_i$  with  $i > k$  into account. Alternatively, one could also choose at each discrete time step a different order for updates according to (ANDRIEU et al., 2003). The assumption of a well-regulated updating policy will be overcome in the continuous-time limit, i.e., in case where the neural dynamics are described as a Markov jump process. In Appendix B, we prove the following central theorem:

**Theorem 1.**  $p(\zeta, z)$  is the unique invariant distribution of operator  $T$ , i.e.,  $T$  is aperiodic and irreducible and satisfies

$$p(\zeta, z) = \sum_{\zeta', z'} T(\zeta, z | \zeta', z') \cdot p(\zeta', z') \quad . \quad (2.10)$$

## 2.5 Neural sampling with a relative refractory mechanism

The proof of this Theorem is provided by Lemmata 1 – 3 in Appendix B. The statement that  $T$  (which is composed of the operators  $T^k$ ) is irreducible and aperiodic ensures that  $p$  is the *unique* invariant distribution of the Markov chain defined by  $T$ , i.e., that irrespective of the initial network state the successive application of  $T$  explores the whole state space in a non-periodic manner.

This theorem guarantees that after a sufficient “burn-in” time (more precisely in the limit of an infinite “burn-in” time), the dynamics of the network, which are given by the transition operator  $T$ , produce samples from the distribution  $p(\zeta, z)$ . As by construction  $\sum_{\zeta} p(\zeta, z) = p(z)$ , the Markov chain provides samples from the given distribution  $p(z)$ . Furthermore, the network  $\mathcal{N}$  can carry out probabilistic inference for this distribution. For example,  $\mathcal{N}$  can be used to sample from the posterior distribution  $p(z_1 \dots, z_l | z_{l+1}, \dots, z_K)$  over  $z_1 \dots, z_l$  given  $z_{l+1}, \dots, z_K$ . One just needs to clamp those neurons  $v_{l+1}, \dots, v_K$  to the corresponding observed values. This could be implemented by injecting a strong positive (negative) current into the units with  $z_j = 1$  ( $z_j = 0$ ). Then, as soon as the stochastic dynamics of  $\mathcal{N}$  has converged to its invariant distribution, the averaged firing rate of neuron  $v_1$  is proportional to the following desired marginal probability

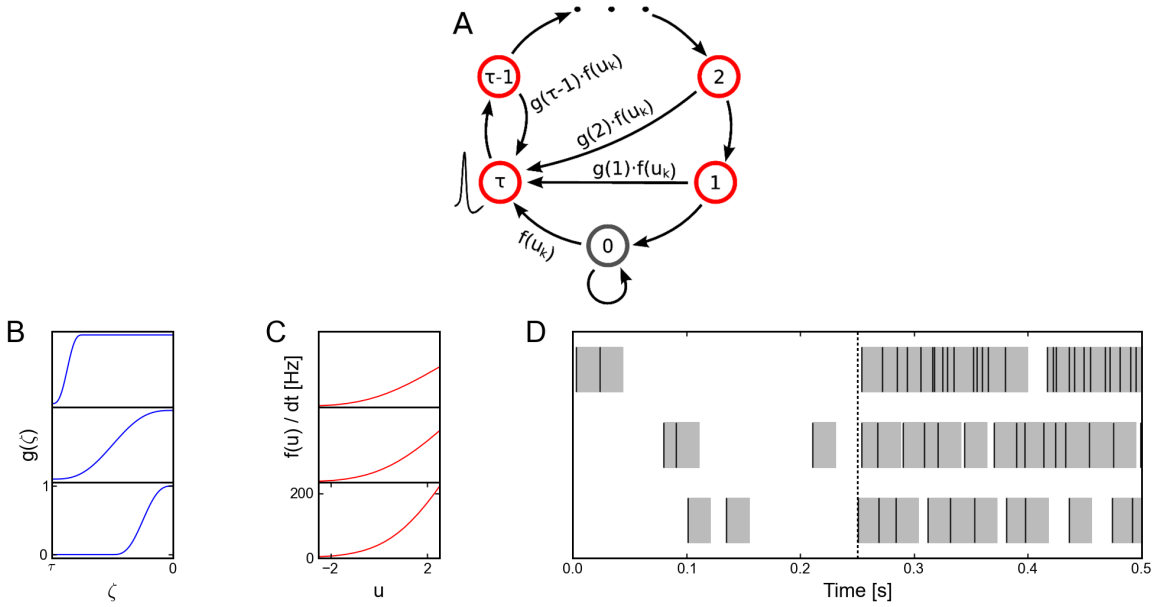
$$p(z_1 = 1 | z_{l+1}, \dots, z_K) = \sum_{z_2, \dots, z_l} p(z_1 = 1, z_2, \dots, z_l | z_{l+1}, \dots, z_K) \quad .$$

In a biological neural system this result of probabilistic inference could for example be read out by an integrator neuron that counts spikes from this neuron  $v_1$  within a behaviorally relevant time window of a few hundred milliseconds, similarly as the experimentally reported integrator neurons in area LIP of monkey cortex (T. YANG and SHADLEN, 2007; GOLD and SHADLEN, 2007). Another readout neuron that receives spike input from  $v_k$  could at the same time estimate  $p(z_k = 1 | z_{l+1}, \dots, z_K)$  for another RV  $z_k$ . But valuable information for probabilistic inference is not only provided by firing rates or spike counts, but also by spike correlations of the neurons  $v_1, \dots, v_l$  in  $\mathcal{N}$ . For example, the probability  $p(z_1 = 1, z_2 = 1 | z_{l+1}, \dots, z_K)$  can be estimated by a readout neuron that responds to superpositions of EPSPs caused by near-coincident firing of neurons  $v_1$  and  $v_2$  within a time interval of length  $\tau$ . Thus, a large number of different probabilistic inferences can be carried out efficiently in parallel by readout neurons that receive spike input from different subsets of neurons in the network  $\mathcal{N}$ .

## 2.5 Neural sampling with a relative refractory mechanism

For the previously described simple neuron model, the refractory process was assumed to last for  $\tau$  time steps, exactly as long as the postsynaptic potentials caused by each spike. In this section we relax this assumption by introducing a more complex and biologically more realistic neuron model, where the duration of

## 2 Neural dynamics as sampling



**Fig. 2.2: Neuron model with relative refractory mechanism.** The figure shows the transition operator  $T^k$ , refractory functions  $g$  and activation functions  $f$  for the neuron model with relative refractory mechanism. (A) Transition probabilities of the internal variable  $\zeta_k$  given by  $T^k$ . (B) Three examples of possible refractory functions  $g$ . They assume value 0 when the neuron cannot spike, and return to value 1 (full readiness to fire again) with different time courses. The value of  $g$  at intermediate time points regulates the current probability of firing of neuron  $v_k$  (see A). The x-axis is equivalent to the number of time steps since last spike (running from 0 to  $\tau$  from left to right). (C) Associated activation functions  $f$  according to (2.11). (D) Spike trains produced by the resulting three different neuron models with (hypothetical) membrane potentials that jump at time 0.25s from a constant low value to a constant high value. Black horizontal bars indicate spikes, and the active states  $z_k = 1$  are indicated by gray shaded areas of duration  $\tau \cdot dt = 20\text{ms}$  after each spike. It can be seen from this example that different refractory mechanisms give rise to different spiking dynamics.

the refractory process is decoupled from the duration  $\tau$  of a postsynaptic potential. Thus, this model can for example also fire bursts of spikes with an interspike interval  $< \tau$ . The introduction of this more complex neuron model comes at the price that one can no longer prove that a network of such neurons samples from the desired distribution  $p$ . Nevertheless, if the sigmoidal activation function  $\sigma$  is replaced by a different activation function  $f$ , one can still prove that the sampling is “locally correct”, as specified in equation (2.12) below. Furthermore, our computer simulations suggest that also globally the error introduced by the more complex neuron model is not functionally significant, i.e. that statistical dependencies between the RVs  $z$  are still faithfully captured.

The neuron model with a relative refractory period is defined in the following way. Consider some arbitrary refractory function  $g : [0, \dots, \tau] \rightarrow \mathbb{R}$  with  $g(\tau) = 0$ ,  $g(0) = 1$ , and  $g(l) \geq 0$  for  $l = 1, \dots, \tau - 1$ . The idea is that  $g(\zeta_k)$  models the readiness of the neuron to fire in its state  $\zeta_k$ . This readiness has value 0 when the neuron has fired at the preceding time step (i.e.,  $\zeta_k = \tau$ ), and assumes the resting state 1 when  $\zeta_k$  has dropped to 0. In between, the readiness may take on any non-negative value according to the function  $g(\zeta_k)$ . The function  $g$  does



## 2.5 Neural sampling with a relative refractory mechanism

not need to be monotonic, allowing for example that it increases to high values in between, yielding a preferred interspike interval of an oscillatory neuron. The firing probability of neuron  $v_k$  in state  $\zeta_k$  is given by  $g(\zeta_k) \cdot f(u_k)$ , where  $f(u_k)$  is an appropriate function of the membrane potential as described below. Thus this function  $g$  is closely related to the function  $\eta$  (called afterpotential) in the spike response model (GERSTNER and KISTLER, 2002) as well as to the self-excitation kernel in Generalized Linear Models (PILLOW et al., 2008). In general, different neurons in the network may have different refractory profiles, which can be modeled by a different refractory function for each neuron  $v_k$ . However for the sake of notational simplicity we assume a single refractory function in the following.

In the presence of this refractory function  $g$  one needs to replace the sigmoidal activation function  $\sigma(u_k - \log \tau)$  by a suitable function  $f(u_k)$  that satisfies the condition

$$\exp(u) = f(u) \frac{\sum_{\eta=1}^{\tau} \prod_{\zeta=\eta+1}^{\tau} (1 - g(\zeta) \cdot f(u))}{\prod_{\zeta=1}^{\tau} (1 - g(\zeta) \cdot f(u))} \quad (2.11)$$

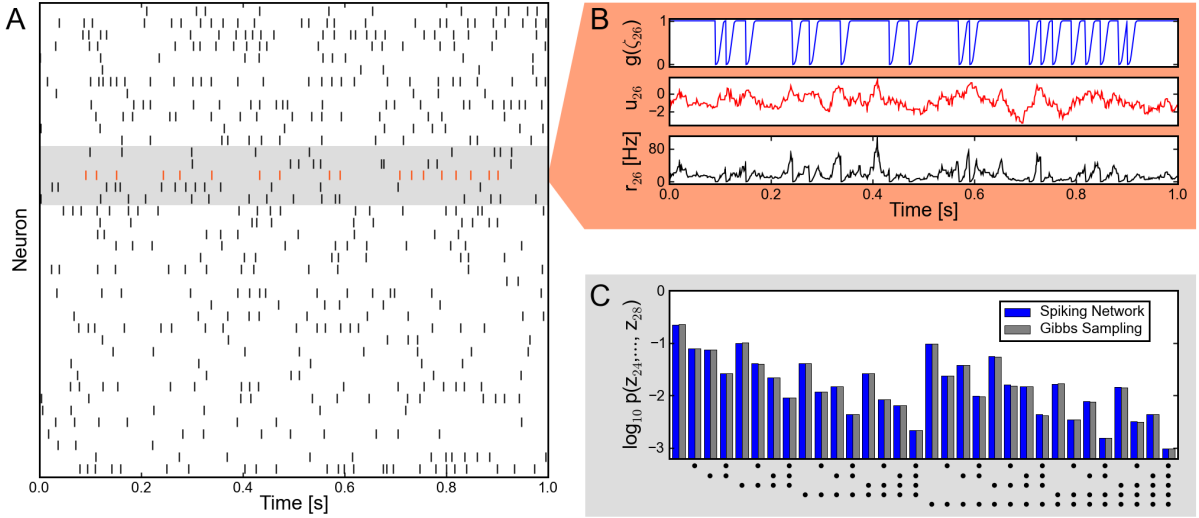
for all real numbers  $u$ . This equation can be derived (see Appendix B, Lemma 5) if one requires each neuron  $v_k$  to represent the correct distribution  $p(z_k | z_{\setminus k})$  over  $z_k$  conditioned the variables  $z_{\setminus k}$ . One can show that, for any  $g$  as above, there always exists a continuous, monotonic function  $f$  which satisfies this equation (see Lemma 4 in Appendix B). Unfortunately (2.11) cannot be solved analytically for  $f$  in general. Hence, for simulations we approximate the function  $f$  for a given  $g$  by numerically solving (2.11) on a grid and interpolating between the grid points with a constant function. Examples for several functions  $g$  and the associated  $f$  are shown in Fig. 2.2B and Fig. 2.2C respectively. Furthermore, spike trains emitted by single neurons with these refractory functions  $g$  and the corresponding functions  $f$  are shown in Fig. 2.2D for the case of piecewise constant membrane potentials. This figure indicates, that functions  $g$  that define a shorter refractory effect lead to higher firing rates and more irregular firing. It is worth noticing that the standard activation function  $\sigma(u_k - \log \tau)$  is the solution of equation (2.11) for the absolute refractory function, i.e., for  $g(0) = g(1) = 1$  and  $g(l) = 0$  for  $1 < l \leq \tau$ .

The transition operator  $T^k$  is defined for this model in a very similar way as before. However, for  $1 < \zeta'_k \leq \tau$ , when the variable  $\zeta'_k$  was deterministically reduced by 1 in the simpler model (yielding  $\zeta_k = \zeta'_k - 1$ ), this reduction occurs now only with probability  $1 - g(\zeta'_k) \cdot f(u_k)$ . With probability  $g(\zeta'_k) \cdot f(u_k)$  the operator  $T^k$  sets  $\zeta_k = \tau$ , modeling the firing of another spike of neuron  $v_k$  at this time point. The neural computability condition (2.4) remains unchanged, e.g.,  $u_k = b_k + \sum_{i=1}^K W_{ki} z_i$  for a Boltzmann distribution. A schema of the stochastic dynamics of this local state transition operator  $T^k(\zeta_k | \zeta'_k, z_{\setminus k})$  is shown in Fig. 2.2A.

This transition operator  $T^k$  has the following properties. In Lemma 5 in Appendix B it is proven that the unique invariant distribution of  $T^k$ , denoted as  $q_k^*(\zeta_k, z_k | \zeta_{\setminus k}, z_{\setminus k})$ , gives rise to the correct marginal distribution over  $z_k$ , i.e.

$$\sum_{\zeta_k=0}^{\tau} q_k^*(\zeta_k, z_k | \zeta_{\setminus k}, z_{\setminus k}) = p(z_k | z_{\setminus k}) \quad . \quad (2.12)$$

## 2 Neural dynamics as sampling



**Fig. 2.3: Sampling from a Boltzmann distribution by spiking neurons with relative refractory mechanism.** (A) Spike raster of the network. (B) Traces of internal state variables of a neuron (# 26, indicated by orange spikes in A). The rich interaction of the network gives rise to rapidly changing membrane potentials and instantaneous firing rates. (C) Joint distribution of 5 neurons (gray shaded area in A) obtained by the spiking neural network and Gibbs sampling from the same distribution. Active states  $z_i = 1$  are indicated by a black dot, using one row for each neuron  $v_i$ , the columns list all  $2^5 = 32$  possible states ( $z_{24}, \dots, z_{28}$ ) of these 5 neurons. The tight match between both distributions suggests that the spiking network represents the target probability distribution  $p$  with high accuracy.

This means that a neuron whose dynamics is described by  $T^k$  samples from the correct distribution  $p(z_k | z_{\setminus k})$  if it receives a static input from the other neurons in the network, i.e., as long as its membrane potential  $u_k$  is constant. Hence the “local” computation performed by such neuron can be considered as correct. If however, several neurons in the network change their states in a short interval of time, the joint distribution over  $z$  is in general not the desired one, i.e.,  $\sum_{\zeta} q^*(\zeta, z) \neq p(z)$ , where  $q^*(\zeta, z)$  denotes the invariant distribution of  $T = T^1 \circ \dots \circ T^K$ . In Appendix B, we present simulation results that indicate that the error of the approximation to the desired Boltzmann distributions introduced by neural sampling with relative refractory mechanism is rather minute. It is shown that the neural sampling approximation error is orders of magnitudes below the one introduced by a fully factorized distribution (which amounts to assuming correct marginal distributions  $p(z_k)$  and independent neurons).

To illustrate the sampling process with the relative refractory mechanism, we examine a network of  $K = 40$  neurons. We aim to sample from a Boltzmann distribution (2.5) with parameters  $W_{ij}$ ,  $b_i$  being randomly drawn from normal distributions. For the neuron model, we use the relative refractory mechanism shown in the mid row of Fig. 2.2B. A detailed description of the simulation and the parameters used is given in Appendix B. A spike pattern of the resulting sampling network is shown in Fig. 2.3A. The network features a sparse, irregular spike response with average firing rate of 13.9Hz. For one neuron  $v_{26}$ , indicated with orange spikes, the internal dynamics are shown in Fig. 2.3B. After each action potential

the neuron’s refractory function  $g(\zeta_{26})$  drops to zero and reduces the probability of spiking again in a short time interval. The influence of the remaining network  $z_{\setminus 26}$  is transmitted to neuron  $v_{26}$  via PSPs of duration  $\tau \cdot dt = 20$  ms and sums up to the fluctuating membrane potential  $u_{26}$ . As reflected in the highly variable membrane potential even this small network exhibits rich interactions. To represent the correct distribution  $p(z_{26}|z_{\setminus 26})$  over  $z_{26}$  conditioned on  $z_{\setminus 26}$ , the neuron  $v_{26}$  continuously adapts its instantaneous firing rate. To quantify the precision with which the spiking network draws samples from the target distribution (2.5), Fig. 2.3C shows the joint distribution of 5 neurons. For comparison we accompany the distribution of sampled network states with the result obtained from the standard Gibbs sampling algorithm (considered as the ground truth). Since the number of possible states  $z$  grows exponentially in the number of neurons, we restrict ourselves for visualization purposes to the distribution  $p(z_{24}, \dots, z_{28})$  of the gray shaded units and marginalize over the remaining network. The probabilities are estimated from  $10^7$  samples, i.e., from  $10^7$  successive states  $z$  of the Markov chain. Stochastic deviations of the estimated probabilities due to the finite number of samples are quite small (typical errors  $\Delta p(z) / \sqrt{p(z)} \approx 10^{-3}$ ) and are comparable to systematic deviations due to the only locally correct computation of neurons with relative refractory mechanism. In Appendix B, we present further simulation results showing that the proposed networks consisting of neurons with relative refractory mechanism approximate the desired target distributions faithfully over a large range of distribution parameters.

In order to illustrate that the proposed sampling networks feature biologically quite realistic spiking dynamics, we present in Appendix B several neural firing statistics (e.g., the inter-spike interval histogram) of the network model. In general, the statistics computed from the model match experimentally observed statistics well. The proposed network models are based on the assumption of rectangular-shaped, renewal PSPs. More precisely, we define renewal (or non-additive) PSPs in the following way. Renewal PSPs evoked by a single synapse do not add up but are merely prolonged in their duration (according to equation (2.6)); renewal PSPs elicited at different synapses nevertheless add up in the normal way. In Appendix B, we investigate the impact of replacing the theoretically ideal rectangular-shaped, renewal PSPs with biologically more realistic alpha-shaped, additive PSPs. Simulation results suggest that the network model with alpha-shaped PSPs does not capture the target distribution as accurately as with the theoretically ideal PSP shapes, statistical dependencies between the RVs  $z$  are however still approximated reasonably well.

## 2.6 Neural sampling in continuous time

The neural sampling model proposed above was formulated in discrete time of step size  $dt$ , inspired by the discrete time nature of MCMC techniques in statistics and machine learning as well as to make simulations possible on digital computers.

## 2 Neural dynamics as sampling

However, models in continuous time (e.g., ordinary differential equations) are arguably more natural and “realistic” descriptions of temporally varying biological processes. This gives rise to the question whether one can find a sensible limit of the discrete time model in the limit  $dt \rightarrow 0$ , yielding a sampling network model in continuous time. Another motivation for considering continuous time models for neural sampling is the fact that many mathematical models for recurrent networks are formulated in continuous time (GERSTNER and KISTLER, 2002), and a comparison to these existing models would be facilitated. Here we propose a stochastically spiking neural network model in continuous time, whose states still represent correct samples from the desired probability distribution  $p(\mathbf{z})$  at any time  $t$ . These types of models are usually referred to as Markov jump processes. It can be shown that discretizing this continuous time model yields the discrete time model defined earlier, which thus can be regarded as a version suitable for simulations on a digital computer.

We define the continuous time model in the following way. Let  $t_k^l$ , for  $l = 0, 1, \dots$ , denote the firing times of neuron  $\nu_k$ . The refractory process of this neuron, in analogy to Fig. 2.1 and equation (2.8)-(2.9) for the case of discrete time, is described by the following differential equation for the auxiliary variable  $\zeta_k$ , which may now assume any nonnegative real number  $0 \leq \zeta_k \leq 1$ :

$$\frac{d}{dt}\zeta_k(t) = \begin{cases} -\frac{1}{\tau} & \text{for } \zeta_k > 0 \\ \sum_l \delta(t - t_k^l) & \text{for } \zeta_k = 0 \end{cases} . \quad (2.13)$$

Here  $\delta(t - t_k^l)$  denotes Dirac’s Delta centered at the spike time  $t_k^l$ . This differential equation describes the following simple dynamics. The auxiliary variable  $\zeta_k(t)$  decays linearly with time constant  $\tau$  when the neuron is refractory, i.e.,  $\zeta_k(t) > 0$ . Once  $\zeta_k(t)$  arrives at its resting state 0 it remains there, corresponding to the neuron being ready to spike again (more precisely, in order to avoid point measures we set it to a random value in  $[-2\epsilon, -\epsilon]$ , see Appendix B). In the resting state, the neuron has the probability density  $\frac{1}{\tau} \exp(u_k(t))$  to fire at every time  $t$ . If it fires at  $t_k^l$ , this results in setting  $\zeta_k(t_k^l) = 1$ , which is formalized in equation (2.13) by the sum of Dirac Delta’s  $\sum_l \delta(t - t_k^l)$ . Here the current membrane potential  $u_k(t)$  at time  $t$  is defined as in the discrete time case, e.g., by  $u_k = b_k + \sum_{i=1}^K W_{ki}z_i(t)$  for the case of a Boltzmann distribution (2.5). The binary variable  $z_k(t)$  is defined to be 1 if  $\zeta_k(t) > 0$  and 0 if the neuron is in the resting state  $\zeta_k(t) = 0$ . Biologically, the term  $W_{ki}z_i(t)$  can again be interpreted as the value at time  $t$  of a rectangular-shaped PSP (with a duration of  $\tau$ ) that neuron  $\nu_i$  evokes in neuron  $\nu_k$ . As the spikes are discrete events in continuous time, the probability of two or more neurons spiking at the same time is zero. This allows for updating all neurons in parallel using a differential equation.

In analogy to the discrete time case, the neural network in continuous time can be shown to sample from the desired distribution  $p(\mathbf{z})$ , i.e.,  $p(\mathbf{z})$  is an invariant distribution of the network dynamics defined above. However, to establish this fact, one has to rely on a different mathematical framework. The probability distribution  $p_t(\boldsymbol{\zeta})$  of the auxiliary variables  $\zeta_1(t), \dots, \zeta_K(t)$  as a function of time  $t$ , which

describes the evolution of the network, obeys a partial differential equation, the so-called Differential-Chapman-Kolmogorov equation (see GARDINER (2009)):

$$\partial_t p_t(\zeta) = (T p_t)(\zeta), \quad (2.14)$$

where the operator  $T$ , which captures the dynamics of the network, is implicitly defined by the differential equations (2.13) and the spiking probabilities. This operator  $T$  is the continuous time equivalent to the transition operator  $T$  in the discrete time case. The operator  $T$  consists here of two components. The *drift term* captures the deterministic decay process of  $\zeta_k(t)$ , stemming from the term  $-1/\tau$  in equation (2.13). The *jump term* describes the non-continuous aspects of the path  $\zeta_k(t)$  associated with “jumping” from  $\zeta_k(t_k^l - dt) = 0$  to  $\zeta_k(t_k^l) = 1$  at the time  $t_k^l$  when the neuron fires.

In Appendix B, we prove that the resulting time invariant distribution, i.e., the distribution that solves  $\partial_t p_t(\zeta) = 0$ , now denoted  $p(\zeta)$  as it is not a function of time, gives rise to the desired marginal distribution  $p(z)$  over  $z$ :

$$\int d\zeta \delta(z, \zeta^{>0}) p(\zeta) = p(z), \quad (2.15)$$

where  $\delta(z, \zeta^{>0}) = (\delta(z_1, \zeta_1^{>0}), \dots, \delta(z_K, \zeta_K^{>0}))$  and  $\zeta_k^{>0} = 1$  if  $\zeta_k > 0$  and  $\zeta_k^{>0} = 0$  otherwise.  $\delta(z_k, \zeta_k^{>0}) = 1$  denotes Kronecker’s Delta with  $\delta(z_k, \zeta_k^{>0}) = 1$  if  $z_k = \zeta_k^{>0}$  and  $\delta(z_k, \zeta_k^{>0}) = 0$  otherwise. Thus, the function  $\delta(z, \zeta^{>0})$  simply reflects the definition that  $z_k(t) = 1$  if  $\zeta_k(t) > 0$  and 0 otherwise. For an explicit definition of  $T$ , a proof of the above statement, and some additional comments see Appendix B.

The neural samplers in discrete and continuous time are closely related. The model in discrete time provides an increasingly more precise description of the inherent spike dynamics when the duration  $dt$  of the discrete time step is reduced, causing an increase of  $\tau$  (such that  $\tau \cdot dt$  is constant) and therefore a reduced firing probability of each neuron at any discrete time step (see the term  $\log \tau$  in equation (2.8)). In the limit of  $dt$  approaching 0, the probability that two or more neurons will fire at the same time approaches 0, and the discrete time sampler becomes equal to the continuous time system defined above, which updates all units in parallel.

It is also possible to formulate a continuous time version of the neural sampler based on neuron models with relative refractory mechanisms. In Appendix B, the resulting continuous time neuron model with a relative refractory mechanism is defined. Theoretical results similar to the discrete time case can be derived for this sampler (see Lemmata 9 and 10 in Appendix B): It is shown that each neuron “locally” performs the correct computation under the assumption of static input from the remaining neurons. However one can no longer prove in general that the global network samples from the target distribution  $p$ .

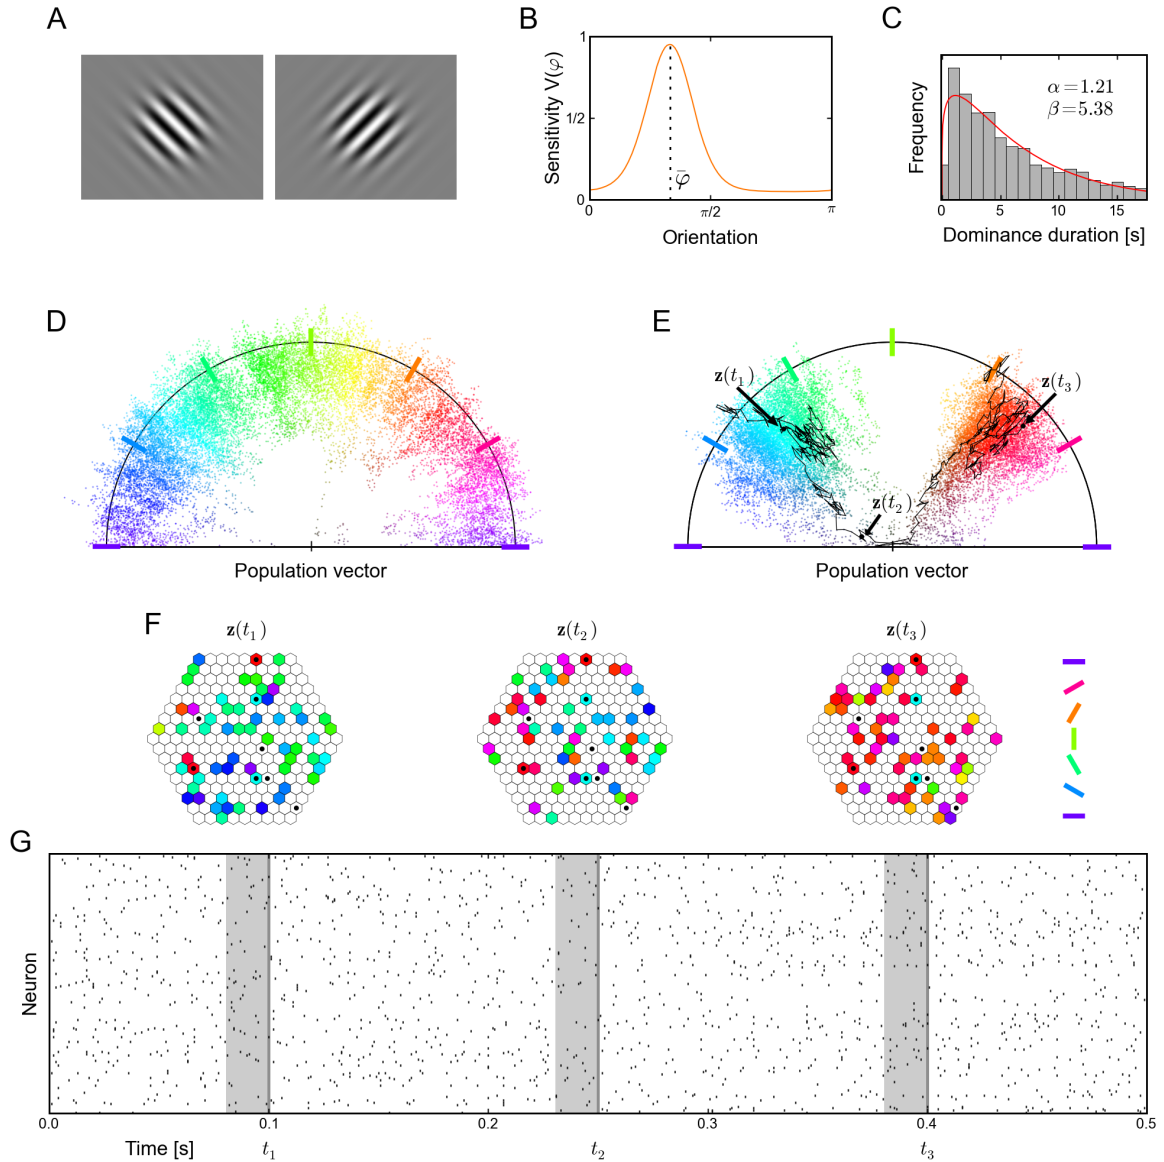
### 2.7 Demonstration of perceptual multistability in sampling networks

In the following we present a network model for perceptual multistability based on the neural sampling framework introduced above. This simulation study is aimed at showing that the proposed network can indeed sample from a desired distribution and also perform inference, i.e., sample from the correct corresponding posterior distribution. It is not meant to be a highly realistic or exhaustive model of perceptual multistability nor of biologically plausible learning mechanisms. Such models would naturally require considerably more modeling work.

Perceptual multistability evoked by ambiguous sensory input, such as a 2D drawing (e.g., Necker cube) that allows for different consistent 3D interpretations, has become a frequently studied perceptual phenomenon. The most important finding is that the perceptual system of humans and nonhuman primates does not produce a superposition of different possible percepts of an ambiguous stimulus, but rather switches between different self-consistent global percepts in a spontaneous manner. Binocular rivalry, where different images are presented to the left and right eye, has become a standard experimental paradigm for studying this effect (LEOPOLD et al., 2002; BLAKE and LOGOTHETIS, 2002; ALAIS and BLAKE, 2005; BARTELS and LOGOTHETIS, 2010). A typical pair of stimuli are the two images shown in Fig. 2.4A. Here the percepts of humans and nonhuman primates switch (seemingly stochastically) between the two presented orientations. Multiple studies (HOYER and HYVÄRINEN, 2003; SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009) proposed that several aspects of experimental data on perceptual multistability can be explained if one assumes that percepts correspond to samples from the conditional distribution over interpretations (e.g., different 3D shapes) given the visual input (e.g., the 2D drawing). Furthermore, the experimentally observed fact that percepts tend to be stable on the time scale of seconds suggests that perception can be interpreted as probabilistic inference that is carried out by MCMC sampling which produces successively correlated samples. In (GERSHMAN et al., 2009) it is shown that this MCMC interpretation is also able to qualitatively reproduce the experimentally observed distribution of dominance durations, i.e., the distribution of time intervals between perceptual switches. However, in lack of an adequate model for sampling by a recurrent network of spiking neurons, these studies could describe this approach only on a rather abstract level, and pointed out the open problem to relate this algorithmic approach to neural processes. We have demonstrated in a computer simulation that the previously described model for neural sampling could in principle fill this gap, providing a modeling framework that is on the one hand consistent with the dynamics of networks of spiking neurons, and which can on the other hand also be clearly understood from the perspective of probabilistic inference through MCMC sampling.

In the following we model some essential aspects of an experimental setup for binocular rivalry with grating stimuli (see Fig. 2.4A) in a recurrent network of

## 2.7 Demonstration of perceptual multistability in sampling networks



**Fig. 2.4: Perceptual multistability as probabilistic inference with neural sampling.** (A) Typical visual stimuli for the left and right eye in binocular rivalry experiments. (B) Tuning curve of a neuron with preferred orientation  $\bar{\varphi}$ . (C) Distribution of dominance durations in the network under ambiguous input. The red curve shows a Gamma distribution fitted the data. (D) 2-dim. projection (population vector) of the distribution  $p(z)$  encoded in the spiking network. The network favors coherent global states of arbitrary orientation. (E) 2-dim. projection of the bimodal posterior distribution under an ambiguous input (two different orientations mimicking the stimuli shown in A). Black trace: 500 ms network trajectory  $z(t)$  around a perceptual switch. (F) Network states at 3 time points  $t_1$ ,  $t_2$ ,  $t_3$  marked in E. Neurons that fired in the preceding 20 ms (see gray bar in G) are plotted in the color of their preferred orientation. Inactive neurons are shown in white. While states  $z(t_1)$  and  $z(t_3)$  represent coherent orientations,  $z(t_2)$  shows an incoherent state close to a perceptual switch. Clamped neurons (which the posterior is conditioned on) are marked by a black dot. (G) Spike raster of the unclamped neurons during the 500 ms epoch marked by the black trace in E. Gray bars indicate the 20 ms time intervals that define the network states shown in F.

## 2 Neural dynamics as sampling

spiking neurons with the previously described relative refractory mechanism. We assigned to each of the 217 neurons in the network  $\mathcal{N}$  a tuning curve  $V_k(\varphi)$ , centered around its preferred orientation  $\bar{\varphi}_k$  as shown in Fig. 2.4B. The preferred orientations  $\bar{\varphi}_k$  of the neurons were chosen to uniformly cover the entire interval  $[0, \pi)$  of possible orientations and were randomly assigned to the neurons. The neurons were arranged on a hexagonal grid as depicted in Fig. 2.4F. Any two neurons with distance  $\leq 8$  were synaptically connected (neighboring units had distance 1). We assume that these neurons represent neurons in the visual system that have roughly the same or neighboring receptive field, and that each neuron receives visual input from either the left or the right eye. The network connections were chosen such that neurons that have similar (very different) preferred orientations are connected with positive (negative) weights (for details see Appendix B).

We examined the resulting distribution  $p(z)$  over the 217 dimensional network states. To provide an intuitive visualization of these high dimensional network states  $z$ , we resort to a 2-dimensional projection, the population vector of a state  $z$  (see Appendix B for details of the applied population vector decoding scheme). Only the endpoints of the population vectors are drawn (as colored points) in Fig. 2.4D,E. The orientation of the population vector is assumed to correspond to the dominant orientation of the percept, and its distance from the origin encodes the strength of this percept. We also, somewhat informally, call the strength of a percept its coherence and a network state which represents a coherent percept a coherent network state. A coherent network state hence results in a population vector of large magnitude. Each direction of a population vector is color coded in Fig. 2.4D,E, using the color code for directions shown on the right hand side of Fig. 2.4F. In Fig. 2.4D the distribution  $p(z)$  of the network is illustrated by sampling of the network for 20s, with samples  $z$  taken every millisecond. Each dot equals a sampled network state  $z$ . In a biological interpretation the spike response of the freely evolving network reflects spontaneous activity, since no observations, i.e., no external input, was added to the system. Fig. 2.4D shows that the spontaneous activity of this simple network of spiking neurons moves preferably through coherent network states for all possible orientations due to the chosen recurrent network connections (being positive for neurons with similar preferred orientation and negative otherwise). This can directly be seen from the rare occurrence of population vectors with small magnitude (vectors close to the “center”) in Fig. 2.4D.

To study percepts elicited by ambiguous stimuli, where inputs like in Fig. 2.4A are shown simultaneously to the left and right eye during a binocular rivalry experiment, we provided ambiguous input to the network. Two cells with preferred orientation  $\bar{\varphi}_k \approx 45^\circ$  and two cells with  $\bar{\varphi}_k \approx 135^\circ$  were clamped to 1. Additionally four neurons with  $\bar{\varphi}_k \approx 0^\circ$  resp.  $90^\circ$  were muted by clamping to 0. This ambiguous input is incompatible with a coherent percept, as it corresponds to two orthogonal orientations presented at the same time. The resulting distribution over the state of the 209 remaining neurons is shown for a time span of 20s of simulated biological time (with samples taken every millisecond) in Fig. 2.4E. One clearly sees that



## 2.7 Demonstration of perceptual multistability in sampling networks

the network spends most of the time in network states that correspond to one of the two simultaneously presented input orientations ( $45^\circ$  and  $135^\circ$ ), and virtually no time on orientations in between. This implements a sampling process from a bimodal conditional distribution. The black line marks a 500 ms trace of network states  $z$  around a perceptual switch: The network remained in one mode of high probability – corresponding to one percept – for some period of time, and then quickly traversed the state space to another mode – corresponding to a different percept.

Three of the states  $z$  around this perceptual switch ( $z(t_1)$ ,  $z(t_2)$  and  $z(t_3)$  in Fig. 2.4E) are explicitly shown in Fig. 2.4F. Neurons  $\nu_k$  that fired during the preceding interval of 20 ms (marked in gray in Fig. 2.4G) are drawn in the respective color of their preferred orientation. Inactive neurons are drawn in white, and clamped neurons are marked by a black dot ( $\bullet$ ).

Fig. 2.4G shows the action potentials of the 209 non-clamped neurons during the same 500 ms trace around the perceptual switch. One sees that the sampling process is expressed in this neural network model by a sparse, asynchronous and irregular spike response. It is worth mentioning that the average firing rate when sampling from the posterior distribution is only slightly higher than the average firing rate of spontaneous activity (16.1 Hz and 15.4 Hz respectively), which is reminiscent of related experimental data (FISER et al., 2004). Thus on the basis of the overall network activity it is indistinguishable whether the network carries out an inference task or freely samples from its prior distribution. It is furthermore notable, that a focus of the network activity on the two orientations that are given by the external input can be achieved in this model, in spite of the fact that only two of the 217 neurons were clamped for each of them. This numerical relationship is reminiscent of standard data on the weak input from LGN to V1 that is provided in the brain (BINZEGGER et al., 2004; BINZEGGER et al., 2009), and indicates that the proposed neural sampling model could provide a possible mechanism (under the modeling assumptions made above) for cortical processing of such numerically weak external inputs.

The distribution of the resulting dominance durations, i.e., the time between perceptual switches, for the previously described setup with ambiguous input is shown for a continuous run of  $10^4$  s in Fig. 2.4C (a similar method as in GERSHMAN et al. (2009) was used to measure dominance durations, see Appendix B). This distribution can be approximated quite well by a Gamma distribution, which also provides a good fit to experimental data (see the discussion in GERSHMAN et al. (2009)). We expect that also other features of the more abstract MCMC model for biological vision of (SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009), such as contextual biases and traveling waves, will emerge in larger and more detailed implementations of the MCMC approach through the proposed neural sampling method in networks of spiking neurons.

### 2.8 Discussion

We have presented a spiking neural network that samples from a given probability distribution via its inherent network dynamics. In particular the network is able to carry out probabilistic inference through sampling. The model, based on assumptions about the underlying probability distribution (formalized by the neural computability condition) as well as on certain assumptions regarding the underlying MCMC model, provides one possible neural implementation of the “inference-by-sampling paradigm” emerging in computational neuroscience.

During inference the observations (i.e., the variables which we wish to condition on) are modeled in this study by clamping the corresponding neurons by strong external input to the observed binary value. Units which receive no input or input with vanishing contrast (stimulus intensity) are treated as unobserved. Using this admittedly quite simplistic model of the input, we observed in simulations that our network model exhibits the following property: The onset of a sensory stimulus reduces the variability of the firing activity, which represents (after stimulus onset) a conditional distribution, rather than the prior distribution (see the difference between panels **D** and **E** of Fig. 2.4). It is tempting to compare these results to the experimental finding of reduced firing rate variability after stimulus onset observed in several cortical areas (CHURCHLAND et al., 2010). We wish to point out however, that a consistent treatment of zero contrast stimuli requires more thorough modeling efforts (e.g., by explicitly adding a random variable for the stimulus intensity (FISER et al., 2010; BERKES et al., 2011)), which is not the focus of the presented work.

Virtually all high-level computational tasks that a brain has to solve can be formalized as optimization problems, that take into account a (possibly large) number of soft or hard constraints. In typical applications of probabilistic inference in science and engineering (see e.g. BISHOP (2006); KOLLER and FRIEDMAN (2009)) such constraints are encoded in e.g., conditional probability tables or factors. In a biological setup they could possibly be encoded through the synaptic weights of a recurrent network of spiking neurons. The solution of such optimizations problems in a probabilistic framework via sampling, as implemented in our model, provides an alternative to deterministic solutions, as traditionally implemented in neural networks (see, e.g., HOPFIELD and TANK (1985) for the case of constraint satisfaction problems). Whereas an attractor neural network converges to *one* (possibly approximate) solution of the problem, a stochastic network may alternate between different approximate solutions and stay the longest at those approximate solutions that provide the best fit. This might be advantageous, as given more time a stochastic network can explore more of the state space and avoid shallow local minima. Responses to ambiguous sensory stimuli (LEOPOLD et al., 2002; BLAKE and LOGOTHETIS, 2002; ALAIS and BLAKE, 2005; BARTELS and LOGOTHETIS, 2010) might be interpreted as an optimization with soft constraints. The interpretation of human thinking as sampling process solving an inference task, recently proposed in cognitive science (VUL and PASHLER, 2008; DENISON et al., 2009; GRIFFITHS

and TENENBAUM, 2006), further emphasizes that considering neural activity as an inferential process via sampling promises to be a fruitful approach.

Our approach builds on, and extends, previous work where recurrent networks of non-spiking stochastic neurons (commonly considered in artificial neural networks) were shown to be able to carry out probabilistic inference through Gibbs sampling (ACKLEY et al., 1985). In G. HINTON and BROWN (2000) a first extension of this approach to a network of recurrently connected spiking neurons had been presented. The dynamics of the recurrently connected spiking neurons are described as stepwise sampling from the posterior of a temporal Restricted Boltzmann Machine (tRBM) by introducing a clever interpretation of the temporal spike code as time varying parameters of a multivariate Gaussian distribution. Drawing one sample from the posterior of a RBM is, by construction, a trivial one-step task. In contrast to our model, the model of G. HINTON and BROWN (2000) does not produce multiple samples from a fixed posterior distribution, given the fixed input, but produces exactly one sample consisting of the temporal sequence of the hidden nodes, given a temporal input sequence. Similar temporal models, sometimes called Bayesian filtering, also underlie the important contributions of (R. ZEMEL et al., 2005) and (DENEVE, 2008b). In (DENEVE, 2008b) every single neuron is described as hidden Markov Model (HMM) with two states. Instead of drawing samples from the instantaneous posterior distribution using stochastic spikes, (DENEVE, 2008b) presents a deterministic spike generation with the intention to convey the analog probability value rather than discrete samples. The approach presented here can be interpreted as a biologically more realistic version of Gibbs sampling for a specific class of probability distributions by taking into account a spike-based communication, finite duration PSPs and refractory mechanisms. Other implementations based on different distributions (e.g., directed graphical models) and different sampling methods (e.g., reversible MCMC methods) are of course conceivable and worth exploring.

In a computer experiment (see Fig. 2.4), we used our proposed network to model aspects of biological vision as probabilistic inference along the lines of argumentation put forward in (HOYER and HYVÄRINEN, 2003; SUNDARESWARA and SCHRATER, 2008; GERSHMAN et al., 2009). Our model was chosen to be quite simplistic, just to demonstrate that a number of experimental data on the dynamics of spontaneous activity (KENET et al., 2003; FOX and RAICHLE, 2007; BERKES et al., 2011) and binocular rivalry (ALAIS and BLAKE, 2005; BLAKE and LOGOTHETIS, 2002; LEOPOLD et al., 2002; BARTELS and LOGOTHETIS, 2010) can in principle be captured by this approach. The main point of the modeling study is to show that rather realistic neural dynamics can support computational functions rigorously formalized as inference via sampling.

We have also presented a model of spiking dynamics in continuous time that performs sampling from a given probability distribution. Although computer simulations of biological networks of neurons often actually use discrete time, it is desirable to also have a sound approach for understanding and describing the network sampling dynamics in continuous time, as the latter is arguable a natural

## 2 Neural dynamics as sampling

framework for describing temporal processes in biology. Furthermore comparison to many existing continuous time neuron and network models of neurons is facilitated.

We have made various simplifying assumptions regarding neural processes, e.g., simple symbolic postsynaptic potentials in the form of step-functions (reminiscent of plateau potentials caused by dendritic NMDA spikes (ANTIC et al., 2010)). More accurate models for neurons have to integrate a multitude of time constants that represent different temporal processes on the physical, molecular, and genetic level. Hence the open problem arises, to which extent this multitude of time constants and other complex dynamics can be integrated into theoretical models of neural sampling. We have gone one first step in this direction by showing that in computer simulations the two temporal processes that we have considered (refractory processes and postsynaptic potentials) can approximately be decoupled. Furthermore, we have presented simulation results suggesting that more realistic alpha-shaped, additive EPSPs are compatible with the functionality of the proposed network model.

Finally, we want to point out that the prospect of using networks of spiking neurons for probabilistic inference via sampling suggests new applications for energy-efficient spike-based and massively parallel electronic hardware that is currently under development (MEROLLA et al., 2007; BRUEDERLE et al., 2010).

## Homeostatic and Hebbian plasticity united under a Bayesian account

### Contents

---

3.1	Introduction . . . . .	38
3.2	Homeostatic plasticity in winner-take-all networks . . . . .	40
3.3	Theory for the WTA model . . . . .	42
3.4	Spiking network dynamics with homeostatic plasticity . . . . .	45
3.5	Homeostatic plasticity in recurrent spiking networks . . . . .	46
3.6	Discussion . . . . .	48

---

**Abstract.** Recent spiking network models of Bayesian inference and unsupervised learning frequently assume either inputs to arrive in a special format or employ complex computations in neuronal activation functions and synaptic plasticity rules. Here we show in a rigorous mathematical treatment how homeostatic processes, which have previously received little attention in this context, can overcome common theoretical limitations and facilitate the neural implementation and performance of existing models. In particular, we show that homeostatic plasticity can be understood as the enforcement of a ‘balancing’ posterior constraint during probabilistic inference and learning with Expectation Maximization. We link homeostatic dynamics to the theory of variational inference, and show that nontrivial terms, which typically appear during probabilistic inference in a large class of models, drop out. We demonstrate the feasibility of our approach in a spiking Winner-Take-All architecture of Bayesian inference and learning. Finally, we sketch how the mathematical framework can be extended to richer recurrent network architectures. Altogether, our theory provides a novel perspective on the interplay of homeostatic processes and synaptic plasticity in cortical microcircuits, and points to an essential role of homeostasis during inference and learning in spiking networks.

**Acknowledgments and author contributions.** This chapter is based on the publication

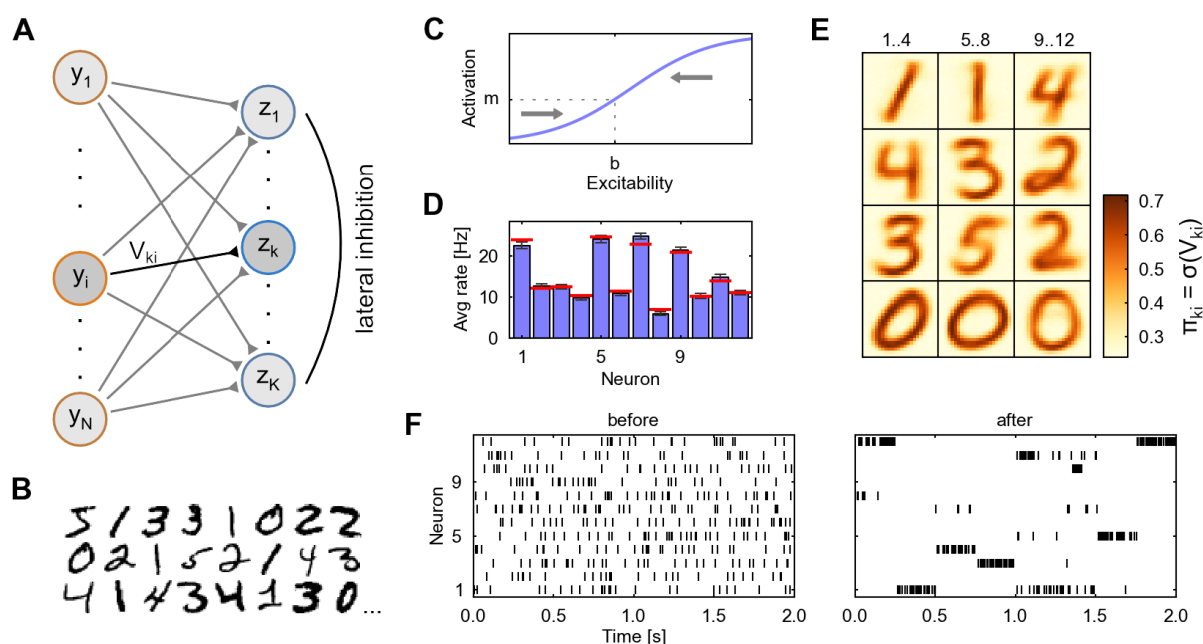
STEFAN HABENSCHUSS, JOHANNES BILL, BERNHARD NESSLER (2012). “Homeostatic plasticity in Bayesian spiking networks as Expectation Maximization with posterior constraints.” In: *Advances in Neural Information Processing Systems*.

To this study, I contributed as a functional first author (equal contribution with SH). The theory was jointly developed by SH and JB, with support of BN for the extension of the theory to recurrent networks. The computer experiments were designed by SH and JB, and were performed by JB. All authors contributed to writing the manuscript.

## 3.1 Introduction

Experimental findings from neuro- and cognitive sciences have led to the hypothesis that humans create and maintain an internal model of their environment in neuronal circuitry of the brain during learning and development (KÖRDING and WOLPERT, 2004; ORBAN et al., 2008; FISER et al., 2010; BERKES et al., 2011), and employ this model for Bayesian inference in everyday cognition (GRIFFITHS and TENENBAUM, 2006; ANGELAKI et al., 2009). Yet, how these computations are carried out in the brain remains largely unknown. A number of innovative models has been proposed recently which demonstrate that in principle, spiking networks can carry out quite complex probabilistic inference tasks (DENEVE, 2008b; STEIMER et al., 2009; BUESING et al., 2011; PECEVSKI et al., 2011), and even learn to adapt to their inputs near optimally through various forms of plasticity (DENEVE, 2008a; NESSLER et al., 2009a; BREA et al., 2012; REZENDE et al., 2012; KECK et al., 2012). Still, in network models for concurrent online inference and learning, most approaches introduce distinct assumptions: Both (NESSLER et al., 2009a) in a spiking Winner-take-all (WTA) network, and (KECK et al., 2012) in a rate based WTA network, identified the limitation that inputs must be normalized before being presented to the network, in order to circumvent an otherwise nontrivial (and arguably non-local) dependency of the intrinsic excitability on all afferent synapses of a neuron. NESSLER et al. (2009a) relied on population coded input spike trains; KECK et al. (2012) proposed feed-forward inhibition as a possible neural mechanism to achieve this normalization. A theoretically related issue has been encountered by Deneve (DENEVE, 2008b; DENEVE, 2008a), in which inference and learning is realized in a two-state Hidden Markov Model by a single spiking neuron. Although synaptic learning rules are found to be locally computable, the learning update for intrinsic excitabilities remains intricate. In a different approach, Brea et al. (BREA et al., 2012) have recently proposed a promising model for Bayes optimal sequence learning in spiking networks in which a global reward signal, which is computed from the network state and synaptic weights, modulates otherwise purely local learning rules. Also the recent innovative model for variational learning in recurrent spiking networks by Rezende et al. (REZENDE et al., 2012) relies on sophisticated updates of variational parameters that complement otherwise local learning rules.

There exists great interest in developing Bayesian spiking models which require minimal non-standard neural mechanisms or additional assumptions on the input distribution: such models are expected to foster the analysis of biological circuits



**Fig. 3.1: Homeostatic plasticity in WTA circuits.** (A) Spiking WTA network model. (B) Input templates from MNIST database (digits 0-5) are presented in random order to the network as spike trains (the input template switches after every 250ms, black/white pixels are translated to high/low firing rates between 20 and 90 Hz). (C) Sketch of intrinsic homeostatic plasticity maintaining a certain target average activation. (D) Homeostatic plasticity induces average firing rates (blue) close to target values (red). (E) After a learning period, each WTA neuron has specialized on a particular input motif. (F) WTA output spikes during a test phase before and after learning. Learning leads to a sparse output code.

from a Bayesian perspective (TENENBAUM et al., 2011), and to provide a versatile computational framework for novel neuromorphic hardware (SCHEMMELE et al., 2010). With these goals in mind, we introduce here a novel theoretical perspective on homeostatic plasticity in Bayesian spiking networks that complements previous approaches by constraining statistical properties of the *network response* rather than the input distribution. In particular we introduce ‘balancing’ posterior constraints which can be implemented in a purely local manner by the spiking network through a simple rule that is strongly reminiscent of homeostatic intrinsic plasticity in cortex (DESAI et al., 1999; WATT and DESAI, 2010). Importantly, it turns out that the emerging network dynamics eliminate a particular class of nontrivial computations that frequently arise in Bayesian spiking networks.

First we develop the mathematical framework for Expectation Maximization (EM) with homeostatic posterior constraints in an instructive Winner-Take-all network model of probabilistic inference and unsupervised learning. Building upon the theoretical results of (GRACA et al., 2008), we establish a rigorous link between homeostatic intrinsic plasticity and variational inference. In a second step, we sketch how the framework can be extended to recurrent spiking networks; by introducing posterior constraints on the correlation structure, we recover local plasticity rules for recurrent synaptic weights.

### 3.2 Homeostatic plasticity in winner-take-all networks

We first introduce, as an illustrative and representative example, a generative mixture model  $p(\mathbf{z}, \mathbf{y}|\mathbf{V})$  with hidden causes  $\mathbf{z}$  and binary observed variables  $\mathbf{y}$ , and a spiking WTA network  $\mathcal{N}$  which receives inputs  $\mathbf{y}(t)$  via synaptic weights  $\mathbf{V}$ . As shown in (NESSLER et al., 2009a), such a network  $\mathcal{N}$  can implement probabilistic inference  $p(\mathbf{z}|\mathbf{y}, \mathbf{V})$  through its spiking dynamics, and maximum likelihood learning through local synaptic learning rules (see Fig. 3.1A). The mixture model comprises  $K$  binary and mutually exclusive components  $z_k \in \{0, 1\}$ ,  $\sum_{k=1}^K z_k = 1$ , each specialized on a different  $N$ -dimensional input pattern:

$$p(\mathbf{y}, \mathbf{z}|\mathbf{V}) = \prod_{k=1}^K e^{\hat{b}_k z_k} \prod_{i=1}^N \left[ (\pi_{ki})^{y_i} \cdot (1 - \pi_{ki})^{1-y_i} \right]^{z_k} \quad (3.1)$$

$$\Leftrightarrow \log p(\mathbf{y}, \mathbf{z}|\mathbf{V}) = \sum_k z_k \left( \sum_i V_{ki} y_i - A_k + \hat{b}_k \right) , \quad (3.2)$$

$$\text{with } \sum_k e^{\hat{b}_k} = 1 \text{ and } \pi_{ki} = \sigma(V_{ki}) \text{ and } A_k = \sum_i \log(1 + e^{V_{ki}}) , \quad (3.3)$$

where  $\sigma(x) = (1 + \exp(-x))^{-1}$  denotes the logistic function, and  $\pi_{ki}$  the expected activation of input  $i$  under the mixture component  $k$ . For simplicity and notational convenience, we will treat the prior parameters  $\hat{b}_k$  as constants throughout the paper. Probabilistic inference of hidden causes  $z_k$  based on an observed input  $\mathbf{y}$  can be implemented by a spiking WTA network  $\mathcal{N}$  of  $K$  neurons which fire with the instantaneous spiking probability (for  $\delta t \rightarrow 0$ ),

$$p(z_k \text{ spikes in } [t, t + \delta t]) = \delta t \cdot r_{\text{net}} \cdot \frac{e^{u_k(t)}}{\sum_j e^{u_j(t)}} \propto p(z_k = 1|\mathbf{y}, \mathbf{V}) , \quad (3.4)$$

with the input potential  $u_k(t) = \sum_i V_{ki} y_i(t) - A_k + \hat{b}_k$ . Each WTA neuron  $k$  receives spiking inputs  $y_i$  via synaptic weights  $V_{ki}$  and responds with an instantaneous spiking probability which depends exponentially on its input potential  $u_k$  in accordance with biological findings (JOLIVET et al., 2006). Stochastic winner-take-all (soft-max) competition between the neurons is modeled in (3.4) via divisive normalization (E. SIMONCELLI and HEEGER, 1998). The input is defined as  $y_i(t) = 1$  if input neuron  $i$  emitted a spike within the last  $\tau$  milliseconds, and 0 otherwise, corresponding to a rectangular post-synaptic potential (PSP) of length  $\tau$ . We define  $z_k(t) = 1$  at spike times  $t$  of neuron  $k$  and  $z_k(t) = 0$  otherwise.

In addition to the spiking input, each neuron's potential  $u_k$  features an intrinsic excitability  $-A_k + \hat{b}_k$ . Note that, besides the prior constant  $\hat{b}_k$ , this excitability depends on the normalizing term  $A_k$ , and hence on all afferent synaptic weights through (3.3): WTA neurons which encode strong patterns with high probabilities  $\pi_{ki}$  require lower intrinsic excitabilities, while neurons with weak patterns require larger excitabilities. In the presence of synaptic plasticity, i.e., time-varying  $V_{ki}$ , it is unclear how biologically realistic neurons could communicate ongoing changes



### 3.2 Homeostatic plasticity in winner-take-all networks

in synaptic weights from distal synaptic sites to the soma. This critical issue was apparently identified in (NESSLER et al., 2009a) and (KECK et al., 2012); both papers circumvent the problem (in similar probabilistic models) by constraining the input  $\mathbf{y}$  (and also the synaptic weights in (KECK et al., 2012)) in order to maintain constant and uniform values  $A_k$  across all WTA neurons.

Here, we propose a different approach to cope with the nontrivial computations  $A_k$  during inference and learning in the network. Instead of assuming that the inputs  $\mathbf{y}$  meet a normalization constraint, we constrain the *network response* during inference, by applying homeostatic dynamics to the intrinsic excitabilities. This approach turns out to be beneficial in the presence of time-varying synaptic weights, i.e., during ongoing changes of  $V_{ki}$  and  $A_k$ . The resulting interplay of intrinsic and synaptic plasticity can be best understood from the standard EM lower bound (BISHOP, 2006),

$$\mathcal{F}(\mathbf{V}, q(z|\mathbf{y})) = \mathcal{L}(\mathbf{V}) - \langle D_{\text{KL}}(q(z|\mathbf{y}) || p(z|\mathbf{y}, \mathbf{V})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{E-step, (3.5)}$$

$$= \langle \log p(\mathbf{y}, \mathbf{z}|\mathbf{V}) \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} + \langle H(q(z|\mathbf{y})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{M-step, (3.6)}$$

where  $\mathcal{L}(\mathbf{V}) = \langle \log p(\mathbf{y}|\mathbf{V}) \rangle_{p^*(\mathbf{y})}$  denotes the log-likelihood of the input under the model,  $D_{\text{KL}}(\cdot || \cdot)$  the Kullback-Leibler divergence, and  $H(\cdot)$  the entropy. The decomposition holds for arbitrary distributions  $q$ . In hitherto proposed neural implementations of EM (DENEVE, 2008a; NESSLER et al., 2009a; KECK et al., 2012; SATO, 1999), the network implements the current posterior distribution in the E-step, i.e.,  $q = p$  and  $D_{\text{KL}}(q || p) = 0$ . In contrast, by applying homeostatic plasticity, the network response will be constrained to implement a variational posterior from a class of “homeostatic” distributions  $\mathcal{Q}$ : the long-term average activation of each WTA neuron  $z_k$  is constrained to an a priori defined target value. Notably, we will see that the resulting network response  $q^*$  describes an optimal variational E-Step in the sense that  $q^*(z|\mathbf{y}) = \arg \min_{q \in \mathcal{Q}} D_{\text{KL}}(q(z|\mathbf{y}) || p(z|\mathbf{y}, \mathbf{V}))$ . Importantly, homeostatic plasticity fully regulates the intrinsic excitabilities, and as a side effect eliminates the non-local terms  $A_k$  in the E-step, while synaptic plasticity of the weights  $V_{ki}$  optimizes the underlying probabilistic model  $p(\mathbf{y}, \mathbf{z}|\mathbf{V})$  in the M-step.

In summary, the network response implements  $q^*$  as the variational E-step, the M-Step can be performed via gradient ascent on (3.6) with respect to  $V_{ki}$ . As derived in section 3.3, this gives rise to the following temporal dynamics and plasticity rules in the spiking network, which instantiate a stochastic version of the variational EM scheme:

$$u_k(t) = \sum_i V_{ki} y_i(t) + b_k, \quad \dot{b}_k(t) = \eta_b \cdot (r_{\text{net}} \cdot m_k - \delta(z_k(t) - 1)), \quad (3.7)$$

$$\dot{V}_{ki}(t) = \eta_V \cdot \delta(z_k(t) - 1) \cdot (y_j(t) - \sigma(V_{ki})), \quad (3.8)$$

where  $\delta(\cdot)$  denotes the Dirac delta function, and  $\eta_b, \eta_V$  are learning rates (which were kept time-invariant in the simulations with  $\eta_b = 10 \cdot \eta_V$ ). Note that (3.8) is a spike-timing dependent plasticity rule (cf. (NESSLER et al., 2009a)) and is non-zero only at post-synaptic spike times  $t$ , for which  $z_k(t) = 1$ . The effect of

### 3 Homeostatic and Hebbian plasticity united under a Bayesian account

the homeostatic intrinsic plasticity rule (3.7) is illustrated in Fig. 3.1C: it aims to keep the long-term average activation of each WTA neuron  $k$  close to a certain target value  $m_k$ . More precisely, if  $r_k$  is a neuron’s long-term average firing rate, then homeostatic plasticity will ensure that  $r_k/r_{\text{net}} \approx m_k$ . The target activations  $m_k \in (0, 1)$  can be chosen freely with the obvious constraint that  $\sum_k m_k = 1$ . Note that (3.7) is strongly reminiscent of homeostatic intrinsic plasticity in cortex (DESAI et al., 1999; WATT and DESAI, 2010).

We have implemented these dynamics in a computer simulation of a WTA spiking network  $\mathcal{N}$ . Inputs  $\mathbf{y}(t)$  were defined by translating handwritten digits 0-5 (Fig. 3.1B) from the MNIST dataset (LECUN et al., 1998) into input spike trains. Fig. 3.1D shows that, at the end of a  $10^4$ s learning period, homeostatic plasticity has indeed achieved that  $r_k \approx r_{\text{net}} \cdot m_k$ . Fig. 3.1E illustrates the patterns learned by each WTA neuron after this period (shown are the  $\pi_{ki}$ ). Apparently, the WTA neurons have specialized on patterns of different intensity which correspond to different values of  $A_k$ . Fig. 3.1F shows the output spiking behavior of the circuit before and after learning in response to a set of test patterns. The specialization to different patterns has led to a distinct sparse output code, in which any particular test pattern evokes output spikes from only one or two WTA neurons. Note that homeostasis forces all WTA neurons to participate in the competition, and thus prevents neurons from becoming underactive if their synaptic weights decrease, and from becoming overactive if their synaptic weights increase, much like the original  $A_k$  terms (which are nontrivial to compute for the network). Indeed, the learned synaptic parameters and the resulting output behavior corresponds to what would be expected from an optimal learning algorithm for the mixture model (3.1)-(3.3).<sup>1</sup>

### 3.3 Theory for the WTA model

In the following, we develop the three theoretical key results for the WTA model (3.1)-(3.3):

- Homeostatic intrinsic plasticity finds the network response distribution  $q^*(\mathbf{z}|\mathbf{y}) \in \mathcal{Q}$  closest to the posterior distribution  $p(\mathbf{z}|\mathbf{y}, \mathbf{V})$ , from a set of “homeostatic” distributions  $\mathcal{Q}$ .
- The interplay of homeostatic and synaptic plasticity can be understood from the perspective of variational EM.
- The critical non-local terms  $A_k$  defined by (3.3) drop out of the network dynamics.

---

<sup>1</sup> Without adaptation of intrinsic excitabilities, the network would start performing erroneous inference, learning would reinforce this erroneous behavior, and performance would quickly break down. We have verified this in simulations for the present WTA model: Consistently across trials, a small subset of WTA neurons became dominantly active while most neurons remained silent.

### E-step: variational inference with homeostasis

The variational distribution  $q(\mathbf{z}|\mathbf{y})$  we consider for the model (3.1)-(3.3) is a  $2^N \cdot K$  dimensional object. Since  $q$  describes a conditional probability distribution, it is non-negative and normalized for all  $\mathbf{y}$ . In addition, we constrain  $q$  to be a ‘‘homeostatic’’ distribution  $q \in \mathcal{Q}$  such that the average activation of each hidden variable (neuron)  $z_k$  equals an a-priori specified mean activation  $m_k$  under the input statistics  $p^*(\mathbf{y})$ . This is sketched in Fig. 3.2. Formally we define the constraint set,

$$\mathcal{Q} = \{q : \langle z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} = m_k, \text{ for all } k = 1 \dots K\} , \quad \text{with } \sum_k m_k = 1 . \quad (3.9)$$

The constrained maximization problem  $q^*(\mathbf{z}|\mathbf{y}) = \arg \max_{q \in \mathcal{Q}} \mathcal{F}(\mathbf{V}, q(\mathbf{z}|\mathbf{y}))$  can be solved with the help of Lagrange multipliers (cf. (GRACA et al., 2008)). We find that the  $q^*$  which maximizes the objective function  $\mathcal{F}$  during the E-step (and thus minimizes the KL-divergence to the posterior  $p(\mathbf{z}|\mathbf{y}, \mathbf{V})$ ) has the convenient form  $q^*(\mathbf{z}|\mathbf{y}) \propto p(\mathbf{z}|\mathbf{y}, \mathbf{V}) \cdot \exp(\sum_k \beta_k^* z_k)$  with some  $\beta_k^*$ . Hence, it suffices to consider distributions of the form,

$$q_{\beta}(\mathbf{z}|\mathbf{y}) \propto \exp\left(\sum_k z_k \left(\sum_i V_{ki} y_i + \underbrace{\hat{b}_k - A_k + \beta_k}_{=: b_k}\right)\right) , \quad (3.10)$$

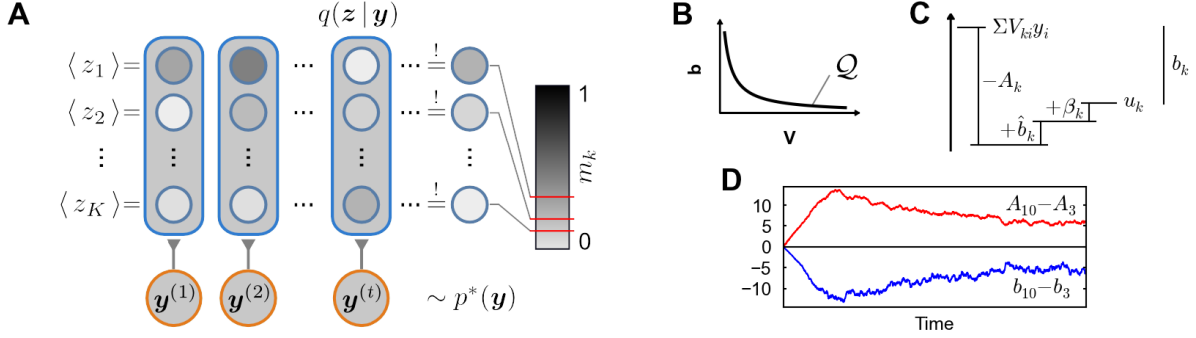
for the maximization problem. We identify  $\beta_k$  as the variational parameters which remain to be optimized. Note that any distribution of this form can be implemented by the spiking network  $\mathcal{N}$  if the intrinsic excitabilities are set to  $b_k = -A_k + \hat{b}_k + \beta_k$ . The optimal variational distribution  $q^*(\mathbf{z}|\mathbf{y}) = q_{\beta^*}(\mathbf{z}|\mathbf{y})$  then has  $\beta^* = \arg \max_{\beta} \Psi(\beta)$ , i.e. the variational parameter vector which maximizes the dual (GRACA et al., 2008),

$$\Psi(\beta) = \sum_k \beta_k m_k - \langle \log \sum_{\mathbf{z}} p(\mathbf{z}|\mathbf{y}, \mathbf{V}) \exp(\sum_k \beta_k z_k) \rangle_{p^*(\mathbf{y})} . \quad (3.11)$$

Due to concavity of the dual, a unique global maximizer  $\beta^*$  exists, and thus also the corresponding optimal intrinsic excitabilities  $b_k^* = -A_k + \hat{b}_k + \beta_k^*$  are unique. Hence, the posterior constraint  $q \in \mathcal{Q}$  can be illustrated as in Fig. 3.2B: For each synaptic weight configuration  $\mathbf{V}$  there exists, under a particular input distribution  $p^*(\mathbf{y})$ , a unique configuration of intrinsic excitabilities  $\mathbf{b}$  such that the resulting network output fulfills the homeostatic constraints. The theoretical relation between the intrinsic excitabilities  $b_k$ , the original nontrivial term  $-A_k$  and the variational parameters  $\beta_k$  is sketched in Fig. 3.2C. Importantly, while  $b_k$  is implemented in the network,  $A_k$ ,  $\beta_k$  and  $\hat{b}_k$  are not explicitly represented in the implementation anymore. Finding the optimal  $\mathbf{b}$  in the dual perspective, i.e. those intrinsic excitabilities which fulfill the homeostatic constraints, amounts to gradient ascent  $\partial_{\beta} \Psi(\beta)$  on the dual, which leads to the following homeostatic learning rule for the intrinsic excitabilities,

$$\Delta b_k \propto \partial_{\beta_k} \Psi(\beta) = m_k - \langle z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} . \quad (3.12)$$

### 3 Homeostatic and Hebbian plasticity united under a Bayesian account



**Fig. 3.2: Theory for the WTA model.** (A) Homeostatic posterior constraints in the WTA model: Under the variational distribution  $q$ , the average activation of each variable  $z_k$  must equal  $m_k$ . (B) For each set of synaptic weights  $\mathbf{V}$  there exists a unique assignment of intrinsic excitabilities  $\mathbf{b}$ , such that the constraints are fulfilled. (C) Theoretical decomposition of the intrinsic excitability  $b_k$  into  $-A_k$ ,  $\hat{b}_k$  and  $\beta_k$ . (D) During variational EM the  $b_k$  predominantly “track” the dynamically changing non-local terms  $-A_k$  (relative comparison between two WTA neurons from Figure 3.1).

Note that the intrinsic homeostatic plasticity rule (3.7) in the network corresponds to a sample-based stochastic version of this theoretically derived adaptation mechanism (3.12). Hence, given enough time, homeostatic plasticity will automatically install near-optimal intrinsic excitabilities  $\mathbf{b} \approx \mathbf{b}^*$  and implement the correct variational distribution  $q^*$  up to stochastic fluctuations in  $\mathbf{b}$  due to the non-zero learning rate  $\eta_b$ . The non-local terms  $A_k$  have entirely dropped out of the network dynamics, since the intrinsic excitabilities  $b_k$  can be arbitrarily initialized, and are then fully regulated by the local homeostatic rule, which does not require knowledge of  $A_k$ .

As a side remark, note that although the variational parameters  $\beta_k$  are not explicitly present in the implementation, they can be theoretically recovered from the network at any point, via  $\beta_k = b_k + A_k - \hat{b}_k$ . Notably, in all our simulations we have consistently found small absolute values of  $\beta_k$ , corresponding to a small KL-divergence between  $q^*$  and  $p$ .<sup>2</sup> Hence, a major effect of the local homeostatic plasticity rule during learning is to dynamically track and effectively implement the non-local terms  $-A_k$ . This is shown in Fig. 3.2D, in which the relative excitabilities of two WTA neurons  $b_k - b_j$  are plotted against the corresponding non-local  $A_k - A_j$  over the course of learning in the first simulation (Fig. 3.1).

#### M-step: interplay of synaptic and homeostatic intrinsic plasticity

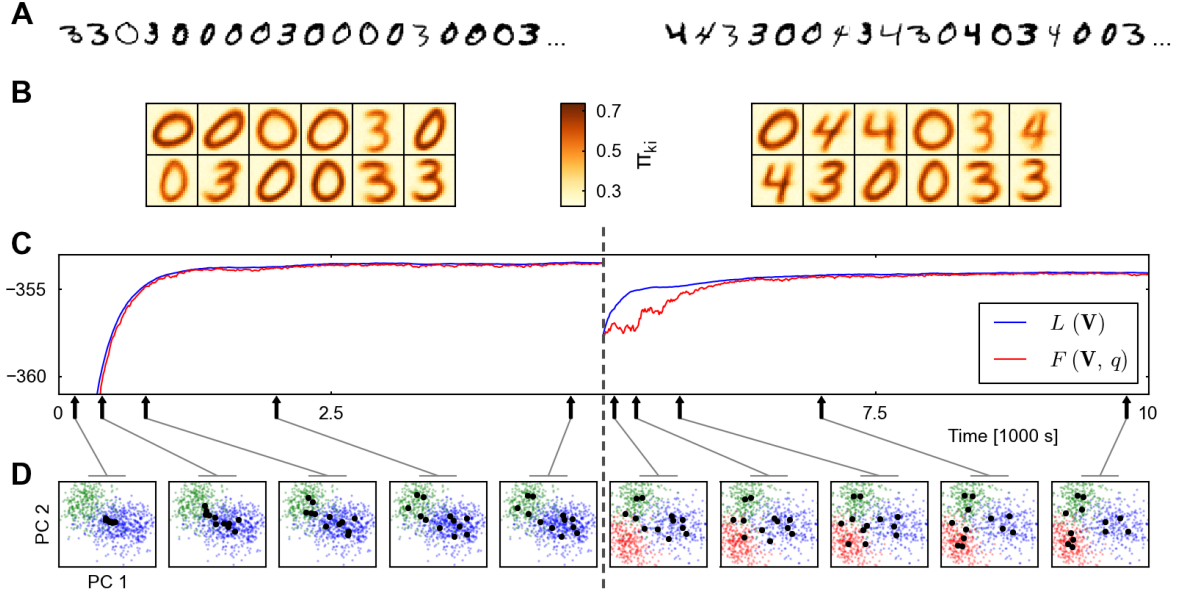
During the M-step, we aim to increase the EM lower bound  $\mathcal{F}$  in (3.6) w.r.t. the synaptic parameters  $\mathbf{V}$ . Gradient ascent yields,

$$\partial_{V_{ki}} \mathcal{F}(\mathbf{V}, q(\mathbf{z}|\mathbf{y})) = \langle \partial_{V_{ki}} \log p(\mathbf{y}, \mathbf{z}|\mathbf{V}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \quad (3.13)$$

$$= \langle z_k \cdot (y_j - \sigma(V_{ki})) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}, \quad (3.14)$$

<sup>2</sup>This is assuming for simplicity uniform prior parameters  $\hat{b}_k$ . Note that a small KL-divergence is in fact often observed during variational EM since  $\mathcal{F}$ , which contains the negative KL-divergence, is being maximized.

### 3.4 Spiking network dynamics with homeostatic plasticity



**Fig. 3.3: Dynamical properties of homeostatically regulated spiking networks.** (A) Input templates from MNIST dataset (digits 0,3 at a ratio 2:1, and digits 0,3,4 at a ratio 1:1:1) used during the first and second learning period, respectively. (B) Learned patterns at the end of each learning period. (C). Network performance converges in the course of learning.  $\mathcal{F}$  is a tight lower bound to  $\mathcal{L}$ . (D) Illustration of pattern learning and re-learning dynamics in a 2-D projection in the input space. Each black dot corresponds to the pattern  $\pi_{ki}$  of one WTA neuron  $k$ . Colored dots are input samples from the training set (blue/green/red  $\leftrightarrow$  digits 0/3/4).

where  $q$  is the variational distribution determined during the E-step, i.e., we can set  $q = q^*$ . Note the formal correspondence of (3.14) with the network synaptic learning rule (3.8). Indeed, if the network activity implements  $q^*$ , it can be shown easily that the expected update of synaptic weights due to the synaptic plasticity (3.8) is proportional to (3.14), and hence implements a stochastic version of the theoretical M-step (cf. (NESSLER et al., 2009a)).

### 3.4 Spiking network dynamics with homeostatic plasticity

To highlight a number of salient dynamical properties emerging from homeostatic plasticity in the considered WTA model, Fig. 3.3 shows a simulation of the same network  $\mathcal{N}$  with homeostatic dynamics as in Fig. 3.1, only with different input statistics presented to the network, and uniform  $m_k = \frac{1}{K}$ . During the first 5000s, different writings of 0's and 3's from the MNIST dataset were presented, with 0's occurring twice as often as 3's. Then the input distribution  $p^*(\mathbf{y})$  abruptly switched to include also 4's, with each digit occurring equally often. The following observations can be made: Due to the homeostatic constraint, each neuron responds on average to  $m_k \cdot T$  out of  $T$  presented inputs. As a consequence, the number of neurons which specialize on a particular digit is directly proportional to the frequency of occurrence of that digit, i.e. 8:4 and 4:4:4 after the first and second

### 3 Homeostatic and Hebbian plasticity united under a Bayesian account

learning period, respectively (Fig. 3.3B). In general, if uniform target activations  $m_k$  are chosen, output resources are allocated precisely in proportion to input frequency. Fig. 3.3C depicts the time course of the EM lower bound  $\mathcal{F}$  as well as the average likelihood  $\mathcal{L}$  (assuming uniform  $\hat{b}_k$ ) under the model during a single simulation run, demonstrating both convergence and tightness of the lower bound. As expected due to the stabilizing dynamics of homeostasis, we found variability in performance among different trials to be small (not shown). Fig. 3.3D illustrates the dynamics of learning and re-learning of patterns  $\pi_{ki}$  in a 2D projection of input patterns onto the first two principal components.

## 3.5 Homeostatic plasticity in recurrent spiking networks

The neural model so far was essentially a feed-forward network, in which every postsynaptic spike can directly be interpreted as one sample of the instantaneous posterior distribution (NESSLER et al., 2009a). The lateral inhibition served only to ensure the normalization of the posterior. We will now extend the concept of homeostatic processes as posterior constraints to the broader class of recurrent networks and sketch the utility of the developed framework beyond the regulation of intrinsic excitabilities.

Recently it was shown in (BUESING et al., 2011; PECEVSKI et al., 2011) that recurrent networks of stochastically spiking neurons can in principle carry out probabilistic inference through a sampling process. At every point in time, the joint network state  $z(t)$  represents one sample of a posterior. However, (BUESING et al., 2011) and (PECEVSKI et al., 2011) did not consider unsupervised learning on spiking input streams.

For the following considerations, we divide the definition of the probabilistic model in two parts. First, we define a Boltzmann distribution,

$$p(\mathbf{z}) = \exp\left(\sum_k \hat{b}_k z_k + \frac{1}{2} \sum_{j \neq k} \hat{W}_{kj} z_k z_j\right) / \text{norm.} \quad , \quad (3.15)$$

with  $\hat{W}_{kj} = \hat{W}_{jk}$  as “prior” for the hidden variables  $z$  which will be represented by a recurrently connected network of  $K$  spiking neurons. For the purpose of this section, we treat  $\hat{b}_k$  and  $\hat{W}_{kj}$  as constants. Secondly, we define a conditional distribution in the exponential-family form (BISHOP, 2006),

$$p(\mathbf{y}|\mathbf{z}, \mathbf{V}) = \exp(f_0(\mathbf{y}) + \sum_{k,i} V_{ki} z_k f_i(\mathbf{y}) - A(\mathbf{z}, \mathbf{V})) \quad , \quad (3.16)$$

that specifies the likelihood of observable inputs  $\mathbf{y}$ , given a certain network state  $\mathbf{z}$ . This defines the generative model  $p(\mathbf{y}, \mathbf{z}|\mathbf{V}) = p(\mathbf{z}) p(\mathbf{y}|\mathbf{z}, \mathbf{V})$ .

We map this probabilistic model to the spiking network and define that for every  $k$  and every point in time  $t$  the variable  $z_k(t)$  has the value 1, if the corresponding

### 3.5 Homeostatic plasticity in recurrent spiking networks

neuron has fired within the time window  $(t - \tau, t]$ . In accordance with the neural sampling theory, in order for a spiking network to sample from the correct posterior  $p(\mathbf{z}|\mathbf{y}, \mathbf{V}) \propto p(\mathbf{z}) p(\mathbf{y}|\mathbf{z}, \mathbf{V})$  given the input  $\mathbf{y}$ , each neuron must compute in its membrane potential the log-odd (BUESING et al., 2011),

$$u_k = \log \frac{p(z_k = 1 | \mathbf{z}_{\setminus k}, \mathbf{V})}{p(z_k = 0 | \mathbf{z}_{\setminus k}, \mathbf{V})} = \underbrace{\sum_i V_{ki} f_i(\mathbf{y})}_{\text{feedforward drive}} \underbrace{-A_k(\mathbf{V}) + \hat{b}_k}_{\text{intr. excitability}} + \sum_{j \neq k} \underbrace{(-A_{kj}(\mathbf{V}) + \hat{W}_{kj})}_{\text{recurrent weight}} z_j - \dots \quad (3.17)$$

where  $\mathbf{z}_{\setminus k} = (z_1, \dots, z_{k-1}, z_{k+1}, \dots, z_K)^\top$ . The  $A_k, A_{kj}, \dots$  are given by the decomposition of  $A(\mathbf{z}, \mathbf{V})$  along the binary combinations of  $\mathbf{z}$  as,

$$A(\mathbf{z}, \mathbf{V}) = A_0(\mathbf{V}) + \sum_k z_k A_k(\mathbf{V}) + \frac{1}{2} \sum_{j \neq k} z_k z_j A_{kj}(\mathbf{V}) + \dots \quad (3.18)$$

Note, that we do not aim at this point to give learning rules for the prior parameters  $\hat{b}_k$  and  $\hat{W}_{kj}$ . Instead we proceed as in the last section and specify a-priori desired properties of the average network response under the input distribution  $p^*(\mathbf{y})$ ,

$$c_{kj} = \langle z_k z_j \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \quad \text{and} \quad m_k = \langle z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} . \quad (3.19)$$

Let us explore some illustrative configurations for  $m_k$  and  $c_{kj}$ . One obvious choice is closely related to the goal of maximizing the entropy of the output code by fixing  $\langle z_k \rangle$  to  $\frac{1}{K}$  and  $\langle z_k z_j \rangle$  to  $\langle z_k \rangle \langle z_j \rangle = \frac{1}{K^2}$ , thus enforcing second order correlations to be zero. Another intuitive choice would be to set all  $\langle z_k z_j \rangle$  very close to zero, which excludes that two neurons can be active simultaneously and thus recovers the function of a WTA. It is further conceivable to assign positive correlation targets to groups of neurons, thereby creating populations with redundant codes. Finally, with a topographical organization of neurons in mind, all three basic ideas sketched above might be combined: one could assign positive correlations to neighboring neurons in order to create local cooperative populations, mutual exclusion at intermediate distance, and zero correlation targets between distant neurons.

With this in mind, we can formulate the goal of learning for the network in the context of EM with posterior constraints: we constrain the E-step such that the average posterior fulfills the chosen targets, and adapt the forward weights  $\mathbf{V}$  in the M-step according to (3.6). Analogous to the first-order case, the variational solution of the E-step under these constraints takes the form,

$$q_{\beta, \omega}(\mathbf{z}|\mathbf{y}) \propto p(\mathbf{z}|\mathbf{y}, \mathbf{V}) \cdot \exp \left( \sum_k \beta_k z_k + \frac{1}{2} \sum_{j \neq k} \omega_{kj} z_k z_j \right) , \quad (3.20)$$

with symmetric  $\omega_{kl} = \omega_{lk}$  as variational parameters. A neural sampling network  $\mathcal{N}$  with input weights  $V_{ki}$  will sample from  $q_{\beta, \omega}$  if the intrinsic excitabilities are set to  $b_k = -A_k + \hat{b}_k + \beta_k$ , and the symmetric recurrent synaptic weights to  $W_{kj} = -A_{kj} +$

### 3 Homeostatic and Hebbian plasticity united under a Bayesian account

$\hat{W}_{kj} + \omega_{kj}$ . The variational parameters  $\beta, \omega$  (and hence also  $\mathbf{b}, \mathbf{W}$ ) which optimize the dual problem  $\Psi(\mathbf{b}, \omega)$  are uniquely defined and can be found iteratively via gradient ascent. Analogous to the last section, this yields the intrinsic plasticity rule (3.12) for  $b_k$ . In addition, we obtain for the recurrent synapses  $W_{kj}$ ,

$$\Delta W_{kj} \propto c_{kj} - \langle z_k z_j \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} , \quad (3.21)$$

which translates to an anti-Hebbian spike-timing dependent plasticity rule in the network implementation.

For any concrete instantiation of  $f_0(\mathbf{y}), f_i(\mathbf{y})$  and  $A(\mathbf{z}, \mathbf{V})$  in (3.16) it is possible to derive learning rules for  $V_{ki}$  for the M-step via  $\partial_{V_{ki}} \mathcal{F}(\mathbf{V}, q)$ . Of course not all models entail local synaptic learning rules. In particular it might be necessary to assume conditional independence of the inputs  $\mathbf{y}$  given the network state  $\mathbf{z}$ , i.e.,  $p(\mathbf{y}|\mathbf{z}, \mathbf{V}) = \prod_i p(y_i|\mathbf{z}, \mathbf{V})$ . Furthermore, in order to fulfill the neural computability condition (3.17) for neural sampling (BUESING et al., 2011) with a recurrent network of point neurons, it might be necessary to choose  $A(\mathbf{z}, \mathbf{V})$  such that terms of order higher than 2 vanish in the decomposition. This can be shown to hold, for example, in a model with conditionally independent Gaussian distributed inputs  $y_i$ . It is ongoing work to find further biologically realistic network models in the sense of this theory and to assess their computational capabilities through computer experiments.

## 3.6 Discussion

Complex and non-local computations, which appear during probabilistic inference and learning, arguably constitute one of the cardinal challenges in the development of biologically realistic Bayesian spiking network models. In this paper we have introduced homeostatic plasticity, which to the best of our knowledge had not been considered before in the context of EM in spiking networks, as a theoretically grounded approach to stabilize and facilitate learning in a large class of network models. Our theory complements previously proposed neural mechanisms and provides, in particular, a simple and biologically realistic alternative to the assumptions on the input distribution made in (NESSLER et al., 2009a) and (KECK et al., 2012). Indeed, our results challenge the hypothesis of (KECK et al., 2012) that feedforward inhibition is critical for correctly learning the structure of the data with biologically plausible plasticity rules. More generally, it turns out that the enforcement of a balancing posterior constraint often simplifies inference in recurrent spiking networks by eliminating nontrivial computations. Our results suggest a crucial role of homeostatic plasticity in the Bayesian brain: to constrain activity patterns in cortex to assist the autonomous optimization of an internal model of the environment.



# Chapter 4

## Statistical learning with compound memristive synapses

### Contents

4.1	Introduction . . . . .	50
4.2	Stochastic memristors as plastic synapses . . . . .	52
4.3	Compound memristive synapses in winner-take-all networks . . . . .	57
4.4	Memristive synapses support inference and online learning . . . . .	58
4.5	Demonstration of unsupervised learning . . . . .	63
4.6	Influence of synaptic resolution . . . . .	64
4.7	Robustness to device variations . . . . .	66
4.8	Discussion . . . . .	69

**Abstract.** Memristors have recently emerged as promising circuit elements to mimic the function of biological synapses in neuromorphic computing. The fabrication of reliable nanoscale memristive synapses, that feature continuous conductance changes based on the timing of pre- and postsynaptic spikes, has however turned out to be challenging. In this article, we propose an alternative approach, the compound memristive synapse, that circumvents this problem by the use of memristors with binary memristive states. A compound memristive synapse employs multiple bistable memristors in parallel to jointly form one synapse, thereby providing a spectrum of synaptic efficacies. We investigate the computational implications of synaptic plasticity in the compound synapse by integrating the recently observed phenomenon of stochastic filament formation into an abstract model of stochastic switching. Using this abstract model, we first show how standard pulsing schemes give rise to spike-timing dependent plasticity (STDP) with a stabilizing weight dependence in compound synapses. In a next step, we study unsupervised learning with compound synapses in networks of spiking neurons organized in a winner-take-all architecture. Our theoretical analysis reveals that compound-synapse STDP implements generalized Expectation-Maximization in the spiking network. Specifically, the emergent synapse configuration represents the most salient features of the input distribution in a Mixture-of-Gaussians generative model. Furthermore, the network’s spike response to spiking input streams approximates a well-defined

## 4 Statistical learning with compound memristive synapses

Bayesian posterior distribution. We show in computer simulations how such networks learn to represent high-dimensional distributions over images of handwritten digits with high fidelity even in presence of substantial device variations and under severe noise conditions. Therefore, the compound memristive synapse may provide a synaptic design principle for future neuromorphic architectures.

**Acknowledgments and author contributions.** This chapter is based on the publication

JOHANNES BILL AND ROBERT LEGENSTEIN (2014). “A compound memristive synapse model for statistical learning through STDP in spiking neural networks.” In: *Frontiers in neuroscience*.

To this study, I contributed as first author. JB and RL conceived the study, designed the experiments and wrote the manuscript. JB developed the theory and performed the computer simulations. The authors thank Radu Berdan, Alex Serb and Themis Prodromakis for valuable discussions on memristors, Giacomo Indiveri for initial discussions on the compound memristive synapse model, Georg Goeri for preparatory studies on robustness to device variability, and David Kappel as well as two anonymous reviewers for helpful comments on the manuscript. JB thanks Bernhard Nessler and Stefan Habenschuss for all the fruitful discussions on the SEM theory and statistical learning in general.

### 4.1 Introduction

A characteristic property of massively parallel computation in biological and artificial neural circuits is the need for intensive communication between neuronal elements. As a consequence, area- and energy consumption of neuromorphic circuits is often dominated by those circuits that implement synaptic transmission between neural elements (SCHEMMEL et al., 2008). Biological synapses are highly dynamic computational entities, exhibiting plasticity on various time scales (MALENKA and BEAR, 2004). In order to capture the most salient of these aspects, silicon synapses thus demand extensive circuitry if implemented in CMOS technology. On this account, novel nanoscale circuit elements have recently gained interest in the field of neuromorphic engineering as a promising alternative solution for the implementation of artificial synaptic connections. In particular, for mixed-signal neuromorphic CMOS architectures, which combine traditional digital circuits with analog components, memristors are considered a promising class of circuit elements due to high integration densities, synapse-like plasticity dynamics, and low power consumption (INDIVERI et al., 2013; S. H. Jo et al., 2010; CHOI et al., 2009; KUZUM et al., 2011). One particularly important feature of memristors is that their electrical resistance (often termed “memristance”) can be altered in a persistent manner by applying a voltage to its terminals, leading to the eponymous perception of memristors as resistors with memory. As an important application of this property, it was

shown that the memristance can be changed based on the spike timings of the pre- and postsynaptic neurons in a manner that approximates spike-timing dependent plasticity (STDP), a plasticity rule that is believed to represent a first approximation for the changes of synaptic efficacies in biological synapses (MARKRAM et al., 1997; CAPORALE and DAN, 2008; MARKRAM et al., 2012). On the level of single synapses, this important property has been confirmed experimentally in real memristors (MAYR et al., 2012) and has been included into computational models of memristive plasticity (SERRANO-GOTARREDONA et al., 2013). On the network level, models of memristive STDP were employed in computer simulations to demonstrate the potential applicability of neuromorphic designs with memristive synapses to pattern recognition tasks (QUERLIOZ et al., 2011).

In practice however, the production of functional memristive synapses with nanoscale dimensions has proven difficult, mainly due to large device variations and their unreliable behavior. It turned out to be particularly challenging to fabricate reliable nanoscale memristive synapses that feature a continuous spectrum of conductance values. As an alternative solution, it was proposed to employ bistable memristors as neuromorphic synapses instead since they exhibit a high degree of uniformity (FANG et al., 2011) and high durability (S. H. Jo et al., 2009b). For switching between their two stable conductance states, bistable memristors can be operated in a deterministic as well as in a stochastic regime (S. H. Jo et al., 2009b).

Using machine learning theory, we show in this article that stochastically switching bistable memristors become computationally particularly powerful in mixed-signal neuromorphic architectures when multiple memristors are combined to jointly form one synapse. Such a joint operation of multiple memristors can be interpreted as the collective function of ion channels in biological synapses (INDIVERI et al., 2013). Concretely, we propose the *compound memristive synapse model* which employs  $M$  bistable memristors operating in parallel to form a single synaptic weight between two neurons. To implement synaptic plasticity, we employ standard STDP pulsing schemes (QUERLIOZ et al., 2011; SERRANO-GOTARREDONA et al., 2013) and exploit the stochastic nature of memristive switching (YU et al., 2013; S. H. Jo et al., 2009b; GABA et al., 2013; SURI et al., 2013). For the analysis of the resulting plasticity dynamics, we perceive individual memristors as binary stochastic switches. This abstract description was previously utilized to capture the most salient features of experimentally observed memristive switching (SURI et al., 2013) and appears compatible with pivotal aspects of the experimental literature (S. H. Jo et al., 2009b; GABA et al., 2013).

We show analytically and through computer simulations that the change of the synaptic efficacy for a given pairing of pre- and postsynaptic spikes follows an STDP-like plasticity rule such that the expected weight change depends on the momentary synaptic weight in a stabilizing manner. The resulting *compound-synapse STDP* enables a synapse to attain many memristive states depending on the history of pre- and postsynaptic activity. A stabilizing weight dependence of synaptic plasticity exists in biological synapses (BI and POO, 1998) and has been shown to facilitate learning and adaptation in neural systems (MORRISON et al., 2007;

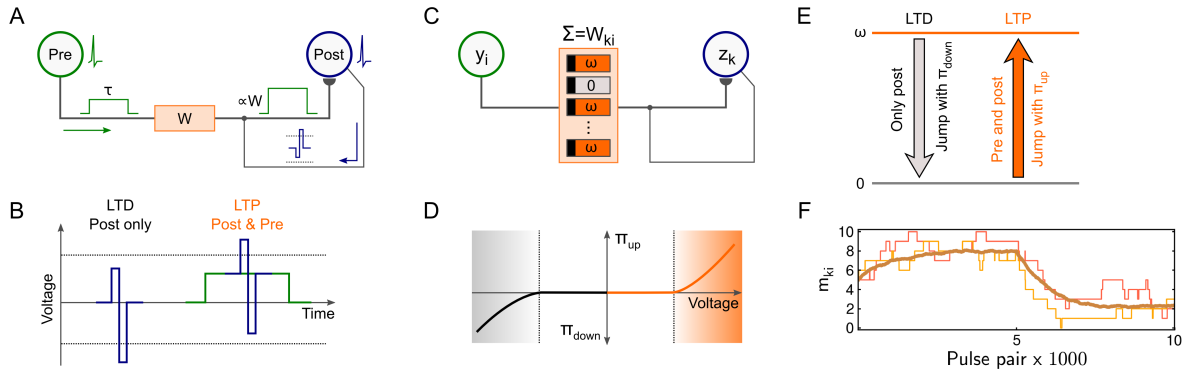
## 4 Statistical learning with compound memristive synapses

M. C. VAN ROSSUM et al., 2000). In particular, it has been shown in (NESSLER et al., 2013) that in stochastic winner-take-all (WTA) architectures, STDP with stabilizing weight dependence implements an online Expectation-Maximization algorithm. When exposed to input examples, neurons in the WTA network learn to represent the hidden causes of the observed input in a well-defined generative model. This adaptation proceeds in a purely unsupervised manner. We adopt a similar strategy here and apply the compound memristive synapse model in a network of stochastically spiking neurons arranged in a WTA architecture. We show analytically that compound-synapse STDP optimizes the synaptic efficacies such that the WTA network neurons in the hidden layer represent the most salient features of the input distribution in a Mixture-of-Gaussians generative model. After training, the network performs Bayesian inference over the hidden causes for the given input pattern. We show in computer simulations that such networks are able to learn to represent high-dimensional distributions over images of handwritten digits. After unsupervised training, the network transforms noisy input spike-patterns into a sparse and reliable spike code that supports classification of images. It turns out that even small compound synapses, consisting of only four bistable constituents per synapse, are sufficient for reliable image classification in our simulations. We furthermore show that the proposed model is able to represent the input distribution with high fidelity even in the presence of substantial device variations and under severe noise conditions. These findings render the compound synapse model a promising design principle for novel high-density, low-power mixed-signal CMOS architectures.

### 4.2 Stochastic memristors as plastic synapses

Memristors have gained increasing attention in neuromorphic engineering as possible substrates for plastic synapses (S. H. JO et al., 2010) due to the possibility to change their electrical conductance without the requirement of extensive supporting circuitry. Recently, ZAMARREÑO-RAMOS et al. (2011) and QUERLIOZ et al. (2011) have proposed a pulsing scheme to realize spike-timing dependent plasticity (STDP) with memristive synapses in response to pre- and postsynaptic activity as sketched in Fig. 4.1A: Pre-synaptic spikes trigger a rectangular voltage pulse of duration  $\tau$  (shown in green) that is sent to the memristor's presynaptic terminal. Similarly, postsynaptic neurons send back a copy of their spikes to the postsynaptic terminal as a brief voltage pulse (shown in blue). The combined effect of these pulses was shown to trigger STDP-type plasticity in the memristor as illustrated in Fig. 4.1B, where we adopted the convention to measure the voltage drop at the memristor as "presynaptic minus postsynaptic potential". If only the postsynaptic neuron spikes (Fig. 4.1B, left), the voltage pulse exceeds the lower threshold of the memristor, leading to long-term depression (LTD). Conversely, if a postsynaptic spike follows a presynaptic spike within duration  $\tau$  (Fig. 4.1B, right), the combined voltage trace exceeds the memristor's upper threshold, thus resulting in long-term

## 4.2 Stochastic memristors as plastic synapses



**Fig. 4.1: Compound synapse model with stochastic memristors.** (A) STDP pulsing scheme. Input spikes elicit a rectangular voltage trace (green, left) that is sent to the presynaptic terminal of the memristor synapse. Post-synaptic spikes elicit a brief voltage pulse (blue) which is sent back to the synapse. (B) Solitary postsynaptic spikes trigger LTD since the voltage exceeds the lower threshold of the memristor. Simultaneous pre- and postsynaptic spikes trigger LTP since the voltage exceeds the upper threshold. (C) Compound memristive synapse model. A synapse is composed of  $M$  bistable memristors operating in parallel. Each memristor can either be active (weight  $\omega$ ) or inactive (weight 0). The total synaptic weight  $W_{ki}$  between input neuron  $y_i$  and network neuron  $z_k$  is the sum of the individual memristor weights. (D) Bistable memristors switch stochastically between the active and inactive state depending on the applied voltage difference across its terminals. Switching to the active state (inactive state) occurs with probability  $\pi_{up}$  ( $\pi_{down}$ ) if a certain threshold voltage (dotted line) is exceeded. (E) Summary of stochastic transitions for compound-synapse STDP. (F) In an STDP pairing experiment, the stabilizing weight dependence of compound synapse plasticity governs convergence to a dynamic equilibrium. 10,000 plasticity pulses were applied to a synapse with  $M = 10$  constituents. During the first half, 80% (20%) of the events were of LTP (LTD) type. During the second half, the probability for LTP (LTD) events was inverted to 20% (80%). Thin lines: number of active memristors  $m_{ki}(t)$  for two example simulation runs. Thick line: Average  $\langle m_{ki}(t) \rangle$  over 100 runs. The average weight converges to a dynamic equilibrium.

potentiation (LTP). Pre-synaptic pulses alone do not trigger any plasticity. Overall, the memristor's conductance obeys an STDP-type plasticity rule.

The direct implementation of the above plasticity rule in mixed-signal neuromorphic architectures, however, faces practical challenges due to the continuous spectrum of the conductance values it relies on. Memristors that support a (quasi-) continuous spectrum of memristive states often suffer from instabilities to maintain their conductance value ("volatility") and typically show unreliable changes under repetitive application of identical pulses.

Here we explore an alternative approach that employs multiple bistable memristors, that support only two distinct conductance states per memristor, to jointly form one synapse. Such devices were reported to exhibit a high degree of uniformity (FANG et al., 2011). Concretely, we propose a synapse model which employs  $M$  bistable memristors operating in parallel to form a single synaptic weight  $W_{ki}$  between the  $i$ -th input neuron  $y_i$  and the  $k$ -th network neuron  $z_k$ . The model which we refer to as *compound memristive synapse* is sketched in Fig. 4.1C. Each memristor is assumed to provide two stable states: a high-conductive (active) state and a low-conductive (inactive) state. Since the dynamic range of memristors typically covers several orders of magnitude, the weight contribution of inactive memristors is almost negligible. In line with this notion, each inactive memristor contributes

## 4 Statistical learning with compound memristive synapses

weight 0 in the synapse model and each active memristor contributes weight  $\omega$ . Since parallel conductances sum up, the total weight of the compound memristive synapse reads

$$W_{ki} = \omega \cdot m_{ki} \quad (4.1)$$

with  $m_{ki} \in \{0, 1, \dots, M\}$  denoting the number of active memristors. As a consequence, a compound memristive synapse supports  $M + 1$  discrete weight levels, ranging from 0 to the maximum weight  $W^{\max} = \omega \cdot M$ .

Plasticity in this synapse model naturally emerges from transitions between the active and inactive state of the bistable constituents. However, deterministic transitions, which are, for instance, desirable in memristor-based memory cells, impair the performance in a neuromorphic online learning setup with compound memristive synapses: If all constituents of a synapse, that experience the same pre- and postsynaptic spikes, change their state simultaneously, the compound weight toggles between  $W_{ki} = 0$  and  $W_{ki} = W^{\max}$  depending on the latest pulse pair, not showing any gradual trace of memory formation as required for STDP.

A possible remedy to this issue can be found in memristors that exhibit stochastic rather than deterministic switching between their stable states. Yu et al. (2013), for instance, reported stochastic transitions in  $\text{HfO}_x/\text{TiO}_x$  memristors (from the class of anion-based memristors) and explored in computer simulations how stochastic bistable memristors could be used in a neuromorphic learning architecture. Similar studies explored the usability of stochastically switching bistable cation-based memristive materials (S. H. Jo et al., 2009b; GABA et al., 2013; SURI et al., 2013). In these nanoscale devices, changes in the memristance were shown to be dominated by the formation of a single conductive filament (S. H. Jo et al., 2009b). Stochastic switching was demonstrated for both directions (active  $\leftrightarrow$  inactive) (SURI et al., 2013) with switching probabilities being adjustable via the duration (SURI et al., 2013; GABA et al., 2013) and amplitude (S. H. Jo et al., 2009b; GABA et al., 2013) of the voltage applied across the terminals. These observations have led to the conclusion that “switching can be fully stochastic” with switching probabilities being almost unaffected by the rate at which consecutive plasticity pulses are delivered (GABA et al., 2013). Furthermore, bistable memristors can be extremely durable, not showing any notable degradation over hundreds of thousands of programming cycles (S. H. Jo et al., 2009b). Owing to these pivotal properties, SURI et al. (2013) proposed an STDP-type plasticity rule that perceives bistable memristors as simple stochastic switches.

Here we generalize this idea to compound memristive synapses and investigate the computational function of the arising plasticity rules in spiking networks from a machine learning perspective. Fig. 4.1D illustrates a simple model of stochastic switching of individual memristors that employs the STDP-pulsing scheme from ZAMARREÑO-RAMOS et al. (2011) and QUERLIOZ et al. (2011) discussed above: If the voltage between the pre- and postsynaptic terminal of a device exceeds a certain threshold (dotted lines) for the duration of the back-propagating spike signal, stochastic switching may occur. Inactive memristors jump to the active state

## 4.2 Stochastic memristors as plastic synapses

with probability  $\pi_{\text{up}}$  provided a sufficiently strong positive voltage, resulting in stochastic LTP. Similarly, active memristors turn inactive with probability  $\pi_{\text{down}}$  given a sufficiently strong negative voltage, leading to stochastic LTD. No switching occurs if only a small (or zero) voltage is applied, or if the memristor is already in the respective target state. Since the applied pre- and postsynaptic pulse amplitudes are free parameters in the model, the jumping probabilities  $\pi_{\text{up}}$  and  $\pi_{\text{down}}$  can be controlled, to a certain extent, by the experimenter. This simple model captures the most salient aspects of stochastic switching in bistable memristive devices as discussed above, cp. also (SURI et al., 2013) and the Discussion section. In order to distinguish the abstract memristor model from physical memristive devices, we will in the following refer to the model memristors as “stochastic switches”, “constituents”, or simply as “switches” for the sake of brevity. Furthermore, we refer to the resulting stochastic plastic behavior of compound memristive synapses in response to pre-post spike pairs, summarized in Fig. 4.1E, as *compound-synapse STDP*.

From the transition probabilities of individual stochastic switches we can calculate the expected temporal weight change  $\left\langle \frac{d}{dt} W_{ki} \right\rangle$  of the compound memristive synapse as a function of pre- and postsynaptic activity. Formally, we denote the presence of a rectangular input pulse (green in Fig. 4.1A) of the  $i$ -th input by  $y_i(t) = 1$  (and the absence by  $y_i(t) = 0$ ). The brief pulses that are sent back from a postsynaptic neuron  $z_k$  to the synapse (blue in Fig. 4.1A) are formally treated as point events at the spike times of the postsynaptic neuron. We denote the spike time of the  $f^{\text{th}}$  spike of neuron  $z_k$  by  $t_k^f$ . The spike train  $s_k(t)$  of a neuron  $z_k$  is formally defined as the sum of Dirac delta pulses  $\delta(\cdot)$  at the spike times:  $s_k(t) = \sum_f \delta(t - t_k^f)$ . When a synaptic efficacy  $W_{ki}$  is subject to a stochastic LTP update, there are  $(M - m_{ki})$  constituents in the compound memristive synapse that are currently inactive and could undergo an LTP transition. Each constituent independently switches to its active state with probability  $\pi_{\text{up}}$ , thereby contributing  $\omega$  to the compound weight  $W_{ki}$ . Hence, the expected weight change for the LTP condition reads  $(M - m_{ki}) \omega \pi_{\text{up}}$ . A similar argumentation applies to the LTD case. In summary, considering that plasticity always requires a postsynaptic spike, and that LTP is induced in the presence of a presynaptic pulse ( $y_i(t) = 1$ ), while LTD is induced in the absence of a presynaptic pulse ( $y_i(t) = 0$ ), the expected weight change of the compound memristive synapse reads:

$$\begin{aligned} \left\langle \frac{d}{dt} W_{ki} \right\rangle &= s_k(t) \cdot \left[ \underbrace{(M - m_{ki}) \omega \pi_{\text{up}} y_i(t)}_{\text{LTP}} - \underbrace{m_{ki} \omega \pi_{\text{down}} (1 - y_i(t))}_{\text{LTD}} \right] \quad (4.2) \\ &= s_k(t) \cdot [M \omega \pi_{\text{up}} y_i(t) - m_{ki} \omega \pi_{\text{up}} y_i(t) - m_{ki} \omega \pi_{\text{down}} + m_{ki} \omega \pi_{\text{down}} y_i(t)] \\ &= s_k(t) \cdot [W^{\text{max}} \pi_{\text{up}} y_i(t) - W_{ki} \pi_{\text{down}}] + s_k(t) W_{ki} y_i(t) \cdot (\pi_{\text{down}} - \pi_{\text{up}}) . \end{aligned}$$

In order to obtain a simple closed form solution of the weight changes, we set  $\pi_{\text{up}} = \pi_{\text{down}}$ , i.e., the probability of potentiation of a single switch under LTP equals its probability of depression under LTD. This choice will facilitate the

#### 4 Statistical learning with compound memristive synapses

theoretical analysis of learning in a spiking network, later on. In a hardware implementation, the switching probabilities could be adjusted via the pre- and postsynaptic amplitudes of the STDP pulses. We obtain the following closed form solution for the expected weight change:

$$\left\langle \frac{d}{dt} W_{ki} \right\rangle = \pi_{\text{up}} W^{\text{max}} \cdot s_k(t) \cdot [y_i(t) - W_{ki}/W^{\text{max}}] \quad . \quad (4.3)$$


Notably, the plasticity rule (4.3) differs from standard additive STDP rules in that it includes the weight dependent term  $W_{ki}/W^{\text{max}}$ . This weight dependence has its origin in the varying number of (in-)active stochastic switches  $m_{ki}$  that could actually undergo plastic changes and is in line with a prominent finding from neurobiology (Bi and Poo, 1998): Relative changes  $\Delta W_{ki}/W_{ki}$  become weaker for strong weights under LTP, while under LTD the relative change is weight-independent. Studies in computational neuroscience, see e.g. (M. C. VAN ROSSUM et al., 2000), found that this type of weight dependence facilitates the formation of stable connections in spiking networks.

In Fig. 4.1F, we illustrate the stochastic convergence of a compound synaptic weight to a stable, dynamic equilibrium in a simple STDP pairing experiment. The synapse consists of  $M = 10$  bistable switches with switching probabilities set to  $\pi_{\text{up}} = \pi_{\text{down}} = 0.001$ . These synapse parameters will also be used in network level simulations, later on. In a small computer simulation, 5 of the 10 constituents are initially set active ( $m_{ki}(t = 0) = 5$ ). Then, 5000 postsynaptic spikes are sent to the postsynaptic terminal of the synapse for triggering stochastic switching in the bistable constituents. 80% of these events are randomly paired with a presynaptic pulse, i.e., 80% of the events are of LTP type and 20% are of LTD type. After 5000 plasticity pulses, the statistics are inverted to 20% LTP and 80% LTD events for another 5000 plasticity pulses. The thin lines in Fig. 4.1F show the evolution of  $m_{ki}(t)$  in this STDP pairing experiment for two independent simulation runs. In the first half, the synaptic weight tends to settle around higher weight values, while in the second half, it stochastically declines towards lower weight values. The gradual convergence of the average weight, as suggested by eq. (4.3), becomes apparent by taking the mean over 100 simulation runs (thick line). The stabilizing weight dependence of the STDP rule leads to convergence to a dynamic equilibrium such that the mean value shows slow, continuous changes as expected from the theory. Individual synapses fluctuate stochastically around this mean.

Eq. (4.3) is reminiscent of a theoretically derived synaptic plasticity rule for statistical model optimization in spiking neural networks proposed by NESSLER et al. (2013). Building upon the theoretical approach developed by NESSLER et al. (2013), we will next turn to the question how to conceive stochastic learning with compound memristive synapses from a Bayesian perspective as unsupervised model optimization via a powerful optimization method that is known as Expectation-Maximization in the machine learning literature.



### 4.3 Compound memristive synapses in winner-take-all networks



**Fig. 4.2: Spiking network for probabilistic inference and online learning.** (A) Winner-take-all network architecture with lateral inhibition and synaptic weights  $W_{ki}$ . (B) Network neurons  $z_k$  implicitly maintain a Gaussian likelihood function for each input  $y_i$  in their afferent synaptic weights  $W_{ki}$ . The mean  $\mu_{ki}$  of the distribution is encoded by the fraction  $W_{ki}/W^{\max} = m_{ki}/M$  of active switches in the compound memristive synapse, i.e. stronger synaptic weights  $W_{ki}$  correspond to higher mean values  $\mu_{ki}$ . Inset: Local implicit graphical model. (C) Illustration of Bayesian inference for two competing network neurons  $z_k, z_j$  and one active input  $y_i(t) = 1$ . Different means  $\mu_{ki}, \mu_{ji}$  encoded in the weights give rise to different values in the likelihood function and shape the posterior distribution according to Bayes rule.

### 4.3 Compound memristive synapses in winner-take-all networks

The winner-take-all (WTA) network structure is a ubiquitous circuit motif in neocortex (DOUGLAS and MARTIN, 2004; LANSNER, 2009) and is often utilized in neuromorphic engineering (MEAD and ISMAIL, 1989; INDIVERI, 2000). Recently, WTA networks attracted increasing attention in theoretical studies on statistical learning (KECK et al., 2012; HABENSCHUSS et al., 2012; HABENSCHUSS et al., 2013b; NESSLER et al., 2013; KAPPEL et al., 2014) because their comparatively simple network dynamics facilitate a comprehensive mathematical treatment. In this section, we introduce the WTA architecture that we will use to study the learning capabilities of spiking neural networks with compound memristive synapses subject to STDP.

The WTA network architecture is sketched in Fig. 4.2A. The network consists of  $N$  spiking input neurons  $y_1, \dots, y_N$  and  $K$  spiking network neurons  $z_1, \dots, z_K$  with all-to-all connectivity in the forward synapses. Lateral inhibition introduces competition among the network neurons. Network neurons are stochastic spike response neurons (GERSTNER and KISTLER, 2002) with the membrane potential  $u_k$  of network neuron  $z_k$  being given by

$$u_k(t) = b_k + \sum_{i=1}^N W_{ki} \cdot y_i(t) . \quad (4.4)$$

The membrane potential  $u_k(t)$  integrates the inputs  $y_i(t)$ , i.e., the rectangular voltage pulses following each input spike, linearly through the synaptic weights  $W_{ki}$ . The parameter  $b_k$  denotes the intrinsic excitability of the neuron and controls its general disposition to fire. In Appendix C, we outline how the linear membrane potential (4.4) can be realized with leaky integrators, a common neuron model

## 4 Statistical learning with compound memristive synapses

in neuromorphic designs. For the spike response of the network neurons  $z_k$ , a stochastic firing mechanism is employed. In the WTA network, neurons  $z_k$  spike in a Poissonian manner with instantaneous firing rate  $\rho_k(t)$  that depends on the membrane potential  $u_k(t)$  and on lateral inhibition  $u_{\text{inh}}(t)$ :

$$\rho_k(t) = r_{\text{net}} \cdot e^{u_k(t) - u_{\text{inh}}(t)} , \quad (4.5)$$

with a constant  $r_{\text{net}} > 0$  that scales the overall firing rate of the network. In other words, the neuron spikes with probability  $\rho_k(t) \cdot \delta t$  in a small time window  $\delta t \rightarrow 0$ . The inhibitory contribution  $u_{\text{inh}}(t) := \log \sum_{j=1}^K \exp(u_j(t))$  summarizes the effect of lateral inhibition in the network and introduces WTA-competition between the network neurons to fire in response to a given stimulus  $y_1(t), \dots, y_N(t)$ . Notably, the exponential relationship (4.5) between an idealized membrane potential and neuronal firing is consistent with biological findings about the response properties of neocortical pyramidal neurons (JOLIVET et al., 2006).

The feed-forward synapses from inputs  $y_i$  to network neurons  $z_k$  are implemented as compound memristive synapses and their synaptic weights  $W_{ki}$  are adapted through stochastic STDP as described above. The intrinsic excitabilities  $b_k$  are adapted according to a homeostatic plasticity rule (HABENSCHUSS et al., 2012) that ensures that all network neurons take part in the network response and thus facilitates the emergence of a rich neural representation that covers the entire input space. Besides its observed stabilizing effect (QUERLIOZ et al., 2011), homeostatic intrinsic plasticity plays a distinct computational role when combined with synaptic learning: Network neurons that maintain many strong synapses  $W_{ki}$  gain an “unrightful advantage” during WTA competition over neurons that are specialized on low-activity input patterns (and therefore maintain weaker weights). A detailed analysis of the learning dynamics in the network shows that, in order to compensate for this advantage, the former must be burdened with a lower excitability  $b_k$  than the latter. A formal definition of the homeostatic plasticity mechanism and a discussion of its computational role from a theory perspective are provided in Appendix C, see also (HABENSCHUSS et al., 2012).

### 4.4 Memristive synapses support inference and online learning

In this section, we thoroughly analyze the learning effects of STDP in compound memristive synapses in the stochastic WTA network model. For the mathematical analysis, we describe the inputs  $y_i(t)$  and the stochastic neuron responses  $s_k(t)$  with the help of probability theory. To this end, we perceive the spiking activity of the input neurons  $y_i$  and network neurons  $z_k$  as samples of random variables (RVs)  $Y_i$  and  $Z_k$  respectively. Consistent with the assumption that spikes from input neurons produce voltage pulses of duration  $\tau$  in the circuit implementation (see Fig. 4.1A), we set  $Y_i = y_i(t)$ , i.e.,  $Y_i = 1$  if input neuron  $y_i$  spiked within  $[t - \tau, t]$

and  $Y_i = 0$  otherwise. The assignment of output spikes  $s_k(t)$  to RVs  $Z_k$  is different: The random variable  $Z_k$  labels the winner of the WTA network at spike times of network neurons. Hence, the value of  $Z_k$  is only defined at the moments when one of the  $K$  network neurons spikes, i.e., when the spike train  $s_j(t) \neq 0$  for some neuron  $z_j$ . In this case,  $Z_k$  encodes which neuron spiked, and we set  $Z_k = 1$  if  $k = j$  and  $Z_k = 0$  if  $k \neq j$ .

Using this interpretation of neural activity as realizations of RVs, the network's stochastic response  $s_k(t)$  to an input configuration  $\mathbf{y}(t) = (y_1(t), \dots, y_N(t))$  gives rise to a conditional probability distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  over the network RVs  $\mathbf{Z} := (Z_1, \dots, Z_K)$  conditioned on the input RVs  $\mathbf{Y} := (Y_1, \dots, Y_N)$ .<sup>1</sup> In line with the definition of the RVs  $Z_k$ , the distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  describes the network response only when one of the network neurons  $z_k$  fires. The response distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y} = \mathbf{y}(t))$  for any fixed input configuration  $\mathbf{y}(t)$  can directly be calculated from equations (4.4) and (4.5) (note that the probability  $p_{\text{net}}(Z_k = 1 | \mathbf{Y} = \mathbf{y}(t))$  for an individual RV  $Z_k$  to be active is proportional to the firing rate  $\rho_k(t)$  of neuron  $z_k$ ):

$$p_{\text{net}}(Z_k = 1 | \mathbf{Y} = \mathbf{y}(t)) = \rho_k(t) / r_{\text{net}} = e^{u_k(t) - u_{\text{inh}}(t)} = \frac{e^{b_k + \sum_{i=1}^N W_{ki} \cdot y_i(t)}}{\sum_{j=1}^K e^{b_j + \sum_{i=1}^N W_{ji} \cdot y_i(t)}} \cdot (4.6)$$

Equation (4.6) fully characterizes the network response distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  for any given input  $\mathbf{Y} = \mathbf{y}(t)$ . We next turn to the question how the response distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  can be understood as the result of a meaningful probabilistic computation. Specifically, we will show that the spike response of the WTA network approximates the Bayesian posterior distribution during inference in a well-defined probabilistic model. This probabilistic model is implicitly encoded in the synaptic weights  $W_{ki}$ , and synaptic plasticity can thus be perceived as an ongoing refinement of the involved probability distributions. Indeed, the STDP rule (4.3) of the compound memristive synapses turns out to be optimal in the sense that it instantiates *generalized Expectation-Maximization* in the WTA network, a powerful algorithm for unsupervised learning from machine learning theory. Our findings build upon theoretical work on synaptic learning in spiking neural networks from NESSLER et al. (2009a); NESSLER et al. (2013); HABENSCHUSS et al. (2012) and HABENSCHUSS et al. (2013b).

The key idea for identifying the response distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  as the result of a Bayesian computation, is to hypothetically reverse the network computation and view the spike response of a network neuron  $z_k$  as the *hidden cause* behind the observed input  $\mathbf{y}(t)$ . In this view, the network is treated as a *generative model* that implicitly defines a prior distribution  $p(\mathbf{Z})$  over hidden causes  $Z_k$  and a set of likelihood distributions  $p(\mathbf{Y} | Z_k = 1)$ , one for each hidden cause  $Z_k$ . The shape of

<sup>1</sup>In Appendix C, the distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  will be identified as the variational distribution  $q(\mathbf{z} | \mathbf{y})$  for generalized EM learning, in the sense of the Introduction of this thesis.

#### 4 Statistical learning with compound memristive synapses

the distributions is defined by the parameters of the network, e.g., the synaptic weights  $W_{ki}$ . An important property of these implicitly encoded distributions – that also motivates the term ‘generative model’ – is that they give rise to a (hypothetical) distribution over the inputs,  $p(\mathbf{Y} = \mathbf{y}(t)) = \sum_{k=1}^K p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1) \cdot p(Z_k = 1)$ . This equation also explains why a RV  $Z_k$  is called a hidden cause: If we observe an input vector  $\mathbf{y}(t)$ , that has high probability only in one of the likelihood distributions  $p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1)$ , then we can consider the RV  $Z_k$  as a likely (but unobservable) cause for the observation according to the generative model. A common objective in machine learning theory, known as *Maximum-Likelihood learning*, is to find parameters that bring the implicit distribution  $p(\mathbf{Y})$  of the generative model as close as possible to the distribution of the actually observed input. Then the hidden causes of the generative model are expected to represent important features of the observed input (e.g., some typical input clusters). Leaving the hypothetical generative perspective again, in the network’s real operation an input  $\mathbf{y}(t)$  is presented to the network and the hidden causes  $Z_k$  need to be inferred (e.g., the cluster the input  $\mathbf{y}(t)$  belongs to). The mathematically correct result of this inference is given by Bayes rule

$$p(Z_k = 1 | \mathbf{Y} = \mathbf{y}(t)) \propto p(Z_k = 1) \cdot p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1) \quad (4.7)$$

which combines the likelihood  $p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1)$  with the prior  $p(Z_k = 1)$ .

NESSLER et al. (2013) showed that in a WTA network architecture, that evolves according to eq. (4.4) and (4.5), the synaptic weights  $W_{ki}$  can be understood as an implicit neural encoding of likelihood distributions  $p(\mathbf{Y} | Z_k = 1)$ , and that the network response  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y})$  approximates the posterior distribution according to eq. (4.7). Hence, WTA networks can be regarded as implicit generative models. Furthermore – and even more importantly from a theoretical perspective – NESSLER et al. (2013) showed that the implicit likelihood model, that is encoded in the weights  $W_{ki}$ , can be optimized in an unsupervised manner by a weight-dependent STDP rule. Indeed, there exists a tight link between the type of weight dependence in the STDP rule and the type of implicit likelihood model it optimizes. The exponential weight dependence in Nessler’s rule, however, differs from the linear weight dependence we identified for the plasticity rule in eq. (4.3). This raises the question what type of implicit likelihood model is encoded and optimized by the compound memristive synapses.

An intuition about an appropriate probabilistic interpretation of compound-synapse STDP can be obtained from the equilibrium points of the plasticity rule (4.3). We first observe, that plasticity is always triggered by a postsynaptic spike, i.e.,  $Z_k = 1$  for some neuron  $z_k$ . In the spirit of spike-triggered averaging (E. P. SIMONCELLI et al., 2004), we can then study the conditional distribution  $p(Y_i = y_i(t) | Z_k = 1)$  of an input  $y_i$  at the moment of the network response since the coincidence of pre- and postsynaptic spiking activity is the driving force behind any weight change in the STDP rule. By assuming that plasticity has converged to a dynamic equilibrium, the average contributions of LTP and LTD cancel each other, i.e.,  $\left\langle \frac{d}{dt} W_{ki} \right\rangle_{p(Y_i | Z_k=1)} = 0$ .

From this condition, we obtain the following relation between the synaptic weight  $W_{ki}$  and the conditional input distribution  $p(Y_i | Z_k = 1)$ :

$$\begin{aligned} 0 &\stackrel{!}{=} \left\langle \frac{d}{dt} W_{ki} \right\rangle_{p(Y_i | Z_k=1)} = \langle \tau_{\text{up}} W^{\text{max}} \cdot [Y_i - W_{ki}/W^{\text{max}}] \rangle_{p(Y_i | Z_k=1)} \\ &\Rightarrow W_{ki} = W^{\text{max}} \cdot \langle Y_i \rangle_{p(Y_i | Z_k=1)} . \end{aligned} \quad (4.8)$$

According to this analysis, the synaptic weight  $W_{ki}$  represents the expected value of input neuron  $y_i$  at the moment of a postsynaptic spike in network neuron  $z_k$  in a linear manner. In the compound memristive synapse, this expectation is encoded in the  $M + 1$  possible weight states  $W_{ki} = 0, \omega, \dots, W^{\text{max}}$  of the synapse. The linear encoding (4.8) is compatible with the convergence points which we observed previously in the small STDP pairing experiment in Fig. 4.1F.

The above analysis only serves as an intuition and is no substitute for a thorough mathematical treatment of the learning process. A rigorous formal derivation that, for instance, also takes into account the dynamically changing response properties of the network neurons due to recurrent interactions and plastic changes in the weights, is provided in Appendix C. It reveals that the above intuition holds. More precisely, the likelihood distributions  $p(\mathbf{Y} | Z_k = 1)$  that are optimized in a WTA circuit with compound-synapse STDP are given by the product of the likelihoods of individual inputs

$$p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1) = \prod_{i=1}^N p(Y_i = y_i(t) | Z_k = 1) , \quad (4.9)$$

and the likelihood for each individual input channel  $y_i$  is given by a Gaussian distribution

$$p(Y_i = y_i(t) | Z_k = 1) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{(y_i(t) - \mu_{ki})^2}{2\sigma^2}} . \quad (4.10)$$

The mean values  $\mu_{ki}$  and the standard deviation  $\sigma$  of the likelihood distributions (4.10) are identified as

$$\mu_{ki} = W_{ki}/W^{\text{max}} = m_{ki}/M \quad \text{and} \quad \sigma = 1/\sqrt{(W^{\text{max}})} . \quad (4.11)$$

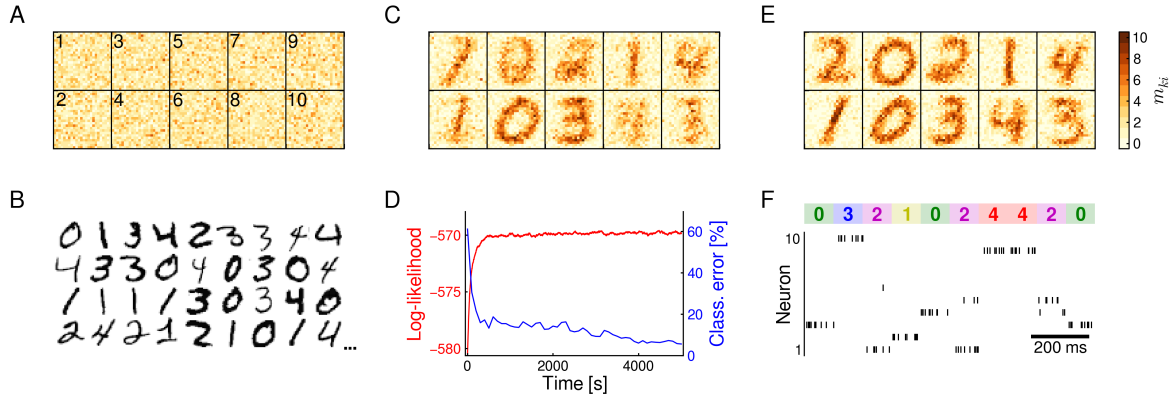
Hence, the mean  $\mu_{ki}$  of the distribution for input channel  $y_i$  is given by the fraction of active constituents in the compound memristive synapse. The width  $\sigma = 1/\sqrt{(W^{\text{max}})}$  of the distribution is determined by the maximum weight  $W^{\text{max}} = M \cdot \omega$  and could, for instance, be controlled by the weight contribution  $\omega$  of an individual stochastic switch. The resulting probabilistic model of the WTA network is a *Mixture-of-Gaussians* generative model (see Appendix C for a formal definition).

In order to illustrate the computational properties of this generative model, the likelihood distribution  $p(Y_i | Z_k = 1)$  for a single input  $y_i$  and a single active hidden cause  $Z_k = 1$  is sketched in Fig. 4.2B. An active hidden cause  $Z_k = 1$  assigns

probabilities to all possible instantiations  $y_i(t)$  of  $Y_i$ . In principle, the Gaussian likelihood distribution supports arbitrary real-valued input instantiations  $y_i(t) \in \mathbb{R}$ . We will come back to this observation in the Discussion section where we address possible extensions of the WTA network to support more complex input types. In this article, we consider only binary inputs that take on the value 0 (input pulse absent) or 1 (input pulse present), see the presynaptic pulses in Fig. 4.1A. The corresponding likelihood values  $p(Y_i = 0 | Z_k = 1)$  and  $p(Y_i = 1 | Z_k = 1)$  are determined by the mean  $\mu_{ki}$  and the variance  $\sigma$  of the likelihood distribution, see eqs. (4.10) and (4.11). The task of the network when presented with an input  $\mathbf{y}(t)$  is to infer the posterior distribution over hidden causes  $p(Z_k = 1 | \mathbf{Y} = \mathbf{y}(t))$  and produce spikes according to this distribution. The optimal solution is given by Bayes rule (4.7) with the likelihood  $p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1)$  given by equations (4.9) and (4.10), and an (input independent) a priori probability  $p(Z_k = 1)$ . As we prove in Appendix C, the response distribution  $p_{\text{net}}(\mathbf{Z} | \mathbf{Y} = \mathbf{y}(t))$  of the spiking WTA network implements a close and well-defined variational approximation of this posterior distribution. A minimal example of such Bayesian inference is sketched in Fig. 4.2C, where we consider a small WTA network with only one input  $y_i$  and two network neurons  $z_k$  and  $z_j$ . For a given input instantiation  $y_i(t)$  the values  $p(Y_i = y_i(t) | Z_k = 1)$  and  $p(Y_i = y_i(t) | Z_j = 1)$  measure the likelihoods of the two competing hypotheses that neuron  $z_k$  or neuron  $z_j$  is the hidden cause of the observed input  $y_i(t)$ . These likelihood values shape the Bayesian posterior distribution (4.7) by contributing one factor to the product in eq. (4.9).

As a consequence, a network neuron  $z_k$  is particularly responsive to those input configurations  $\mathbf{y}(t)$  that are associated with high likelihood values  $p(\mathbf{Y} = \mathbf{y}(t) | Z_k = 1)$ . The likelihood distributions are determined by the means  $\mu_{ki}$ , i.e., by the number  $m_{ki}$  of active switches in the synapses that change according to the compound-synapse STDP rule. Through this mechanism, synaptic plasticity governs the emergence of prototypic patterns that the network neurons are most responsive to, and thereby turns each neuron  $z_k$  into a probabilistic expert for certain input configurations  $\mathbf{y}(t)$ . The aim of learning in the WTA network is to distribute the probabilistic experts  $z_k$  such that the likelihood of the presented input is (on average) as high as possible. This objective is an equivalent formulation of Maximum-Likelihood learning, and the *log-likelihood function*, which measures the average (logarithm of the) input likelihood, is a widely-used measure to determine how well a learning system is adapted to the presented input. Formally, the learning process can be described within the framework of generalized Expectation-Maximization (DEMPSTER et al., 1977; NESSLER et al., 2013; HABENSCHUSS et al., 2012). In Appendix C, we show that the generalized Expectation-Maximization algorithm is implemented in the WTA network via the interplay of the compound-synapse STDP rule (4.3), that adapts the synaptic weights  $W_{ki}$ , and the homeostatic intrinsic plasticity rule, that regulates the intrinsic excitability  $b_k$  of the neurons such that each neuron maintains a long-term average firing rate. The interplay of synaptic and intrinsic plasticity achieves the aim of Maximum-Likelihood learning in the WTA network in the following sense:

## 4.5 Demonstration of unsupervised learning



**Fig. 4.3: Learning of hand-written digits.** (A) Synaptic weights at time  $t = 0$  s after random initialization. Color intensities show the number  $m_{ki} \in [0, 10]$  of active switches for each connection. The indices  $k$  of the postsynaptic neurons are indicated in the top-left corners. (B) Examples from the MNIST data set. Pixel intensities of the digits were encoded as Poisson spike trains and presented to the network. Digits 0 to 4 were presented with a new digit being shown every 100 ms. (C) Weight matrix at time  $t = 200$  s. Memristive synapses start to integrate salient features of the input stream. (D) Over the course of learning the log-likelihood function (red) increases, indicating that the network continuously refines its implicit statistical model of the presented data. This refinement leads to an improved classification performance (blue) on an independent test set. (E) At the end of the learning experiment at  $t = 5000$  s, the synapses have specialized on different prototypes of the presented digits, rendering each network neuron a probabilistic expert for a certain digit. (F) Network response for 1 s at the end of learning. The presented digit is shown at the top. The input is transformed into a sparse and reliable spike code. Some digits invoke a spike response of more than one neuron. This ambiguous response encodes uncertainty in the variational posterior distribution.

*The expected synaptic weight changes  $\left\langle \frac{d}{dt} W_{ki} \right\rangle$  of the compound memristive synapses on average increase a lower bound of the log-likelihood function in a Mixture-of-Gaussians generative model during online learning until a (local) optimum is reached.*

## 4.5 Demonstration of unsupervised learning

We tested the learning capabilities of the compound memristive synapse model in a standard machine learning task for hand-written digit recognition. In a computer simulation, we set up a WTA network with  $N = 24 \times 24$  input neurons and  $K = 10$  network neurons. Each synaptic weight  $W_{ki}$  was composed of  $M = 10$  stochastic bistable switches, each contributing  $\omega = 0.1$  in its active state. The switching probabilities  $\pi_{\text{up}} = \pi_{\text{down}} = 0.001$  were set to a quite low value. This corresponds to a long integration time for gradual memory formation in order to assess the general ability of the synapse model in online learning tasks. Before training, the stochastic switches were initialized randomly as shown in Fig. 4.3A. The network was then exposed to hand-written digits 0 to 4 from the MNIST training data set (Lecun et al., 1998). Examples from the data set are shown in Fig. 4.3B. Each pixel was encoded by one input neuron  $y_i$ . Digits were presented as Poisson spike trains with firing rates depending on pixel intensities. Overall, the network was trained

## 4 Statistical learning with compound memristive synapses

in an unsupervised setup for 5000 s with a new digit being presented every 100 ms. Fig. 4.3C shows the weight matrix in an early stage of learning at  $t = 200$  s. At this stage, the synapses begin to integrate salient statistical features of the input, such as the generally low activity along the frame. Furthermore, a specialization to certain digit classes becomes apparent for some of the network neurons.

Over the course of learning, the synapses continuously improve the network's implicit generative model of the presented input. This refinement is reflected in the log-likelihood function shown in Fig. 4.3D that measures how well the probabilistic model is adapted to the input distribution. The ongoing refinement also becomes apparent by a more intuitive – and practically more relevant – measure, namely the classification performance of the network on an independent test set of handwritten digits. The classification error (blue in Fig. 4.3D; see Appendix C for details) continuously decreases as training progresses. The improved performance on an independent test set furthermore indicates that the network develops a generally well-suited representation of the input and evades the risk of over-fitting.

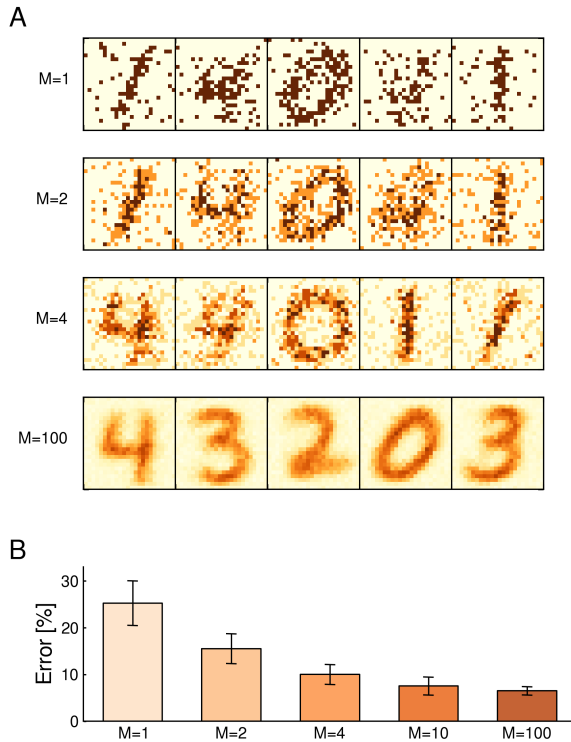
At the end of training, after  $t = 5000$  s, a set of prototypic digits has emerged in the compound memristive synapses as shown in Fig. 4.3E. The well-adapted synapse array turns each network neuron into a probabilistic expert for a certain digit class. As a consequence, the network has learned to transform the  $N$ -dimensional, noisy spike input into a sparse and reliable spike code, as shown in Fig. 4.3F. Typically, exactly one network neuron  $z_k$  fires in response to the input. But also the seemingly unclear cases, when two neurons respond simultaneously, carry meaningful information in a Bayesian interpretation: Since network spikes approximate the posterior distribution  $p(\mathbf{Z} | \mathbf{Y} = \mathbf{y}(t))$  through sampling, an ambiguous spike response encodes the level of uncertainty during probabilistic inference.

### 4.6 Influence of synaptic resolution

In the previous section, we have demonstrated that the compound memristor synapse model is able to learn statistical regularities in the input stream and enables the spiking WTA network to perform probabilistic inference in a well-defined generative model. The demonstration employed compound synapses with  $M = 10$  stochastic switches per synapse. Notably, the number of constituents  $M$  is a free parameter of the model and determines the weight resolution of the compound synapse. Increasing the synaptic resolution by recruiting more bistable switches per synapse is generally expected to improve the accuracy of the input representation, but comes at the cost of reduced integration density in a neuromorphic design. In the following, we therefore explore the opposite direction, i.e., unsupervised learning with a low weight resolution.

For estimating the influence of the weight resolution on the learning capabilities of the WTA network, we repeated the above computer simulation for different values of  $M$ , while holding the maximum weight  $W^{\max} = \omega \cdot M$ , and thus the



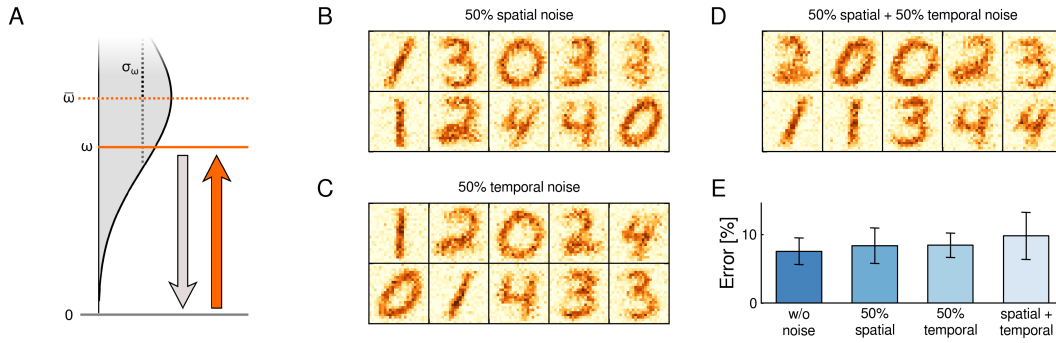


**Fig. 4.4: Influence of the synaptic resolution.** (A) Examples of learned weight matrices with different numbers of constituents  $M$  per synapse. Even binary weights with only one stochastic switch learn to maintain a coarse image of prototypic digits. In the limit of large  $M$ , compound synapses support a quasi-continuous weight spectrum. (B) Classification error for different weight resolutions after 5000 s of learning, based on 20 independently trained networks for each value of  $M$ . Errorbars: standard deviation among networks.

variance  $\sigma = 1/\sqrt{(W^{\max})}$  of the implicit generative model, fixed. Fig. 4.4A shows examples of the digits stored in the synapse array after 5000s of learning for  $M = 1, 2, 4$  and 100 stochastic switches per synapse. Even binary synapses with  $M = 1$  successfully identify noisy archetypes of the input digits. This observation is in line with previous studies on learning with binary weights (FUSI, 2002). The accuracy of the representation quickly increases with higher  $M$ -values. As an (academic) reference, we also included a simulation with  $M = 100$  switches per synapse which support a quasi-continuous state spectrum.

The resulting ability of the WTA network to recognize hand-written digits is shown in Fig. 4.4B in terms of the classification error on a test set. Each bar depicts the mean performance of 20 independently trained networks per  $M$ -value, errorbars show the standard deviation among networks. Taking the academic example with  $M = 100$  as a reference for the performance achievable by the small WTA network, the computer simulations suggest that as few as  $M = 4$  constituents per synapse may be sufficient for practical applications. While individual weights  $W_{ki}$  only store little information (ca. 2.3 bits in case of  $M = 4$ ) about the expected input in channel  $y_i$ , the partial evidence received from each of the  $N = 576$  input channels is integrated by the network in a statistically correct manner to form a sharply peaked posterior, most of the time.

## 4 Statistical learning with compound memristive synapses



**Fig. 4.5: Robustness to spatial and temporal noise.** (A) Noise was added by drawing the weight  $\omega$  of the active state from a normal distribution. Three types of noise were tested: For spatial noise, a value  $\omega \sim \mathcal{N}(\omega; \bar{\omega}, \sigma_{\omega}^2)$  was assigned to each switch and maintained throughout the experiment. For temporal noise, the value  $\omega$  was redrawn after every LTP transition. For spatial + temporal noise, the device specific spatial weight value determined the mean for redrawing  $\omega$  after LTP transitions. (B-D) Number  $m_{ki}$  of active switches after learning at  $t = 5000$  s. Shown are example networks for all three noise types. (E) Classification error after learning, based on 20 independently trained networks per noise type. Errorbars denote standard deviation among networks.

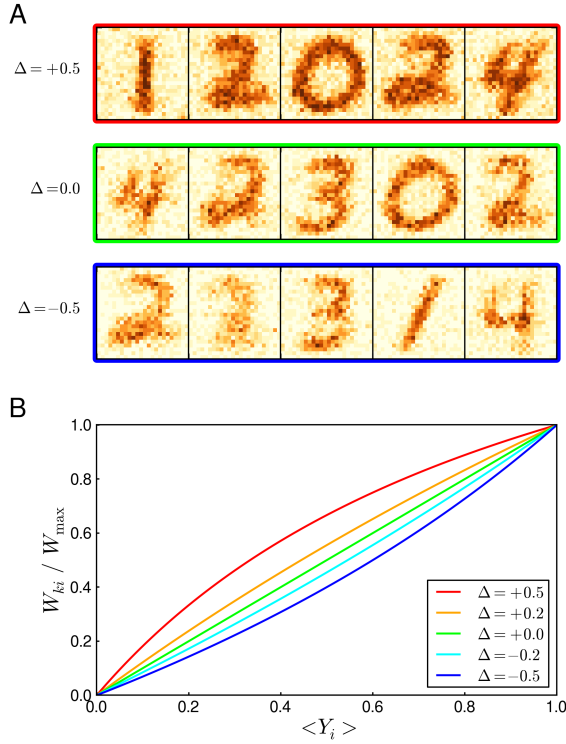
### 4.7 Robustness to device variations

We have demonstrated so far that spiking networks with compound memristive synapses can learn a faithful representation of their input when synapses consist of idealized bistable constituents. Large-scale physical implementations, however, are likely to exhibit substantial device variabilities and imperfections. Plasticity in the compound synapse model depends on two device properties that are likely to be distorted in physical implementations: the conductance of each individual constituent  $\omega$  and the switching probabilities  $\pi_{\text{up}}/\pi_{\text{down}}$ . In the following, we address the impact of distortions in these two properties on the WTA network learning capabilities, separately.

We first turned to the conductance value  $\omega$  of individual switches and investigated the robustness of learning to two fundamentally different types of noise in  $\omega$ , namely spatial noise and temporal noise:

- *Spatial noise* describes device-to-device variations and addresses peculiarities of individual memristors that remain stable over time.
- *Temporal noise* refers to trial-to-trial variations and covers device instabilities over the course of learning.

Both types of noise can be suspected to induce serious disturbances during learning: In case of spatial noise, although device-to-device variations could average out if many memristors are employed, any remaining deviations give rise to sustained systematic errors that may build up over the course of learning. In case of temporal noise, while trial-to-trial variations could average out over time, any synaptic update rests upon a disturbed instantiation of the weight matrix, i.e., on a noisy (and false) assumption. In computer simulations, we accounted for these types of noise separately by disturbing the active-state weight value  $\omega$  of the switches as



**Fig. 4.6: Robustness to unbalanced switching probabilities.** (A) Examples of learned weight matrices when the switching probabilities  $\pi_{\text{up}}$  and  $\pi_{\text{down}}$  are systematically unbalanced. The parameter  $\Delta := (\pi_{\text{up}} - \pi_{\text{down}}) / \pi_{\text{up}}$  measures the relative imbalance between LTP and LTD transitions. The top (bottom) row shows weight matrices resulting from a 50% decrease (increase) of the LTD switching probability. The balanced case is shown in the middle row for comparison. (B) Systematic imbalance ( $\Delta \neq 0$ ) leads to changed equilibrium points in the STDP rule and can be theoretically understood as a non-linear encoding of the expected input by the synaptic weights.

sketched in Fig. 4.5A. To model spatial noise, the weight  $\omega$  was randomly drawn prior to training for each stochastic switch from a normal distribution  $\mathcal{N}(\omega; \bar{\omega}, \sigma_{\omega}^2)$  with mean  $\bar{\omega}$  and standard deviation  $\sigma_{\omega}$ . To capture the effect of temporal noise, in contrast, the weight  $\omega$  was redrawn from  $\mathcal{N}(\omega; \bar{\omega}, \sigma_{\omega}^2)$  whenever the constituent switched to its active state in an LTP transition. Furthermore, we examined the combined effect of both noise types being present simultaneously. In this combined case, the mean value for temporal noise was determined by the device-specific spatially perturbed weight value of each constituent. In any case, the range of perturbed weights was truncated to  $\omega \geq 0$  to rule out negative conductances.

We repeated the experiment of Figure 4.3 under each of these noise conditions. The mean  $\bar{\omega} = 0.1$  was set to the undisturbed weight value of the previous, idealized experiment. The noise level was set to  $\sigma_{\omega} = 0.05$ , i.e., to 50% of the mean. Example cases of weight matrices after learning (shown are the  $m_{ki}$ 's from individual simulation runs) are presented in Fig. 4.5B-D for the three noise conditions “spatial”, “temporal” and “spatial+temporal”, respectively. Surprisingly, hardly any difference to the idealized setup is observable. Nevertheless, under 20 repetitions of the learning simulation the detrimental influence of noise becomes visible in the classification performance (see Fig. 4.5E) as noise appears to slightly increase the mean of the classification error. In summary, these results reveal a remarkable robustness of learning with compound memristive synapses to substantial device variability and severe temporal instability.

## 4 Statistical learning with compound memristive synapses

We next turned to the question how distorted switching probabilities  $\pi_{\text{up}}$  and  $\pi_{\text{down}}$  influence the learning dynamics in the WTA network. In the theory section, we had assumed that  $\pi_{\text{up}} = \pi_{\text{down}}$ , i.e., that the switching probabilities underlying LTP and LTD are balanced. This assumption, which could to some extent be achieved in a calibration step, yielded the elegant learning rule (4.3) and thereby facilitated the theoretical analysis. A physical implementation, however, will likely exhibit unbalanced switching probabilities in the majority of memristors. We examined the influence of unbalanced switching, in two ways. First, we applied spatial noise to the switching probabilities of individual constituents by drawing  $\pi_{\text{up}}$  and  $\pi_{\text{down}}$  (separately) from normal distributions with 50% noise level (truncated to  $0 \leq \pi_{\text{up}}, \pi_{\text{down}} \leq 1$ ). Thus, about half of the stochastic switches were more responsive to LTP pulses, the other half more to LTD pulses; even more, some of the constituents only showed switching in one direction, or were completely unresponsive. Nevertheless, synapses developed a faithful representation of prototypic digits in a repetition of the experiment in Figure 4.3 (data not shown). Also the classification performance was only mildly impaired (classification error:  $8.7\% \pm 2.7\%$  based on 20 networks) compared to ideal, noise-free synapses (class. error:  $7.5\% \pm 1.9\%$ ).

In a second step, we pursued a slightly different – more principled – approach that permits a theoretical interpretation of how the altered synapse dynamics give rise to a different encoding of the expected input by the synaptic weights. Instead of drawing random parameters for each constituent, we systematically chose the LTD switching probability  $\pi_{\text{down}}$  larger (or smaller) than the LTP probability  $\pi_{\text{up}}$  throughout the synapse array. This systematic imbalance displays a worst-case scenario during learning since all synapses either favor (or suppress) LTD over LTP. We denote the relative imbalance between  $\pi_{\text{up}}$  and  $\pi_{\text{down}}$  by  $\Delta := (\pi_{\text{up}} - \pi_{\text{down}}) / \pi_{\text{up}}$ . For instance,  $\Delta = -0.5$  means that the probability for LTD transitions is 50% higher than for LTP transitions. Fig. 4.6A shows examples of weight matrices after 5000s of learning in a repetition of the experiment in Figure 4.3. In the top row,  $\Delta = +0.5$ , LTP transitions are favored over LTD transitions, resulting in generally stronger weights in comparison with balanced STDP ( $\Delta = 0.0$ , middle row). Conversely, in the bottom row,  $\Delta = -0.5$ , the systematic strengthening of LTD leads to weaker weight patterns. Nevertheless, synaptic weight values converged to a dynamic equilibrium in either case since the STDP rule preserves its general stabilizing weight dependence. As can be expected from the prototypic digits that emerged in the weight matrices, the classification performance of the WTA networks was not considerably impaired by the unbalanced switching (classification errors estimated from 20 networks:  $6.7\% \pm 0.8\%$  for  $\Delta = +0.5$ ;  $7.5\% \pm 1.9\%$  for  $\Delta = 0$ ;  $11.8\% \pm 4.0\%$  for  $\Delta = -0.5$ ). Indeed, positive  $\Delta$ -values even performed slightly (but not significantly) better than balanced switching. A conceptual understanding of the altered learning dynamics can be obtained from the equilibrium points of the unbalanced STDP rule. The short calculation, that had led to eq. (4.8) for the balanced case, can be repeated for unbalanced switching probabilities. Fig. 4.6B shows the resulting encoding of the expected input value  $\langle Y_i \rangle_{p(Y_i | Z_k=1)}$  by the synaptic weight  $W_{ki}$  for different  $\Delta$ -values. The example digits shown in panel A correspond to the red, green and blue graph in panel B, respectively. This analysis illustrates how

unbalanced switching probabilities give rise to a non-linear encoding of the input in the WTA network. In particular, it can be seen how the same expected input value  $\langle Y_i \rangle_{p(Y_i|Z_k=1)}$  leads to stronger weights for  $\Delta > 0$ , and weaker weights for  $\Delta < 0$ .

## 4.8 Discussion

We have proposed the compound memristive synapse model for neuromorphic architectures that employs multiple memristors in parallel to form a plastic synapse. A fundamental property of the synapse model is that individual memristors exhibit stochastic switching between two stable memristive states rather than obeying a deterministic update rule. Yet, the expected weight change of the compound memristive synapse, as it arises from the stochastic switching of its constituents, yielded an STDP-type plasticity rule with a stabilizing, linear weight dependence. We examined the computational capabilities of the compound-synapse STDP rule in WTA networks, a common circuit motif in cortical and neuromorphic architectures, by analyzing the network and synapse dynamics from the perspective of probability theory and machine learning. The comprehensive mathematical treatment revealed that compound memristive synapses enable a spiking network to perform Bayesian inference in and autonomous statistical optimization of a Mixture-of-Gaussians generative model via generalized Expectation-Maximization. Accompanying computer simulations demonstrated the practical capability of the synapse model to perform unsupervised classification tasks and, furthermore, revealed a remarkable robustness of the compound synapses to substantial device variations and imperfections.

### Comprehensive learning theory of memristive plasticity

Our work contributes a theoretical foundation for memristive learning in neural networks to the endeavor to employ memristors as plastic synapses in self-calibrating systems. SNIDER (2008); QUERLIOZ et al. (2011) and SERRANO-GOTARREDONA et al. (2013) have investigated how different pre- and postsynaptic waveforms can shape a memristive STDP learning window. S. H. Jo et al. (2010) and MAYR et al. (2012) have demonstrated STDP-type plasticity in Ag/Si and BiFeO<sub>3</sub> memristors. Yu et al. (2013) reported stochastic switching between stable states in oxide-based memristive synapses. GABA et al. (2013) studied the parameter dependence of switching probabilities in metal filament based memristors, indicating a renewal process that is independent of the overall network firing rate. A strategy for integrating nanoscale memristive synapses into a hybrid memristor-CMOS network architecture was proposed by INDIVERI et al. (2013). The beneficial contribution of stochasticity to learning with CMOS synapse circuits was explored by CHICCA et al. (2014). Here we have established a firm link between the emergent synapse configurations observed in such architectures (see e.g., QUERLIOZ et al. (2011))

## 4 Statistical learning with compound memristive synapses

and a rigorous mathematical description of memristive learning on the system level using machine learning theory. Our findings on memristive learning from a Bayesian perspective build upon a series of theoretical contributions on synaptic learning in spiking neural networks: NESSLER et al. (2009a) and NESSLER et al. (2013) identified a general link between STDP-type synaptic plasticity and statistical model optimization for probabilistic inference in WTA networks. HABENSCHUSS et al. (2012) extended this work to incorporate also homeostatic intrinsic plasticity, thereby overcoming several limiting assumptions on the input presentation. Finally, HABENSCHUSS et al. (2013b) investigated how the learning framework can be generalized to support a broad class of probability distributions. We expect that utilizing machine learning theory for describing the effects of specific memristor synapse models can significantly promote our understanding of memristive learning and its computational prospects.

### Heading for a full hardware integration

Plasticity in the compound memristor synapse model relies on stochastic transitions between two stable states. Such bistable devices – or more generally, devices with a clearly discrete state spectrum – were reported to exhibit a high degree of uniformity (D. LEE et al., 2006; FANG et al., 2011) and temporal stability (INDIVERI et al., 2013). Notably, the theoretical approach we have pursued in this work could likely be extended to cover memristors with more than two stable states and to support more complex input and plasticity mechanisms. For instance, the Gaussian likelihood distributions  $p(\mathbf{Y} | \mathbf{Z})$  identified in the present study, in principle support inference over arbitrary real-valued input states  $\mathbf{y}(t)$ . Such states could arise if the input is presented in the form of exponentially decaying or additive postsynaptic potentials. Such more complex input types could afford more versatile STDP pulsing schemes, and the resulting memristor plasticity rules could likely be incorporated in an adapted model of statistical learning. The reason we restricted the input to binary values  $y_i(t)$  is found in the STDP pulsing scheme that employs binary presynaptic waveforms. In this case, the theoretically derived generative model reveals how active and inactive inputs contribute to the network's spike response by means of a Gaussian likelihood distribution  $p(Y_i = y_i(t) | Z_k = 1)$  that is sampled only at  $y_i(t) = 0$  and  $y_i(t) = 1$ .

In this article, we have employed a simple model for stochastic switching in memristive devices where switching occurs with probabilities  $\pi_{\text{up}}$ ,  $\pi_{\text{down}}$  which depend on the applied voltage difference across the memristor terminals. This phenomenological model captures the most salient aspects of switching in real memristive materials (S. H. JO et al., 2009b; GABA et al., 2013) and was used as an abstraction of memristive switching in a recent experimental study (SURI et al., 2013). In future research, it will be important to evaluate the effectiveness of this model either with physical memristors or in simulations based on detailed memristor models. The authors of (SURI et al., 2013) raised the concern that the precise

Parameter name	Hardware	Theory
Learning rate	$\pi_{\text{up}} \cdot W^{\text{max}}$	$\eta_W$
Max. weight $W^{\text{max}}$	$\omega \cdot M$	$1/\sigma^2$
Likelihood mean $\mu_{ki}$	$m_{ki}/M$	$\sigma^2 \cdot W_{ki}$
Synaptic resolution	$\omega$	$1/(M \cdot \sigma^2)$

**Table 4.1: Translation of synapse parameters between the hardware and theory domain.**

switching probabilities of individual devices are potentially hard to control in large-scale systems. In this regard, our simulation results indicate that learning with compound memristive synapses tolerates significant noise levels in the switching probabilities. We expect the origin for the observed robustness to be twofold: Firstly, imbalances in  $\pi_{\text{up}}$ ,  $\pi_{\text{down}}$  between different constituents are expected to partly average out in compound synapses according to the central limit theorem; secondly, the stabilizing weight dependence of compound-synapse STDP ensures that even unbalanced switching leads to stable weight configurations, albeit with slightly shifted convergence points.

Another potential issue for learning with compound memristive synapses is the absolute value of the switching probability. The product  $\pi_{\text{up}} \cdot W^{\text{max}}$  can be linked to a learning rate in the theory domain (see Table 4.1) which controls how many samples from the input history are integrated into the implicit generative model during online learning. A slow and gradual memory formation, which is desirable for developing a representation of large and complex input data sets, relies on small learning rates, i.e., on small switching probabilities. It has to be seen if memristive materials that exhibit stochastic switching provide sufficiently small switching probabilities. A possible remedy in a hardware integration could be to multiplex the back-propagating signals from network neurons such that only a random subset of the memristors is notified of a network spike at a time (FUSI, 2002).

Regarding the physical model neurons of a hybrid memristor-CMOS architecture, two types of currents occur in the WTA network. The input integration via forward-synapses is spike based and could be realized with standard leaky integrators (see Appendix C). Lateral inhibition, in contrast, depends on the neuronal membrane potentials, and the involved inhibitory circuits should ideally transmit potentials instead of spikes. Alternatively, the effect of lateral inhibition could be approximated in a spike-based manner by populations of inhibitory neurons. Independent of the specific implementation of lateral inhibition, the resulting potential  $u_k - u_{\text{inh}}$  controls the stochastic response of the WTA neurons that could either be implemented genuinely with a stochastic firing mechanism or be emulated with integrate-and-fire neurons (PETROVICI et al., 2013).

### Inherently stochastic nature of compound-synapse STDP

The spiking WTA network architecture with compound memristive synapses exploits stochasticity in various ways, in that the stochastic firing of network neurons in response to a transient input trajectory triggers stochastic STDP updates in the synaptic weights. From a learning perspective, the high degree of stochasticity contributes to the network's ongoing exploration for potential improvements in the parameter space. While the learning theory only guarantees convergence to a local optimum of the weight configuration, the stochastic nature of the ongoing exploration enables the network to evade small local optima in the parameter landscape, and thereby improves the robustness of learning (compared to traditional batch Expectation-Maximization).

For the derivation of the learning algorithm, we have focused on the weight-dependent STDP rule (4.3) which describes the expected temporal weight change  $\left\langle \frac{d}{dt} W_{ki} \right\rangle$  of the compound synapse. The stochasticity of memristive switching, however, gives rise to a probability distribution over the weights, as well. Indeed, in equilibrium we expect that the number of active constituents  $m_{ki}$  follows a binomial-type weight distribution. This points to a potential knob for adjusting the amount of stochasticity used during online learning: when many memristors are recruited per synapse, i.e., for large  $M$ , we expect a reduced variance in the weight distribution.

Besides the level of stochasticity, the parameter  $M$  also controls the weight resolution of the compound synapse. In Figure 4.4, we have investigated the impact of the weight resolution on the learning capabilities of the WTA network. Notably, we observed that even with  $M = 4$ , i.e., with synapses that feature only 5 weight levels, the network performed reasonably well in the hand-written digit recognition task. The observation that even a low synaptic weight resolution can yield a satisfactory performance has important practical implications for nanoscale circuit designs, where integration density and power consumption impose crucial constraints, since the area allocated by the synapse array grows linearly in the synaptic size  $M$ . For instance, SRAM cells can be fabricated with a cell size of  $0.127 \mu\text{m}^2$  (S.-Y. Wu et al., 2009), corresponding to a memory density of  $\lesssim 1 \text{ Gb}/\text{cm}^2$ . Importantly, this estimate does not include any additional plasticity circuits for implementing STDP or similar plasticity mechanisms. Functional Ag/Si memristive crossbars with  $2 \text{ Gb}/\text{cm}^2$  memory density were demonstrated by (S. H. Jo et al., 2009a), with densities up to  $10 \text{ Gb}/\text{cm}^2$  being envisioned (S. H. Jo et al., 2009a; S. H. Jo et al., 2009b). In the long term, memristive crossbars are expected to combine the advantages of SRAM and Flash memory regarding energy efficiency, non-volatility and integration density (J. J. YANG et al., 2013).



## Generalization to other materials and future research

In recent years, a plethora of (in a broader sense) memristive materials has been discovered, and the characterization and refinement of their switching dynamics is evolving rapidly. At least four types of stochastically switching memristive devices can be distinguished: Switching in (1) anion-based (e.g.,  $\text{HfO}_x/\text{TiO}_x$  (Yu et al., 2013)) and (2) cation-based (e.g.,  $\text{Ag}/\text{GeS}_2$  (SURI et al., 2013)) devices mainly originates from conductive filament formation (J. J. YANG et al., 2013). In contrast, (3) single-electron latching switches (e.g., CMOS/MOLEcular (CMOL) CrossNets (J. H. LEE et al., 2006)) rely on electronic tunneling effects and, thus, their stochastic switching dynamics arise directly from the underlying physical process. Similarly, (4) magnetoresistive devices (e.g., spin-transfer torque magnetic memory (STT-MRAM) (VINCENT et al., 2014)) can inherit stochastic switching dynamics from fundamental physical properties. Some manufacturing processes related to these ideas (like conductive-bridging RAM and STT-MRAM) reached already an early industrial stage, others are still primarily subject of academic research. While the microscopic origins of plasticity in these memristor types are fundamentally different, they all share stochastic, persistent switching between bistable memory states on a phenomenological level. We therefore believe that the compound memristive synapse model displays a promising concept for future work in diverse research fields.

Independent of the underlying switching mechanism, any nanoscale synaptic crossbar will likely exhibit imperfections and imbalances due to process variations. Here, we have investigated spatial and temporal noise in the weight values as well as deviations in the switching probabilities under unbiased, uniform conditions. However, physical implementations can be expected to also suffer from more systematic imperfections, such as structural imbalances (e.g., one corner of the array being more reactive) or crosstalk between neighboring devices. While our computer simulations indicate a remarkable general robustness against device variations of various types, additional research is required to estimate the influence of such systematic, and potentially coupled, deviations.

## Conclusion

In this article, we have introduced the compound memristive synapse model together with the compound-synapse STDP rule for weight adaptation. Compound-synapse STDP, a stabilizing weight-dependent plasticity rule, naturally emerges under a standard STDP pulsing scheme. In addition, by employing memristors with bistable memristive states, compound memristive synapses may circumvent practical challenges in the design of reliable nanoscale memristive materials. Both, our theoretical analysis and our computer simulations confirmed that compound-synapse STDP endows networks of spiking neurons with powerful learning capabilities. Hence, the compound memristive synapse model may provide a synaptic design principle for future neuromorphic hardware architectures.



# Chapter 5

## Distributed Bayesian computation and learning in spiking neural sheets

### Contents

5.1	Introduction . . . . .	76
5.2	Probabilistic inference in spatially extended spiking networks . .	81
5.3	Emergence of local experts through synaptic plasticity . . . . .	89
5.4	Plastic recurrent synapses integrate structural knowledge . . . . .	95
5.5	Emergence of excitatory subnetworks in neural sheets . . . . .	102
5.6	Discussion . . . . .	104

**Abstract.** During the last decade, Bayesian probability theory has emerged as a framework in cognitive science and neuroscience for describing perception, reasoning and learning of mammals. However, our understanding of how probabilistic computations could be organized in the brain, and how the observed connectivity structure of cortical microcircuits supports these calculations, is rudimentary at best. In this study, we investigate statistical inference and self-organized learning in a spatially extended spiking network model, that accommodates both local competitive and large-scale associative aspects of neural information processing, under a unified Bayesian account. Specifically, we show how the spiking dynamics of a recurrent network with lateral excitation and local inhibition in response to distributed spiking input, can be understood as sampling from a variational posterior distribution of a well-defined implicit probabilistic model. This interpretation further permits a rigorous analytical treatment of experience-dependent plasticity on the network level. Using machine learning theory, we derive update rules for neuron and synapse parameters which equate with Hebbian synaptic and homeostatic intrinsic plasticity rules in a neural implementation. In computer simulations, we demonstrate that the interplay of these plasticity rules leads to the emergence of probabilistic local experts that form distributed assemblies of similarly tuned cells communicating through lateral excitatory connections. The resulting sparse distributed spike code of a well-adapted network carries compressed information on salient input features combined with prior experience on correlations among them. Our theory predicts that the emergence of such efficient representations benefits from network architectures in which the range of local inhibition matches the spatial extent of pyramidal cells that share common afferent input.

**Acknowledgments and author contributions.** This chapter is based on the manuscript

JOHANNES BILL, LARS BUESING, STEFAN HABENSCHUSS, BERNHARD NESSLER, WOLFGANG MAASS, ROBERT LEGENSTEIN (2015). “Distributed Bayesian computation and self-organized learning in sheets of spiking neurons with local lateral inhibition.” In: *(submitted for publication in Feb. 2015)*.

To this study, I contributed as first author. The study was conceived by JB, LB and WM, with the theory being developed by JB, LB, SH and BN. The computer experiments were designed by JB, SH and RL, and were conducted by JB. The manuscript was written by JB, SH and RL. The authors thank David Kappel for his helpful comments on the manuscript.

### 5.1 Introduction

Humans and animals perceive their environment through a stream of data from various high-dimensional sensory modalities. Successful behavior requires that the individual dimensions of this data stream are aligned with one another and integrated into a compact representation that promotes rapid decision making and generalization. Typically, the available sensory information on which decisions have to be based is noisy, unreliable and incomplete. Hence, it is essential that such representations respect the statistical nature of sensory data and that knowledge about statistical and causal relations among events in the external world are taken into account when a representation is generated. In recent years, Bayesian inference has been identified in cognitive science as a powerful normative framework for the description of cognitive processes in face of uncertainty in humans (TENENBAUM et al., 2011; TÉGLÁS et al., 2011) and animals (ANGELAKI et al., 2009). The Bayesian framework has also been successfully employed for a formal description of learning, for instance in perceptual (GOODMAN et al., 2008; ORBAN et al., 2008) and sensorimotor (KÖRDING and WOLPERT, 2004; BERNIKER et al., 2010) learning tasks.

In the Bayesian framework, quantities of interest are formally treated as random variables (RVs), and beliefs about their current values are formalized as probability distributions over these RVs (BISHOP, 2006). Typically, one distinguishes between observations  $y_i$ ,  $i = 1, \dots, N$ , representing directly observable variables, and latent variables  $z_k$ ,  $k = 1, \dots, K$  which cannot be observed directly. Latent variables represent abstract features and concepts that allow to structure and conceive the given input. As an everyday example, the high dimensional vector  $\mathbf{y} = (y_1, \dots, y_N)^\top$  of input RVs could summarize the entire sensory stream from the visual, auditory, and vestibular system while driving by bike in a city during rush hour. Latent variables  $\mathbf{z} = (z_1, \dots, z_K)^\top$  could represent streets, cars, and pedestrians, or even more abstract features, such as estimated velocities, potential threats or anticipated actions. Thus, the latent variables  $\mathbf{z}$  form a compact representation of relevant aspects of the scene

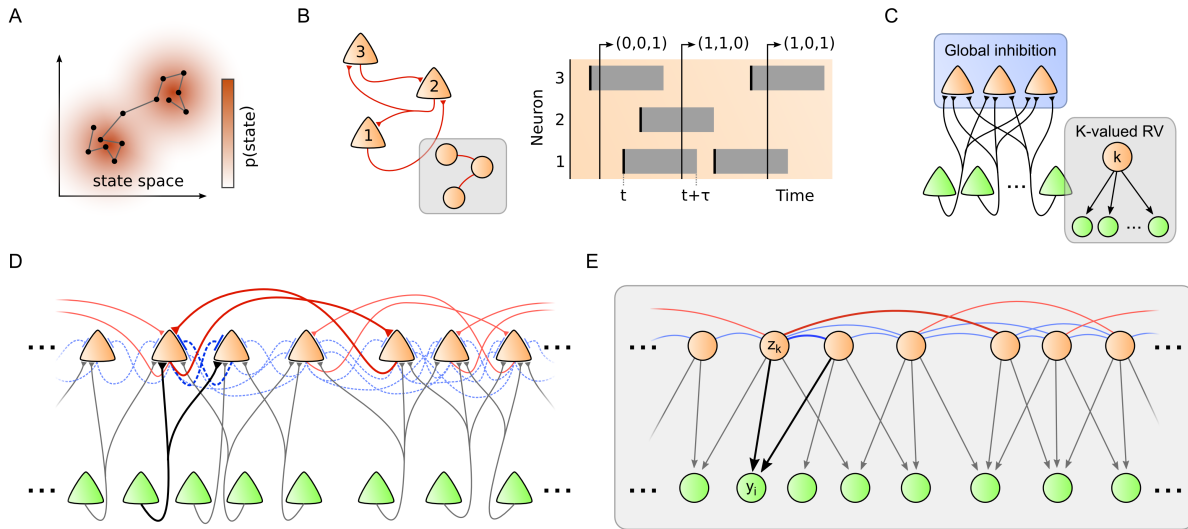
which is typically more stable and informative than individual local observables  $y_i$ . In this manner, local observations  $y_i$ , which are often unreliable, e.g., due to noise or an occluding obstacle, can be integrated into a consistent global interpretation.

In a Bayesian framework, all knowledge about statistical dependencies between RVs is represented in terms of a probability distribution  $p(z, \mathbf{y} | \boldsymbol{\theta})$ , where the parameter vector  $\boldsymbol{\theta}$  shapes the statistical model of the environment. The aim of Bayesian inference is to infer a belief over the possible states of the latent variables  $z$  for the given observations  $\mathbf{y}$ . More precisely, Bayesian inference is the calculation of the posterior distribution  $p(z | \mathbf{y}, \boldsymbol{\theta})$  over the latent variables  $z$  given the input  $\mathbf{y}$ . In ambiguous situations (e.g. when estimating the speed of a car), the posterior  $p(z | \mathbf{y}, \boldsymbol{\theta})$  provides not only the most likely interpretation of the input, but rather the probabilities of the most likely and all alternative interpretations. Finally,  $p(z | \mathbf{y}, \boldsymbol{\theta})$  can also capture correlations among hidden variables (e.g. different drivers will maintain roughly the same speed level).

During recent years, several modeling approaches explored how Bayesian computations can be performed by spiking networks, and how the involved probability distributions can be represented by the neuronal spike response (R. S. ZEMEL et al., 1998; R. P. RAO, 2004; DENEVE, 2005; MA et al., 2006; R. P. RAO, 2007; STEIMER et al., 2009; BUESING et al., 2011; HABENSCHUSS et al., 2013a; PETROVICI et al., 2013; KAPPEL et al., 2014; SAVIN and DENEVE, 2014; LEGENSTEIN and MAASS, 2014). Two major lines of research regarding the neural representation of probability distributions can be distinguished: distributional (probabilistic) population codes (R. S. ZEMEL et al., 1998; MA et al., 2006) and sample-based representations (HOYER and HYVÄRINEN, 2003; FISER et al., 2010). In sample-based representations, spike responses are interpreted as samples  $z$  from the posterior distribution as illustrated in Fig. 5.1A: Through its inherent dynamics, the network trajectory visits states  $z(t)$  proportionally to  $p(z | \mathbf{y}, \boldsymbol{\theta})$ . Thus, neural activity is hypothesized to encode distributions in the sequence of network states. Sample-based representations are particularly appealing for theoretical considerations, since they are highly versatile and naturally support the representation of complex, potentially multimodal distributions over large numbers of variables (FISER et al., 2010; HABENSCHUSS et al., 2013a).

Recently, a generic spiking network model that samples from a known probability distribution was proposed by Buesing et al. (BUESING et al., 2011). The underlying theory describes the dynamics of networks of idealized stochastic spiking neurons as a Markov chain Monte Carlo (MCMC) sampling algorithm. In this model, each binary RV  $z_k \in \{0, 1\}$  is associated with one spiking neuron in the network, and spikes of these neurons are interpreted as realizations of the corresponding RVs, see Fig. 5.1B. After a spike of the  $k$ -th neuron at time  $t^s$ , the associated RV  $z_k$  turns active, i.e.  $z_k = 1$ , for a fixed duration  $\tau$ , typically on the order of 10 milliseconds. At time  $t^s + \tau$ , the RV switches back to the inactive ground state  $z_k = 0$ . This interpretation of neuronal spike patterns as realizations of RVs defines the vector  $\mathbf{z} = (z_1, \dots, z_K)$  of all associated RVs at any time  $t$ :

## 5 Distributed Bayesian computation and learning in spiking neural sheets



**Fig. 5.1: Neural sheet model with local inhibition for distributed Bayesian inference and self-organized learning.** (A) The sampling hypothesis proposes that probability distributions are represented in the brain such that the time the network spends in state  $z$  is proportional to the probability  $p(z)$ . (B) In (BUESING et al., 2011) it was shown that recurrent networks of stochastic spiking neurons can implement Markov chain Monte Carlo sampling in a well-defined graphical model (inset). Each neuron is identified with a binary random variable (RV). The state of the RV at time  $t$  encodes whether the neuron has fired shortly before (right). (C) In (NESSLER et al., 2013) it was shown that a local population of neurons (orange), organized in a Winner-Take-All (WTA) architecture, can learn an implicit probabilistic model of spiking input (green) through STDP-type plasticity. In (NESSLER et al., 2013), competition between the neurons was established via a global inhibitory current. Inset: Corresponding graphical mixture model. (D) We propose a spatially structured neural sheet model with lateral inhibition and recurrent excitation for distributed Bayesian computation and self-organized learning. The network model unites the benefits of (BUESING et al., 2011) and (NESSLER et al., 2013). Strong inhibitory connections (dashed blue) between nearby network neurons establish local competition. Sparse recurrent excitatory synapses (red) connect more distant neurons. In addition, each network neuron integrates spiking input from a local subset of input neurons (green). (E) Graphical model of the neural sheet in D. Nearby binary network RVs  $z_k$  (orange nodes) maintain competitive links (blue) while more distant variables can maintain associative links (red). Bottom-up input synapses in D give rise to generative downward arrows to the input RVs  $y_i$  (green nodes).

$$z_k(t) = 1 \Leftrightarrow \text{Neuron } k \text{ fired in } (t - \tau, t] . \quad (5.1)$$

Based on this link between random variables and neuronal spike responses, the authors of (BUESING et al., 2011) identified a sufficient condition for a population of stochastic spiking neurons to sample from a well-defined probability distribution  $p(z | \theta)$ , i.e., the relative occurrences of states  $z(t)$  visited by the network trajectory are distributed according to  $p(z | \theta)$  in the limit  $t \rightarrow \infty$ . This implementation of MCMC sampling in networks of spiking neurons was termed *neural sampling*. The inset of Fig. 5.1B depicts the graphical model of the distribution  $p(z | \theta)$  which the small neural network in the figure samples from. Notably, the links in the graphical model mirror the structure of synaptic network connections.

The neural sampling theory explains how networks of spiking neurons can sample from a given distribution  $p(z | \theta)$  over latent variables  $z$  for given parameters  $\theta$ . It

does, however, not cover the question how these latent variables enter the network in the first place, i.e., how observed input  $\mathbf{y}$  can be integrated and represented through latent variables  $\mathbf{z}$  and how this representation can be learned from the statistics of observables  $\mathbf{y}$ . This question is addressed in the current article. In particular, we exhibit a network architecture for sheets of stochastically spiking neurons for which

- the spike response can be understood as neural sampling from the Bayesian posterior distribution of a well-defined probabilistic model,
- local synaptic plasticity rules can be derived for self-organized model optimization of the parameters  $\theta$  using machine learning theory, and
- emerging recurrent connections store correlations between latent variables  $\mathbf{z}$  and help to maintain coherent network states for resolving ambiguous input.

For the self-organized adaptation of network parameters  $\theta$  based on the statistics of observations  $\mathbf{y}$ , we adopt a “generative perspective” which is often used in Bayesian modeling (BISHOP, 2006; R. P. RAO and BALLARD, 1999; LOCHMANN and DENEVE, 2011; NESSLER et al., 2013). In the generative perspective, we interpret the latent variables  $\mathbf{z}$  as so-called hidden causes of the inputs  $\mathbf{y}$ . This view permits to identify a conditional distribution  $p(\mathbf{y} | \mathbf{z}, \theta)$  (called “likelihood”) which describes the distribution over inputs if the network would (hypothetically) generate its own spiking inputs based on its current network state  $\mathbf{z}$ . The conditional distribution  $p(\mathbf{y} | \mathbf{z}, \theta)$  describes how likely an observation  $\mathbf{y}$  is under the assumption that it was generated by the hidden cause  $\mathbf{z}$ . In a complementary, reversed perspective, by holding the observation  $\mathbf{y}$  fixed and viewing  $p(\mathbf{y} | \mathbf{z}, \theta)$  as a function of  $\mathbf{z}$ , this allows to assess how well different  $\mathbf{z}$ -configurations could serve as explanations for the given observation. This reversed perspective is formalized in Bayes rule

$$p(\mathbf{z} | \mathbf{y}, \theta) = \frac{p(\mathbf{y} | \mathbf{z}, \theta) p(\mathbf{z} | \theta)}{p(\mathbf{y} | \theta)} . \quad (5.2)$$

Bayes rule (5.2) tells us how to infer the statistically optimal posterior distribution  $p(\mathbf{z} | \mathbf{y}, \theta)$  of network responses  $\mathbf{z}$  to any stimulus  $\mathbf{y}$ . For this inference, two components are essential. The *prior*  $p(\mathbf{z} | \theta)$  encodes that some combinations of latent variables are generally more likely than others (e.g. driving speeds of cars are highly correlated during rush hour). The *likelihood*  $p(\mathbf{y} | \mathbf{z}, \theta)$  formalizes the constraint that the values of latent variables  $\mathbf{z}$  should be such that the current observations  $\mathbf{y}$  are probable under the generative probabilistic model. Note that the shape of both distributions is determined by the network parameters  $\theta$ . The denominator  $p(\mathbf{y} | \theta)$  can often be ignored since it just provides a normalizing factor for the posterior. A key benefit of the generative perspective is that it provides a theoretically well-founded approach to self-organized (unsupervised) learning. This is because inference of hidden causes works best when the parameters  $\theta$  are tuned such that the hypothetical generative distribution of the network  $p(\mathbf{y} | \theta)$  matches the true distribution of observables  $p^*(\mathbf{y})$ . The process of minimizing the

mismatch between the probabilistic model and the input statistics is known as “maximum likelihood learning”.

Adopting the generative perspective, Nessler et al. (NESSLER et al., 2013) recently developed a model for inference and learning in local populations of stochastic spiking neurons with lateral inhibition, see Fig. 5.1C. Network neurons in this model receive all-to-all connections from a set of input neurons, and each network neuron maintains a set of afferent synaptic weights that render it an expert for detecting certain input patterns. A global inhibitory current enforces competition among the network neurons. Nessler et al. showed that the spiking activity of such a network with  $K$  neurons in response to input can be understood as Bayesian inference in an implicit probabilistic mixture model with  $K$  hidden causes (see inset of Fig. 5.1C). Furthermore, it was shown that maximum likelihood learning in this model gives rise to synaptic update rules that appear compatible with experimental data on spike-timing dependent plasticity (STDP) (NESSLER et al., 2013; HABENSCHUSS et al., 2013b).

In this article, we combine the benefits of maximum likelihood learning in networks with lateral inhibition with the general theory of neural sampling. We consider a spatially structured network architecture (see Fig. 5.1D) where network neurons represent latent variables  $z$ . The spatially extended architecture generalizes the network motif considered in (NESSLER et al., 2013) in that network neurons inhibit each other locally through lateral inhibition (DOUGLAS and MARTIN, 2004) and, in addition, may form sparse excitatory connections beyond the range of lateral inhibition (BOSKING et al., 1997; STEPANYANTS et al., 2009). Afferent connections from input neurons, that represent observables  $y$ , branch and synapse locally in the sheet of network neurons. Lateral inhibition structures the network in local WTA-like subcircuits similar to (NESSLER et al., 2013) such that network neurons that receive input from overlapping sets of input neurons are subject to strong lateral inhibition. As we will see, this constraint ensures that theoretically correct inference and learning can be implemented through simple local neural operations. In contrast to (NESSLER et al., 2013) however, which considered only a single WTA circuit motif, each network neuron in the architecture proposed here can participate in several WTA-subcircuits, and multiple network neurons with disjoint input can be recruited in parallel to cooperatively explain the spatially distributed input. Building on the neural sampling theory (BUESING et al., 2011), we show that the response of this network architecture to spiking input can be understood as neural sampling-based Bayesian inference in the structured graphical model shown in Fig. 5.1E. The graphical model has two main components: recurrent links between network nodes, and generative input links pointing from network nodes to input nodes. The recurrent links encode statistical correlations between latent variables  $z$  by shaping the prior distribution  $p(z | \theta)$  in Bayes rule (5.2). The generative input links encode the likelihood model  $p(y | z, \theta)$ . We show that both these components of the probabilistic model can be optimized concurrently through local synaptic plasticity rules in this network architecture. In particular, we derive iterative update rules for maximum likelihood learning which give rise to Hebbian-type synaptic



and homeostatic intrinsic plasticity rules in the neural network. The joint application of these rules can be understood as Stochastic Online Expectation-Maximization (SATO, 1999), a powerful machine learning algorithm for unsupervised model optimization. While a theoretically optimal STDP-type plasticity rule can be derived for afferent connections that define the likelihood model  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})$ , an approximate solution is proposed for recurrent connections that define the prior distribution  $p(\mathbf{z} | \boldsymbol{\theta})$  over latent variables. In computer simulations, we verify that the spiking network can calculate and represent the theoretically correct Bayesian posterior distribution with high accuracy. We demonstrate how synaptic plasticity shapes the network response to extract and convey the most salient features of the input in a sparse distributed spike code, and how recurrent connections capture correlations between latent variables  $\mathbf{z}$  to maintain a coherent network-wide interpretation. These simulations also reveal how the sampling network calculates and represents uncertainty in case of ambiguous or uninformative input. When the presented input appears consistent with multiple possible (but mutually incompatible) explanations, the network response encodes the associated multimodal posterior distribution  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  by switching iteratively between different coherent network states. Finally, our theoretical analysis points to an integral role of lateral inhibition during learning: it is the local inhibition network motif that gives rise to local synaptic plasticity rules and that facilitates the emergence of probabilistic local experts. The resulting well-adapted network transforms high-dimensional spiking input streams into an efficient sparse code.

The remainder of the manuscript is structured as follows. We first introduce the probabilistic model and show how the resulting posterior distribution  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  can be calculated and represented by a spiking neural network with local afferent connections, lateral inhibition and sparse recurrent excitation through neural sampling. We then investigate synaptic learning of afferent connections and recurrent connections in separate subsections. Finally, we apply the complete spiking network architecture to a two-dimensional model of neural tissue in which we observe the emergence of excitatory subnetworks through the interplay of afferent and recurrent synaptic plasticity.

## 5.2 Probabilistic inference in spatially extended spiking networks

In this section, we derive how the spike response of the neural sheet model with the architecture in Fig. 5.1D can be understood as an ongoing sampling process from a Bayesian posterior distribution  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  that arises from the graphical model in Fig. 5.1E.

## Spiking neural network model

The network architecture in Fig. 5.1D comprises network neurons and input neurons. The spiking activity of input neurons is fed externally into the network. For the network neurons, we employ the stochastic spike response neuron model from Buesing et al. (BUESING et al., 2011) that describes the state of each network neuron by a binary variable  $z_k(t) \in \{0, 1\}$  according to eq. (5.1): After a spike of the  $k$ -th network neuron,  $z_k(t)$  turns active for duration  $\tau$ . After that period, the variable switches back to the inactive state  $z_k(t) = 0$ . Similarly, the state of the  $i$ -th input neuron is described by a binary variable  $y_i(t) \in \{0, 1\}$  with  $y_i(t) = 1$  for duration  $\tau$  after a spike of the  $i$ -th input neuron. The membrane potential of the  $k$ -th network neuron is given by

$$u_k(t) = b_k + \sum_{j=1}^K W_{kj}^{\text{exc}} z_j(t) + \sum_{j=1}^K W_{kj}^{\text{inh}} z_j(t) + \sum_{i=1}^N V_{ki} y_i(t) . \quad (5.3)$$

Here,  $b_k$  denotes a neuron's intrinsic excitability and captures, for instance, the influence of the voltage gap between resting and threshold potential in more detailed neuron models (JOLIVET et al., 2006; PETROVICI et al., 2013).  $W_{kj}^{\text{exc}}$  and  $W_{kj}^{\text{inh}}$  are recurrent synaptic weights between network neurons to instantiate sparse recurrent excitation (BOSKING et al., 1997) and local lateral inhibition (MARKRAM et al., 2004). For simplicity, we model the effect of inhibition as direct negative connections  $W_{kj}^{\text{inh}}$  between network neurons, i.e., we do not model interneurons explicitly. The afferent weights  $V_{ki}$  denote the strength of synapses from the  $i$ -th input to the  $k$ -th network neuron. For notational convenience, weight 0 is assigned to non-existing connections. As in (BUESING et al., 2011), neurons communicate via rectangular post-synaptic potentials (PSPs) of duration  $\tau$  and with amplitude  $W_{kj}^{\text{exc}}$ ,  $W_{kj}^{\text{inh}}$  and  $V_{ki}$  respectively. Throughout this work, we chose  $\tau = 10$  ms as an estimate for PSP durations. Thus, the membrane potential  $u_k(t)$  integrates the current value of all presynaptic variables  $z_j(t)$  and  $y_i(t)$  at any time.

Network neurons emit spikes stochastically with instantaneous firing probability

$$\rho_k(t) = \lim_{\delta t \rightarrow 0} \frac{1}{\delta t} p(\text{Neuron } k \text{ fires in } [t + \delta t]) = \frac{1}{\tau} e^{u_k(t)} (1 - z_k(t)) . \quad (5.4)$$

Here,  $(1 - z_k(t))$  describes a refractory period that is inversely related to the state  $z_k(t)$ , i.e., when  $z_k(t)$  is active, the neuron is refractory and cannot emit another spike. The exponential dependence of the firing probability  $\rho_k$  on  $u_k$  was confirmed to be a reasonable modeling assumption for neurons in a noisy environment (JOLIVET et al., 2006; PETROVICI et al., 2013).

### The network response as the result of a meaningful Bayesian computation

We aim to understand the activity of the network neurons in response to spiking input as the result of a Bayesian computation. To this end, we establish a link between the stochastic spike response properties of individual neurons and the joint activity distribution of the entire network. As in the neural sampling theory (BUESING et al., 2011), the states  $z_k(t)$  and  $y_i(t)$  of the neurons are formally treated as random variables  $z_k$  and  $y_i$ . This probabilistic description of stochastic network activity in response to given input amounts to a conditional probability distribution  $p(z | \mathbf{y}, \boldsymbol{\theta})$  for each possible input configuration  $\mathbf{y}$ . In order to assign a computational meaning to the conditional distributions  $p(z | \mathbf{y}, \boldsymbol{\theta})$ , we seek to identify a generative model, consisting of prior  $p(z | \boldsymbol{\theta})$  and likelihood  $p(\mathbf{y} | z, \boldsymbol{\theta})$ , for which the conditional distributions  $p(z | \mathbf{y}, \boldsymbol{\theta})$  arise from Bayes rule (5.2) as the correct posterior. If we find such a generative model, the network's spike response can be understood as probabilistic inference through sampling from the posterior  $p(z | \mathbf{y}, \boldsymbol{\theta})$ .

To decide whether the posterior distribution of a generative model is compatible with the neural dynamics of the spiking network, we make use of a sufficient condition that was identified in (BUESING et al., 2011). The so-called neural computability condition connects the membrane potentials  $u_k$  of individual spike response neurons (5.4) with the activity distribution of the recurrent network. For our case of a posterior distribution  $p(z | \mathbf{y}, \boldsymbol{\theta})$  with arbitrary (but fixed) input  $\mathbf{y}$ , the neural computability condition reads:

$$u_k \stackrel{!}{=} \log \frac{p(z_k = 1 | z_{\setminus k}, \mathbf{y}, \boldsymbol{\theta})}{p(z_k = 0 | z_{\setminus k}, \mathbf{y}, \boldsymbol{\theta})} \quad \text{for all } k = 1, \dots, K \quad (5.5)$$

with  $z_{\setminus k} = (z_1, \dots, z_{k-1}, z_{k+1}, \dots, z_K)$  denoting the state of the remaining network. If the condition holds for all possible input states  $\mathbf{y}$  and all possible network states  $z$ , the neural sampling theory guarantees that the trajectory of network states  $z(t)$  is distributed according to  $p(z | \mathbf{y}, \boldsymbol{\theta})$  in the limit  $t \rightarrow \infty$  for any fixed input  $\mathbf{y}$ , i.e., the spiking network samples from  $p(z | \mathbf{y}, \boldsymbol{\theta})$ .

To establish the equivalence between the spiking network in Fig. 5.1D and the Bayesian model in Fig. 5.1E, we proceed as follows. First, we introduce prior  $p(z | \boldsymbol{\theta})$  and likelihood distributions  $p(\mathbf{y} | z, \boldsymbol{\theta})$  that match the graphical model in Fig. 5.1E. These distributions are shaped by a set of parameters that preserve additional degrees of freedom. From these distributions, we can then calculate the posterior  $p(z | \mathbf{y}, \boldsymbol{\theta})$  according to Bayes rule (5.2). Finally, we apply the sufficient condition (5.5) to decide under what conditions the spiking network will sample from the correct posterior distribution. In particular, this will allow us to determine the connectivity structure of the network in detail, and furthermore, to precisely

## 5 Distributed Bayesian computation and learning in spiking neural sheets

map the abstract parameters of the generative model to the synaptic efficacies and intrinsic excitabilities of the neural sheet model.

We set out with a class of prior distributions  $p(z | \theta)$ , namely Boltzmann distributions, that have been shown (BUESING et al., 2011) to be compatible with the neural sampling dynamics of the recurrent spiking network model (without input):

$$p(z | \theta) = \frac{1}{Z} \exp \left[ \frac{1}{2} z^T \hat{W}^{\text{exc}} z + \frac{1}{2} z^T \hat{W}^{\text{inh}} z + z^T \hat{b} \right], \quad (5.6)$$

where  $Z$  is a normalizing factor. The distribution assigns probabilities to binary random vectors  $z = (z_1, \dots, z_K)^T$  with  $z_k \in \{0, 1\}$ . The parameters  $\hat{b}$ ,  $\hat{W}^{\text{exc}}$ ,  $\hat{W}^{\text{inh}}$  shape the distribution and consist of a real-valued bias vector  $\hat{b} = (\hat{b}_1, \dots, \hat{b}_K)^T$  and symmetric, zero-diagonal  $K \times K$  coupling matrices  $\hat{W}^{\text{exc}}$  and  $\hat{W}^{\text{inh}}$ . The matrix notations  $z^T \hat{W} z$  denote bilinear forms. Distributions of the form (5.6) are the maximum entropy distributions for given first and second moments, i.e., this class of distributions captures all mean activations  $\langle z_k \rangle$  and covariances  $\langle z_k z_j \rangle - \langle z_k \rangle \langle z_j \rangle$  of network variables. We endow the prior (5.6) with the spatial structure sketched in the upper row of Fig. 5.1E. Orange circles depict the random variables  $z_k$  that correspond to network neurons. The random variables  $z_k$  maintain sparse excitatory recurrent connections  $\hat{W}_{kj}^{\text{exc}}$  (red links) on an intermediate range. Excitatory links between variables make their coactivation more probable in the prior (5.6) and thus encode network-wide state combinations  $z$  that are more likely to occur than others. This associative memory aspect will turn out to be particularly powerful in the context of recurrent learning where structural knowledge is integrated by the prior and will allow the network to maintain coherent network states even in face of ambiguous input. The sparse excitatory links on intermediate distances are complemented with strong negative connections  $\hat{W}_{kj}^{\text{inh}}$  (theoretically  $\hat{W}_{kj}^{\text{inh}} \rightarrow -\infty$ , blue links) on a local scale among nearby network neurons. These negative connections will correspond to local lateral inhibition in the network and, as we will see in the context of learning, facilitate self-organized statistical model optimization through synaptic plasticity.

We next turn to the likelihood distribution  $p(y | z, \theta)$  that establishes the connection between the network variables  $z_k$  and the input variables  $y_i$ . We adopt a generative perspective in order to identify the likelihood distribution  $p(y | z, \theta)$  and view the network variables  $z_k$  as “hidden causes” of their inputs  $y_i$ . Taking the generative perspective, we assume that each input  $y_i$  is (hypothetically) generated by the corresponding subset of connected network variables  $z_k$ . This amounts to the downward arrows in Fig. 5.1E where each input  $y_i$  (green circles) receives converging arrows only from nearby network variables  $z_k$ , the so-called parents of  $y_i$ . For our idealized architecture, we further assume that no two parents  $z_k$  and  $z_j$  of an individual input  $y_i$  can be active simultaneously. This is ensured by  $z_k$  and  $z_j$  sharing strong negative connections  $\hat{W}_{kj}^{\text{inh}}$ , i.e., the parents of  $y_i$  are assumed to lie within the range of lateral inhibition (blue links in Fig. 5.1E). As a consequence,

each input  $y_i$  is generated by at most one of its parents at a time. We denote the probability of  $y_i$  to be active by  $\pi_{ki}$  (with  $0 < \pi_{ki} < 1$ ) given that it is generated by the active hidden cause  $z_k$ . If none of its parents is active, the chance of activity is assigned a default value  $\pi_{0i}$  which we treat as a constant. This is summarized in the following Bernoulli distribution:

$$p(y_i = 1 | \mathbf{z}, \boldsymbol{\theta}) = \begin{cases} \pi_{ki} & \text{if } z_k = 1 \text{ for a parent } z_k \text{ of } y_i \\ \pi_{0i} & \text{if } z_k = 0 \text{ for all parents of } y_i. \end{cases} \quad (5.7)$$

The extension to the full likelihood distribution  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})$ , that obeys the graphical model in Fig. 5.1E, is straightforward: Since  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})$  is conditioned on the entire random vector  $\mathbf{z}$ , which trivially comprises the parents of all inputs, the  $y_i$ 's are conditionally independent and we obtain  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) = \prod_{i=1}^N p(y_i | \mathbf{z}, \boldsymbol{\theta})$  for the likelihood of all  $N$  input variables. This completes the definition of prior and likelihood distributions that match the graphical model sketched in Fig. 5.1E. A formal definition of the probabilistic model is provided in *Generative model* in Appendix D. There we also describe extensions of the likelihood distribution to support natural and real-valued inputs, i.e., to  $y_i \in \mathbb{N}$  and  $y_i \in \mathbb{R}$ .

### The neural network can sample from the posterior of the generative model

From the prior (5.6) and the likelihoods (5.7), the posterior distribution  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  can be calculated in closed-form using Bayes rule. By applying the neural computability condition (5.5), a straightforward derivation reveals that the membrane potential (5.3) is indeed compatible with the posterior  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  of the generative model. The closed-form posterior and the derivation are provided in *Inference in the generative model* in Appendix D. Furthermore, the calculation yields a translation between the network parameters  $b_k, W_{kj}^{\text{exc}}, W_{kj}^{\text{inh}}, V_{ki}$  and the abstract parameters  $\hat{b}_k, \hat{W}_{kj}^{\text{exc}}, \hat{W}_{kj}^{\text{inh}}, \pi_{ki}$  (and the constants  $\pi_{0i}$ ), and results in the following corollary:

**Corollary 1 (Inference).** *Let network parameters and abstract parameters be identified via*

$$b_k = \hat{b}_k - A_k \quad (5.8)$$

$$V_{ki} = \log \left( \frac{\pi_{ki}}{1 - \pi_{ki}} \right) - \log \left( \frac{\pi_{0i}}{1 - \pi_{0i}} \right) \quad (5.9)$$

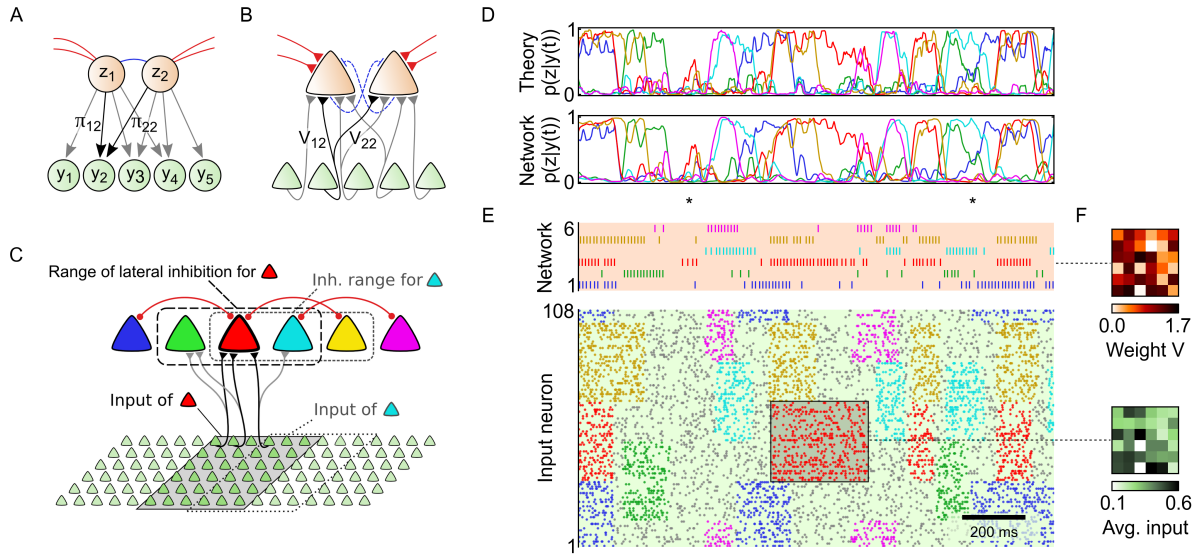
$$W_{kj}^{\text{exc}} = \hat{W}_{kj}^{\text{exc}}, \quad W_{kj}^{\text{inh}} = \hat{W}_{kj}^{\text{inh}} \quad (5.10)$$

with  $A_k := \sum_{i=1}^N \log [1 + \pi_{0i} \cdot (e^{V_{ki}} - 1)]$ , and let the recurrent weights  $\hat{W}^{\text{exc}}$  and  $\hat{W}^{\text{inh}}$  be symmetric matrices with zero diagonal. Then, for any fixed input instantiation  $\mathbf{y}(t) = \mathbf{y}$ , the response  $\mathbf{z}(t)$  of a recurrent spiking network consisting of stochastic neurons (5.4) with linear membrane potential (5.3) is distributed according to  $\mathbf{z}(t) \sim p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  in the limit  $t \rightarrow \infty$ , i.e. the network samples from the posterior distribution of the probabilistic model defined by the prior (5.6) and the likelihoods (5.7).

Corollary 1 explains the structural similarity of the graphical model in Fig. 5.1E and the network architecture in Fig. 5.1D: Each red (blue) recurrent link between latent variables  $z_k$  and  $z_j$  in the graph corresponds to a symmetric reciprocal excitatory (inhibitory) synaptic connection between the  $k$ -th and  $j$ -th network neuron; and each downward arrow from  $z_k$  to  $y_i$  in the graph gives rise to a synapse  $V_{ki}$  from the  $i$ -th input neuron to the  $k$ -th network neuron. In particular, the assumption in the abstract model, that at most one parent  $z_k$  explains a dependent input variable  $y_i$  at a time, translates to a local lateral inhibition motif in the spiking network: Any two network neurons, that share common input, inhibit each other through lateral inhibition. This theoretically derived network motif is reminiscent of cortical lateral inhibition frequently reported across areas and species (DOUGLAS and MARTIN, 2004). In the neural sheet model, local lateral inhibition introduces Winner-take-all (WTA) competition among nearby network neurons which thus play the role of local feature detectors. However, the local WTA circuits are not separated in the sheet, but rather interwoven such that each network neuron can participate in multiple overlapping WTA sub-circuits. In addition, lateral excitatory connections  $W_{kj}^{\text{exc}}$  encode associations between spatially distant feature combinations. This generalized concept of WTA circuits in the continuous sheet contrasts with existing models of interacting WTAs (LANSNER, 2009) which investigated disjoint non-overlapping WTA sub-circuits that operate in parallel. The full probabilistic model  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$  of the neural sheet can thus be understood as a spatially extended reservoir of contiguous competing local feature detectors (corresponding to the likelihood  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})$ ) and an associative memory over the feature set (corresponding to the prior  $p(\mathbf{z} | \boldsymbol{\theta})$ ).

The implementation of Bayes rule (5.2) by the spiking network is illustrated in Fig. 5.2A and B for the minimal example of two neighboring neurons with an overlapping, but not identical, subset of inputs  $y_i$ . Both neurons  $z_1$  and  $z_2$  maintain an implicit likelihood model for their respective local inputs (top-down arrows in A, bottom-up synapses in B). Additionally, the prior installs competition between the neurons via negative reciprocal connections  $W^{\text{inh}}$  (blue edge in A, blue synapses in B). Consequently, the neurons  $z_1$  and  $z_2$  preferentially fire to different input patterns  $\mathbf{y}$ . The response of  $z_1$  and  $z_2$  could be further modulated by the activity of other latent variables  $z_j$  via recurrent excitatory connections (red edges in A, red synapses in B). To illustrate how Bayesian inference is realized in the network, let us consider this two-neuron example and assume that there are no other network neurons. For the small graphical model sketched in Fig. 5.2A and B, suppose that both neurons have the same excitability  $b_1 = b_2 = -2$  and that afferent weights are given by  $(V_{11}, \dots, V_{15}) = (2.0, 0.1, 1.0, 2.5, -)$  and  $(V_{21}, \dots, V_{25}) = (-, 2.0, 0.2, 0.5, 1.5)$  with ‘-’ indicating a non-existing connection. According to eq. (5.9), these afferent weight values correspond to expected activity patterns  $(\pi_{11}, \dots, \pi_{15}) = (0.65, 0.22, 0.40, 0.75, -)$  and  $(\pi_{21}, \dots, \pi_{25}) = (-, 0.65, 0.23, 0.29, 0.53)$  where we assumed  $\pi_{0i} = 0.2$  for the default constant. Then, if only inputs  $y_1$  and  $y_4$  are active,  $z_1$  will typically win the competition, and  $z_2$  will be suppressed due to inhibition and remain silent most of the time. In contrast, if only  $y_2$  and  $y_5$  are active,  $z_2$  will be predominantly active. Finally, if  $y_1$  and  $y_2$  are active,  $z_1$  and  $z_2$

## 5.2 Probabilistic inference in spatially extended spiking networks



**Fig. 5.2: Sheets of spiking neurons can perform Bayesian inference on distributed spiking input.** (A) Local generative model with two competing hidden causes and five inputs. Each hidden cause stores a specific input pattern in the top-down parameters  $\pi_{ki}$ . (B) Corresponding local neural network. Top-down parameters  $\pi_{ki}$  translate to bottom-up synaptic weights  $V_{ki}$ , turning each network neuron into a probabilistic expert for a specific local input pattern. (C) Example network with six network neurons. Neighboring neurons with overlapping input inhibit each other (dashed line: range of lateral inhibition for red neuron). The spiking network is linked to a generative model  $p(\mathbf{y}, \mathbf{z} | \theta)$  according to Corollary 1. (D) The network in C performs sample-based inference of hidden causes  $\mathbf{z}$  under time-varying input  $\mathbf{y}(t)$  in the associated generative model. Comparison of posterior marginals  $p(z_k = 1 | \mathbf{y}(t), \theta)$  between the analytically calculated exact posterior (top) and the average network response (bottom, estimated from 1000 simulation runs) under the time-varying input  $\mathbf{y}(t)$  in panel E. All traces were smoothed with a 20ms box kernel for visual clarity. (E) Top: Network spike response in a single simulation run. Bottom: Input spike trains (colored: structured input, gray: background activity). The network response is an ongoing sampling process from the posterior  $p(\mathbf{z} | \mathbf{y}(t), \theta)$ . Time points marked by a '\*' exemplify characteristic properties of sample-based Bayesian information processing. (F) Top: Bottom-up weights  $V_{ki}$  of the red neuron. Bottom: Average input activity while the matching pattern is presented (during the 300ms period marked in E).

will respond approximately equally often. In this case, the network activity will alternate between the posterior states  $(z_1, z_2) = (0, 0), (0, 1)$  and  $(1, 0)$ . In general, according to Corollary 1, the frequency of occurrence of each network state  $(z_1, z_2)$  will be proportional to the Bayesian posterior probability  $p(z_1, z_2 | \mathbf{y}, \theta)$  for any input instantiation  $\mathbf{y}$ .

### Demonstration of sample-based inference on transient spiking input

Corollary 1 ensures that the network will sample from the correct posterior distribution  $p(\mathbf{z} | \mathbf{y}, \theta)$  in arbitrarily large network architectures for any fixed input configuration  $\mathbf{y}(t) = \mathbf{y}$ . However, Corollary 1 offers no strict guarantees in case of time-varying input since it only ensures that the network activity converges to the equilibrium distribution  $p(\mathbf{z} | \mathbf{y}, \theta)$  in the limit  $t \rightarrow \infty$ . For transient spiking input, the underlying Markov chain cannot fully converge to its equilibrium distribution

and the sampling network will “lag behind” the true posterior. Therefore, the network can only be expected to provide approximate samples from the analytically correct posterior distribution in response to transient spiking input.

We quantitatively assessed the network’s ability to perform inference on time-varying input  $\mathbf{y}(t)$  with the help of computer simulations in the small, analytically tractable network architecture shown in Fig. 5.2C. This 6-neuron network will furthermore serve to illustrate salient response properties of the neural sheet during neural sampling-based inference. In the network, the prior features both sparse excitatory connections and local lateral inhibition such that neighboring neurons compete with each other through inhibitory connections. More distant neurons maintain excitatory connections (red links:  $W_{kj}^{\text{exc}} = 1$ ), representing knowledge that these two hidden causes are likely to co-occur. Furthermore, each neuron has a (randomly generated) preferred local input pattern stored in its synaptic weights  $V_{ki}$ . For instance, the top panel of Fig. 5.2F shows the afferent weights of the red neuron. A demonstration of neural sampling by this circuit in response to time-varying input  $\mathbf{y}(t)$  is shown in Fig. 5.2E: High-dimensional ( $N = 108$ ) input spike trains are presented to the network (bottom). Input spikes consist of uninformative background spikes (gray), interleaved with periods of more structured inputs (colored spikes). Uninformative background spikes are Poisson spike trains with a uniform rate chosen such that the average activation of input neurons is  $\langle y_i \rangle = 0.2$  (ca. 20 Hz). This value also determined the constant  $\pi_{0i}$ . Structured input spikes are Poisson spike trains with rates that were chosen to match the activity patterns stored in the afferent weights  $V_{ki}$  of the network neurons. For example, the red neuron is particularly good at detecting the red spike pattern. Fig. 5.2E (top) shows the network response to the input spike pattern (bottom) during a single simulation run. Whenever a structured local input pattern occurs, the corresponding network neuron starts firing. Yet, the network response is not deterministic: sometimes network neurons elicit spikes even when presented with unstructured input (see e.g. at the 1<sup>st</sup> time point marked by a ‘\*’); in other cases, a competing neighboring neuron emits a spike (see e.g. at the 2<sup>nd</sup> time point marked by a ‘\*’). Due to the stochastic nature of the network response, the exact spike pattern of the network will be slightly different in each simulation run. This apparent trial-to-trial variability is an inherent feature of any sample-based inference process.

We next turn to the question how well the network response approximates the correct Bayesian posterior at any time  $t$ . As we have seen, Corollary 1 suggests to interpret the network states  $\mathbf{z}(t)$  as approximate samples from the posterior distribution  $p(\mathbf{z} | \mathbf{y}(t), \boldsymbol{\theta})$ . Owing to the transient input signal  $\mathbf{y}(t)$ , however, only few samples  $\mathbf{z}(t)$  can be drawn under stable conditions, thereby impeding full convergence to the equilibrium. A quantitative comparison addressing this question is provided in Fig. 5.2D. From many repetitions of the experiment – all with the same input spike pattern  $\mathbf{y}(t)$  – we can estimate the distribution the network actually samples from. The top row of Fig. 5.2D shows the exact marginal probabilities  $p_{\text{theo}}(z_k = 1 | \mathbf{y}(t), \boldsymbol{\theta})$  at any time  $t$ , analytically calculated from Bayes rule. The



### 5.3 Emergence of local experts through synaptic plasticity

bottom row shows the average network response  $p_{\text{net}}(z_k = 1 | \mathbf{y}(t), \boldsymbol{\theta})$ , estimated from the samples from 1000 repetitions of the experiment. The comparison indicates that the sampling network approximates the correct posterior probabilities with high accuracy at almost any time, capturing not only qualitative aspects of the transient posterior distribution but also the quantitative composition of the distribution in face of input fluctuations. Only for particularly rapid input fluctuations, which lead to sharp peaks in the posterior, the network shows a slightly delayed and sometimes inaccurate response. A more detailed statistical evaluation of the sampling quality is provided in *Details to the computer simulations* in Appendix D, along with a brief discussion on the origins of stochastic and systematic deviations in the sampled distribution. In conclusion, the quantitative comparison in Fig. 5.2D shows that the stochastic network response  $z(t)$  can be understood as an ongoing Bayesian inference process, even in case of time-varying input  $\mathbf{y}(t)$ .

Fig. 5.2D and E also exemplify characteristic properties of sample-based Bayesian information processing. At times, input instantiations  $\mathbf{y}(t)$  may be noisy or ambiguous such that the hidden causes of the presented input cannot be inferred with certainty. Two typical examples are marked with a '\*' symbol: At the first time point marked, the unstructured background input accidentally bears some resemblance to the red and the yellow input patterns, such that the posterior probability for inferring the red/yellow patterns temporarily jumps to  $p(z_k = 1 | \mathbf{y}(t), \boldsymbol{\theta}) \approx 1/2$  for  $k = 3$  and  $k = 5$ . This brief moment of uncertainty due to stochastic fluctuations in the input is represented by the network via a small number of spikes of the red ( $k=3$ ) and the yellow ( $k=5$ ) network neurons. While the exact timing of the spikes is a stochastic process, the instantaneous spiking probability in the sampling network is well in line with the analytically calculated posterior. At the second time point marked, two competing neurons (dark blue and green) could explain their local input well, as can be seen from the marginal posterior  $p(z_k = 1 | \mathbf{y}, \boldsymbol{\theta})$ . In such ambiguous situations, both neurons are eager to fire, yet they compete due to lateral inhibition. As a result, network activity switches between two local interpretations where either one of the two hidden causes is active. Due to the stochastic nature of the sampling process, the particular switching times between the competing network states change with each repetition of the experiment, leading to trial-to-trial variability from the perspective of an external observer.

### 5.3 Emergence of local experts through synaptic plasticity

Animals and humans possess the ability not only to infer hidden causes of their perceptions, but also to adapt their internal model, that underlies these inferences, to their specific environment. From a Bayesian perspective, this amounts to reshaping the parameters of the internal model  $p(\mathbf{y}, z | \boldsymbol{\theta})$  to better suit the true input statistics. For our analysis of plasticity, we keep the lateral inhibition structure, i.e, the effect of interneurons, fixed, and focus on plasticity of excitatory synapses. The guiding questions for the following two subsections are:

## 5 Distributed Bayesian computation and learning in spiking neural sheets

- How can the afferent weights  $V_{ki}$ , recurrent weights  $W_{kj}^{\text{exc}}$  and neuronal excitabilities  $b_k$  be adapted to support inference in a statistically optimal manner?
- To what extent can this network-wide optimization be accomplished with only local plasticity rules?

To address these questions we employed a standard objective function (BISHOP, 2006) for statistical model optimization, namely the log-likelihood of the input under the model:

$$\mathcal{L}(\boldsymbol{\theta}) = \langle \log p(\mathbf{y} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})} , \quad (5.11)$$

where  $p^*(\mathbf{y})$  denotes the true distribution of inputs  $\mathbf{y}$  actually presented to the network. Eq. (5.11) makes use of a conceptual advantage of the Bayesian approach: So far we were only interested in the posterior distribution  $p(\mathbf{z} | \mathbf{y}, \boldsymbol{\theta})$  which describes the stochastic network response  $\mathbf{z}$  under a given stimulus  $\mathbf{y}$ ; now we switch to the complementary generative view and examine the distribution  $p(\mathbf{y} | \boldsymbol{\theta}) = \sum_{\mathbf{z}} p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$  of data hypothetically generated by the model. The distribution  $p(\mathbf{y} | \boldsymbol{\theta})$  can be viewed as the outcome when the probabilistic model would “dream” its own environment. Adopting the complementary view is only possible because a full probabilistic model  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$  of the network is available. The function  $\mathcal{L}(\boldsymbol{\theta})$  then measures the likelihood of the actually presented input  $\mathbf{y} \sim p^*(\mathbf{y})$  to occur in  $p(\mathbf{y} | \boldsymbol{\theta})$ . Since  $\mathcal{L}(\boldsymbol{\theta}) = -D_{\text{KL}}(p^* || p) + \text{const.}$ , increasing  $\mathcal{L}(\boldsymbol{\theta})$  is equivalent to reducing the Kullback-Leibler divergence  $D_{\text{KL}}(p^* || p)$  that measures the dissimilarity between the two distributions. Therefore, maximizing  $\mathcal{L}(\boldsymbol{\theta})$  means to align the internal model  $p(\mathbf{y} | \boldsymbol{\theta})$  with the true input distribution  $p^*(\mathbf{y})$ . This objective is commonly known as “maximum likelihood learning” in the machine learning literature (BISHOP, 2006).

In a first step, we investigated the maximization of  $\mathcal{L}(\boldsymbol{\theta})$  with respect to  $\mathbf{V}$ , i.e., we examined the role of plastic afferent synapses  $V_{ki}$ . Plastic recurrent connections  $W_{kj}^{\text{exc}}$  will be addressed in a separate subsection, and we set  $\mathbf{W}^{\text{exc}} = \mathbf{0}$  for now. Update rules for the afferent weights  $V_{ki}$  can be directly derived from the probabilistic model using the mathematical framework of generalized online Expectation Maximization (NESSLER et al., 2013; HABENSCHUSS et al., 2012). The derivation is provided in *Model optimization via Generalized Expectation Maximization* in Appendix D, and yields the following plasticity rules for synaptic weights  $V_{ki}$  and intrinsic excitabilities  $b_k$ :

$$\frac{\partial}{\partial t} V_{ki} = \eta_V \cdot z_k(t) \cdot (y_i(t) - \sigma(V_{ki} + V_{0i})) \quad (5.12)$$

$$\frac{\partial}{\partial t} b_k = \eta_b \cdot (m_k - z_k(t)) . \quad (5.13)$$

Here,  $\eta_V$  and  $\eta_b$  denote small learning rates,  $\sigma(x) := 1/(1 + \exp(-x))$  is the logistic function,  $V_{0i} := \log(\pi_{0i}/(1 - \pi_{0i}))$  is a constant in order to respect the

### 5.3 Emergence of local experts through synaptic plasticity

default activity  $\pi_{0i}$ , and  $m_k$  is a long-term average target response  $\langle z_k \rangle$  of the  $k$ -th network neuron. Importantly, both plasticity rules only rely on information that is locally and instantaneously available to the neurons and synapses: Each afferent synapse  $V_{ki}$  adapts its weight in a Hebbian-type update based on pre- ( $y_i$ ) and post- ( $z_k$ ) synaptic activity through the weight-dependent plasticity rule (5.12); each neuron changes its intrinsic excitability  $b_k$  in a homeostatic fashion based on its current spike response  $z_k$ . The joint application of the rules (5.12) and (5.13) can be shown to implement unsupervised statistical model optimization in the sense of the following corollary:

**Corollary 2 (Learning of afferent synapses).** *The implicit probabilistic model defined by (5.6) and (5.7) can be optimized with respect to the afferent synaptic parameters  $V_{ki}$  by a Generalized Expectation Maximization algorithm which continuously increases a lower bound*

$$\mathcal{F} \leq \mathcal{L}(\theta)$$

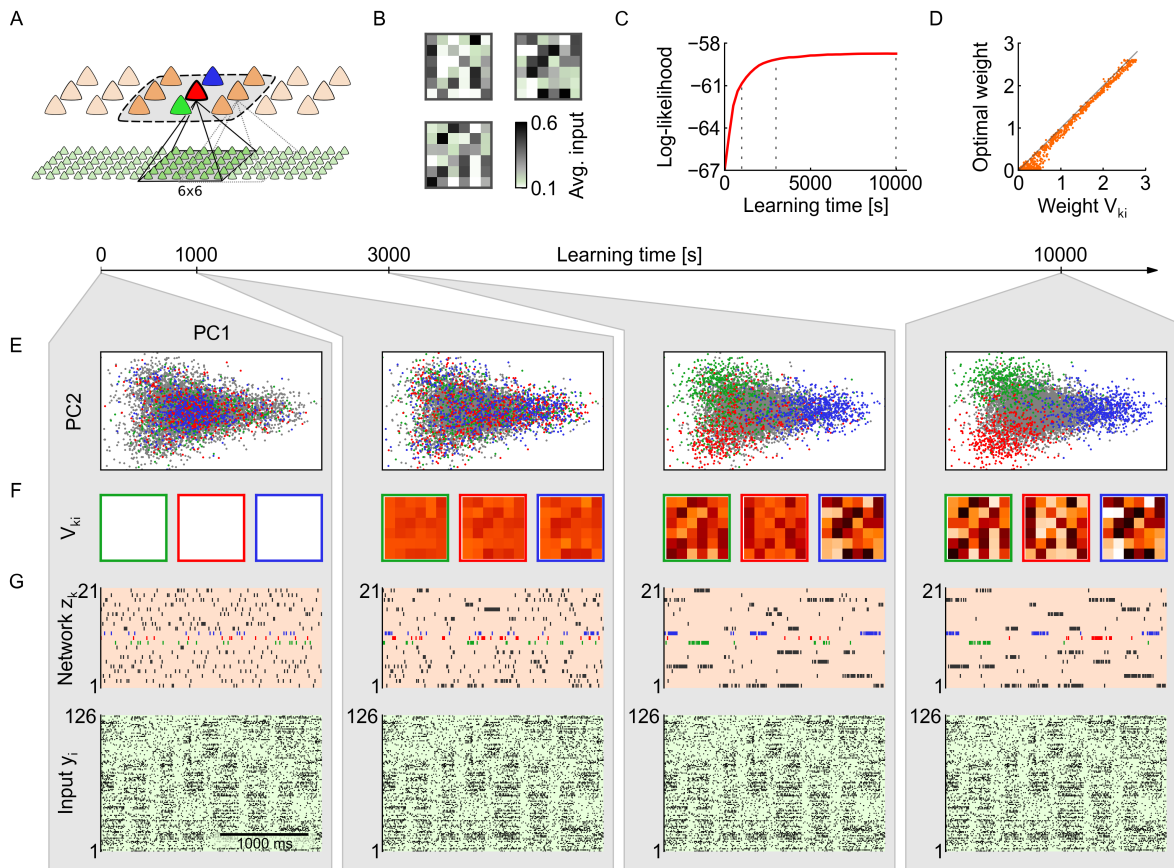
*of the log-likelihood  $\mathcal{L}(\theta)$  until a local optimum of  $\mathcal{F}$  is reached. The concurrent application of the learning rules (5.12) and (5.13) implements an online approximation of this optimization algorithm.*

The approximate character of the network implementation only arises from incomplete convergence of the network's Markov chain during inference and from the non-infinitesimally small learning rates  $\eta_V$  and  $\eta_b$ . Ignoring these effects, i.e., in the limit of small learning rates ( $\eta_b \rightarrow 0$  and  $\eta_V/\eta_b \rightarrow 0$ ) and assuming instantaneous convergence of the Markov chain, we find that the plasticity rules become exact. Then, the direction of the expected learning update is given by,

$$\langle z_k \cdot (y_i - \sigma(V_{ki} + V_{0i})) \rangle = \frac{\partial \mathcal{F}}{\partial V_{ki}} \quad \text{for all } k, i, \quad (5.14)$$

where the expectation  $\langle \cdot \rangle$  is taken with respect to the presented input  $\mathbf{y}(t)$  and the network response  $\mathbf{z}(t)$ . In other words, plastic changes in the synaptic weights  $V_{ki}$  point on average in the direction of increasing  $\mathcal{F}$ . Maximizing a lower bound  $\mathcal{F}$  on  $\mathcal{L}$ , instead of direct optimization of  $\mathcal{L}$ , is a common trick (GRACA et al., 2008) in machine learning to obtain a tractable learning problem. In our model this approach serves to obtain a spiking network implementation of maximum likelihood learning in which all required information is available locally at the neurons and synapses. The local availability of information during learning comes at a cost. The network does not sample from the exact posterior anymore, but from a well-defined variational posterior (HABENSCHUSS et al., 2012): The variational posterior, the network samples from, is the closest distribution to the analytically exact posterior (measured in terms of the Kullback-Leibler divergence) that satisfies the homeostatic long-term target activations  $m_k$ . A full derivation of the plasticity rules (5.12) and (5.13), including a precise definition of the lower bound  $\mathcal{F}$  and the variational posterior distribution, is provided in *Model optimization via Generalized Expectation Maximization* in Appendix D.

## 5 Distributed Bayesian computation and learning in spiking neural sheets



**Fig. 5.3: Emergence of probabilistic local experts through synaptic plasticity.** (A) Network architecture with  $21 \times 6$  inputs and  $7 \times 3$  network neurons. The green, red and blue neuron receive input from the same  $6 \times 6$  subset of input neurons. The input subsets of neighboring groups are shifted by three. The dashed line indicates the range of lateral inhibition. (B) At each of the overlapping  $6 \times 6$  locations, three randomly drawn activity patterns can occur. Shown are the activity patterns for the location highlighted in A. (C) Synaptic and intrinsic plasticity shape the probabilistic model that is encoded by the network. The log-likelihood function  $\mathcal{L}(\theta)$  measures how well the network is adapted to the presented input. (D) One-to-one comparison of the synaptic weights  $V_{ki}$  to analytically calculated optimal weight values, at the end of learning ( $T = 10,000$  s). (E) 2-dimensional projection of the local input distribution  $p^*(\mathbf{y})$ . Each dot is one input instantiation  $\mathbf{y}(t)$  at the  $6 \times 6$  input field in A. Dots are colored according to which of the three neurons (green/red/blue) fired in response (gray: none fired). The neural plasticity rules achieve a clustering of local inputs into local categories. (F) Evolution of synaptic weights  $V_{ki}$  of the green/red/blue neuron over the course of learning. Each neuron becomes a probabilistic local expert for a certain input pattern (cp. panel B). (G) The plastic network develops a sparse, structured spike code that conveys compressed information about the presented input. Bottom: input spike trains. Top: Network response at different stages of learning.

### Demonstration of self-organized learning in the neural sheet

We tested the derived learning rules for afferent weights  $V_{ki}$  and intrinsic excitabilities  $b_k$  in a computer simulation of a sheet of  $7 \times 3$  network neurons (Fig. 5.3A). The sheet model receives synaptic input from a total  $21 \times 6$  afferent cells. Each network neuron receives input from a subset of  $6 \times 6$  inputs. The network neurons are arranged such that in each column there are three neurons sharing the same  $6 \times 6$  input (such as the green/red/blue neuron in Fig. 5.3A). Neighboring columns of network neurons receive inputs from an overlapping subset of afferent cells ( $6 \times 6$  input subset shifted by three columns of input neurons to the left and right respectively). Neurons in the same column as well as neurons in neighboring columns inhibit each other. The input data was generated in a similar manner as in Fig. 5.2, by interleaving a background Poisson spike train with periods of structured Poisson spikes generated from a small number of stereotypical rate patterns. At each  $6 \times 6$  input field there are three such recurring activity patterns. At any moment, at most one such pattern is presented at any spatial position in the input. The resulting spike trains have complex spatio-temporal structure, as shown in Fig. 5.3G (bottom).

Initially all afferent weights are set to  $V_{ki} = 0$ . Afferent weights are plastic and follow eq. (5.12). For the purpose of neuroscientific modeling, synaptic weights  $V_{ki}$  were restricted to positive values. The theory for inference and learning would support positive and negative weights, including sign changes. Excitabilities are uniformly initialized at  $b_k = -2$  and follow eq. (5.13) with the average target response  $\langle z_k \rangle$  of each network neuron set to  $m_k = 6.5\%$ . Whenever a network neuron spikes, synaptic plasticity is triggered and the current activity of the local  $6 \times 6$  input field, that is connected to the active network neuron, leaves a small trace in the afferent synaptic weights of that neuron. As a result, the same neuron is more likely to fire again when a similar input pattern occurs in the future, which leads to further strengthening of the synaptic weights. Due to the combined effect of synaptic plasticity and local competition among neurons, network neurons start specializing on different salient input patterns of their respective  $6 \times 6$  input fields (Fig. 5.3F). Homeostatic intrinsic plasticity ensures during this learning process that neurons which specialize on weak patterns (weak synaptic weights) are not disadvantaged compared to neurons which focus on strong patterns (strong synaptic weights), by regulating intrinsic excitabilities such that all network neurons maintain their long-term average activity  $m_k$ .

A direct consequence of the gradual specialization of network neurons on local salient input patterns is that the network response becomes increasingly more structured and reliable during learning (Fig. 5.3G, top): each network neuron becomes a *probabilistic local expert* for one of the salient activity patterns in its local  $6 \times 6$  input field. This assigns a particular meaning to each of the random variables  $z_k$  of the network: The activity of neuron  $z_k$  represents the presence or absence of the salient local input feature which is encoded in its afferent weights  $V_{ki}$ . Nearby network neurons (such as the blue/red/green neurons in Fig. 5.3A) are in

competition due to lateral inhibition. Therefore, whenever an input is presented to the network, the activity in the local  $6 \times 6$  input field is effectively categorized by similarity to the preferred patterns of the three neurons. This is visualized in Fig. 5.3E: Each dot in the 2D-projection represents one instantiation of the local  $6 \times 6$  input field highlighted in Fig. 5.3A. Grey dots indicate that none of the three neurons (green/red/blue) fired in response to the local input ( $\approx 80\%$  of the time in accordance with the targets  $m_k$ ). Colored dots indicate that the respective network neuron fired (colors as in Fig. 5.3A). After learning, the input space is segmented into three regions (Fig. 5.3E, right). Ambiguous input instances at borders between regions evoke probabilistic responses (e.g. between blue and red region). In this case, the network stochastically responds with one of the two (or three) possible interpretations in order to approximate the posterior probabilities of each hidden cause. The resulting representation on the network level is a sparse structured spike response that conveys highly compressed information about the input.

The qualitative changes in network behavior described above are paralleled by a quantitative improvement of network performance measured by the log-likelihood  $\mathcal{L}(\theta)$  (Fig. 5.3C), as predicted by Corollary 2: From the perspective of statistical model optimization, the learning dynamics due to synaptic and homeostatic plasticity guide a local stochastic search in parameter space which on average increases the lower bound  $\mathcal{F}$  of the log-likelihood  $\mathcal{L}(\theta)$ . At the end of learning, after  $T = 10,000$  s, the network has identified a faithful representation of the actually presented local activity patterns (see Fig. 5.3D). For the shown comparison, theoretically optimal weights were calculated from the presented activity patterns according to eq. (5.9). The reason for the small but systematic differences (learned weights are a bit stronger than predicted) can be found in the facts that the network samples from a variational posterior distribution and that the plasticity rules (5.12) and (5.13) optimize a low bound  $\mathcal{F}$  instead of the log-likelihood  $\mathcal{L}$ .

In summary, we have demonstrated in this subsection how a neural-sampling network can adapt its internal parameters to perform probabilistic inference on distributed spiking input streams. The theoretically derived plasticity rules (5.12) and (5.13) enable the sampling network to develop a sparse and reliable spike code that carries the most salient information of the input stream. This statistical optimization process evolves in a fully self-organized manner by turning network neurons into probabilistic local experts that compete in explaining the presented spike input according to the rules of probability theory.

### The distinct role of lateral inhibition for synaptic learning

Before we address plasticity of recurrent synapses, an important contribution of inhibition to synaptic learning deserves a brief discussion. The simplicity of the synaptic plasticity rule (5.12) arises from the salient lateral inhibition network motif. To identify the role of lateral inhibition for synaptic learning, it is instructive

## 5.4 Plastic recurrent synapses integrate structural knowledge

to review the derivation of (5.12) in the absence of inhibition. By repeating the derivation with  $\mathbf{W}^{\text{inh}} = \mathbf{0}$ , we obtain:

$$\frac{\partial}{\partial t} V_{ki} = \eta_V \cdot z_k(t) \cdot \left[ y_i(t) - \sigma \left( V_{0i} + \sum_{j=1}^K V_{ji} z_j(t) \right) \right]. \quad (5.15)$$

The sum on the right-hand side in (5.15) depends on the activity of neighboring cells  $z_j$ ,  $j \neq k$ , as well as on their afferent weights  $V_{ji}$ , thereby rendering statistically optimal learning non-local. In contrast, local inhibition introduces competition among nearby neurons such that each input variable  $y_i$  is explained by at most one hidden cause  $z_k$  at a time, and, as a consequence, the complex non-local term in eq. (5.15) vanishes. Notably, this outcome is not an artifact of the specific probabilistic model we use, but rather is a general consequence of explaining away effects in any graphical model with converging arrows. This finding suggests that lateral inhibition among nearby neurons assists synaptic learning in a Bayesian framework of model optimization. On the other hand, when lateral inhibition extends beyond neurons with shared afferent input, the expressive power of the probabilistic model is reduced since less hidden causes  $z_k$  are allowed to be active simultaneously. These theoretical considerations suggest that the emergence of efficient representations benefits from network architectures in which the range of local inhibition matches the spatial extent of excitatory cells that share common afferent input.

## 5.4 Plastic recurrent synapses integrate structural knowledge

We have demonstrated how synaptic plasticity can guide statistically optimal learning of afferent connections  $V_{ki}$ . This learning process led to the emergence of probabilistic local experts. As a result, the configuration  $\mathbf{z}(t)$  of active network neurons indicates the subset of currently present local features in the spiking input  $\mathbf{y}(t)$ . In most biologically relevant scenarios, these local input features are unlikely to be statistically independent. For early visual areas, for instance, we can expect the input in nearby spatial receptive fields to exhibit some degree of correlation. Similarly, across sensory modalities certain visual, auditory or tactile stimuli will often occur together. These statistical correlations among local features give rise to non-vanishing covariances  $\langle z_k z_j \rangle - \langle z_k \rangle \langle z_j \rangle$  in the network. In the probabilistic model  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$ , the covariance between network neurons is determined by the recurrent weights  $W_{kj}^{\text{exc}}$  which shape the prior distribution  $p(\mathbf{z} | \boldsymbol{\theta})$ . Since a probabilistic model supports inference best if its prior  $p(\mathbf{z} | \boldsymbol{\theta})$  reflects the input-evoked correlation structure (BERKES et al., 2011), we extend our investigation of statistically optimal learning to plastic recurrent synapses  $W_{kj}^{\text{exc}}$ . For the derivation, we can follow the same approach that has already afforded the plasticity rule (5.12) for afferent synapses. The key idea is to identify a synaptic plasticity rule that points

## 5 Distributed Bayesian computation and learning in spiking neural sheets

on average in the direction of  $\partial \mathcal{F} / \partial W_{kj}^{\text{exc}}$ , i.e., in the direction of increasing  $\mathcal{F}$ . The derivation is provided in *Model optimization via Generalized Expectation Maximization* in Appendix D and yields the following theoretically optimal plasticity rule for recurrent weights:

$$\frac{\partial}{\partial t} W_{kj}^{\text{exc}} = \eta_W \cdot \left( z_k(t) z_j(t) - \phi_{kj}^{\text{opt}}(\hat{W}^{\text{exc}}, \hat{W}^{\text{inh}}, \hat{\mathbf{b}}) \right) \quad (5.16)$$

with learning rate  $\eta_W$  and  $\phi_{kj}^{\text{opt}}(\hat{W}^{\text{exc}}, \hat{W}^{\text{inh}}, \hat{\mathbf{b}}) = \langle z_k z_j \rangle_{p(z|\boldsymbol{\theta})}$ . The plasticity rule (5.16) features a long-term potentiation (LTP) and a long-term depression (LTD) term: Concurrent activation  $z_k(t) z_j(t)$  of network neurons in response to the input strengthens the synapse in a local Hebbian LTP update. Depression, however, turns out to be non-local since the term  $\phi_{kj}^{\text{opt}}$  depends on all parameters  $\hat{W}^{\text{exc}}$ ,  $\hat{W}^{\text{inh}}$  and  $\hat{\mathbf{b}}$  of the prior distribution  $p(z|\boldsymbol{\theta})$ . This increased complexity of learning in recurrent systems is well-known in machine learning theory (SALAKHUTDINOV and G. HINTON, 2009), and we can employ (non-local) machine learning techniques to determine the value of  $\phi_{kj}^{\text{opt}}$  in a so-called sleep phase. This algorithmic approach is known as wake-sleep learning in the literature (G. E. HINTON and SEJNOWSKI, 1986). Since the calculation of  $\phi_{kj}^{\text{opt}}$  is often computationally costly, the development of approximate solutions (such as contrastive divergence (G. HINTON, 2002; CARREIRA-PERPINAN and G. E. HINTON, 2005)), which are tailored to particular network architectures and learning tasks, turned out beneficial in machine learning. The question, whether the brain makes use of similar learning strategies, is subject of ongoing theoretical and experimental research (see e.g. (PEYRACHE et al., 2009; Z.-W. LIU et al., 2010), or (TONONI and CIRELLI, 2014) for a recent review), and it appears indeed conceivable that nature found ways to estimate  $\phi_{kj}^{\text{opt}}$ .

In the present study, we contribute to this intriguing hypothesis by exploring to what extent even simple plasticity rules could be sufficient to approximate the non-local plasticity dynamics of eq. (5.16) in the neural sheet model. Specifically, we are interested in plasticity rules which (i) rely on only local information, (ii) can be applied uniformly to all recurrent synapses, and (iii) are shaped by only a small set of parameters. In case of the architectures often considered in the machine learning literature (stacked Restricted Boltzmann Machines (RBMs) and variants thereof) this would likely be a hopeless endeavor. However, the network architecture of the neural sheet is fundamentally different from stacked RBM architectures, and the local integration of input may further ease the complexity of recurrent learning. The conceptual separation of feature detection (through plastic afferent weights  $V_{ki}$ ) on the one hand, and feature structure (through plastic recurrent weights  $W_{kj}^{\text{exc}}$ ) on the other hand, can be expected to facilitate learning in that – once the essential features have been identified – the complexity of learning the prior shows some resemblance to the reduced complexity of training a fully visible Boltzmann machine (HYVÄRINEN, 2006). Thus, the emergence of probabilistic local experts may provide a guidance for learning of recurrent connections. Inspired by



## 5.4 Plastic recurrent synapses integrate structural knowledge

the structure of eq. (5.16) and the weight dependence of the LTD term in eq. (5.12), we make the following ansatz for a local synaptic plasticity rule:

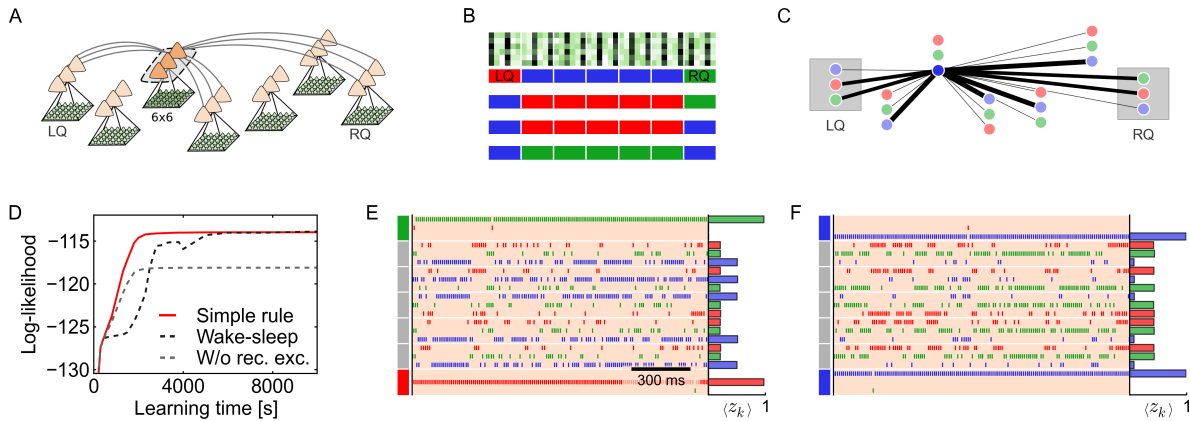
$$\frac{\partial}{\partial t} W_{kj}^{\text{exc}} = \eta_W \cdot \left( z_k(t) z_j(t) - \phi_{\vartheta}(W_{kj}^{\text{exc}}) \right) \quad (5.17)$$

where the LTD-term  $\phi_{\vartheta}(W_{kj}^{\text{exc}})$  only depends on the weight of the respective synapse, and the parameters  $\vartheta$  of  $\phi_{\vartheta}$  are to be determined for the given learning task. Ideally, the LTD function  $\phi_{\vartheta}$  should be chosen such that eq. (5.16) and (5.17) lead to the same weight values  $W_{kj}^{\text{exc}}$  under given neuronal spike patterns. In the simple rule, plasticity is governed by two antagonistic terms: LTP is proportional to the average coactivation  $\langle z_k z_j \rangle$  of a synapse's pre- and post-synaptic neurons; LTD is proportional to  $\phi_{\vartheta}(W_{kj}^{\text{exc}})$ . Hence, if  $\phi_{\vartheta}$  is a continuous and monotonically increasing function in  $W_{kj}^{\text{exc}}$ , there exists an equilibrium weight  $W_{kj}^{\text{exc}}$  for each value of the coactivation  $\langle z_k z_j \rangle$ . Based on this observation, a principled approach to identify a suitable LTD function  $\phi_{\vartheta}$  is to first determine optimal weights  $W_{kj}^{\text{opt}}$  for a given learning problem using the theoretically optimal wake-sleep rule (5.16), and then to fit  $\phi_{\vartheta}(W_{kj}^{\text{exc}})$  such that the identified optimal weights are approximately reproduced under given coactivations  $\langle z_k z_j \rangle$  by the simple rule (5.17). In *Approximate plasticity rule for recurrent synapses* in Appendix D, we show how for a set of optimal weights  $W_{kj}^{\text{opt}}$  and corresponding coactivations  $\langle z_k z_j \rangle$  an LTD function  $\phi_{\vartheta}(W_{kj}^{\text{exc}})$  can be constructed that features the desired convergence points  $W_{kj}^{\text{opt}}$ . It must be noted that the local plasticity rule (5.17) can only serve as a heuristic for approximating the plasticity dynamics of the theory-based wake-sleep rule (5.16). As such, the parameters  $\vartheta$  need to be adjusted for every learning scenario and there exists no strict guarantee of stable plasticity dynamics. Therefore, we explored the suitability of the local recurrent plasticity rule in computer simulations by comparing the learning results of the local rule with the outcome of wake-sleep learning.

### Demonstration of self-organized integration of structural knowledge

We tested the learning capabilities of the simple plasticity rule (5.17) and compared it with the theoretically optimal wake-sleep rule (5.16) in a computer simulation with seven spatially separate network populations (see Fig. 5.4A). Each population consisted of three neurons and received spiking input from a group of  $6 \times 6$  inputs. Within each population neurons are subject to lateral inhibition. Neurons from different populations are linked via all-to-all recurrent excitatory connections. At each spatially separate  $6 \times 6$  input location one of three local activity patterns may occur. The local activity patterns are simple stripe patterns (see Fig. 5.4B top) that can easily be learned by the afferent weights  $V_{ki}$ . For brevity, we refer to the three activity patterns as the 'red', 'green' and 'blue' pattern in the following.

## 5 Distributed Bayesian computation and learning in spiking neural sheets



**Fig. 5.4: Plastic recurrent synapses integrate structural knowledge.** (A) Network architecture with seven recurrently connected local populations (all-to-all beyond inhibition). Input weights  $V_{ki}$  and recurrent excitatory weights  $W_{kj}^{\text{exc}}$  are plastic. One of three stripe activity patterns (visualized as red/green/blue) is presented as spiking input to each local population. (B) Examples of the input structure. Top: input activity of the strip patterns (250 ms average). Bottom: The two outer locations serve as left and right cue (LQ,RQ). Cue input patterns are chosen independently. The cues determine the type (“color”) of the inner patterns in that the inner inputs are always different from both cues. In addition, inner patterns are always consistent. Hence, inner network neurons must consider both cues and the state of the other inner network neurons to infer their own state. (C) Recurrent weights of the highlighted blue-tuned neuron after 25,000 s of learning with the simple local plasticity rule (5.17). Network neurons are colored according to the local pattern they have become experts for. Line width encodes the synaptic weight  $W_{kj}^{\text{exc}}$  (min. line width 0.2 for  $W_{kj}^{\text{exc}} = 0$ ). In accordance with the input structure, the inner blue neuron has developed strong excitatory connections to the other inner blue neurons and moderately strong connections to the red and green cue neurons. (D) Comparison of the log-likelihood over the first 10,000 s of learning for three recurrent plasticity conditions: with the simple recurrent plasticity rule (red), with the wake-sleep algorithm (dashed black), and without recurrent excitation (dashed gray). Recurrent plasticity significantly enhances the network’s learning capabilities. (E) The knowledge on the input structure, that was learned by the recurrent synapses  $W_{kj}^{\text{exc}}$ , enables inference of correct global network states in face of incomplete input. Only the outer cues are presented while all inner inputs show an uninformative (“gray”) pattern. With red and green cues, the network correctly infers that the inner hidden causes should be blue, most of the time. Horizontal bars: mean activity  $\langle z_k \rangle$  of network neurons (average over 100 s). (F) When both outer cues are blue, the already incomplete input is furthermore ambiguous. During inference, the network switches stochastically between the two consistent global interpretations.

In order to assess the network’s ability to detect and integrate highly interdependent correlation structures among features at different locations, we introduced complex dependencies between the input presented to the seven populations: While the two *outer* input locations served as a left cue (LQ) and right cue (RQ), and were chosen independently, the five *inner* locations were chosen to be (a) different from the cues, and (b) consistent with each other. For instance, if the outer cue patterns were red and green, all inner input patterns were blue. If, however, both cues showed the same pattern (e.g., both blue), an ambiguous situation arose: In this case, the inner patterns were still chosen to be different from the cue, but additionally all inner patterns were consistently of the same type (either all red or all green). Several examples illustrating this correlation structure are sketched in Fig. 5.4B. As a consequence, an inner network neuron, that is tuned to one of the three local patterns, has to consider both cues and the state of all other inner network neurons in order to infer its own state correctly.

## 5.4 Plastic recurrent synapses integrate structural knowledge

We tested the network’s ability to recover the statistical structure of the input in a simulation with concurrent learning of afferent weights and recurrent weights (initial values:  $V_{ki} = W_{kj}^{\text{exc}} = 0$  and restricted to positive values). In a first simulation run, recurrent synaptic plasticity followed the theoretically optimal wake-sleep plasticity rule (5.16). As expected from the theory, all network neurons developed a tuning to one of the local input patterns, and recurrent weights correctly reflected the correlation structure of the task after 25,000 s of learning: Similarly tuned inner network neurons formed strong excitatory recurrent links  $W_{kj}^{\text{exc}}$ ; in addition, inner neurons developed excitatory connections with compatible cue neurons (e.g. with the red and green cue for a blue-tuned inner neuron). The ongoing adaptation of the network is reflected in the log-likelihood function  $\mathcal{L}(\theta)$  shown in Fig. 5.4D. The black dashed line shows the learning progress with the theoretically optimal learning rule. For comparison, we performed an independent simulation with recurrent plasticity switched off (dashed gray, all  $W_{kj}^{\text{exc}}$  fixed at zero). Without recurrent excitation, the log-likelihood settled at a significantly lower value. The gap in  $\mathcal{L}$  between the two simulations corresponds to the structural information stored in the prior  $p(z | \theta)$ .

Building on the theory-based learning result, we constructed a suitable LTD function  $\phi_{\vartheta}$  to obtain a simple plasticity rule. The set of target weights  $W_{kj}^{\text{opt}}$ , that emerged during the first simulation run, and the set of coactivations  $\langle z_k z_j \rangle$ , that had led to these weights, are provided in *Details to the computer simulations* in Appendix D. Based on the observed functional dependence between  $W_{kj}^{\text{opt}}$  and  $\langle z_k z_j \rangle$ , we chose the following LTD function:

$$\phi_{\vartheta}(W_{kj}^{\text{exc}}) = m_k \cdot m_j + \frac{1}{\gamma} \tan \left( \frac{\pi}{2} \frac{W_{kj}^{\text{exc}}}{W^{\text{max}}} \right) \quad (5.18)$$

with parameters  $\vartheta = (W^{\text{max}} = 1.41, \gamma = 31.6)$  fitted to the data. The LTD function (5.18) has two components: The term  $m_k \cdot m_j$  describes the expected coactivation  $\langle z_k z_j \rangle$  if the neurons were statistically independent. This case is associated with weight  $W_{kj}^{\text{exc}} = 0$  when synaptic plasticity follows eq. (5.17). The second term accounts for positive correlations between  $z_k$  and  $z_j$  through a stabilizing weight dependence. While the specific tangent-shape is a heuristic, the functional form has some properties that are generally expected for Hebbian-type plasticity: The LTD function is strictly monotonically increasing such that higher coactivations settle at stronger efficacies; and, the maximum efficacy of a synapse is bounded since the LTD contribution goes to infinity as  $W_{kj}^{\text{exc}}$  approaches  $W^{\text{max}}$ . The scaling parameter  $\gamma$  sets the strength of LTD.

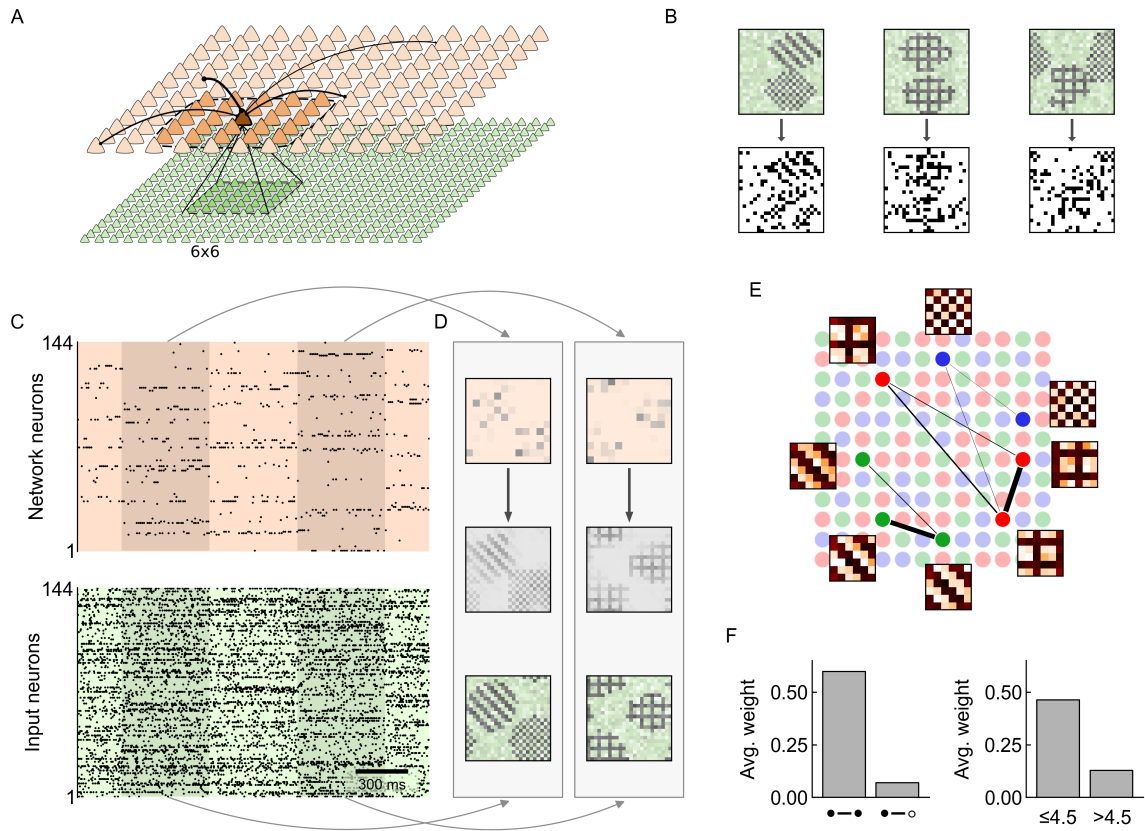
In a second simulation run, all recurrent and afferent weights were reset to zero, and recurrent synaptic plasticity followed the simple local plasticity rule given by eq. (5.17) and (5.18). As seen in Fig. 5.4D, the simple rule is virtually indistinguishable from the theoretically optimal rule in terms of the log-likelihood.

The resulting recurrent connectivity structure after 25,000 s of learning is shown in Fig. 5.4C for the example of an inner network neuron that has specialized on the blue input pattern: The neuron developed strong recurrent weights with all other inner neurons of similar tuning, and moderately strong weights with red- and green-tuned cue neurons. This connectivity matrix mirrors the connectivity structure obtained with theoretically optimal wake-sleep learning (see *Details to the computer simulations* in Appendix D for numerical values).

The learned structural knowledge can be exploited by the sampling network during inference. Most distinctly, the benefits of well-adapted recurrent excitatory connections become apparent in response to incomplete, or even ambiguous, stimuli (Fig. 5.4E and F). To emulate a scenario of incomplete observations, only outer cues were presented to the network while all inner inputs showed uninformative, uniform activity (indicated as gray). For example, when the cues are set to red and green, as shown in Fig. 5.4E, the network correctly infers that all inner features should be blue most of the time. In the given experimental protocol, this knowledge can only be communicated via recurrent excitatory synapses. In a second example shown in Fig. 5.4F, the already incomplete input is furthermore chosen to be ambiguous by presenting a both-blue cue. Hence, the inner features should be either all red or all green. In this case, the network activity stochastically switches between the two inferred consistent interpretations (note that the switching of network neurons at inner locations is synchronized), as expected for sampling from a bi-modal posterior distribution  $p(z | y, \theta)$ .

In conclusion, the computer experiment in Fig. 5.4 demonstrates that the sampling network model can concurrently identify salient features of its input stream and recover complex correlations among spatially distributed features through synaptic plasticity. Self-organized learning of recurrent excitatory connections can be understood as an ongoing refinement of the prior distribution in the network's internal model of presented input. The obtained structural knowledge on the input distribution significantly enhances the network's ability to maintain globally coherent network states in face of incomplete and ambiguous observations. For the examined inference task, which required integration of two independent cues as well as the state of multiple other network neurons, we observed that even the simple local plasticity rule (5.17)+(5.18) endowed the network with close-to-perfect learning capabilities. As a word of caution, owing to theoretical considerations it is unlikely, that such simple plasticity rules will be sufficient to solve arbitrarily sophisticated learning problems. Yet, as we have seen, the architecture of the neural sheet model eases the complexity of recurrent learning in that the local input integration through probabilistic experts can guide learning of recurrent connections.

## 5.4 Plastic recurrent synapses integrate structural knowledge



**Fig. 5.5: Emergence of excitatory subnetworks in neural sheets.** (A) Network architecture with  $24 \times 24$  inputs and  $12 \times 12$  network neurons. Each network neuron receives input from a subset of  $6 \times 6$  inputs, as illustrated for the highlighted neuron. Input subsets of neighboring neurons are shifted by two. The dashed line marks the range of lateral inhibition for the highlighted neuron. Additionally, pairs of network neurons, that do not inhibit each other, maintain sparse recurrent excitatory synapses  $W_{kj}^{\text{exc}}$  (25% prob. for a reciprocal connection). All bottom-up weights  $V_{ki}$  and recurrent weights  $W_{kj}^{\text{exc}}$  are plastic. (B) Three exemplary input activity patterns (top; 500 ms average) and spiking input instantiations  $\mathbf{y}(t)$  therefrom (bottom). Two of three local input templates (grid/diagonals/checker) are presented at different random locations at a time. (C) Input spike trains (bottom) and network response (top) after 10,000 s of learning. Input patterns switch every 500 ms. Only a random subset of 144 inputs is shown to facilitate the comparison of spike density. The network response is considerably sparser than the input. (D) Reconstruction of the input from the sparse network response. Top: Average network response  $\langle \mathbf{z}(t) \rangle$  during the same periods. Middle: Reconstruction  $\langle \mathbf{y}^{\text{gen}} \rangle$  of the input from the network response by means of the learned generative model. (E) Emergence of distributed assemblies of local experts. During learning each network neuron becomes a probabilistic expert for one of the local activity patterns (indicated by red/green/blue color). Input weights  $V_{ki}$  are shown for highlighted neurons. In addition, the network developed recurrent excitatory connections  $W_{kj}^{\text{exc}}$ . Neurons, that are experts for similar patterns, form excitatory subnetworks after learning (line width proportional to  $W_{kj}^{\text{exc}}$ ). (F) Recurrent excitatory connections between similarly tuned neurons are considerably stronger than connections between neurons with different specialization (left). In addition, connections between nearby pairs of network neurons are stronger than between distant pairs (Euclidean distance; based on 4198 recurrent exc. synapses).

## 5.5 Emergence of excitatory subnetworks in neural sheets

In previous subsections, we have examined how small sheets of spiking neurons with local lateral inhibition and recurrent excitation can perform probabilistic inference by sampling from the posterior distribution  $p(z | \mathbf{y}, \theta)$  of a well-defined probabilistic model, and how the parameters  $\theta$  of the underlying probabilistic model can be adapted to the presented input by plastic afferent and recurrent synapses. For illustration purposes, network architectures were small and computer simulations were tailored to highlight specific aspects of sample-based inference and statistical learning.

In this final subsection, we combine all mechanisms described above and explore Bayesian information processing and self-organized learning in a spatially extended neural sheet model. The employed network architecture is shown in Fig. 5.5A. Network neurons are organized in a two-dimensional lamina and integrate spiking input locally via plastic afferent synapses (initial weight  $V_{ki} = 0$ ). Lateral inhibition is spatially confined for any point of reference (dashed line for the highlighted cell) yet omnipresent in the homogeneous continuous tissue, giving rise to a plethora of interleaved and overlapping competitive subcircuits. Beyond the range of inhibition, network neurons maintain sparse, plastic excitatory connections (per pair: 25% chance of a reciprocal connection, initial weight  $W_{kj}^{\text{exc}} = 0$ ). Note that, due to the sheer amount of possible network states  $z(t)$ , a traditional analytical calculation of the posterior distribution  $p(z | \mathbf{y}, \theta)$  becomes intractable in this architecture.

We presented spiking input to the network as shown in Fig. 5.5B: Three local activity motifs (grid, checkerboard, diagonal stripes; top row of Fig. 5.5B) can appear at different locations in the input space. Each local motif is larger than the diameter of afferent connections  $V_{ki}$  to individual network neurons (highlighted green neurons in Fig. 5.5A). At any time, two randomly drawn motifs are presented at different non-overlapping random locations. Inputs  $y_i$  that are not part of a local motif maintain a low-activity background firing rate. Poisson spike trains were generated from the resulting rate patterns (top row of Fig. 5.5B), leading to versatile (and seemingly noisy) input  $\mathbf{y}(t)$  that was presented to the network (bottom row of Fig. 5.5B).

This input mimics some important aspects of cortical information processing: First, input neurons are tuned to certain input features (the presence/absence of the three input motifs), but their spiking activity is highly stochastic and the same input neuron  $y_i$  can be responsive to multiple different input patterns. Consequently, inputs  $y_i$  are not very informative when observed individually. Second, input motifs are too large to be explained by a single network neuron, and thus multiple network neurons must be recruited in parallel for the explanation. Finally, multiple independent streams of information can be presented to the network at different locations simultaneously.

We tested the capability of the neural sheet to integrate and adapt to this input in a 10,000 s learning experiment. All parameters ( $W_{kj}^{\text{exc}}, V_{ki}, b_k$ ) were plastic. As

## 5.5 Emergence of excitatory subnetworks in neural sheets

before, we first trained the network with the theoretically optimal recurrent plasticity rule (5.16). Fitting the parameters  $\vartheta$  for the simple rule (5.17)+(5.18) yielded  $W^{\max} = 2.70$  and  $\gamma = 734$ . Note that, even though the two parameters  $\vartheta$  are tailored to the learning task and are thus assumed to be given to the network, the resulting plasticity rule  $\phi_{\vartheta}$  is of drastically reduced complexity compared with wake-sleep learning. In particular, all recurrent synapses share the same LTD function  $\phi_{\vartheta}(W_{kj}^{\text{exc}})$ . Using the fully-local plasticity rules for  $W_{kj}^{\text{exc}}$ ,  $V_{ki}$ , and  $b_k$ , the network was then simulated again “from scratch”. Fig. 5.5C shows spike trains of the input and the network neurons after learning. To facilitate the comparison of the network response to the input, only a random subset of 144 input neurons is shown. The network response is considerably sparser than the input, with only few network neurons firing simultaneously at a long-term average firing rate of ca. 2.5 Hz. To examine what aspects of the input are conveyed by the sparse response, we can adopt the generative perspective, again. From the learned afferent weights  $V_{ki}$ , the translation to the input domain as established by eq. (5.9) and the likelihood model  $p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta})$ , we can reconstruct the average input  $\langle \mathbf{y}^{\text{gen}} \rangle$  expected from the network response  $\mathbf{z}(t)$  during a short time window:

$$\langle \mathbf{y}_i^{\text{gen}} \rangle = \frac{1}{\Delta} \int_{t_0}^{t_0+\Delta} \sigma \left( V_{0i} + \sum_k z_k(t) V_{ki} \right) dt , \quad (5.19)$$

where we chose  $\Delta = 500$  ms. This reconstruction is shown in Fig. 5.5D for the two highlighted time windows of Fig. 5.5C. The reconstructed  $\langle \mathbf{y}^{\text{gen}} \rangle$  are plotted in gray scale alongside the true average inputs  $\langle \mathbf{y}(t) \rangle$  plotted in green. While small differences, especially at the spatial feature boundaries, are visible, the network response still conveys the most salient information of the presented input. In other words, the network has developed a code which is not only sparse but also very efficient, carrying highly compressed information about the presented input spike patterns in only few spikes.

Fig. 5.5E provides insight into how this code has been established: Each network neuron has become a local expert for one of the three local inputs motifs, and together, network neurons cover the entire input and feature space. Excitatory connections among network neurons store positive associations, i.e., a strong excitatory weight between two network neurons indicates that the corresponding local features tend to co-occur in the input. In line with this notion, we observe the emergence of recurrent excitatory subnetworks among those neurons that receive input from contiguous spatial domains and exhibit tuning to similar input features. These learned associations can be used during inference for ensuring that more likely configurations of network states are preferentially visited by the network in case of uncertainty. The qualitative observation of excitatory subnetworks is further supported by the statistical analysis shown in Fig. 5.5F: Synaptic weights between network neurons with similar tuning are on average stronger than weights between neurons with different tuning properties (left). And, synapses between

nearby cells are stronger than connections between distant cells (right), reflecting the local nature of the presented input.

### 5.6 Discussion

We have proposed a spatially structured spiking network model for distributed Bayesian inference and self-organized learning through synaptic plasticity. Building on the theory of MCMC sampling, we have shown how the transient spike response of the recurrent neural network can be understood as an ongoing sampling process from a well-defined Bayesian posterior distribution. Our study extends work by Buesing et al. in that it endows the generic network architecture of (BUESING et al., 2011) with a spatial structure (namely, local lateral inhibition and sparse recurrent excitation) and the ability to integrate distributed spiking input. It turned out that the local integration of distributed spiking input streams assigns a particular meaning to the previously abstract random variables of the neural circuit: Network neurons encode the presence or absence of salient input features which are stored in the network's afferent synapses. This leads to the conception of network neurons as probabilistic local experts which are organized in a tissue of interwoven local winner-take-all circuit motifs. Beyond the range of lateral inhibition, network neurons communicate via sparse recurrent excitatory connections on an intermediate spatial scale. From a theoretical perspective, the recurrent transfer of information is linked to the prior distribution in Bayes rule and captures structural knowledge on statistical correlations among spatially separate input features. This ability to align instantaneous observations with previously obtained structural knowledge according to the rules of probability theory enables the neural sheet to maintain coherent global network states even when the presented input is incomplete or ambiguous. Moreover, having a full probabilistic description  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$ , that covers both the input and the network response, at hand permitted a rigorous mathematical treatment of self-organized learning. Extending work of (HABENSCHUSS et al., 2012) and (NESSLER et al., 2013), we demonstrated that the interplay of STDP-type synaptic and homeostatic intrinsic plasticity can approximate stochastic online Expectation Maximization, a powerful machine learning algorithm. This is a remarkable finding: Global statistical model optimization can be achieved in a spatially extended network through only local information exchange (pre- and post-synaptic spiking activity) in a fully unsupervised manner.

#### Theoretical aspects of information processing in the neural sheet

The theoretical analysis of inference and learning in the neural sheet model revealed a distinct functional role of individual network components. These are briefly addressed in the following and comprise the functional contribution of precise relative spike timing, the lateral inhibition connectivity motif, and the dynamics of self-organized synaptic learning.



**The relative spike timing carries essential information.** In the proposed neural sheet model, multivariate posterior distributions  $p(z | y, \theta)$  are represented by interpreting network spikes as realizations of random variables  $z_k$  according to eq. (5.1): After a spike of the  $k$ -th network neuron, the associated RV  $z_k$  turns active for a short time  $\tau$  (10 ms in this study). This gives rise to vector-valued network states  $z(t)$  in which multiple RVs  $z_k$  will typically be active, simultaneously. As a direct consequence, the relative spike timing of neuronal subgroups carries important information on a millisecond time scale: overlapping on-times encode the probability of coactivation of the associated RVs in the multivariate distribution  $p(z | y, \theta)$ . The computational importance of incorporating coactivations during inference becomes particularly evident in ambiguous situations that support multiple coherent – but mutually exclusive – explanations. The sampling network’s spike response in such an ambiguous scenario was demonstrated in Fig. 5.4 where the network switched stochastically between two coherent perceptual modes in a synchronized manner. The information density conveyed by this structured, bimodal response goes far beyond the quality of conclusions that could be drawn from just observing the average firing of individual neurons. Figuratively, considering only marginal responses would disregard the insight that an obstacle could be circumvented either on the left hand or the right hand side, but would instead suggest to steer a happy medium.

**Lateral inhibition facilitates local synaptic learning.** This information-rich neural representation of globally coherent network states was shown to emerge fully autonomously through the interplay of local synaptic and intrinsic plasticity rules. The derivation of the weight-dependent plasticity rule (5.12) for afferent synapses  $V_{ki}$  furthermore revealed an essential role of the local lateral inhibition network motif during learning. The contribution of lateral inhibition becomes apparent when the derivation of eq. (5.12) is repeated in the absence of the inhibition motif, resulting in the update rule (5.15). The latter rule (5.15) requires information not only on the pre- and post-synaptic spiking activity of a plastic synapse, but also on the specific weight values of other nearby synapses, thereby rendering learning non-local. This increased complexity of parameter learning is a general problem in graphical models with converging arrows, known as explaining away, and thus, the locality of synaptic learning displays a conceptual challenge for a wide range of Bayesian network architectures. In the neural sheet model, lateral inhibition establishes a particularly strong form of explaining away by ensuring that each input  $y_i$  is explained by at most one network variable  $z_k$ , at a time. The resulting competition among network neurons restores the locality of information required for synaptic learning. We therefore suspect that the competition introduced by local lateral inhibition could assist synaptic learning in a wide range of network models in that it facilitates global statistical model optimization through local synaptic plasticity rules. In particular, in order to maximize the expressive power of the system while preserving local synaptic learning, our theory suggests that self-organized learning of efficient representations benefits from network architectures in which

the range of lateral inhibition matches the spatial extent of network neurons that share common afferent input.

**Recurrent plasticity integrates structural knowledge.** Regarding plasticity of recurrent connections  $W_{kj}^{\text{exc}}$  in the neural sheet, we found optimal learning rules to be fundamentally non-local; a hardly surprising finding for recurrent systems. We therefore explored in computer simulations to what extent even simple local plasticity rules with approximately matching convergence properties could be sufficient to recover complex structural correlations in the presented input. In the small computer experiments of Fig. 5.4 and Fig. 5.5, we observed no significant impairment in the emergent weight configuration when simple local plasticity rules were used (compared with the theoretically optimal non-local rule). This observation indicates that, for tasks of “not too high” complexity, the prior  $p(z | \theta)$  could indeed be adjusted by only local synaptic plasticity. The specific shape of the employed plasticity rule, however, needs to be tailored to statistical properties of the given task. For biological systems, it is thus conceivable that evolution forged tailored simple plasticity dynamics for certain neuronal populations and brain areas that interact with more elaborate plasticity types.

**Robust evasion of suboptimal solutions.** Finally, the learning dynamics of the underlying stochastic online Expectation Maximization algorithm, that performs local gradient ascent in the synaptic weights, deserves a brief discussion. Like every local optimization algorithm, learning in the neural sheet is only guaranteed to converge to a local optimum of the objective function  $\mathcal{F}$ . This is a general issue in unsupervised learning since the likelihood functions of complex probabilistic models may possess many local optima. However, two properties of the neural sheet model are expected to mitigate this issue: First, the network employs stochasticity in two ways, via the random presentation of input samples  $\mathbf{y} \sim p^*(\mathbf{y})$  and via the stochastic nature of the network response  $z \sim p(z | \mathbf{y}, \theta)$ . This induces stochastic fluctuations in the synaptic weights and facilitates the evasion of local optima (compared with a fully deterministic batch algorithm as often used in machine learning). Second, homeostatic intrinsic plasticity forces all neurons to participate in explaining the input. Thereby, many “particularly bad” local optima, which recruit only a small fraction of hidden causes  $z_k$  in the average posterior, are automatically evaded. In accordance with these properties, we observed in the computer simulations that learning was generally very robust and that the network reliably identified near-optimal parameter settings most of the time.

### Idealized modeling assumptions and conceptual limitations

Connectivity structure and neural dynamics of the proposed sheet model have their origin in top-down principles of Bayesian information processing and machine learning theory. In the spiking network implementation of the abstract algorithms,

the time scales involved during inference and learning span multiple orders of magnitude, ranging from milliseconds for the instantaneous sampling process, over tens of seconds for homeostatic plasticity, up to minutes and hours for synaptic plasticity (see also *Interaction of time scales* in Appendix D). In order to keep these complex interactions in the recurrent system theoretically tractable, several idealized modeling assumptions had to be made. We employed a simple spike response neuron model that elicits action potentials stochastically based on an idealized membrane potential. The membrane potential integrates synaptic input linearly, and synaptic transmission is mediated via rectangular non-additive post-synaptic potentials without delay. Homeostasis entered the system only in a single process, namely as homeostatic intrinsic plasticity, with the aim to maintain a predefined average firing activity. In order to stay close to the derived learning rules, synaptic plasticity was implemented on the level of neuronal states, i.e., on the value of  $y_i(t)$  and  $z_k(t)$ , and was not further mapped to the level of spikes. However, in (NESSLER et al., 2013) and (HABENSCHUSS et al., 2013b) it has been demonstrated that (variants of) the plasticity rule (5.12) can be mapped down to the spike level, and that the resulting plasticity dynamics are tightly linked to spike-timing dependent plasticity. In the computer simulations, lateral inhibition was transmitted via strong reciprocal synapses between directly connected network neurons, thereby integrating out the dynamics of putative interneurons. Regarding the connectivity structure, lateral inhibition strictly obeyed the condition that a reciprocal inhibitory connection exists between two network neurons if the neurons share common inputs. In line with the generic probabilistic model of (BUESING et al., 2011), that was employed for the network population, recurrent excitatory connections were symmetric. On a more conceptual level, it is noteworthy that, while the sampling process of the network evolves in time, the underlying probabilistic model  $p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$  is inherently non-temporal, i.e., the theory makes no predictions on the temporal structure of network trajectories or the integration of salient temporal features of the input. Possible extensions of the theory in order to support more complex recurrent (and also asymmetric) connectivity structures, soft lateral inhibition, as well as the integration of temporal sequences are outlined below.

### Relation to cortical microcircuits

The proposed neural sheet model shares some striking similarities with cortical microcircuits. These similarities range from salient connectivity motifs, to microscopic neural dynamics of single cells and synapses, up to population-level response characteristics in living animals. Most clearly, the neural sheet model can be linked to salient aspects of cortical layer 2/3, with network neurons being associated with pyramidal cells and lateral inhibition being mediated disynaptically by fast-spiking interneurons (e.g., basket cells). Certainly, the idealized neural dynamics of the model cannot be expected to find any precise counterpart in biology. And, of course, the abstract network model does not (and could not) intend to provide any complete description of all the subtleties found in cortical microcircuits. Yet,

we believe that the evident similarities, as briefly reviewed in the following, could contribute to sharpening our conception of the complex neural dynamics observed in living tissue.

**Disynaptic inhibition establishes local competition among pyramidal cells.** In vitro experiments indicate that central aspects of the ubiquitous lateral inhibition network motif in the neural sheet model are established in layer 2/3 by somatargeting, fast-spiking (FS) interneurons. FS interneurons preferably form synapses locally (STEPANYANTS et al., 2009) and show particularly high connection probability with nearby pyramidal cells (PCs) (AVERMANN et al., 2012). These GABAergic connections typically involve many (ca. 15) contacts (MARKRAM et al., 2004) per cell-pair and inhibit the target close to its cell body (MARKRAM et al., 2004). In addition, FS-PC connections were reported to feature very short transmission delays of only approx. 1ms (AVERMANN et al., 2012). The dense and fast inhibition of PCs by FS interneurons is complemented with an increased probability for reciprocal connections (YOSHIMURA and CALLAWAY, 2005), i.e., bidirectional links between nearby PC-FS pairs. Notably, reciprocal PC-FS connections were reported to be especially strong in either direction (see (YOSHIMURA and CALLAWAY, 2005; MATEO et al., 2011) for PC→FS, and (YOSHIMURA and CALLAWAY, 2005) for PC←FS). Furthermore, also the PC→FS connection was found to have below-average transmission delays (AVERMANN et al., 2012). This led the authors of (AVERMANN et al., 2012) to the conclusion that, “[t]aken together, our *in vitro* data (this study) and our related *in vivo* data (MATEO et al., 2011) suggest that disynaptic inhibition driven by FS GABAergic neurons in the neocortex mediates competition among excitatory neurons such that perhaps only a small fraction of excitatory layer 2/3 neurons can be active at any given time.” The finding of lateral competition among PCs in layer 2/3 was also confirmed *in vivo* by (ADESNIK and SCANZIANI, 2010).

**Pyramidal cells are clearly tuned while local lateral inhibition is unspecific.** In vivo studies with awake animals characterized typical response properties of PCs and interneurons. During quiet wakefulness, PCs show a much lower firing activity than FS cells (GENTET et al., 2010). This was explained by the stronger synaptic drive required for exciting PCs compared with FS cells, which maintain an average membrane potential close below threshold (GENTET et al., 2010). Furthermore, detailed patch-clamp recordings, that separated individual conductance contributions, pointed to significantly sharper spatial tuning of excitation than inhibition (HAIDER et al., 2013). This finding is supported by studies with anesthetized mice that reported clear orientation selectivity of excitatory neurons, while GABAergic cells showed only little tuning (SOHYA et al., 2007; ZARIWALA et al., 2011), but see (RUNYAN et al., 2010). Indeed, experimental data suggests that the stimulus dependence of interneuron responses could be explained by the integrated activity of surrounding neurons (KERLIN et al., 2010).

**Sparse coding and structured connectivity of excitatory neurons.** Experimental data on the spike response of PCs in layer 2/3 appears to be compatible with a sparse coding scheme. For instance, it was observed in vivo by (SMITH and HÄUSSER, 2010) that even neurons with overlapping receptive fields showed only low correlation. Also (ECKER et al., 2010) reported generally weak noise correlations between neighboring neurons of similar tuning, suggesting a mechanism of active decorrelation among nearby neurons. A recent review (BARTH and POULET, 2012), that assessed experimental evidence for sparse coding, emphasized the particularly sparse response of layer 2/3 and pointed to the generally low firing activity in superficial layers compared with neurons from layer 4, which are considered to be a major input to layer 2/3. The functional characterization of PC responses is complemented by studies on physical connections between excitatory neurons. While inhibition was shown to act mostly locally, recurrent excitation spans larger distances (STEPANYANTS et al., 2009). Furthermore, neurons with positively correlated responses were found to also have an increased probability to maintain reciprocal connections (Ko et al., 2011). More specifically, in (BOSKING et al., 1997) an increased connection probability was reported between neurons with similar orientation tuning and cell body distances larger than  $500 \mu\text{m}$ .

**Bottom-up models propose similar network architectures and function.** Based on the rich experimental data on cortical layer 2/3 both in vivo and in vitro, experimental neurobiologists and computational neuroscientists have proposed functional models of information processing in layer 2/3 that appear compatible with the architecture of the neural sheet model in this article. Douglas and Martin (DOUGLAS and MARTIN, 2004) as well as Lansner (LANSNER, 2009) sketched functional models of layer 2/3 that rely on local competitive circuits which communicate via associative excitatory links, although these studies did not emphasize the interwoven structure of local inhibition. The local competitive aspect of information processing in layer 2/3 was examined in detail by (AVERMANN et al., 2012) and (MATEO et al., 2011). Finally, the architecture of the neural sheet model is highly reminiscent of the connectivity structures within a cortical column model of barrel cortex, as proposed in a recent review by Petersen and Crochet (PETERSEN and CROCHET, 2013).

### Experimentally testable predictions

In our model, the correlation structure between stimulus features is reflected in the lateral connectivity of the network after learning. In the generative model, the plastic modification of lateral connections corresponds to reshaping the prior distribution  $p(z | \theta)$ . Whether such correlation structure of features can be observed in the prior could be tested in an experiment that extends the setup of Berkes et al. (BERKES et al., 2011). There, it was observed in ferret V1 that the distribution of spontaneous activity (dark stimulus condition, identified with the prior  $p(z | \theta)$ ) became increasingly similar with the average evoked activity (natural stimulus conditions, identified with the average posterior  $\langle p(z | \mathbf{y}, \theta) \rangle$ ) during development.

## 5 Distributed Bayesian computation and learning in spiking neural sheets

Our model predicts that a change in the correlation structure of environmental features should be mirrored in the generative prior and thus in the correlation structure of cortical spontaneous activity after learning.

In Fig. 5.4, we have demonstrated synchronized perceptual switching between two modes of activity under ambiguous stimulus conditions. Such stochastic switching displays a characteristic property of sample-based representations when the posterior distribution is bimodal. Note that in our model, perceptual switching is a result of the learned prior  $p(z | \theta)$  that assigns very low probability to network states that seem inconsistent with previous experience. According to our model, the prior is subject to ongoing learning. Hence, for initial presentations of an ambiguous stimulus, the prediction is that network activity will alternate between response patterns that are consistent with the alternative stimulus interpretations rather than evoking a stable intermediate response (see Fig. 5.4F). After sufficient learning time, however, it is expected that a distinct neural representation emerges for the previously ambiguous stimulus, on which the network activity settles.

### Related theoretical work and integration into larger networks

During the last decade, several theoretical studies examined how Bayesian computations could be performed and represented by spiking neural networks. In the following, we discuss how the neural sheet model relates to existing work with a focus on two key aspects: (1) The fundamental (and unanswered) question how probability distributions are encoded in neuronal activity patterns, and (2) the compatibility of the proposed sheet model, that focused on spatial aspects of distributed inference and learning, with spiking network models that addressed other aspects of Bayesian information processing, such as more complex causal relations or temporal integration. The two aspects are discussed separately.

**Neural representations of Bayesian computation.** The algorithms that underlie Bayesian computations in neural network models are diverse. However, regarding the representation of the arising posterior distributions  $p(z | y, \theta)$  two general lines of research can be identified, namely sample-based codes and distributional codes. In sample-based codes (as employed by the neural sheet model) the observed network state is interpreted as an instantiation of one or more random variables  $z$ . By observing the sequence  $z(t)$  of network states over time, the distribution  $p(z | y, \theta)$  is represented with increasing precision through the relative frequency of state occurrences (HOYER and HYVÄRINEN, 2003; FISER et al., 2010). Examples of a direct (one-to-one) mapping between neurons and random variables are found in (BERKES et al., 2011), (BUESING et al., 2011), (NESSLER et al., 2013) and (KAPPEL et al., 2014). A conceptual separation between network neurons and the represented RVs was explored by (PECEVSKI et al., 2011) where only a subset of so-called principle neurons carry meaningful information for downstream populations. The idea to rigorously separate neuronal spike patterns from the represented RVs was recently explored

by (SAVIN and DENEVE, 2014), thereby allowing the simultaneous operation of multiple entangled sampling chains within a single network. In a different direction, recent theoretical work (HABENSCHUSS et al., 2013a) has shown that the notion of “network states” can be extended to also cover entire population trajectories. On a general account, it has been emphasized (FISER et al., 2010) that sample-based codes offer several conceptual advantages, such as high representational flexibility, easy marginalization, natural emergence of response variability, and general suitability for learning.

Distributional codes provide a complementary (and at first glance, irreconcilable) neural representation of probability distributions. A characteristic property of distributional codes is the (almost) instantaneous representation of either the entire posterior distribution or, at least, pivotal statistical properties thereof (e.g. marginals or mean values). Just as the case for sampling networks, the exact inference algorithms implemented by distributional code network models are manifold. One line of research (R. P. RAO, 2004; DENEVE, 2005; R. P. RAO, 2007; STEIMER et al., 2009), builds upon the belief propagation algorithm that aims to calculate the marginal posteriors  $\langle z_k \rangle_{p(z|y,\theta)}$  for all variables. This approach enables inference in complex graphical models, including temporal integration. However, since correlations among RVs in the posterior are not accommodated, the precise relative spike timing between different neurons carries no relevant information: what matters is the spike count in a given time interval, not the spike timing. A second line of research is established by (probabilistic) population codes (see e.g., (R. S. ZEMEL et al., 1998; MA et al., 2006; POUGET et al., 2013)). These models aim to infer the posterior of a hidden (“true”) stimulus parameter (e.g., bar orientation) from the observed input activity in a known generative model. The inference process relies on integrating the spike count of many input neurons with known tuning properties, simultaneously. Neuronal trial-to-trial variability arises in this deterministic inference scheme solely from stochasticity in the spiking input.

Despite their different origins, sample-based codes and probabilistic population codes can provide mutually compatible interpretations, at least in some scenarios. Consider, for instance, a local column in the neural sheet model with multiple network neurons within the range of lateral inhibition. This local network is very similar to the competitive networks examined in (NESSLER et al., 2013). Over the course of learning, each neuron will develop a tuning, e.g., a preferred bar orientation (cp. Fig. 5 in (NESSLER et al., 2013)), such that the local network neurons will jointly cover the entire local input space. Due to the width of the likelihood distribution associated with each hidden cause, the response curves of roughly similarly tuned neurons will partially overlap. Formally, the resulting spike response of the local population to a given stimulus displays a sample-based code of a multinomial (mixture) posterior distribution. However, by knowing the tuning curves of each network neuron, the response of the sampling network could equally well be interpreted as a distributional code of the stimulus parameter (cp. Fig. 2 in (POUGET et al., 2013)). Conversely, the continuous hidden stimulus parameter in PPC models could always be mapped to a set of locally competing network neurons

in a sampling sheet. Thus, it appears that, at least in some cases, sample-based codes and probabilistic population codes are mutually compatible, and further experimental research will be required to investigate spike response characteristics in more complex scenarios.

**Integration into larger Bayesian spiking networks.** In this work, we have employed a rather basic prior distribution  $p(z|\theta)$  for associative memory formation, namely a single-layer Boltzmann distribution. The theory of the neural sheet model, however, also supports more complex network structures such as deep learning architectures (SALAKHUTDINOV and G. HINTON, 2009) that constitute one of the most powerful tools for unsupervised learning in machine learning theory. Such deep network architectures would add hierarchical information exchange to the system (e.g., top-down or contextual information). Furthermore, arbitrary Bayesian networks, i.e., directed graphical models, could likely be used for the prior distribution as long as the graphical model features the local lateral inhibition network motif (mutually exclusive activity of hidden variables with shared input). Several spiking network implementations for sample-based inference in general graphical models can be found in (PECEVSKI et al., 2011). Notably, these implementations overcome the constraint of symmetric recurrent connections in the network. An asymmetric recurrent connectivity structure is also a salient property of temporal models. To endow the spatial model with the ability of temporal integration, it would therefore be intriguing to combine the proposed sheet model with the Hidden Markov Model network implementation by (KAPPEL et al., 2014) or the neural particle filtering approach by (LEGENSTEIN and MAASS, 2014). In a different research direction, we observed in computer simulations that relaxing the assumption of strong lateral inhibition still leads to reasonable learning results in many scenarios. This indicates that even soft inhibition could be sufficient to govern the emergence of probabilistic experts. Yet, our current theory does not provide a proper interpretation of the resulting network response since the arising coactivation of neighboring neurons likely demands a generative model that supports soft explaining away. Finally, it would be interesting to combine the fully self-organized network model with reinforcement learning signals from the environment (e.g., via top-down feedback or third-factor plasticity rules). This could endow the spiking network with the ability of Bayesian decision making and action selection.

### Operation paradigm for novel computing platforms

Beyond neuroscience, the proposed neural sheet model displays an intriguing design principle for neuromorphic architectures (MEAD and ISMAIL, 1989). Neuromorphic systems rely on rigorous parallelization of information processing by implementing physical models of neurons and synapses in microscale electronic circuitry (SCHEMMEL et al., 2008). A key intention behind the development of these systems is the construction of fault tolerant, self-organized computing devices that



overcome the traditional strictly serial and deterministic design of von Neumann architectures.

The proposed neural sheet model appears to be ideally suited for neuromorphic implementations: information exchange is fully asynchronous without a central clock; the emergent sparse spike code reduces the load on the interneuronal bus system; and the local confinement of lateral inhibition facilitates compact mapping and efficient routing. Furthermore, idealized modeling assumptions, such as stochastic neurons, rectangular PSP shapes or tailored inhibitory connections, could likely be accounted for in engineered systems.

One of the main challenges in neuromorphic engineering lies in the required high-density integration of plastic synapses (SCHEMMEL et al., 2006). On the way to a principled solution, material scientists have made great progress in recent years in using memristors as plastic synapses (INDIVERI et al., 2013; J. J. YANG et al., 2013). Memristors (CHUA, 1971; STRUKOV et al., 2008) are novel nanoscale materials that adapt their electrical conductance based on the history of the current and voltage flux through the device. Thus, memristors are expected to accommodate both synaptic transmission and synaptic plasticity without the need of extensive supporting circuitry. Experimental (SURI et al., 2013; MAYR et al., 2012) and theoretical (SNIDER, 2008; SERRANO-GOTARREDONA et al., 2013) studies have explored how the plasticity dynamics of memristors can be utilized in spiking neural networks. In particular, it was shown in (QUERLIOZ et al., 2011) and (BILL and LEGENSTEIN, 2014) that conductance changes in certain memristive materials appear compatible with (variants of) the weight dependent plasticity rule (5.12). Therefore, we expect that the proposed neural sheet model can provide a promising paradigm for the efficient operation of neuromorphic hardware as massively parallel computing devices for probabilistic inference and self-organized learning.

## Conclusion

We have proposed a spiking neural sheet model for sample-based probabilistic inference and self-organized learning. The spatially structured network combines aspects of local competitive learning and large-scale associative memory formation under a unified Bayesian account. While the model has been derived from mathematical principles of Markov chain Monte Carlo sampling and Machine Learning theory, the resulting network architecture and neural dynamics bear striking resemblance with salient connectivity motifs and neuronal response properties observed experimentally in cortical microcircuits. Therefore, we believe that the theoretical findings presented in this article can significantly contribute to our understanding of the intricate dynamics observed in cortical networks.



# Appendix



# Appendix **A**

## List of publications

1. JOHANNES BILL, LARS BÜSING, STEFAN HABENSCHUSS, BERNHARD NESSLER, WOLFGANG MAASS, ROBERT LEGENSTEIN (2015). "Distributed Bayesian computation and self-organized learning in sheets of spiking neurons with local lateral inhibition." In: (*under review*).
2. MIHAI PETROVICI\*, JOHANNES BILL\*, ILJA BYTSCHOK, JOHANNES SCHEMMEL, KARLHEINZ MEIER (2015). "Stochastic inference with deterministic spiking neurons." In: (*under review*).
3. DIMITRI PROBST, MIHAI PETROVICI, ILJA BYTSCHOK, JOHANNES BILL, DEJAN PECEVSKI, JOHANNES SCHEMMEL, KARLHEINZ MEIER (2015). "Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons." In: *Frontiers in Computational Neuroscience*.
4. JOHANNES BILL, ROBERT LEGENSTEIN (2014). "A compound memristive synapse model for statistical learning through STDP in spiking neural networks." In: *Frontiers in Neuroscience*.
5. STEFAN HABENSCHUSS\*, JOHANNES BILL\*, BERNHARD NESSLER (2012). "Homeostatic plasticity in Bayesian spiking networks as Expectation Maximization with posterior constraints." In: *Advances in Neural Information Processing Systems (NIPS)*.
6. LARS BÜSING, JOHANNES BILL, BERNHARD NESSLER, WOLFGANG MAASS (2011). "Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons." In: *PLoS Computational Biology*.
7. JOHANNES BILL, KLAUS SCHUCH, DANIEL BRÜDERLE, JOHANNES SCHEMMEL, WOLFGANG MAASS, KARLHEINZ MEIER (2010). "Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity." In: *Frontiers in Computational Neuroscience*.
8. DANIEL BRÜDERLE, JOHANNES BILL, BERNHARD KAPLAN, JENS KREMKOW, KARLHEINZ MEIER, ERIC MÜLLER AND JOHANNES SCHEMMEL (2010). "Simulator-Like Exploration of Cortical Network Architectures with a Mixed-Signal VLSI System." In: *IEEE International Symposium on Circuits and Systems (ISCAS)*.
9. JOHANNES BILL\*, MARTIN-I. TRAPPE\*, IGOR LESANOVSKY, PETER SCHMELCHER (2006). "Resonant quantum dynamics of neutral spin-1 particles in a magnetic guide." In: *Physical Review A*.

\* = Equal contribution



## Appendix to chapter 2: Neural dynamics as sampling

In this appendix, we provide details and proofs for the neural sampling theory, followed by details for the computer simulations. Then we investigate typical firing statistics of individual neurons during neural sampling and examine the approximation quality of neural sampling with different neuron and synapse models.

### B.1 Derivation of the neural sampling theory

#### Notation

To keep the derivations in a compact form, we introduce the following notations. We define the function  $\zeta_k^{>0}$  of  $\zeta_k$  to be 1 if  $\zeta_k > 0$  and 0 otherwise. Analogously we define  $\zeta_{\setminus k}^{>0} = (\zeta_1^{>0}, \dots, \zeta_{k-1}^{>0}, \zeta_{k+1}^{>0}, \dots, \zeta_K^{>0})$ . Let  $\delta(\cdot, \cdot)$  denote Kronecker's Delta, i.e.,  $\delta(x, y) = 1$  if  $x = y$  and 0 otherwise whereas  $\delta(\cdot)$  denotes Dirac's Delta, i.e.,  $\int f(x)\delta(x)dx = f(0)$ . Furthermore  $\chi_I(x)$  is the indicator function of the set  $I$ , i.e.,  $\chi_I(x) = 1$  if  $x \in I$  and  $\chi_I(x) = 0$  if  $x \notin I$ .

#### Details to neural sampling with absolute refractory period in discrete time

The following Lemmata 1 – 3 provide a proof of Theorem 1. For completeness we begin this paragraph with a recapitulation of the definitions stated in Chapter 2. We then identify some central properties of the joint probability distribution  $p(\zeta, z)$  and prove that the proposed network samples from the desired invariant distribution.

For a given distribution  $p(z)$  over the binary variables  $z \in \{0, 1\}^K$  with  $\forall z \in \{0, 1\}^K p(z) \neq 0$ , the joint distribution over  $(\zeta, z)$  with  $\zeta \in \{0, 1, \dots, \tau\}^K$  is defined

## B Appendix to chapter 2: Neural dynamics as sampling

in the following way (see equation 2.7):

$$p(\zeta_k|z_k) := \begin{cases} \tau^{-1} & \text{for } z_k = 1 \wedge \zeta_k > 0 \\ 1 & \text{for } z_k = 0 \wedge \zeta_k = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$p(\zeta|z) := \prod_{k=1}^K p(\zeta_k|z_k)$$

$$p(\zeta, z) := p(\zeta|z)p(z).$$

The assumption  $p(z) \neq 0$  for all  $z$  is required to show the irreducibility of the Markov chain, a prerequisite to ensure the uniqueness of the invariant distribution of the MCMC dynamics. Furthermore, for the given distribution  $p(z)$  we define the functions  $u_k : \{0, 1\}^{K-1} \rightarrow \mathbb{R}$  for  $k \in \{1, \dots, K\}$  which map  $z_{\setminus k} \mapsto u_k(z_{\setminus k})$ :

$$u_k(z_{\setminus k}) := \text{logit}(p(z_k = 1|z_{\setminus k})) = \log \frac{p(z_k = 1|z_{\setminus k})}{p(z_k = 0|z_{\setminus k})}.$$

Instead of  $u_k(z_{\setminus k})$  we simply write  $u_k$  in the following.

**Lemma 1.** *The distribution  $p(\zeta, z)$  has conditional distributions of the following form:*

$$p(\zeta_k|\zeta_{\setminus k}, z_{\setminus k}) = p(\zeta_k|z_{\setminus k}) = \begin{cases} \frac{\sigma(u_k)}{\tau} & \text{for } \zeta_k > 0 \\ 1 - \sigma(u_k) & \text{otherwise} \end{cases}$$

$$p(z_k|\zeta, z_{\setminus k}) = p(z_k|\zeta_k) = \begin{cases} 1 & \text{for } \zeta_k > 0 \wedge z_k = 1 \\ 1 & \text{for } \zeta_k = 0 \wedge z_k = 0 \\ 0 & \text{otherwise} \end{cases}.$$

These results can also be written more compactly in the following form:  $p(\zeta_k|z_{\setminus k}) = \sigma(u_k)\chi_{\{1, \dots, \tau\}}(\zeta_k)\frac{1}{\tau} + (1 - \sigma(u_k))\delta(\zeta_k, 0)$  and  $p(z_k|\zeta_k) = \delta(z_k, \zeta_k^{>0})$ .

*Proof.* Here we use the fact that the logistic function  $\sigma$  is the inverse of the logit function, i.e.,  $p(z_k = 1|z_{\setminus k}) = \sigma(u_k)$ .

$$p(\zeta_k|\zeta_{\setminus k}, z_{\setminus k}) = \sum_{z_k=0}^1 \frac{p(\zeta, z)}{p(\zeta_{\setminus k}, z_{\setminus k})} = \sum_{z_k=0}^1 \frac{p(\zeta, z)}{p(\zeta_{\setminus k}|z_{\setminus k})p(z_{\setminus k})} = \sum_{z_k=0}^1 \frac{(\prod_{l \neq k} p(\zeta_l|z_l)) p(\zeta_k|z_k)p(z)}{(\prod_{l \neq k} p(\zeta_l|z_l)) p(z_{\setminus k})}$$

$$= \sum_{z_k=0}^1 p(\zeta_k|z_k)p(z_k|z_{\setminus k}) = \sigma(u_k)\chi_{\{1, \dots, \tau\}}(\zeta_k)\frac{1}{\tau} + (1 - \sigma(u_k))\delta(\zeta_k, 0).$$

This also shows that  $\zeta_k$  is independent from  $\zeta_{\setminus k}$  given  $z_{\setminus k}$ , i.e.,  $p(\zeta_k|\zeta_{\setminus k}, z_{\setminus k}) =$



$p(\zeta_k | z_{\setminus k})$ . Now we show the second relation using Bayes' rule:

$$\begin{aligned}
 p(z_k | \zeta, z_{\setminus k}) &= \frac{p(\zeta_k | \zeta_{\setminus k}, z)}{p(\zeta_k | \zeta_{\setminus k}, z_{\setminus k})} p(z_k | \zeta_{\setminus k}, z_{\setminus k}) \\
 &= \frac{z_k \chi_{\{1, \dots, \tau\}}(\zeta_k)^{\frac{1}{\tau}} + (1 - z_k) \delta(\zeta_k, 0)}{\sigma(u_k) \chi_{\{1, \dots, \tau\}}(\zeta_k)^{\frac{1}{\tau}} + (1 - \sigma(u_k)) \delta(\zeta_k, 0)} p(z_k | z_{\setminus k}) \\
 &= \begin{cases} z_k & \text{for } \zeta_k > 0 \\ 1 - z_k & \text{for } \zeta_k = 0 \end{cases} \\
 &= \delta(z_k, \zeta_k^{>0}).
 \end{aligned}$$

□

In order to facilitate the verification of the next two Lemmata, we first restate the definition of the operators  $T^k$  in a more concise way:

$$\begin{aligned}
 T &:= T^1 \circ \dots \circ T^K \\
 T^k(\zeta, z | \zeta', z') &:= T^k(\zeta_k, z_k | \zeta'_k, z'_k) \delta(\zeta_{\setminus k}, \zeta'_{\setminus k}) \delta(z_{\setminus k}, z'_{\setminus k}) \\
 T^k(\zeta_k, z_k | \zeta', z') &:= \delta(z_k, \zeta_k^{>0}) \cdot T^k(\zeta_k | \zeta'_k, z'_k) \\
 T^k(\zeta_k | \zeta'_k, z'_k) &:= \begin{cases} \sigma(u'_k - \log \tau) & \text{for } \zeta_k = \tau \wedge \zeta'_k = 0, 1 \\ 1 - \sigma(u'_k - \log \tau) & \text{for } \zeta_k = 0 \wedge \zeta'_k = 0, 1 \\ 1 & \text{for } \zeta_k = \zeta'_k - 1 \wedge \zeta'_k > 1 \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

where  $u'_k := u_k(z'_k) = \text{logit}(p(z_k = 1 | z'_{\setminus k}))$ .

**Lemma 2.** For all  $k = 1, \dots, K$  the operator  $T^k(\zeta_k | \zeta'_k, z'_k)$  leaves the conditional distribution  $p(\zeta_k | z'_{\setminus k})$  invariant.

*Proof.* For sake of simplicity, denote  $T^k(\zeta_k = i | \zeta'_k = j, z'_{\setminus k}) = T_{ij}^k$  for  $i, j \in \{0, 1, \dots, \tau\}$  and  $p(\zeta_k = i | z'_{\setminus k}) = p_i$ . We have to show  $p_i \stackrel{!}{=} \sum_{j=0}^{\tau} T_{ij}^k p_j$  for  $i \in \{0, 1, \dots, \tau\}$ .

First we show  $p_\tau = \sum_{j=0}^{\tau} T_{\tau j}^k p_j$  using  $p_0 = 1 - \sigma(u_k)$  and  $p_1 = p_2 = \dots = p_\tau = \sigma(u_k) \tau^{-1}$  (which results from Lemma 1):

$$\begin{aligned}
 \sum_{j=0}^{\tau} T_{\tau j}^k p_j &= T_{\tau 0}^k p_0 + T_{\tau 1}^k p_1 = \sigma(u_k - \log \tau) (1 - \sigma(u_k)) + \sigma(u_k - \log \tau) \sigma(u_k) \tau^{-1} \\
 &= \sigma(u_k - \log \tau) \sigma(u_k) \tau^{-1} (\tau \exp(-u_k) + 1) = \sigma(u_k - \log \tau) \sigma(u_k) \tau^{-1} (\sigma(u_k - \log \tau))^{-1} \\
 &= \sigma(u_k) \tau^{-1} \stackrel{!}{=} p_\tau.
 \end{aligned}$$

Here we used the definition of the logistic function  $\sigma(x) = (1 + \exp(-x))^{-1}$  and  $\sigma(x)(1 - \sigma(x))^{-1} = \exp(x)$ . Now we show  $p_0 = \sum_{j=0}^{\tau} T_{0j}^k p_j$ :

## B Appendix to chapter 2: Neural dynamics as sampling

$$\begin{aligned}
\sum_{j=0}^{\tau} T_{0j}^k p_j &= T_{00}^k p_0 + T_{01}^k p_1 = (1 - \sigma(u_k - \log \tau))(1 - \sigma(u_k)) + (1 - \sigma(u_k - \log \tau))\sigma(u_k)\tau^{-1} \\
&= (1 - \sigma(u_k - \log \tau))(1 - \sigma(u_k)) \left(1 + \exp(u_k)\tau^{-1}\right) \\
&= \sigma(-u_k + \log \tau)(1 - \sigma(u_k))(\sigma(-u_k + \log \tau))^{-1} \\
&= 1 - \sigma(u_k) \stackrel{!}{=} p_0.
\end{aligned}$$

Here we used  $1 - \sigma(x) = \sigma(-x)$ .

It is trivial to show  $p_i = \sum_{j=0}^{\tau} T_{ij}^k p_j$  for  $i = 1, \dots, \tau - 1$  as  $\sum_{j=0}^{\tau} T_{ij}^k p_j = T_{i,i+1}^k p_{i+1} = p_{i+1} = p_i$ . Here we used the facts that  $T_{i,i+1}^k = 1$  and  $p_i = p_{i+1}$  for  $i = 1, \dots, \tau - 1$  by definition. □

**Lemma 3.** For all  $k = 1, \dots, K$  the operator  $T^k(z, \zeta | \zeta', z')$  leaves the distribution  $p(\zeta, z)$  invariant.

*Proof.* We start from Lemma 2, which states that  $T^k(\zeta_k | \zeta'_k, z'_{\setminus k})$  leaves the conditional distribution  $p(\zeta_k | z'_{\setminus k})$  invariant:

$$\begin{aligned}
&\sum_{\zeta'_k} T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(\zeta'_k | z'_{\setminus k}) = p(\zeta_k | z'_{\setminus k}) \\
\Leftrightarrow &\sum_{\zeta'_k, z'_{\setminus k}} \delta(z_{\setminus k}, z'_{\setminus k}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(\zeta'_k | z'_{\setminus k}) = \sum_{z'_{\setminus k}} \delta(z_{\setminus k}, z'_{\setminus k}) p(\zeta_k | z'_{\setminus k}) = p(\zeta_k | z_{\setminus k}) \\
\Leftrightarrow &\sum_{\zeta'_k, z'_{\setminus k}} \delta(z_k, \zeta_k^{>0}) \delta(z_{\setminus k}, z'_{\setminus k}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(\zeta'_k | z'_{\setminus k}) = \delta(z_k, \zeta_k^{>0}) p(\zeta_k | z_{\setminus k}) = p(z_k | \zeta_k) p(\zeta_k | z_{\setminus k}) \\
\Leftrightarrow &\sum_{\zeta'_k, z'} \delta(z_k, \zeta_k^{>0}) \delta(z_{\setminus k}, z'_{\setminus k}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(z'_k, \zeta'_k | z'_{\setminus k}) = p(z_k, \zeta_k | z_{\setminus k}) \\
\Leftrightarrow &\sum_{\zeta'_k, z'} \delta(z_k, \zeta_k^{>0}) \delta(z_{\setminus k}, z'_{\setminus k}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(z'_k, \zeta'_k | z'_{\setminus k}) p(\zeta_{\setminus k} | z'_{\setminus k}) p(z'_{\setminus k}) = p(z_k, \zeta_k | z_{\setminus k}) p(\zeta_{\setminus k} | z_{\setminus k}) p(z_{\setminus k}) \\
\Leftrightarrow &\sum_{\zeta'_k, z'} \delta(z_k, \zeta_k^{>0}) \delta(z_{\setminus k}, z'_{\setminus k}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) p(\zeta'_k, \zeta_{\setminus k}, z') = p(\zeta, z) \\
\Leftrightarrow &\sum_{\zeta', z'} T^k(z_k, \zeta_k | \zeta', z') \delta(\zeta_{\setminus k}, \zeta'_{\setminus k}) \delta(z_{\setminus k}, z'_{\setminus k}) p(z', \zeta') = p(\zeta, z) \\
\Leftrightarrow &\sum_{\zeta', z'} T^k(z, \zeta | \zeta', z') p(z', \zeta') = p(\zeta, z).
\end{aligned}$$

Here we used the relations  $\delta(z_k, \zeta_k^{>0}) = p(z_k | \zeta_k)$  and  $p(\zeta_k, z_k | z_{\setminus k}) = p(z_k | \zeta_k) p(\zeta_k | z_{\setminus k})$  as well as  $p(\zeta_k | z_{\setminus k}) = p(\zeta_k | \zeta_{\setminus k}, z_{\setminus k})$  which directly follow from the definitions of  $T^k(\zeta, z, | \zeta', z')$  and  $p(\zeta, z)$ .

□

Finally, we can verify that the composed operator  $T = T^1 \circ \dots \circ T^K$  samples from the given distribution  $p$ .

**Theorem 1.**  $p(\zeta, z)$  is the unique invariant distribution of operator  $T$ .

*Proof.* As all  $T^k$  leave  $p(\zeta, z)$  invariant, so does the concatenation  $T = T^1 \circ \dots \circ T^K$ . To ensure that  $p(\zeta, z)$  is the *unique* invariant distribution, we have to show that  $T$  is irreducible and aperiodic.  $T$  is aperiodic as the transition probabilities  $T_{00}^k = 1 - \sigma(u_k - \log \tau) > 0$  and  $T_{00}^k < 1$  (this follows from the assumption  $\forall z p(z) \neq 0$  made above).

The operator  $T$  is also irreducible for the following reason. First we see that from any state  $(\zeta', z')$  in at most  $\tau$  steps we can get to the zero-state  $(\zeta, z) = 0^{2K}$  (and stay there) with non-zero probability, as  $T_{i,i+1}^k = 1$  for  $i = 1, \dots, \tau - 1$  and  $T_{01}^k = 1 - \sigma(u_k - \log \tau) > 0$ . Furthermore, it can be seen that any state  $(\hat{\zeta}, \hat{z})$  can be reached from the zero-state  $(\zeta, z) = 0^{2K}$  in at most  $\tau$  steps since  $T_{N0}^k = \sigma(u_k - \log \tau) > 0$  for any value of  $u_k$ . Hence every final state  $(\hat{\zeta}, \hat{z})$  can be reached from every starting state  $(\zeta', z')$  in at most  $2\tau$  steps with non-vanishing probability.

□

### Details to neural sampling with a relative refractory period in discrete time

We augment the neuron model with a relative refractory period described by a function  $g(\zeta_k)$ . We first ensure existence of the corresponding function  $f(u_k)$ . Based on these functions we then introduce the transition operator  $T$  of the Markov chain. This operator is shown to entail correct “local” computations.

**Lemma 4.** Let  $(g_1, \dots, g_\tau) \in (\mathbb{R}_0^+)^{\tau}$  be a tuple of non-negative real numbers, with  $g_\tau = 0$  and at least one element  $g_i \geq 1$ . This defines the refractory function via  $g(\zeta_k) := g_{\zeta_k}$ . There exists a unique  $C^\infty$  function  $f : \mathbb{R} \rightarrow (0, 1)$  with the following property  $\forall u \in \mathbb{R}$ :

$$f(u) \frac{\sum_{i=1}^{\tau} \prod_{j=i+1}^{\tau} (1 - g_j f(u))}{\prod_{j=1}^{\tau} (1 - g_j f(u))} = \exp(u). \quad (\text{B.1})$$

Furthermore, the function  $f$  has the property:

$$\begin{aligned} \forall i \in \{1, \dots, \tau\} \forall u \in \mathbb{R} : 0 \leq g_i f(u) < 1 \\ \exists i \in \{1, \dots, \tau\} \forall u \in \mathbb{R} : 0 < g_i f(u) < 1. \end{aligned}$$

## B Appendix to chapter 2: Neural dynamics as sampling

*Proof.* Let  $g_{\max} := \max_{j \in \{1, \dots, \tau\}} g_j$ ; we know that  $g_{\max} \geq 1$ . We define the function  $F : (0, 1/g_{\max}) \rightarrow \mathbb{R}^+$ :

$$F(x) := x \sum_{i=1}^{\tau} \left( \frac{1}{\prod_{j=1}^i (1 - g_j x)} \right)$$

We can see that  $F$  is a positive  $\mathcal{C}^\infty$  function on  $(0, 1/g_{\max})$ . Furthermore,  $F(x)/x$  is defined as a sum of functions of the form  $\frac{1}{\prod_{j=1}^i (1 - g_j x)}$ . Each factor  $1/(1 - g_j x)$  is positive and strictly monotonous. Therefore,  $F$  is strictly monotonous on  $(0, 1/g_{\max})$  with the limits:

$$\begin{aligned} \lim_{x \rightarrow 0} F(x) &= 0 \\ \lim_{x \rightarrow 1/g_{\max}} F(x) &= \infty. \end{aligned}$$

Hence the equation  $F(x) = \exp(u)$  has a unique solution for  $x$  called  $f(u) \in (0, 1/g_{\max})$  for all  $u \in \mathbb{R}$ . From applying the implicit function theorem to  $F(x, u) := F(x) - \exp(u)$  it follows that  $f$  is  $\mathcal{C}^\infty$ .

□

From here on, with the letter  $f$  we will denote the function characterized by the above Lemma for the given tuple  $g$  (which denotes the chosen refractory function).

**Definition 1.** Define  $g_0 = 1$ . The transition operator  $T^k$  is defined in the following way for all  $k = 1, \dots, K$ :

$$T^k(\zeta_k, z_k | \zeta', z') := \delta(z_k, \zeta_k^{>0}) T^k(\zeta_k | \zeta'_k, z'_{\setminus k})$$

$$T^k(\zeta_k | \zeta'_k, z'_{\setminus k}) := \begin{cases} g_{\zeta'_k} f(u_k) & \text{for } \zeta_k = \tau \\ 1 - g_{\zeta'_k} f(u_k) & \text{for } \zeta_k = \zeta'_k - 1 \wedge \zeta'_k > 0 \\ 1 - f(u_k) & \text{for } \zeta_k = 0 \wedge \zeta'_k = 0 \\ 0 & \text{otherwise} \end{cases},$$

with  $u_k = u_k(z'_{\setminus k})$ .

**Lemma 5.** For all  $k = 1, \dots, K$  the unique invariant distribution  $q^*(z_k, \zeta_k | \zeta'_{\setminus k}, z'_{\setminus k})$  of the operator  $T^k(z_k, \zeta_k | \zeta', z')$  fulfills  $\sum_{\zeta_k} q^*(z_k, \zeta_k | \zeta'_{\setminus k}, z'_{\setminus k}) = p(z_k | z'_{\setminus k})$ . This means, for a constant configuration  $z'_{\setminus k}$ , the operator  $T^k$  produces samples  $z_k^*$  from the correct conditional distribution  $p(z_k | z'_{\setminus k})$ .

*Proof.* We define:

## B.1 Derivation of the neural sampling theory

$$q^*(z_k, \zeta_k | \zeta'_{\setminus k}, z'_{\setminus k}) := \delta(z_k, \zeta_k^{>0}) q(\zeta_k | z'_{\setminus k}) := \delta(z_k, \zeta_k^{>0}) \left( \sigma(u_k) h(\zeta_k | z'_{\setminus k}) + (1 - \sigma(u_k)) \delta(\zeta_k, 0) \right),$$

where the function  $h(\zeta_k | z'_{\setminus k})$  is defined as:

$$h(\zeta_k | z'_{\setminus k}) := \begin{cases} \frac{\prod_{j=\zeta_k+1}^{\tau} (1 - g_j f(u_k))}{\sum_{\alpha=1}^{\tau} \prod_{j=\alpha+1}^{\tau} (1 - g_j f(u_k))} & \text{for } \zeta_k > 0 \\ 0 & \text{otherwise} \end{cases}.$$

It is trivial to see that  $q^*$  has the correct marginal distribution over  $z_k$ :

$$\begin{aligned} \sum_{\zeta_k} q^*(z_k, \zeta_k | \zeta'_{\setminus k}, z'_{\setminus k}) &= \sum_{\zeta_k} \delta(z_k, \zeta_k^{>0}) \left( \sigma(u_k) h(\zeta_k | z'_{\setminus k}) + (1 - \sigma(u_k)) \delta(\zeta_k, 0) \right) \\ &= \sigma(u_k)^{z_k} (1 - \sigma(u_k))^{1-z_k} = p(z_k | z'_{\setminus k}). \end{aligned}$$

We now show that  $q^*$  is the unique invariant distribution of  $T^k$ . Because of the definition of  $T^k$ , we only have to show that  $q^*(\zeta_k | z'_{\setminus k})$  is the unique invariant distribution of  $T^k(\zeta_k | \zeta'_{\setminus k}, z'_{\setminus k})$ . We denote  $q^*(\zeta_k = i | z'_{\setminus k}) =: q_i$  and  $T^k(\zeta_k = i | \zeta'_k = j, z'_{\setminus k}) =: T_{ij}$ , i.e., we have to show  $\forall i \in \{0, 1, \dots, \tau\} : q_i = \sum_j T_{ij} q_j$ .

It is trivial to show  $q_i = \sum_j T_{ij} q_j$  for  $1 \leq i \leq \tau - 1$ , as there is only one non-vanishing element of transition operator, namely  $T_{i,i+1}$ :

$$\begin{aligned} \sum_{j=0}^{\tau} T_{ij} q_j &= T_{i,i+1} q_{i+1} = (1 - g_{i+1} f(u_k)) q_{i+1} \\ &= (1 - g_{i+1} f(u_k)) h(\zeta_k = i + 1 | z_{\setminus k}) \sigma(u_k) \\ &= h(\zeta_k = i | z_{\setminus k}) p(z_k = 1 | z_{\setminus k}) \stackrel{!}{=} q_i. \end{aligned}$$

Here we used  $q_i = h(\zeta_k = i | z_{\setminus k}) \sigma(u_k)$  for  $i > 0$  and the definition of  $h(\zeta_k | z_{\setminus k})$ .

Now we show  $q_0 = \sum_j T_{0j} q_j$  starting from equation (B.1) and additionally using the relations  $\exp(u_k) = \sigma(u_k) / (1 - \sigma(u_k))$  and  $q_0 = 1 - \sigma(u_k)$  as well as the definition of  $q_1$ . We define for the sake of simplicity  $\psi := \sum_{\alpha=1}^{\tau} \prod_{j=\alpha+1}^{\tau} (1 - g_j f(u_k))$ :

$$\begin{aligned} \sum_{j=0}^{\tau} T_{0j} q_j &= (1 - f(u_k)) q_0 + (1 - g_1 f(u_k)) q_1 \\ &= (1 - f(u_k)) (1 - \sigma(u_k)) + \frac{\sigma(u_k)}{\psi} \prod_{j=1}^{\tau} (1 - g_j f(u_k)) \\ &= (1 - f(u_k)) (1 - \sigma(u_k)) + \sigma(u_k) f(u_k) \exp(-u_k) \\ &= (1 - f(u_k)) (1 - \sigma(u_k)) + f(u_k) (1 - \sigma(u_k)) \stackrel{!}{=} q_0. \end{aligned}$$

We finally show  $q_{\tau} = \sum_j T_{\tau j} q_j$ , using the definition of  $q_{\tau} = \sigma(u_k) h(\zeta_k = \tau | z_{\setminus k}) = \frac{\sigma(u_k)}{\psi}$ :

## B Appendix to chapter 2: Neural dynamics as sampling

$$\begin{aligned}
\sum_{i=0}^{\tau} T_{\tau i} q_i &= \sum_{i=1}^{\tau} g_i f(u_k) q_i + f(u_k) q_0 \\
&= \sum_{i=1}^{\tau} g_i f(u_k) \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) \frac{\sigma(u_k)}{\psi} + f(u_k) q_0 \\
&= \frac{\sigma(u_k)}{\psi} \left( \sum_{i=1}^{\tau} g_i f(u_k) \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) + (1 - g_1 f(u_k)) \prod_{j=2}^{\tau} (1 - g_j f(u_k)) \right) \\
&= \frac{\sigma(u_k)}{\psi} \left( - \sum_{i=1}^{\tau} (1 - g_i f(u_k)) \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) + \sum_{i=1}^{\tau} \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) + \prod_{j=1}^{\tau} (1 - g_j f(u_k)) \right) \\
&= \frac{\sigma(u_k)}{\psi} \left( - \sum_{i=1}^{\tau} \prod_{j=i}^{\tau} (1 - g_j f(u_k)) + \sum_{i=0}^{\tau} \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) \right) \\
&= \frac{\sigma(u_k)}{\psi} \left( - \sum_{i=0}^{\tau-1} \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) + \sum_{i=0}^{\tau} \prod_{j=i+1}^{\tau} (1 - g_j f(u_k)) \right) \\
&= \frac{\sigma(u_k)}{\psi} \left( \prod_{j=\tau+1}^{\tau} (1 - g_j f(u_k)) \right) = \frac{\sigma(u_k)}{\psi} \stackrel{!}{=} q_{\tau}.
\end{aligned}$$

The argument that the transition operator  $T^k$  is aperiodic and irreducible is similar to the one presented in Lemma 1.

□

### Details to neural sampling with an absolute refractory period in continuous time

In contrast to the discrete time model we define the state space of  $\zeta_k$  to be  $\mathbb{R}^+ \cup [-2\epsilon, -\epsilon]$  for  $\epsilon > 0$ , i.e., as the union of the positive real numbers and a small interval  $[-2\epsilon, -\epsilon]$ . We will define the sampling operator in such a way that after neuron  $k$  was refractory for exactly its refractory period  $\tau$ , its refractory variable  $\zeta_k$  is uniformly placed in the small interval  $[-2\epsilon, -\epsilon]$ , which represents now the resting state and replaces  $\zeta_k = 0$ . This avoids point measures (Dirac's Delta) on the value  $\zeta_k = 0$ . This system is still exactly equivalent to the system discussed in the main paper, as all spike-transition probabilities of  $T$  for  $\zeta_k < 0$  are constant. Hence, it does not matter which values  $\zeta_k$  assumes with respect to the spike mechanism during its non-refractory period as long as  $\zeta_k < 0$ .

**Definition 2.** For a given distribution  $p(z)$  over the binary variables  $z \in \{0, 1\}^K$  with  $\forall z \in \{0, 1\}^K p(z) \neq 0$ , we define a joint distribution over  $(\zeta, z)$  with  $\zeta \in \mathbb{R}^K$  in the

following way:

$$p(\zeta_k|z_k) := \begin{cases} 1 & \text{for } 1 \geq \zeta_k > 0 \wedge z_k = 1 \\ \epsilon^{-1} & \text{for } \zeta_k \in I_\epsilon \wedge z_k = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$p(\zeta|z) := \prod_{k=1}^K p(\zeta_k|z_k)$$

$$p(\zeta, z) := p(\zeta|z)p(z),$$

where  $I_\epsilon := [-2\epsilon, -\epsilon]$  is the refractory resting state interval. In accordance with this definition we can also write  $p(\zeta_k|z_k) = z_k \chi_{[0,1]}(\zeta_k) + (1 - z_k) \epsilon^{-1} \chi_{I_\epsilon}(\zeta_k)$ .

**Lemma 6.** The distribution  $p(\zeta, z)$  has the following marginal distribution:

$$p(\zeta_k|\zeta_{\setminus k}) = \sigma(u_k) \chi_{[0,1]}(\zeta_k) + (1 - \sigma(u_k)) \epsilon^{-1} \chi_{I_\epsilon}(\zeta_k)$$

$$= \begin{cases} \sigma(u_k) & \text{for } 1 \geq \zeta_k > 0 \\ (1 - \sigma(u_k)) \epsilon^{-1} & \text{for } \zeta_k \in I_\epsilon \end{cases},$$

where  $u_k := u_k(\zeta_{\setminus k}^{>0})$ .

**Definition 3.** For  $k \in \{1, \dots, K\}$  and  $x \in \mathbb{R}$  the operator  $T_x^k$  is defined in the following way for a function  $q : \mathbb{R} \rightarrow \mathbb{R}$ :

$$(T_x^k q)(\zeta_k) := \tau^{-1} \left( \partial_{\zeta_k} (q(\zeta_k) \chi_{\mathbb{R}^+}(\zeta_k)) - \delta(\zeta_k) F(q) + \exp(x) \delta(\zeta_k - 1) \int_{I_\epsilon} q(\zeta'_k) d\zeta'_k \right. \\ \left. + \chi_{I_\epsilon}(\zeta_k) \left( \epsilon^{-1} F(q) - \exp(x) q(\zeta_k) \right) \right).$$

where the functional  $F$  is defined as the one-sided limit from above at 0:

$$F(q) := \lim_{x \rightarrow 0^+} q(x).$$

The operator  $T$  is defined in the following way for a probability distribution  $q(\zeta)$  on  $\mathbb{R}^K$ :

$$(Tq)(\zeta) := \sum_{k=1}^K (T_{u_k}^k q(\zeta_1, \dots, \zeta_{k-1}, \cdot, \zeta_{k+1}, \zeta_K))(\zeta_k),$$

where  $q(\zeta_1, \dots, \zeta_{k-1}, \cdot, \zeta_{k+1}, \zeta_K) : \mathbb{R} \rightarrow \mathbb{R}$  denotes the function  $q(\zeta)$  of  $\zeta_k$  where  $\zeta_{\setminus k}$  is held constant and  $u_k := u_k(\zeta_{\setminus k}^{>0})$ .

The transition operator  $T$  defines the following Fokker-Planck equation for a time-dependent distribution  $q_t(\zeta)$ :

## B Appendix to chapter 2: Neural dynamics as sampling

$$\partial_t q_t(\zeta) = (Tq_t)(\zeta).$$

The jump and drift functions  $W^k(\zeta|\zeta')$  and  $A^k(\zeta)$  associated to the operator  $T$  are given by:

$$\begin{aligned} W^k(\zeta|\zeta') &= \left( (\epsilon\tau)^{-1} \chi_{I_\epsilon}(\zeta_k) \delta(\zeta'_k) + \delta(\zeta_k - 1) \exp(u_k(\zeta'_k) - \log \tau) \chi_{I_\epsilon}(\zeta'_k) \right) \delta(\zeta_{\setminus k} - \zeta'_{\setminus k}) \\ A^k(\zeta) &= -\tau^{-1} \chi_{\mathbb{R}^+}(\zeta_k) \\ \Rightarrow (Tq_t)(\zeta) &= -\sum_{k=1}^K \partial_{\zeta_k} (A^k(\zeta) q_t(\zeta)) + \sum_{k=1}^K \int \left( W^k(\zeta|\zeta') p(\zeta') - W^k(\zeta'|\zeta) p(\zeta) \right) d\zeta'. \end{aligned}$$

**Lemma 7.** *The operator  $T_{u_k}^k$  leaves the conditional distribution  $p(\zeta_k|\zeta_{\setminus k})$  invariant with  $u_k = u_k(\zeta_{\setminus k}^{>0})$ , i.e.:*

$$(T_{u_k}^k p(\cdot|\zeta_{\setminus k}))(\zeta_k) = 0.$$

*Proof.* This is easy to proof using calculus and the relations  $\partial_{\zeta_k} \chi_{\mathbb{R}^+}(\zeta_k) = \delta(\zeta_k)$  and  $F(p(\cdot|\zeta_{\setminus k})) = \sigma(u_k) = \exp(u_k)(1 - \sigma(u_k))$ .

□

**Lemma 8.**  *$p(\zeta)$  is an invariant distribution of  $T$ , i.e., it is a solution to the invariant Fokker-Planck equation:*

$$\partial_t p(\zeta) = (Tp)(\zeta) = 0.$$

*Proof.* We observe that  $T^k(\alpha p) = \alpha T^k p$  for a constant  $\alpha \in \mathbb{R}$  (which is not a function of  $\zeta_k$ ). Hence:

$$\begin{aligned} T_{u_k}^k p(\zeta_1, \dots, \zeta_{k-1}, \cdot, \zeta_{k+1}, \dots, \zeta_K) &= T_{u_k}^k (p(\cdot|\zeta_{\setminus k}) p(\zeta_{\setminus k})) \\ &= p(\zeta_{\setminus k}) (T_{u_k}^k p(\cdot|\zeta_{\setminus k})) \\ &= 0. \end{aligned}$$

The Lemma follows then from the definition of  $T := \sum_k T_{u_k}^k$ .

□



### Details to neural sampling with a relative refractory period in continuous time

As already assumed in the case of the absolute refractory sampler in continuous time, we define the state space of  $\zeta_k$  to be  $\mathbb{R}^+ \cup [-2\epsilon, -\epsilon]$  for  $\epsilon > 0$ .

**Lemma 9.** *Let  $g$  be a continuous, non-negative function  $g : [0, 1] \rightarrow \mathbb{R}_0^+$  with  $g(\zeta_k) = 1$  for  $\zeta_k \leq 0$ . There exists a unique  $C^\infty$  function  $f : \mathbb{R} \rightarrow \mathbb{R}^+$  with the following property  $\forall u \in \mathbb{R}$ :*

$$f(u) \int_0^1 \exp\left(f(u) \int_0^{\zeta_k} g(\zeta'_k) d\zeta'_k\right) d\zeta_k = \exp(u). \quad (\text{B.2})$$

*Proof.* We define the function  $F : \mathbb{R}_0^+ \rightarrow \mathbb{R}$  in the following way:

$$F(x) := x \int_0^1 \exp(x\alpha(\zeta_k)) d\zeta_k,$$

where  $\alpha(r) := \int_0^r g(\zeta'_k) d\zeta'_k$ . From  $g(\zeta_k) \geq 0$  we can follow that  $\alpha : [0, 1] \rightarrow \mathbb{R}_0^+$  is non-negative.  $F(x)$  is differentiable with the derivative:

$$\begin{aligned} F'(x) &= \int_0^1 \exp(x\alpha(\zeta_k)) d\zeta_k + x \int_0^1 \exp(x\alpha(\zeta_k)) \alpha(\zeta_k) d\zeta_k \\ &\Rightarrow F'(x) > 0. \end{aligned}$$

Hence  $F$  is strictly monotonously increasing. Furthermore, the following relations hold:

$$\begin{aligned} F(0) &= 0 \\ F(x) &\geq x. \end{aligned}$$

Therefore the equation:

$$F(x) = \exp(u),$$

has exactly one solution  $f(u)$  with  $F(f(u)) = \exp(u)$  in  $\mathbb{R}^+$ . From applying the implicit function theorem to  $F(x, u) := F(x) - \exp(u)$  it follows that  $f$  is  $C^\infty$ .

□

**Definition 4.** *For all  $k \in \{1, \dots, K\}$  and  $x \in \mathbb{R}$  the operator  $T_x^k$  is defined in the following way for a function  $q : \mathbb{R} \rightarrow \mathbb{R}$ :*

$$\begin{aligned} (T_x^k q)(\zeta_k) &:= \tau^{-1} \left( \partial_{\zeta_k} (q(\zeta_k) \chi_{\mathbb{R}^+}(\zeta_k)) - \delta(\zeta_k) q(\zeta_k) + f(x) \delta(\zeta_k - 1) \int_{\mathbb{R}} g(\zeta'_k) q(\zeta'_k) d\zeta'_k \right. \\ &\quad \left. + \chi_{I_\epsilon}(\zeta_k) \epsilon^{-1} F(q) - f(x) q(\zeta_k) g(\zeta_k) \right). \end{aligned}$$

The transition operator  $T_x^k$  defines the following Fokker-Planck equation for a time-dependent distribution  $q_t(\zeta_k)$ :

$$\partial_t q_t(\zeta_k) = (T_x^k q_t)(\zeta_k).$$

The jump and drift functions  $W^k(\zeta_k|\zeta'_k)$  and  $A^k(\zeta_k)$  associated to the operator  $T_x^k$  are given by:

$$\begin{aligned} W^k(\zeta_k|\zeta'_k) &= (\epsilon\tau)^{-1}\chi_{I_\epsilon}(\zeta_k)\delta(\zeta'_k) + \tau^{-1}\delta(\zeta_k - 1)f(x)g(\zeta'_k) \\ A^k(\zeta_k) &= -\tau^{-1}\chi_{\mathbb{R}^+}(\zeta_k) \\ \Rightarrow (T_x^k q_t)(\zeta_k) &= -\partial_{\zeta_k}(A^k(\zeta_k)q_t(\zeta_k)) + \int \left( W^k(\zeta_k|\zeta'_k)p(\zeta'_k) - W^k(\zeta'_k|\zeta_k)p(\zeta_k) \right) d\zeta'_k. \end{aligned}$$

**Lemma 10.** For all  $k = 1, \dots, K$  the invariant distribution  $q^*(\zeta_k|z_{\setminus k})$  of the operator  $T_{u_k}^k$  fulfills  $\int \delta(z_k, \zeta_k^{>0})q^*(\zeta_k|z_{\setminus k})d\zeta_k = p(z_k|z_{\setminus k})$ .

*Proof.* We define the distribution  $q^*(\zeta_k|z_{\setminus k})$  as:

$$q^*(\zeta_k|z_{\setminus k}) = (1 - \sigma(u_k)) \left( f(u_k)\chi_{[0,1]}(\zeta_k) \exp(f(u_k)\alpha(\zeta_k)) + \epsilon^{-1}\chi_{I_\epsilon}(\zeta_k) \right),$$

where  $\alpha(\zeta_k) := \int_0^1 g(\zeta'_k)d\zeta'_k$ . By applying the operator  $T_{u_k}^k$  to  $q^*$  one can verify that  $T_{u_k}^k q^* = 0$  holds using the definition of  $f(u_k)$  given in (B.2). Furthermore we can compute the ratio:

$$\begin{aligned} \frac{\int_0^1 q^*(\zeta_k|z_{\setminus k})d\zeta_k}{\int_{I_\epsilon} q^*(\zeta_k|z_{\setminus k})d\zeta_k} &= \frac{p(z_k = 1|z_{\setminus k})}{p(z_k = 0|z_{\setminus k})} \\ &= f(u_k) \int_0^1 \exp\left(f(u_k) \int_0^{\zeta_k} g(\zeta'_k)d\zeta'_k\right) d\zeta_k = \exp(u_k). \end{aligned}$$

□

## B.2 Details to the computer simulations

The simulation results shown in Fig. 2.2, Fig. 2.3 and Fig. 2.4 used the biologically more realistic neuron model with the relative refractory mechanism. During all experiments the first second of simulated time was discarded as burn-in time. The full list of parameters defining the experimental setup is given in Table B.1. All

## B.2 Details to the computer simulations

Description	Variable	Value	Figure	Comment
<i>Simulation Time</i>				
Simulation step size	$dt$	1 ms	2-7	interpretation of an MCMC step
Burn-in time	$t_{\text{burn}}$	1 s	2-7	before recording spikes
Simulation time	$t_{\text{sim}}$	0.5 s	2	
		$10^4$ s	3,5-7	
		20 s	4	$10^4$ s for Figure 2.4C
<i>Network</i>				
Number of neurons	$K$	3	2	unconnected
		40	3,5,6	randomly connected
		217	4	
Connection radius		10	7	100 networks
		0	2	
		$\infty$	3,5-7	
		8	4	
Recurrent weights	$W_{ki}$	$\mathcal{N}(0, 0.3^2)$	3,5-7	from Gaussian distribution
Falling edge	$\tau_+$	20 ms	6,7	for realistic PSP shapes
Rising edge	$\tau_-$	3 ms	6,7	
Scaling factor	$\lambda$	20/17	6,7	
<i>Neuron Model</i>				
Number recovery steps	$\tau$	20	2-7	PSP duration = $\tau \cdot dt = 20$ ms
Refractory function	$g(\zeta)$	$\left[4(1 - \zeta) + \frac{1}{2\pi} \sin(8\pi\zeta)\right]$	2↑	normalized to $\zeta \in [0, 1]$ ,
		$\left[1 - \zeta + \frac{1}{2\pi} \sin(2\pi\zeta)\right]$	2-7	$[x] := \min\{1, \max\{0, x\}\}$
		$\left[1 - 2\zeta + \frac{1}{2\pi} \sin(4\pi\zeta)\right]$	2↓,7	
Excitability	$b_k$	-1 or 2	2	defines membrane potential $u_k$
		$\mathcal{N}(-1.5, 0.5^2)$	3,5-7	from Gaussian distribution
		0	4	initial value
<i>Tuning Function, Training and Inference (Figure 4)</i>				
Peakedness	$\kappa$	3	4	measured: $1.78 \pm 0.15$
Base sensitivity	$v_0$	0.05	4	measured: $0.017 \pm 0.009$
Sensitivity contrast	$C$	0.9	4	measured: $0.760 \pm 0.020$
Training samples	$N_{\text{train}}$	$10^5$	4	
Decorrelation steps		20	4	for contrastive divergence
Learning rate	$\eta$	$10^{-4}$	4	
Number of neurons clamped on/off		4/4	4	

**Table B.1: List of parameters of the computer simulations.**

## B Appendix to chapter 2: Neural dynamics as sampling

occurring joint probability distributions are Boltzmann distributions of the form given in equation (2.5). Example Python (G. VAN ROSSUM and DRAKE JR, 1995) scripts for neural sampling from Boltzmann distributions are available on request and will be provided on our webpage. The example code comprises networks with both absolute and relative refractory mechanism. It requires standard Python packages only and is readily executable.

### Details to Figure 2: Neuron model with relative refractory mechanism

The three refractory functions  $g(\zeta)$  of panel (B) as well as all other simulation parameters are listed in Table B.1. Panel (C) shows the corresponding functions  $f(u)$ , which result from numerically solving equation (2.11). The spike patterns in panel (D) show the response of the neurons when the membrane potential is low ( $u_k = -1$  for  $0 < t < 250$  ms) or high ( $u_k = +2$  for  $250 \text{ ms} < t < 500$  ms). These membrane potentials encode  $p(z_k = 1) = 0.269$  and  $p(z_k = 1) = 0.881$ , respectively according to (2.3) and (2.4). The binary state  $z_k = 1$  is indicated by gray shaded areas of duration  $\tau \cdot dt = 20$  ms after each spike.

### Details to Figure 3: Sampling from a Boltzmann distribution by spiking neurons with relative refractory mechanism

We examined the spike response of a network of 40 randomly connected neurons which sampled from a Boltzmann distribution. The excitabilities  $b_k$  as well as the synaptic weights  $W_{ki}(= W_{ik})$  were drawn from Gaussian distributions (with diagonal elements  $W_{ii} = 0$ ). For the full list of parameters please refer to Table B.1. One second of the arising spike pattern is shown in panel (A). The average firing rate of the network was 13.9 Hz. To highlight the internal dynamics of the neuron model, the values of the refractory function  $g(\zeta_{26})$ , the membrane potential  $u_{26}$  and the instantaneous firing rate  $r_{26}$  of neuron  $v_{26}$  (indicated with red spikes) are shown in panel (B). Here, the instantaneous firing rate  $r_{26}$  is defined for the discrete time Markov chain as

$$r_{26} = p(\text{spike})/dt = T^{26}(\tau|\zeta_{26}, z_{\setminus 26})/dt = g(\zeta_{26}) \cdot f(u_{26})/dt . \quad (\text{B.3})$$

As stated before, the neuron model with relative refractory mechanism  $g_k(\zeta)$  does not entail the correct overall invariant distribution  $p(\mathbf{z})$ . To estimate the impact of this approximation on the joint network dynamics, we compared the distribution  $p(z_{24}, \dots, z_{28})$  over five neurons (indicated by gray background in A) in the spiking network with the correct distribution obtained from Gibbs sampling. The probabilities were estimated from  $10^7$  samples. A more quantitative analysis of the approximation quality of neural sampling with a relative refractory mechanism is provided below.

### Details to Figure 4: Modeling perceptual multistability as probabilistic inference with neural sampling

We demonstrate probabilistic inference and learning in a network of orientation selective neurons. As a simple model we consider a network of 217 neurons on a hexagonal grid as shown in panel (F). Any two neurons with distance  $\leq 8$  were synaptically connected (neighboring units had distance 1). For the remaining parameters of the network and neuron model please refer to Table B.1. Each neuron featured a  $\pi$ -periodic tuning curve as depicted in panel (B):

$$V_k(\varphi) = v_0 + C \cdot \exp[\kappa \cdot \cos(2(\varphi - \bar{\varphi}_k)) - \kappa] \quad (\text{B.4})$$

with base sensitivity  $v_0$ , contrast  $C$ , peakedness  $\kappa$  and preferred orientation  $\bar{\varphi}_k$ . The preferred orientations  $\bar{\varphi}_k$  of the neurons were chosen to cover the entire interval  $[0, \pi)$  of possible orientations with equal spacing and were randomly assigned to the neurons.

For simplicity we did not incorporate the input dynamics in our probabilistic model, but rather trained the network directly like a fully visible Boltzmann machine. We used for this purpose a standard Boltzmann machine learning rule known as contrastive divergence (G. HINTON, 2002; G. E. HINTON, 2010). This learning rule requires posterior samples  $\tilde{z}$ , i.e., network states under the influence of the present input, and approximate prior samples  $z^*$ , which reflect the probability distribution of the network in the absence of stimuli. The update rules for synaptic weights and neuronal excitabilities read:

$$\begin{aligned} \Delta W_{ki} &= \eta_{ki} \cdot (\tilde{z}_k \tilde{z}_i - z_k^* z_i^*) \\ \Delta b_k &= \eta \cdot (\tilde{z}_k - z_k^*) \\ \eta_{ki} &= \begin{cases} \eta & \text{if } v_k \text{ and } v_i \text{ are connected} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (\text{B.5})$$

While more elaborate policies can speed up convergence, we simply used a global learning rate  $\eta$  which was constant in time. The values of  $W_{ki}$  and  $b_k$  were initialized at 0. We generated binary training patterns in the following way:

1. A global orientation  $\varphi$  was drawn uniformly from  $[0, \pi)$ ,
2. each neuron was independently set to be active with probability  $p(z_k = 1) = V_k(\varphi)$ ,
3. the resulting network state  $\tilde{z}$  was taken as posterior sample.

To obtain an approximate prior sample  $z^*$  we let the network run for a short time freely starting from  $(\tilde{\zeta}, \tilde{z})$ . The variables  $\tilde{\zeta}$  were also assumed to be observed with  $\tilde{\zeta}_k \sim \text{iid. uniformly in } \{1, \dots, \tau\}$  if  $\tilde{z}_k = 1$  and  $\tilde{\zeta}_k = 0$  otherwise. After evolving freely for 20 time steps, the resulting network state  $z^*$  was taken as approximate prior sample and  $\mathbf{W}$  and  $\mathbf{b}$  were updated according to (B.5). This process was repeated  $N_{\text{train}} = 10^5$  times. As a result, neurons with similar preferred orientations featured excitatory synaptic connections ( $W_{ki} = 6.4 \cdot 10^{-3} \pm 6.7 \cdot 10^{-3} = \text{mean}$

## B Appendix to chapter 2: Neural dynamics as sampling

$\pm$  standard deviation of weight distribution), those with dissimilar orientations maintained inhibitory synapses ( $W_{ki} = -4.9 \cdot 10^{-3} \pm 5.2 \cdot 10^{-3}$ ). Here, preferred orientations  $\bar{\varphi}_i$  and  $\bar{\varphi}_j$  are defined as similar if  $V_i(\bar{\varphi}_j) - v_0 = V_j(\bar{\varphi}_i) - v_0 > 0.5 C$ , otherwise they are dissimilar. Neuronal biases converged to  $b_k = -0.08 \pm 0.03$ .

We illustrate the learned prior distribution  $p(z)$  of the network through sampled states when the network evolved freely. As seen in panel (D), the population vector – a 2-dimensional projection of the high dimensional network state – typically reflected an arbitrary, yet coherent, orientation (for the definition of the population vector see below). Each dot represents a sampled network state  $z$ .

To apply an ambiguous cue, we clamped 8 out of 217 neurons: Two units with  $\bar{\varphi}_k \approx \pi/4$  and two with  $\bar{\varphi}_k \approx 3\pi/4$  were set active, two units with  $\bar{\varphi}_k \approx 0$  and two with  $\bar{\varphi}_k \approx \pi/2$  were set inactive. This led to a bimodal posterior distribution as shown in panel (E). The sampling network represented this distribution by encoding either global perception separately: The trace of network states  $z(t)$  roamed in one mode for multiple steps before quickly crossing the state space towards the opposite percept.


We define the population vector  $\mathbf{x}$  of a network state  $z$  as a function of the preferred orientations of all active units:

$$\mathbf{x} = (x_0, x_{\pi/4}) = \sum_{k=1}^K z_k \cdot (\cos 2\tilde{\varphi}_k, \sin 2\tilde{\varphi}_k) . \quad (\text{B.6})$$

This definition of  $\mathbf{x}$  is not based on the preferred orientations  $\bar{\varphi}_k$  which are used for generating external input to the network from a given stimulus with orientation  $\varphi$ . It is rather based on the preferred orientations  $\tilde{\varphi}_k$  measured from the network response. We used population vector decoding based on the measured values  $\tilde{\varphi}_k$ , as they are conceptually closer to experimentally measurable preferred orientations, and this decoding hence does not require knowledge of the (unobservable)  $\bar{\varphi}_k$ . For every neuron  $\nu_k$  the preferred orientation  $\tilde{\varphi}_k$  was measured in the following way. We estimated a tuning curve  $\tilde{V}_k(\varphi)$  by a van-Mises fit (of the form (B.4)) to data from stimulation trials in which neuron  $\nu_k$  was not clamped, i.e., where  $\nu_k$  was only stimulated by recurrent input (feedforward input was modeled by clamping 8 out of 217 neurons as a function of stimulus orientation  $\varphi$  as before). Due to the structured recurrent weights, the experimentally measured tuning curves  $\tilde{V}_k(\varphi)$  were found to be reasonably close to the tuning curves  $V_k(\varphi)$  used for external stimulation.  $\tilde{\varphi}_k$  was set to the preferred orientation of  $\tilde{V}_k(\varphi)$  (localization parameter of the van-Mises fit). The measured values  $\tilde{\varphi}_k$  turned out to be consistent with the preferred orientations  $\bar{\varphi}_k$  ( $\bar{\varphi}_k - \tilde{\varphi}_k = 6 \cdot 10^{-4} \pm 8.3 \cdot 10^{-3}$  averaged over all  $K$  neurons). The mean and standard deviation of the remaining parameter values  $v_0$ ,  $C$  and  $\kappa$  of the fitted tuning curves  $\tilde{V}_k(\varphi)$  are listed in Table B.1 next to the ones used for stimulation.

The population vector  $\mathbf{x}$  was defined in (B.6) with the argument  $2\tilde{\varphi}_k$  (instead of  $\tilde{\varphi}_k$ ) as orthogonal orientations should cancel each other and neighborhood relations should be respected. For example neurons with  $\tilde{\varphi}_k = \epsilon$  and  $\tilde{\varphi}_k = \pi - \epsilon$  contribute

### B.3 Firing statistics of neural sampling networks



**Fig. B.1: Firing statistics of neural sampling networks.** (A) Shown is the membrane potential histogram of a typical neuron during sampling. The data is that of neuron  $v_{26}$  from the simulation shown in Fig. 2.3 (the membrane potential and spike trace of  $v_{26}$  are highlighted in Fig. 2.3). (B) The plot shows the ISI distribution of a typical neuron (again  $v_{26}$  from Fig. 2.3) during sampling. The distribution is roughly gamma-shaped, reminiscent of experimentally observed ISI distributions. (C) A scatter plot of the coefficient of variation (CV) versus the average interspike interval (ISI) of each neuron taken from the simulation shown in Fig. 2.3. The value of neuron  $v_{26}$  from Fig. 2.3 is marked by a cross. The simulated data is in accordance with experimentally observed data.

similarly to the population vector for small  $\epsilon$ . But counter to intuition the population vector of a state  $z$  with dominant orientation  $\varphi_z$  will point into direction  $\varphi_x = 2\varphi_z$ . For visualization in panel (D) and (E) we therefore rescaled the population vector: If  $(x_0, x_{\pi/4}) \mapsto (r_x, \varphi_x)$  in polar coordinates, then the dot is located at  $(r_x, \varphi_x/2)$  in accord with intuition. The black semicircles equal  $|\mathbf{x}| = r_x = 45$ .

The population vector  $(x_0, x_{\pi/4}) \in \mathbb{R}^2$  was also used for measuring the dominance durations shown in panel (C). To this  $\mathbb{R}^2$  was divided into 3 areas: (a)  $x_{\pi/4} < -35$ , (b)  $-35 \leq x_{\pi/4} \leq 35$ , (c)  $35 < x_{\pi/4}$ . We detected a perceptual switch when the network state entered area (a) or (c) while the previous perception was (c) or (a), respectively.

In panel (F) neurons  $v_k$  with  $z_k = 1$  are plotted with their preferred orientation color code, inactive neurons are displayed in white. Cells marked by a dot ( $\bullet$ ) were part of the observed variables  $\mathcal{o}$ . The three network states correspond to  $z(t_i)$  with  $t_1 = 100$  ms,  $t_2 = 250$  ms and  $t_3 = 400$  ms in the spike pattern in panel (G). The spike pattern shows the response of the freely evolving units around a perceptual switch during sampling from the posterior distribution. The corresponding trace of the population vector is drawn as black line in panel (E). The width of the light-gray shaded areas in the spike pattern equals the PSP duration  $\tau \cdot dt$ , i.e., neurons that spiked in these intervals were active in the corresponding state in (F).

### B.3 Firing statistics of neural sampling networks

In previous sections it was shown that a spiking neural network can draw samples from a given joint distribution which is in a well-defined class of probability distributions (see the neural computability condition (2.4)). Here, we examine some

statistics of individual neurons in a sampling network which are commonly used to analyze experimental data from recordings. The spike trains and membrane potential data are taken from the simulation presented in Fig. 2.3.

Fig. B.1A,B exemplarily show the distribution of the membrane potential  $u_k$  and the interspike interval (ISI) histogram of a single neuron, namely neuron  $v_{26}$  which was already considered in Fig. 2.3B. The responses of other neurons yield qualitatively similar statistics. The bell-shaped distribution of the membrane potential is commonly observed in neurons embedded in an active network (POSPISCHIL et al., 2009). The ISI histogram reflects the reduced spiking probability immediately after an action potential due the refractory mechanism. Interspike intervals larger than the refractory time constant  $\tau \cdot dt = 20$  ms roughly follow an exponential distribution. Similar ISI distributions were observed during in-vivo recordings in awake, behaving monkeys (SHINOMOTO et al., 2009).

Fig. B.1C shows a scatterplot of the coefficient of variation (CV) of the ISIs versus the average ISI for each neuron in the network. The neurons exhibited a variety of average firing rates between 3.5 Hz and 31.5 Hz. Most of the neurons responded in a highly irregular manner with a  $CV \approx 1$ . Neurons with high firing rates had a slightly lower CV due to the increased influence of the refractory mechanism. The dashed line marks the CV of a Poisson process, i.e., a memoryless spiking behavior. The CV of neuron  $v_{26}$  is marked by a cross. The structure of this plot resembles, e.g., data from recordings in behaving macaque monkeys (SOFTKY and KOCH, 1993) (but note the lower average firing rate).

### B.4 Approximation quality of different neuron and synapse models

The theory of the neuron model with absolute refractory mechanism guarantees sampling from the correct distribution. In contrast, the theory for the neuron model with a relative refractory mechanism only shows that the sampling process is “locally correct”, i.e., that it would yield correct conditional distributions  $p(z_k | z_{\setminus k})$  for each individual neuron if the state of the remaining network  $z_{\setminus k}$  stayed constant. Therefore, the stationary distribution of the sampling process with relative refractory mechanism only provides an approximation to the target distribution. In the following we examine the approximation quality and robustness of sampling networks with different refractory mechanisms for target Boltzmann distributions with parameters randomly drawn from different distributions. Furthermore, we investigate the effect of additive PSP shapes with more realistic time courses.

We generated target Boltzmann distributions with randomly drawn weights  $W_{ki}$  and biases (excitabilities)  $b_k$  and computed the similarity between these reference distributions and the corresponding neural sampling approximations. The setup of these simulations is the same as for the simulation presented in Fig. 2.3. As we aimed to compare the distribution  $q^*(z)$  sampled by the network with the



#### B.4 Approximation quality of different neuron and synapse models

$\sigma$	Absolute refractory	Rel. late recovery	Rel. moderate recovery	Prod. of marginals
0.03	$(3.10 \pm 0.18) \cdot 10^{-4}$	$(3.21 \pm 0.15) \cdot 10^{-4}$	$(3.33 \pm 0.17) \cdot 10^{-4}$	$(4.65 \pm 1.28) \cdot 10^{-4}$
0.3	$(2.98 \pm 0.19) \cdot 10^{-4}$	$(3.20 \pm 0.15) \cdot 10^{-4}$	$(3.58 \pm 0.3) \cdot 10^{-4}$	$(4.94 \pm 1.91) \cdot 10^{-2}$
3.0	$(1.32 \pm 0.45) \cdot 10^{-4}$	$(4.20 \pm 8.70) \cdot 10^{-3}$	$(1.00 \pm 1.82) \cdot 10^{-2}$	$(5.36 \pm 6.71) \cdot 10^{-1}$

**Table B.2: Approximation quality of networks with different refractory mechanisms.** Mean and standard deviation of the Kullback-Leibler divergence  $D_{\text{KL}}(p||q^*)$  between reference Boltzmann distributions  $p$  and neural sampling approximations  $q^*$  for three different neuron models (corresponding to columns) and three different values for the reference distribution hyperparameter  $\sigma$  (corresponding to rows). The parameter  $\sigma$  controls the standard deviation of the weights of the reference distributions  $p(z)$ . In case of very strong synaptic interactions (leading to sharply peaked distributions,  $\sigma = 3$ ) the approximation quality of the spiking network degrades, if the neurons feature a relative refractory mechanism. The data was computed from 100 randomly generated Boltzmann distributions and their neural approximations for each value of  $\sigma$ .

exact Boltzmann distribution  $p(z)$ , we reduced the number of neurons per network to  $K = 10$ . This resulted in a state space of  $2^{10}$  possible network states  $z$  for which the normalization constant for the target Boltzmann distribution could be computed exactly. The weight matrix  $W$  was constraint to be symmetric with vanishing diagonal. Off-diagonal elements were drawn from zero-mean normal distributions with three different standard deviations  $\sigma = 0.03$ ,  $\sigma = 0.3$  and  $\sigma = 3$ , whereas the  $b_k$  were sampled from the same distribution as in Fig. 2.3. For every value of the hyperparameter  $\sigma$  we generated 100 random distributions. For Boltzmann distributions with small weights ( $\sigma = 0.03$ ), the RVs are nearly independent, whereas distributions with intermediate weights ( $\sigma = 0.3$ ) show substantial statistical dependencies between RVs. For very large weights ( $\sigma = 3$ ), the probability mass of the distributions is concentrated on very few states (usually 90% on less than 10 out of the  $2^{10}$  states). Hence, the range of the hyperparameter  $0.03 \leq \sigma \leq 3$  considered here covers a range a very different distributions.


The approximation quality of the sampled distribution was measured in terms of the Kullback-Leibler divergence between the target distribution  $p$  and the neural approximation  $q^*$

$$D_{\text{KL}}(p||q^*) = \sum_z p(z) \log \frac{p(z)}{q^*(z)} . \quad (\text{B.7})$$

We estimated  $q^*$  from  $10^7$  samples for each simulation trial using a Laplace estimator, i.e., we added a priori 1 to the number of occurrences of each state  $z$ .

Table B.2 shows the means and the standard deviations of the Kullback-Leibler divergences between the target Boltzmann distributions and the estimated approximations stemming from neural sampling networks with three different neuron and synapse models: the exact model with absolute refractory mechanism and two models with different relative refractory mechanisms shown in the bottom and middle row in Fig. 2.2B. Additionally, as a reference, we provide the (analytically calculated) Kullback-Leibler divergences for fully factorized distributions, i.e.,  $q^*(z) = \prod_k q^*(z_k)$  with correct marginals  $q^*(z_k) = p(z_k)$  but independent variables  $z_i, z_j$  for  $i \neq j$ .

## B Appendix to chapter 2: Neural dynamics as sampling



**Fig. B.2: Comparison of neural sampling with different neuron and synapse models.** The figure shows a histogram of the Kullback-Leibler divergence between 100 different Boltzmann distributions over  $K = 10$  variables (with parameters randomly drawn, see setup of Figure 2.3) and approximations stemming from different neural sampling networks. Networks with absolute refractory mechanism provide the best approximation (as expected from theoretical guarantees). Networks consisting of neurons with relative refractory mechanisms, with only “locally” correct sampling, also provide a close fit to the true distribution (see inset) compared to a fully factorized approximation (assuming correct marginals and independent variables). Furthermore, it can be seen that sampling networks with more realistic, alpha-shaped, additive PSPs still fit the true distribution reasonably well.

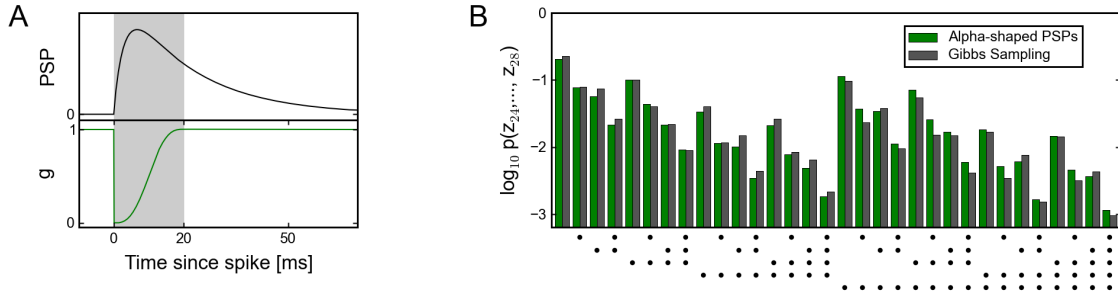
The absolute refractory model provides the best results as we expected due to the theoretical guarantee to sample from the correct distribution (the non-zero Kullback-Leibler divergence is caused by the estimation from a finite number of samples). The models with relative refractory mechanism provide faithful approximations for all values of the hyperparameter  $\sigma$  considered here. These relative refractory models are characterized by the theory to be “locally correct” and turn out to be much more accurate approximations than fully factorized distributions if substantial statistical dependencies between the RVs are present (i.e.,  $\sigma = 0.3$ ,  $\sigma = 3$ ). As expected, a late recovery of the refractory function  $g(\zeta)$  is beneficial for the approximation quality of the model as it is closer to an absolute refractory mechanism. Fig. B.2 explicitly shows the full histograms of the Kullback-Leibler divergences for the intermediate weights group ( $\sigma = 0.3$ ). Systematic deviations due to the relative refractory mechanism are on the same order as the effect of estimating from finite samples (as can be seen, e.g., from a comparison with the absolute refractory model which has 0 systematic error). For completeness, we mention that the divergences of the fully factorized distributions of 2 out of the 100 networks with  $D_{\text{KL}} > 0.1$  are not shown in the plot.

The theorems presented in this article assumed renewed (i.e., non-additive), rectangular PSPs. In the following we examine the effect of additive PSPs with more realistic time courses. We define additive, alpha-shaped PSPs in the following way. The influence  $\Delta u_{ki}$  of each presynaptic neuron  $v_i$  on the postsynaptic membrane potential  $u_k$  is modeled by convolving the input spikes with a kernel  $\kappa$ :

$$\Delta u_{ki}(t) = W_{ki} \cdot \sum_f \kappa(t - t_i^f) \quad (\text{B.8})$$

where  $\kappa(s) = \lambda \cdot (e^{-s/\tau_+} - e^{-s/\tau_-})$  for  $s \geq 0$  and  $\kappa(s) = 0$  for  $s < 0$ , and  $t_i^f$

## B.4 Approximation quality of different neuron and synapse models



**Fig. B.3: Sampling from a Boltzmann distribution with more realistic PSP shapes.** (A) The upper panel shows the shape of a single PSP elicited at time  $t = 0$ . The lower panel shows the time course of the refractory function  $g(\zeta_k(t))$  caused by a single spike of neuron  $v_k$  at  $t = 0$ . The gray-shaded area of length  $\tau \cdot dt = 20 \text{ ms}$  indicates the interval of neuron  $v_k$  being active (i.e.,  $z_k = 1$ ) due to a single spike of neuron  $v_k$  at time  $t = 0$ . (B) Shown is the probability distribution of 5 out of 40 neurons. The plot is similar to Fig. 2.3C, however it is generated with a sampling network that features alpha-shaped, additive PSPs. It can be seen that the network still produces a reasonable approximation to the true Boltzmann distribution (determined by Gibbs sampling).

for  $f \in \mathbb{N}$  are the spike times of the presynaptic neuron  $v_i$ . The time constant governing the rising edge of the PSPs was set to  $\tau_- = 3 \text{ ms}$ . The time constant controlling the falling edge was chosen equal to the duration of rectangular PSPs,  $\tau_+ = \tau \cdot dt = 20 \text{ ms}$ . The scaling parameter  $\lambda$  was set such that the time integral over a single PSP matches the time integral over the theoretically optimal rectangular PSP, i.e.,  $\lambda = \tau \cdot dt / (\tau_+ - \tau_-) = 20/17$ . These parameters display a simple and reasonable choice for the purpose of this study (an optimization of  $\lambda$ ,  $\tau_+$  and  $\tau_-$  is likely to yield an improved approximation quality). Fig. B.3A shows the resulting shape of the non-rectangular PSP. Furthermore the time course of the function  $g(\zeta_k(t))$  caused by a single spike of neuron  $v_k$  is shown in order to illustrate that the time constants of  $g$  and of a PSP are closely related due to the assumption  $\tau_+ = \tau \cdot dt$  made above. Preliminary and non-exhaustive simulations seem to suggest that the choice  $\tau_+ = \tau \cdot dt$  yields better approximation quality than setting  $\tau_+ \gg \tau \cdot dt$  or  $\tau_+ \ll \tau \cdot dt$ ; however it is very well possible that a mismatch between  $\tau_+$  and  $\tau \cdot dt$  can be compensated for by adapting other parameters, e.g., the PSP magnitude or a specific choice of the refractory function  $g$ . Fig. B.3B shows the results of an experiment, similar to the one presented in Fig. 2.3C, with additive, alpha-shaped PSPs and relative refractory mechanism. While differences to Gibbs sampling results are visible, the spiking network still captures dependencies between the binary random variables quite well.

For a quantitative analysis of the approximation quality, we repeated the experiment of Fig. B.2 with additive, alpha-shaped PSPs (shown as green bars). The Kullback-Leibler divergence  $D_{\text{KL}}(p||q^*)$  to the true distribution is clearly higher compared to the case of renewed, rectangular PSPs. Still networks with this more realistic synapse model account for dependencies between the random variables  $z$  and yield a better approximation of  $p(z)$  than fully factorized distributions.



# Appendix C

## Appendix to chapter 4: Statistical learning with compound memristive synapses

In this appendix, we provide the full theory for inference and learning with compound memristive synapses. The theory covers the formal definition of the joint probabilistic model, the analytical calculation of the posterior distribution, and the derivation of updated rules within the framework of generalized Expectation Maximization. We then describe, how the required computations can be calculated and represented by a spiking network, and give details to the computer simulations. Finally, we outline how the abstract stochastic spike response neuron model could be implemented with leaky integrator neurons, a standard neuron model in neuromorphic systems.

### C.1 Probabilistic model definition

The probabilistic model that corresponds to the spiking network is a mixture model with  $K$  mixture components and Gaussian likelihood function. Formally, we define a joint distribution  $p(\mathbf{Y} = \mathbf{y}(t), \mathbf{Z} = \mathbf{z}(t) | \boldsymbol{\theta})$  over  $K$  hidden binary random variables  $\mathbf{Z} = (Z_1, \dots, Z_K)^\top$  with values  $z_k(t) \in \{0, 1\}$ , and  $N$  real-valued visible random variables  $\mathbf{Y} = (Y_1, \dots, Y_N)^\top$  with values  $y_i(t) \in \mathbb{R}$ . The parameter set  $\boldsymbol{\theta} = \{\hat{\mathbf{b}}, \mathbf{W}\}$  consists of a real-valued bias vector  $\hat{\mathbf{b}} = (\hat{b}_1, \dots, \hat{b}_K)^\top$  and a real-valued  $K \times N$  weight matrix  $\mathbf{W}$ . The hidden RVs  $Z_k$  display an unrolled representation of a multinomial RV  $\tilde{Z} \in \{1, \dots, K\}$  that enumerates the mixture components, and we identify  $\tilde{Z} = k \Leftrightarrow Z_k = 1$ , i.e. exactly one binary RV  $Z_k$  is active in the random vector  $\mathbf{Z}$ . In the following, we stick to the unrolled vector notation  $\mathbf{Z}$ , and, for readability, omit the time-dependent notation and further shorten the notation by identifying the RVs with their values, e.g., we write  $p(\mathbf{z} | \boldsymbol{\theta})$  for  $p(\mathbf{Z} = \mathbf{z}(t) | \boldsymbol{\theta})$ .

The likelihood distribution fulfills the naïve Bayes property, i.e., all statistical dependencies between visible RVs  $y_i, y_j$  are explained by the hidden state  $\mathbf{z}$ . More precisely, the generative model has the following structure:

$$p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) = p(\mathbf{z} | \boldsymbol{\theta}) \cdot \prod_{k=1}^K \prod_{i=1}^N p(y_i | z_k = 1, \boldsymbol{\theta})^{z_k} \quad , \quad (\text{C.1})$$

with the prior

$$p(z | \theta) = \frac{e^{z^T \hat{\mathbf{b}}}}{\sum_{j=1}^K e^{\hat{b}_j}} \quad (\text{C.2})$$

and the likelihood

$$p(y_i | z_k = 1, \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{(y_i - \mu_{ki})^2}{2\sigma^2}}, \quad (\text{C.3})$$

with  $\sigma^2$  denoting an arbitrary (but fixed) variance which displays a constant in the model. Eq. (C.3) defines a Gaussian likelihood model for each input RV  $y_i$  with mean  $\mu_{ki}$  which is selected by the active hidden cause  $z_k = 1$  in eq. (C.1). For theoretical considerations, it is convenient to reorganize eq. (C.3) according to its dependency structure:

$$p(y_i | z_k = 1, \theta) = \frac{e^{-y_i^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}} \cdot e^{\frac{\mu_{ki}}{\sigma^2} \cdot y_i} \cdot e^{-\mu_{ki}^2/(2\sigma^2)} \quad (\text{C.4})$$

$$=: h(y_i) \cdot e^{W_{ki} \cdot y_i} \cdot e^{-A_{ki}}, \quad (\text{C.5})$$

where we set  $W_{ki} = \mu_{ki}/\sigma^2$  and  $A_{ki} = \mu_{ki}^2/(2\sigma^2) = \sigma^2 W_{ki}^2/2$ . The first factor does neither depend on the hidden causes  $z_k$  nor on the weight  $W_{ki}$  and will play no role during inference and learning. The second factor describes the coupling between the visible RV  $y_i$  and the active latent RV  $z_k$  through the mean  $W_{ki} = \mu_{ki}/\sigma^2$ . Finally, the third factor solely depends on the weight  $W_{ki}$  (and not on  $y_i$ ) and ensures correct normalization of the distribution.

## C.2 Inference

The posterior distribution given an observation  $\mathbf{y}$  follows directly from Bayes rule:

$$p(z | \mathbf{y}, \theta) = p(z | \theta) \cdot p(\mathbf{y} | z, \theta) / \text{Norm.} \quad (\text{C.6})$$

$$= e^{z^T \hat{\mathbf{b}}} \prod_{k=1}^K \prod_{i=1}^N h(y_i)^{z_k} \cdot e^{z_k \cdot W_{ki} \cdot y_i} \cdot e^{-z_k \cdot A_{ki}} / \text{Norm.} \quad (\text{C.7})$$

$$= e^{z^T \cdot [\hat{\mathbf{b}} - \mathbf{A} + \mathbf{W} \cdot \mathbf{y}]} \cdot \prod_{i=1}^N h(y_i) / \text{Norm.} \quad (\text{C.8})$$

with  $\mathbf{A} := (A_1, \dots, A_k, \dots, A_K)^T$  and  $A_k := \sum_{i=1}^N A_{ki}$ . We evaluate the posterior for a specific hidden RV  $z_k$  to be active and provide the normalization constant explicitly:

$$p(z_k = 1 | \mathbf{y}, \theta) = \frac{e^{\hat{b}_k - A_k + \sum_{i=1}^N W_{ki} \cdot y_i} \cdot \prod_{i=1}^N h(y_i)}{\sum_{j=1}^K e^{\hat{b}_j - A_j + \sum_{i=1}^N W_{ji} \cdot y_i} \cdot \prod_{i=1}^N h(y_i)} \quad (\text{C.9})$$

$$= \frac{e^{\hat{u}_k}}{\sum_{j=1}^K e^{\hat{u}_j}} \quad (\text{C.10})$$

where we defined  $\hat{u}_k = \hat{b}_k - A_k + \sum_{i=1}^N W_{ki} \cdot y_i$ . The quantities  $\hat{u}_k$  are reminiscent of neuronal membrane potentials which consist of bias terms  $\hat{b}_k - A_k$  and synaptic input  $\sum_{i=1}^N W_{ki} y_i$ . However, implicitly the bias terms depend on all afferent synaptic weights since  $\hat{b}_k - A_k = \hat{b}_k - \frac{\sigma^2}{2} \sum_i W_{ki}^2$  and, thus, rely on information not locally available to the neurons. This issue will be resolved in the context of learning: We will identify update rules for both biases and synapses which only use information available locally and thereby make a neural network implementation feasible.

### C.3 Learning via generalized Expectation-Maximization

We investigate unsupervised learning of the probabilistic model based on generalized online Expectation-Maximization (EM) (DEMPSTER et al., 1977), an optimization algorithm from machine learning theory. To this end, we impose additional constraints on the posterior distribution (GRACA et al., 2008) which will enable a neural network implementation via homeostatic intrinsic plasticity (HABENSCHUSS et al., 2012) and STDP-type synaptic plasticity (NESSLER et al., 2013; HABENSCHUSS et al., 2013b). Since the derivation is almost identical to (HABENSCHUSS et al., 2012), we only outline the key steps and main results in the following and refer to (HABENSCHUSS et al., 2012) for the details.

The algorithmic approach rests upon the generalized EM decomposition:

$$\mathcal{F}(\mathbf{W}, q(z|\mathbf{y})) = \mathcal{L}(\boldsymbol{\theta}) - \langle D_{\text{KL}}(q(z|\mathbf{y}) \| p(z|\mathbf{y}, \boldsymbol{\theta})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{E-step} \quad (\text{C.11})$$

$$= \langle \log p(\mathbf{y}, z|\boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} + \langle H(q(z|\mathbf{y})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{M-step} \quad (\text{C.12})$$

with the log-likelihood  $\mathcal{L}(\boldsymbol{\theta}) = \langle \log p(\mathbf{y}|\boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})}$ , the Kullback-Leibler divergence  $D_{\text{KL}}(q(z) \| p(z)) = \sum_z q(z) \cdot \log(q(z)/p(z))$  and the entropy  $H(q(z)) = -\sum_z q(z) \cdot \log q(z)$ . The distribution  $p^*(\mathbf{y})$  denotes the input distribution actually presented to the system. The distribution  $p(\cdot|\boldsymbol{\theta})$  is the probabilistic model defined above. The distribution  $q(z|\mathbf{y})$  is called variational posterior and will ultimately be implemented by the spiking network. The short hand notation  $\langle \cdot \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})}$  denotes the concatenated average  $\langle \langle \cdot \rangle_{q(z|\mathbf{y})} \rangle_{p^*(\mathbf{y})}$  with respect to the input distribution and the resulting variational posterior. In principle, the above decomposition holds for any choice of  $q$ , and since the Kullback-Leibler divergence in eq. (C.11) is strictly non-negative, the objective function  $\mathcal{F}$  is a lower bound of the log-likelihood  $\mathcal{L}$ . During optimization the algorithm will pursue two goals: to increase  $\mathcal{L}$ , i.e. to better adapt the probabilistic model to the data, and to keep  $\langle D_{\text{KL}}(q \| p) \rangle$  small, i.e. to maintain a reliable approximation  $q(z|\mathbf{y})$  of the exact posterior  $p(z|\mathbf{y}, \boldsymbol{\theta})$ .

We first impose a homeostatic constraint on the variational posterior  $q(z|\mathbf{y})$ , namely that the long term average activation of any hidden RV  $z_k$  matches a predefined target value  $c_k$  (with  $\sum_k c_k = 1$ ). Formally, we define a set of constrained distributions  $\mathcal{Q} = \{q : \langle z_k \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} = c_k \quad \forall 1 \leq k \leq K\}$  and demand  $q(z|\mathbf{y}) \in \mathcal{Q}$ . The optimization algorithm then relies on the joint application of an E(expectation)-step

and an M(aximization)-step: During the E-step, we aim to minimize the Kullback-Leibler divergence with respect to  $q \in \mathcal{Q}$  in eq. (C.11); during the M-step, we perform gradient ascent on  $\langle \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle$  with respect to the weights  $W_{ki}$  in eq. (C.12). The E- and M-step will be discussed separately.

The E-step is a constrained optimization problem, namely the minimization of  $\langle D_{\text{KL}}(q || p) \rangle$  such that  $q \in \mathcal{Q}$ , that can be solved through Lagrange multipliers. Since we imposed  $K$  constraints (one per RV  $z_k$ ), we need  $K$  Lagrange multipliers  $\beta_k$ . It turns out that the solution to this optimization problem simply adds the multipliers  $\beta_k$  to the biases  $\hat{b}_k - A_k$  in eq. (C.9). This convenient result gives rise to the definition of intrinsic excitabilities  $b_k := \hat{b}_k - A_k + \beta_k$  which unify biases and multipliers in a single quantity. Furthermore, it turns out that the optimal values of the  $\beta_k$ 's (and thus the  $b_k$ 's) can be determined via iterative update rules that solely rely on the hidden RVs  $z_k$  under the variational response  $q(\mathbf{z} | \mathbf{y})$  and overwrite the non-local terms  $A_k$ . In summary, we obtain the variational posterior distribution  $q$  that solves the E-step:

$$q(\mathbf{z} | \mathbf{y}) = \frac{e^{u_k}}{\sum_{j=1}^K e^{u_j}} \quad \text{with} \quad u_k = b_k + \sum_{i=1}^N W_{ki} \cdot y_i \quad (\text{C.13})$$

$$\Delta b_k \propto \langle c_k - z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z} | \mathbf{y})} \quad . \quad (\text{C.14})$$

The variational posterior in eq. (C.13) is described in terms of membrane potentials  $u_k$  which consist of synaptic input  $\sum_i W_{ki} y_i$  and intrinsic excitabilities  $b_k$ . Eq. (C.14) regulates the intrinsic excitabilities  $b_k$  in a homeostatic fashion: When the average response exceeds the target  $c_k$ , the excitability is reduced, and vice versa.

The M-step can be solved via gradient ascent on  $\mathcal{F}$  with respect to the weights  $W_{ki}$ . The variational posterior  $q$  is a constant during the M-step in EM, and thus the log-joint distribution  $\log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta})$  remains as the only  $W_{ki}$ -dependent term in eq. (C.12). By taking the derivative of the log-joint defined by eq. (C.1), (C.2) and (C.5) with respect to  $W_{ki}$ , we obtain:

$$\frac{d\mathcal{F}}{dW_{ki}} = \langle \partial_{W_{ki}} \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z} | \mathbf{y})} \quad (\text{C.15})$$

$$= \langle \partial_{W_{ki}} z_k \cdot \log p(y_i | z_k = 1, \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z} | \mathbf{y})} \quad (\text{C.16})$$

$$= \langle z_k \cdot (y_i - \sigma^2 \cdot W_{ki}) \rangle_{p^*(\mathbf{y})q(\mathbf{z} | \mathbf{y})} \quad . \quad (\text{C.17})$$

The gradient with respect to the weights  $W_{ki}$  yields Hebbian-type update rules that use pre- $(y_i)$  and post- $(z_k)$  synaptic activity and the current weight  $W_{ki}$  given the input  $p^*(\mathbf{y})$  and the variational response  $q(\mathbf{z} | \mathbf{y})$ . Importantly, only local information is required during the E- and M-step.

## C.4 Spiking network implementation

The spiking neural network model instantiates eq. (C.13), (C.14) and (C.17), i.e., it represents the variational posterior  $q(\mathbf{z} | \mathbf{y})$  for probabilistic inference through its



spike response and implements the derived update rules for generalized online EM learning through intrinsic and synaptic plasticity.

Each of the RVs  $z_k$  is represented by one of the  $K$  network neurons, and each spike in the network is a sample from the variational posterior  $q(\mathbf{z}|\mathbf{y})$  by identifying  $z_k = 1$  for a spike of the  $k$ -th network neuron. By setting the instantaneous firing rate  $\rho_k$  to be

$$\rho_k = \lim_{\delta t \rightarrow 0} p(\text{spike in } [t, t + \delta t]) / \delta t = r_{\text{net}} \cdot e^{u_k - u_{\text{inh}}} \quad (\text{C.18})$$

with  $u_{\text{inh}} := \log \sum_{j=1}^K \exp(u_j)$  the network thus implements eq. (C.13) for any choice of  $r_{\text{net}}$ , i.e.,  $p_{\text{net}} = q$ .

The learning rules (C.14) and (C.17) rely on expected values  $\langle \cdot \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}$ . The expectations can be approximated from input samples  $\mathbf{y} \sim p^*(\mathbf{y})$  and posterior samples  $\mathbf{z} \sim q(\mathbf{z}|\mathbf{y})$  in response to this input. The input vector  $\mathbf{y}$  is defined at any time  $t$  in the network as it measures the instantaneous presence or absence of rectangular input pulses. Samples of the latent variable  $\mathbf{z}$ , in contrast, are only defined at the spike times of the network. Hence integrating expected values  $\langle z_k \rangle$  from the spike response can be expressed most conveniently in terms of the spike train function  $s_k(t) = \sum_f \delta(t - t_k^f)$  of the network neurons. We obtain the following plasticity rules:

$$\frac{db_k}{dt} = \eta_b \cdot [r_{\text{net}} \cdot c_k - s_k(t)] \quad (\text{C.19})$$

$$\frac{dW_{ki}}{dt} = \eta_W \cdot s_k(t) \cdot [y_i - \sigma^2 \cdot W_{ki}] \quad (\text{C.20})$$

with small learning rates  $\eta_b$  and  $\eta_W$ . The homeostatic rule (C.19) regulates the intrinsic excitabilities  $b_k$  such that the average target activations  $\langle z_k(t) \rangle \approx c_k$  are maintained over the presentation of many different input patterns  $\mathbf{y}(t) \sim p^*(\mathbf{y})$  in accordance with eq. (C.14), and thereby implements the E-step. Building on the network response shaped by the E-step, the synaptic rule (C.20) on average increases the objective function  $\mathcal{F}$  since synaptic changes  $\frac{d}{dt} W_{ki}$  on average point in the direction of the  $\mathbf{W}$ -gradient of  $\mathcal{F}$  given by eq. (C.17), thereby implementing the M-step. Since synaptic updates rely on a (sufficiently) precise E-step which, in turn, needs to integrate any changes in the network response due to synaptic plasticity, homeostatic intrinsic plasticity is required to act on faster time scales than synaptic plasticity. As a consequence, the learning rate  $\eta_b$  will typically exceed the learning rate  $\eta_W$  in the spiking network implementation.

The homeostatic intrinsic plasticity rule (C.19) can readily be implemented by the spiking neurons: The intrinsic excitability  $b_k$  of each neuron increases linearly in time with a slow drift  $\eta_b \cdot r_{\text{net}} \cdot c_k$  and is lowered abruptly by  $\eta_b$  at the spike times of

neuron  $z_k$ . Similarly, mapping the synaptic plasticity rule (C.20) to the compound-synapse STDP rule (4.3) is straight-forward due to the structural equivalence of equations (4.3) and (C.20): The learning rate  $\eta_W$  in the theory domain corresponds to  $\pi_{\text{up}} \cdot W^{\text{max}}$  in the hardware domain, e.g., high jumping probabilities  $\pi_{\text{up}}$  and large weight contributions  $\omega = W^{\text{max}}/M$  of individual stochastic switches lead to high learning rates  $\eta_W$ . Furthermore, the maximum weight  $W^{\text{max}}$  can directly be identified with the precision  $1/\sigma^2$  of the likelihood distribution (C.3). In the theory domain, we know that  $\mu_{ki} = \sigma^2 \cdot W_{ki}$ , and hence,  $\mu_{ki} = W_{ki}/W^{\text{max}} = m_{ki}/M$  for the compound synapses. Finally, due to the structural equivalence of equations (4.3) and (C.20) we find that the compound memristor plasticity rule (4.3) inherits the convergence properties from the theoretically derived plasticity rule (C.20) during online learning. The resulting translation of memristor synapse parameters to the abstract model is summarized in Table 4.1.

## C.5 Details to the computer simulations

All computer simulations were performed with customized Python scripts. For the computer simulations, we employed a network architecture with  $K = 10$  network neurons and  $N = 24 \cdot 24 = 576$  inputs. The simulation time step was  $\delta t = 1$  ms and the PSP time constant  $\tau = 10$  ms. The overall network firing rate was set to  $r_{\text{net}} = 100$  Hz, the homeostatic target activation uniformly to  $c_k = 1/K$ . Synapses were composed of  $M = 10$  constituents with weight  $\omega = 0.1$  each. Switching probabilities were set to  $\pi_{\text{up}} = \pi_{\text{down}} = 10^{-3}$ . This corresponds to a Gaussian likelihood model with variance  $\sigma^2 = 1$  and learning rate  $\eta_W = 10^{-3}$ . The learning rate for homeostatic intrinsic plasticity was set to  $\eta_b = 20 \cdot \eta_W$ . For the simulations in Figure 4.4-4.6, certain parameters deviated from the above, depending on the simulation setup. For Figure 4.6, the switching probability  $\pi_{\text{up}} = 10^{-3}$  was kept fixed, and  $\pi_{\text{down}}$  was adapted for different  $\Delta$ -values. All other changes are described directly in Chapter 4.

For learning experiments, digits 0, 1, 2, 3, 4 were extracted in equal proportion from the MNIST training data set (LECUN et al., 1998). A frame of two pixels width was removed, leaving images of size  $24 \times 24$ . The images (indexed by  $s$ ) were scaled linearly to activity patterns  $x_i^s \in [0.05, 0.9]$ , with  $i = 1, \dots, N$ , which were presented to the network as follows. For given activity pattern  $\mathbf{x}^s = (x_1^s, \dots, x_N^s)$ , each input  $i$  spiked with probability  $p_i^{\text{spike}} = 1 - (1 - x_i^s)^{\delta t/\tau}$  per time step  $\delta t$ . During training, a new activity pattern  $\mathbf{x}^s$  was randomly drawn from the training set every 100 ms. Each network was trained for 5000 s. To obtain the unweighted PSP values  $y_i$ , the resulting spike patterns were convolved with a box kernel of duration  $\tau$  and amplitude 1, and then clipped to values  $[0, 1]$ . This defined the input  $\mathbf{y}(t)$ , and thus (implicitly) the data distribution  $\mathbf{y}(t) \sim p^*(\mathbf{y})$ . Notably, the spiking probability  $p_i^{\text{spike}}$  is chosen such that  $\langle y_i \rangle = x_i^s$ .

## C.6 Implementation with leaky integrator neurons

The estimate of the log-likelihood in Figure 4.3D was based on 5000 input samples  $\mathbf{y}(t)$ , which were randomly drawn from the training data, and assumed a uniform prior  $p(z_k = 1 | \boldsymbol{\theta}) = 1/K$  in accordance with the homeostatic target activation. The classification performance in Figure 4.3D, Figure 4.4B and Figure 4.5E was determined as follows. For given configuration of the synapses and intrinsic excitabilities, 100 versions of each digit from the training data set were presented to the network for 1s each. Each neuron was labeled to be tuned to the digit class it was most responsive to. Then 500 versions of each digit from the MNIST test data set were presented to the network for 1s each. The network neuron that spiked most during the 1s period determined the network's classification of the input digit. The classification error is the fraction of wrongly classified digits.

## C.6 Implementation with leaky integrator neurons

The idealized stochastic neurons in the WTA network model feature abstract membrane potentials  $u_k$  that integrate the input  $\mathbf{y}(t)$  through the weights  $W_{ki}$  linearly. In a hardware integration, however, synaptic weights arise from the conductance of memristors, and neurons are physical implementations based on capacitors and various other circuit elements. Here, we outline one possible hardware integration and consider a leaky integrator with membrane potential  $U_k$  that obeys the following dynamics:

$$\tau_m \cdot \frac{dU_k}{dt} = -(U_k - B_k) + I_k / G_L \quad , \quad (\text{C.21})$$

with membrane time constant  $\tau_m$ , leak conductance  $G_L$ , resting potential  $B_k$  and synaptic input current  $I_k$ . In the setup of Fig. 4.1A, input spikes trigger a rectangular voltage pulse of duration  $\tau$  and with amplitude  $U_{\text{pre}}$ . Denoting the conductance of the memristive synapse by  $G$ , this generates a synaptic current  $I = U_{\text{pre}} \cdot G$ . The equilibrium membrane potential (i.e.  $\frac{d}{dt}U_k = 0$ ) under this current is  $U_k = B_k + (U_{\text{pre}}/G_L) \cdot G$ . For small membrane time constant  $\tau_m \rightarrow 0$ , e.g., for a small neuron capacitance, the fast membrane will closely resemble the rectangular presynaptic pulse shape, and the PSP amplitude  $(U_k - B_k)$  will be proportional to the weight  $G$ . This linear integration property of  $U_k$  also holds in case of multiple memristive synapses acting in parallel, and we find

$$U_k = B_k + \frac{U_{\text{pre}}}{G_L} \cdot \sum_i G_{ki} \cdot y_i \quad . \quad (\text{C.22})$$

Consequently, the membrane potential  $U_k$  of the leaky integrator matches the idealized membrane potential  $u_k$  employed in Chapter 4 up to a linear function that serves to translate the voltage-based potential  $U_k$  to the unitless potential  $u_k$ . Using the membrane potential  $U_k$ , the exponential firing behavior (C.18) of the neurons could either be realized with an inherently stochastic firing mechanism.

## C Appendix to chapter 4: Statistical learning with compound memristive synapses

Alternatively, deterministic leaky integrate-and-fire neurons could be operated in a stochastic regime by adapting, for instance, the approach taken in (PETROVICI et al., 2013).

# Appendix **D**

## Appendix to chapter 5: Distributed Bayesian computation and learning in spiking neural sheets

In the following, we provide the full definition and derivation of the neural sheet model for inference and unsupervised learning. We first describe the probabilistic model and show that the spiking network model can sample from the posterior distribution. Then we derive parameter updates for unsupervised model optimization and link them to plasticity rules for sample-based online learning. The heuristic recurrent plasticity rule and a brief discussion on the interacting time scales are presented in separate subsections. Finally, we provide details to the computer simulations and figures.

### D.1 Generative model

#### Definition of variables

We introduce a generative model

$$p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) = p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) \cdot p(\mathbf{z} | \boldsymbol{\theta}) \quad (\text{D.1})$$

over  $N$  observed variables  $y_1, \dots, y_N$ , subsumed in the vector  $\mathbf{y}$ , and  $K$  latent variables  $z_1, \dots, z_K$ , subsumed in the vector  $\mathbf{z}$ . The latent variables are binary,  $z_k \in \{0, 1\}$ , and said to be “active” iff  $z_k = 1$ . The possible values of the observed variables  $y_i$  depend on the employed likelihood model. The full model is governed by parameters  $\boldsymbol{\theta} = (\mathbf{V}, \mathbf{V}_0, \hat{\mathbf{W}}, \hat{\mathbf{b}})$  with real-valued  $K \times N$  afferent weight matrix  $\mathbf{V}$ ,  $N$ -dimensional default vector  $\mathbf{V}_0$  (which will be a constant during learning),  $K \times K$  recurrent weight matrix  $\hat{\mathbf{W}}$ , and  $K$ -dimensional bias vector  $\hat{\mathbf{b}}$ . Furthermore, each latent variable  $z_k$  is connected to a subset of the input variables (its afferent field) which is described by an index set  $\mathcal{I}_k \subseteq \{1, 2, \dots, N\}$ . Likewise, we define the projection field for each input variable  $y_i$  as the index set  $\mathcal{P}_i = \{k | i \in \mathcal{I}_k\}$ . We refer to all-but-one variables of a vector by the shorthand notation  $\mathbf{z}_{\setminus k} := (z_1, \dots, z_{k-1}, z_{k+1}, \dots, z_K)$ .

### Generative Model: Likelihood

Provided the state of the hidden units  $\mathbf{z}$ , the inputs are defined to be independent (local “Naive Bayes”):

$$p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) = \prod_{i=1}^N p(y_i | \mathbf{z}, \boldsymbol{\theta}) \quad . \quad (\text{D.2})$$

For each  $y_i$ , the prior will ensure that at most one hidden unit  $z_k$  in  $\mathcal{P}_i$  is active. We assume that, if present, this single active unit  $z_k$  governs the distribution of  $y_i$ . In case all hidden variables in  $\mathcal{P}_i$  are 0, a “default hypothesis” is used. The resulting input distribution is assumed to be in the natural exponential family. For instance, Bernoulli, Poisson, or Gaussian distributions are in this class. In natural exponential family form,  $p(y_i | \mathbf{z}, \boldsymbol{\theta})$  reads:

$$p(y_i | z_k = 0, \forall k \in \mathcal{P}_i, \boldsymbol{\theta}) = h_i(y_i) e^{V_{0i} y_i - A_{0i}} \quad (\text{D.3})$$

$$p(y_i | z_k = 1, k \in \mathcal{P}_i, \boldsymbol{\theta}) = h_i(y_i) e^{(V_{ki} + V_{0i}) y_i - (A_{ki} + A_{0i})} \quad (\text{D.4})$$

where  $h_i(y_i)$  denotes the base measure, and the normalization constants  $A_{ki}$  and  $A_{0i}$  depend on the parameters  $V_{ki}$  and  $V_{0i}$ . In anticipation of (D.6), the parameters in (D.4) were written as  $(V_{ki} + V_{0i})$ , i.e. relative to the default hypothesis. We adopt the convention that  $V_{ki} = 0$  for  $k \notin \mathcal{P}_i$  to obtain a closed form expression for (D.3) and (D.4):

$$p(y_i | \mathbf{z}, \boldsymbol{\theta}) = h_i(y_i) e^{V_{0i} y_i - A_{0i}} \prod_{k \in \mathcal{P}_i} \left[ e^{V_{ki} y_i - A_{ki}} \right]^{z_k} \quad (\text{D.5})$$

$$= h_i(y_i) e^{V_{0i} y_i - A_{0i}} \exp \left[ \sum_{k=1}^K z_k V_{ki} y_i - z_k A_{ki} \right] \quad (\text{D.6})$$

By combining (D.2) and (D.6) we obtain the full likelihood

$$p(\mathbf{y} | \mathbf{z}, \boldsymbol{\theta}) = h(\mathbf{y}) \exp \left[ \mathbf{z}^\top \mathbf{V} \mathbf{y} - \mathbf{z}^\top \mathbf{A} \right] \quad (\text{D.7})$$

with  $\mathbf{A} = (A_1, \dots, A_K)^\top$  and  $A_k = \sum_{i=1}^N A_{ki}$ . Here we use the shorthands  $\mathbf{z}^\top \mathbf{V} \mathbf{y} = \sum_k \sum_i z_k V_{ki} y_i$  and  $\mathbf{z}^\top \mathbf{A} = \sum_k z_k A_k$ . The function  $h(\mathbf{y}) := \prod_i h_i(y_i) e^{V_{0i} y_i - A_{0i}}$  comprises only terms that do not depend on  $\mathbf{z}$ , and hence, it will play no role in the inference. In Chapter 5, we employed a Bernoulli likelihood model:

$$p(y_i = 1 | z_k = 0, \forall k \in \mathcal{P}_i, \boldsymbol{\theta}) = \pi_{0i} \quad , \quad p(y_i = 1 | z_k = 1, k \in \mathcal{P}_i, \boldsymbol{\theta}) = \pi_{ki} \quad (\text{D.8})$$

or in closed form:

$$p(y_i | \mathbf{z}, \boldsymbol{\theta}) = \pi_{0i}^{y_i} (1 - \pi_{0i})^{1-y_i} \prod_k \left[ \frac{\pi_{ki}^{y_i} (1 - \pi_{ki})^{1-y_i}}{\pi_{0i}^{y_i} (1 - \pi_{0i})^{1-y_i}} \right]^{z_k} \quad . \quad (\text{D.9})$$

By rewriting (D.9) in the exponential family form (D.6), we identify

$$\begin{aligned} V_{ki} &= \log \frac{\pi_{ki}}{1 - \pi_{ki}} - V_{0i} , & A_{ki} &= \log(1 + e^{V_{ki} + V_{0i}}) - A_{0i} , \\ V_{0i} &= \log \frac{\pi_{0i}}{1 - \pi_{0i}} , & A_{0i} &= \log(1 + e^{V_{0i}}) . \end{aligned} \quad (\text{D.10})$$

In particular, the cluster centers can be recovered via  $\pi_{ki} = \sigma(V_{ki} + V_{0i})$ .

Likewise, Poisson and Gaussian distributions can be written in the exponential family form (D.6). This extends the input domain to  $y_i \in \mathbb{N}$  and  $y_i \in \mathbb{R}$  respectively. For a Poisson model with expected values  $\lambda_{0i}$  and  $\lambda_{ki}$  (for the default hypothesis and the hidden causes resp.),

$$p(y_i | z_k = 0, \forall k \in \mathcal{P}_i, \boldsymbol{\theta}) = (y_i!)^{-1} \lambda_{0i}^{y_i} e^{-\lambda_{0i}} , \quad p(y_i | z_k = 1, k \in \mathcal{P}_i, \boldsymbol{\theta}) = (y_i!)^{-1} \lambda_{ki}^{y_i} e^{-\lambda_{ki}} , \quad (\text{D.11})$$

we obtain

$$\begin{aligned} V_{ki} &= \log(\lambda_{ki}) - V_{0i} , & A_{ki} &= \exp(V_{ki} + V_{0i}) - A_{0i} , \\ V_{0i} &= \log(\lambda_{0i}) , & A_{0i} &= \exp(V_{0i}) . \end{aligned} \quad (\text{D.12})$$

For Gaussians with centers  $\mu_{0i}$  and  $\mu_{ki}$  and fixed variance  $\sigma^2$ ,

$$p(y_i | z_k = 0, \forall k \in \mathcal{P}_i, \boldsymbol{\theta}) = \mathcal{N}(y_i; \mu_{0i}, \sigma^2) , \quad p(y_i | z_k = 1, k \in \mathcal{P}_i, \boldsymbol{\theta}) = \mathcal{N}(y_i; \mu_{ki}, \sigma^2) , \quad (\text{D.13})$$

we obtain

$$\begin{aligned} V_{ki} &= \frac{\mu_{ki}}{\sigma^2} - V_{0i} , & A_{ki} &= \frac{\sigma^2}{2} (V_{ki} + V_{0i})^2 - A_{0i} , \\ V_{0i} &= \frac{\mu_{0i}}{\sigma^2} , & A_{0i} &= \frac{\sigma^2}{2} V_{0i}^2 . \end{aligned} \quad (\text{D.14})$$

### Generative model: Prior

The prior  $p(\mathbf{z} | \boldsymbol{\theta})$  guarantees that no two (or more) hidden units  $z_k, z_j$  with overlapping input fields,  $\mathcal{I}_k \cap \mathcal{I}_j \neq \emptyset$ , are active simultaneously, i.e., it ensures that at most one unit  $z_k$  generates each input variable  $y_i$  at the same time. For this work, we choose a Boltzmann machine prior which can introduce dependencies between units with non-overlapping receptive fields through symmetric parameters  $\hat{W}_{kj}^{\text{exc}} = \hat{W}_{jk}^{\text{exc}}$ . Furthermore each variable  $z_k$  has a bias value  $\hat{b}_k$  which directly affects its prior probability of activity:

$$p(\mathbf{z} | \boldsymbol{\theta}) = \frac{1}{Z} \exp \left[ \frac{1}{2} \mathbf{z}^T \hat{\mathbf{W}}^{\text{exc}} \mathbf{z} + \mathbf{z}^T \hat{\mathbf{b}} \right] \cdot \prod_{k=1}^K \prod_{j=1}^K \left( \delta_{\mathcal{I}_k \cap \mathcal{I}_j, \emptyset} \right)^{z_k \cdot z_j} , \quad (\text{D.15})$$

with  $\delta$  denoting the Kronecker delta and  $0^0 := 1$ . The double-product factor ensures, that the assumptions on  $z$  made in the likelihood (D.7) are satisfied, and can be approximated with arbitrary precision by using strong inhibitory weights  $\hat{W}_{jk}^{\text{inh}}$

$$\mathcal{I}_k \cap \mathcal{I}_j \neq \emptyset \Rightarrow \hat{W}_{kj}^{\text{inh}} \rightarrow -\infty \quad \text{for } k \neq j, \quad (\text{D.16})$$

and setting  $\prod_{k=1}^K \prod_{j=1}^K \left( \delta_{\mathcal{I}_k \cap \mathcal{I}_j, \emptyset} \right)^{z_k \cdot z_j} = \exp\left[\frac{1}{2} z^\top \hat{W}^{\text{inh}} z\right]$ . Note that the range of strong inhibition could also extend beyond the minimal required range (D.16). For notational brevity, we subsume the excitatory and inhibitory recurrent weight matrices in a single matrix  $\hat{W} = \hat{W}^{\text{exc}} + \hat{W}^{\text{inh}}$  for the derivation.  $\hat{W}$  is symmetric ( $\hat{W} = \hat{W}^\top$ ) and has zero diagonal ( $\hat{W}_{kk} = 0$ ). Using this notation, the prior simply reads

$$p(z | \theta) = \frac{1}{Z} \exp\left[\frac{1}{2} z^\top \hat{W} z + z^\top \hat{b}\right]. \quad (\text{D.17})$$

## D.2 Inference in the generative model (Corollary 1)

By applying Bayes rule  $p(z | \mathbf{y}, \theta) \propto p(z | \theta) \cdot p(\mathbf{y} | z, \theta)$  on eq. (D.17) and eq. (D.7) we obtain the posterior in closed form:

$$p(z | \mathbf{y}, \theta) = \exp\left[\frac{1}{2} z^\top \hat{W} z + z^\top \mathbf{V} \mathbf{y} + z^\top (\hat{b} - \mathbf{A})\right] / \text{Norm}. \quad (\text{D.18})$$

where the normalization sums the exponential over all possible states of the posterior. To establish the link to the neural sampling theory via the sufficient conditions (5.5), we solve (D.18) for the logit of a single unit  $z_k$ :

$$u_k \stackrel{!}{=} \log \frac{p(z_k = 1 | z_{\setminus k}, \mathbf{y}, \theta)}{p(z_k = 0 | z_{\setminus k}, \mathbf{y}, \theta)} = \log \frac{p(z_k = 1, z_{\setminus k} | \mathbf{y}, \theta)}{p(z_k = 0, z_{\setminus k} | \mathbf{y}, \theta)} \quad (\text{D.19})$$

$$= \log \frac{\exp\left[\frac{1}{2} \sum_{j \neq k} (1 \cdot \hat{W}_{kj} z_j + z_j \hat{W}_{jk} \cdot 1) + \sum_i (1 \cdot V_{ki} y_i) + 1 \cdot (\hat{b}_k - A_k)\right]}{\exp\left[\frac{1}{2} \sum_{j \neq k} (0 \cdot \hat{W}_{kj} z_j + z_j \hat{W}_{jk} \cdot 0) + \sum_i (0 \cdot V_{ki} y_i) + 0 \cdot (\hat{b}_k - A_k)\right]} \quad (\text{D.20})$$

$$+ \log \frac{\exp\left[\frac{1}{2} \sum_{j, l \neq k} z_j \hat{W}_{jl} z_l + \sum_{j \neq k} \sum_i (z_j V_{ji} y_i) + \sum_{j \neq k} z_j \cdot (\hat{b}_j - A_j)\right]}{\exp\left[\frac{1}{2} \sum_{j, l \neq k} z_j \hat{W}_{jl} z_l + \sum_{j \neq k} \sum_i (z_j V_{ji} y_i) + \sum_{j \neq k} z_j \cdot (\hat{b}_j - A_j)\right]} \quad (\text{D.21})$$

$$= \log \frac{\exp\left[\sum_{j \neq k} \hat{W}_{kj} z_j + \sum_i V_{ki} y_i + (\hat{b}_k - A_k)\right]}{\exp[0]} + \log 1 \quad (\text{D.22})$$

$$= \sum_{j=1}^K \hat{W}_{kj} z_j + \sum_{i=1}^N V_{ki} y_i + (\hat{b}_k - A_k). \quad (\text{D.23})$$



### D.3 Model optimization via Generalized Expectation Maximization

Here we made use of the symmetry of  $\hat{W}$  and  $\hat{W}_{kk} = 0$ . Hence, we can map the neuronal membrane potential (5.3) to the parameters of the Bernoulli likelihood model. We find  $\hat{W}_{kj} = \hat{W}_{kj}^{\text{exc}} + \hat{W}_{kj}^{\text{inh}} = W_{kj}^{\text{exc}} + W_{kj}^{\text{inh}}$  for recurrent weights,  $V_{ki} = \log \frac{\pi_{ki}}{1-\pi_{ki}} - \log \frac{\pi_{0i}}{1-\pi_{0i}}$  for afferent weights according to eq. (D.10), and  $b_k = \hat{b}_k - A_k$  for excitabilities, with  $A_k = \sum_i A_{ki}$  given by eq. (D.10). This proves Corollary 1.

## D.3 Model optimization via Generalized Expectation Maximization

In the presence of plastic input synapses  $V_{ki}$ , the terms  $A_k$  change over time and, in a spiking network implementation, add to the intrinsic excitability of the cells. Arguably, the information required for calculating  $A_k$  is not locally available to the neurons since  $A_k$  depends on all afferent synaptic efficacies  $V_{ki}$ . We follow the approach of (HABENSCHUSS et al., 2012) who showed that homeostatic intrinsic plasticity can enable a spiking network to account for time varying  $A_k$ 's. More precisely, the interplay of homeostatic intrinsic plasticity and synaptic plasticity can be understood in the generalized Expectation Maximization framework, and a spiking network can implement a variational posterior distribution  $q(z | \mathbf{y})$  which maintains a long-term average target activity. In close analogy to the derivation in (HABENSCHUSS et al., 2012), we transfer this approach to the spatially extended sheet model in the following.

We impose a posterior constraint on the latent variables  $z$  and investigate learning in the generalized online EM framework. The EM decomposition (NEAL and G. HINTON, 1998) reads

$$\mathcal{F}(\boldsymbol{\theta}, q(z|\mathbf{y})) = \mathcal{L}(\boldsymbol{\theta}) - \langle D_{\text{KL}}(q(z|\mathbf{y}) || p(z|\mathbf{y}, \boldsymbol{\theta})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{E-step} , \quad (\text{D.24})$$

$$= \langle \log p(\mathbf{y}, z | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} + \langle H(q(z|\mathbf{y})) \rangle_{p^*(\mathbf{y})} \quad \rightarrow \text{M-step} , \quad (\text{D.25})$$

with the log-likelihood  $\mathcal{L}(\boldsymbol{\theta}) = \langle \log p(\mathbf{y} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})}$  of the input under the model, the Kullback-Leibler divergence  $D_{\text{KL}}(\cdot || \cdot)$ , and the entropy  $H(\cdot)$ . The decomposition holds for any probability distribution  $q$ , and  $q(z | \mathbf{y})$  defines a variational posterior for every input state. For this work, we constrain  $q$  to a class of ‘‘homeostatic’’ distributions,  $q \in \mathcal{Q}$ , such that each variable  $z_k$  maintains a long-term average activity  $m_k$ ,

$$\mathcal{Q} = \{q : \langle z_k \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})} = m_k \text{ for all } k = 1, \dots, K\} . \quad (\text{D.26})$$

The desired target activations  $\mathbf{m} = (m_1, \dots, m_K)$  are assumed to be compatible with the inhibition structure (D.16), e.g. by choosing  $m_k$  sufficiently small.

### E-step: Homeostatic intrinsic plasticity

During the E-step (D.24), we seek the distribution  $q^* \in \mathcal{Q}$  that minimizes the Kullback-Leibler divergence to the model posterior  $p(z | \mathbf{y}, \boldsymbol{\theta})$ , and thus, maximizes the lower bound  $\mathcal{F}$  on the likelihood  $\mathcal{L}$ . This constrained optimization problem can be solved with Lagrange multipliers. We examine the Lagrange function

$$\Lambda(q) = \langle \text{D}_{\text{KL}}(q(z' | \mathbf{y}') || p(z' | \mathbf{y}', \boldsymbol{\theta})) \rangle_{p^*(\mathbf{y}')} - \sum_{k=1}^K \beta_k \left( \langle z'_k \rangle_{q(z' | \mathbf{y}') p^*(\mathbf{y}')} - m_k \right) - \lambda \left( \langle 1 \rangle_{q(z' | \mathbf{y}')} - 1 \right) \quad (\text{D.27})$$

where the apostrophes indicate that  $\mathbf{y}'$  and  $\mathbf{z}'$  are summation variables, and  $\beta_k$  are the Lagrange multipliers for the  $K$  constraints. The additional Lagrange multiplier  $\lambda$  ensures correct normalization of  $q$ . The root of the derivative with respect to  $q(z | \mathbf{y})$ , with any particular choice of  $z$  and  $\mathbf{y}$ , fulfills  $\partial_{q(z | \mathbf{y})} \Lambda(q) = p^*(\mathbf{y}) [\log(q(z | \mathbf{y}) / p(z | \mathbf{y}, \boldsymbol{\theta})) + 1 - \lambda - \sum_k \beta_k z_k] \stackrel{!}{=} 0$  and, thus, the optimal solution  $q^*$  has the form

$$q^*(z | \mathbf{y}) = p(z | \mathbf{y}, \boldsymbol{\theta}) \exp \left[ \lambda - 1 + \sum_k \beta_k z_k \right] \propto \exp \left[ \frac{1}{2} z^T \hat{\mathbf{W}} z + z^T \mathbf{V} \mathbf{y} + z^T (\underbrace{\hat{\mathbf{b}} - \mathbf{A} + \boldsymbol{\beta}}_{:= \mathbf{b} = (b_1, \dots, b_K)}) \right]. \quad (\text{D.28})$$

Note that the variational distribution in (D.28) is already correctly normalized through the free constant  $\exp(\lambda - 1)$ .

However, the optimal multipliers  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_K)$  are still to be determined. Analogous to (GRACA et al., 2008; HABENSCHUSS et al., 2012), gradient ascent on the dual function

$$\partial_{\beta_k} \Psi(\boldsymbol{\beta}) = \partial_{\beta_k} \boldsymbol{\beta}^T \mathbf{m} - \left\langle \partial_{\beta_k} \log \sum_z p(z | \mathbf{y}, \boldsymbol{\theta}) \exp(\boldsymbol{\beta}^T z) \right\rangle_{p^*(\mathbf{y})} \quad (\text{D.29})$$

$$= m_k - \langle z_k \rangle_{p^*(\mathbf{y}) q(z | \mathbf{y})} \quad (\text{D.30})$$

yields an iterative update rule to determine the optimal Lagrange multipliers  $\beta_k$  in  $q^*$  for the E-step (D.24). During the E-step, the synaptic weights  $V_{ki}$  remain constant (synaptic weight updates are the M-step). Thus, optimizing  $\beta_k$  is equivalent to optimizing  $b_k := \hat{b}_k - A_k + \beta_k$  since  $b_k$  and  $\beta_k$  differ only by an additive constant. In particular, the update rule

$$\partial_{b_k} \Psi = \partial_{\beta_k} \Psi = m_k - \langle z_k \rangle_{p^*(\mathbf{y}) q(z | \mathbf{y})} \quad (\text{D.31})$$

remains unchanged and describes a form of homeostatic plasticity. It compares the average activation  $\langle z_k \rangle_{p^*(\mathbf{y}) q(z | \mathbf{y})}$  with the target activation  $m_k$  and adapts the intrinsic excitability accordingly: When the average activity  $\langle z_k \rangle$  is too low, the excitability  $b_k$  will be increased; if the activity is too high, the excitability will be reduced. Importantly, the homeostatic rule (D.31) requires only local information and “overwrites” the non-local terms  $A_k$  in eq. (D.28) and in  $b_k$ .

### M-step in $V_{ki}$ : Weight-dependent plasticity of afferent weights

During the M-step (D.25), we perform gradient ascent on  $\langle \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}$  with respect to the parameters  $\boldsymbol{\theta}$ . This increases the lower bound  $\mathcal{F}$  on the likelihood  $\mathcal{L}(\boldsymbol{\theta})$  since the entropy  $\langle H(q(\mathbf{z}|\mathbf{y})) \rangle_{p^*(\mathbf{y})}$  does not depend on  $\boldsymbol{\theta}$ . For the afferent weights  $\mathbf{V}$ , the derivative of the log-joint model (D.7)  $\times$  (D.17) reads

$$\partial_{V_{ki}} \mathcal{F} = \partial_{V_{ki}} \langle \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} = \langle z_k \cdot (y_i - \partial_{V_{ki}} A_{ki}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \quad (\text{D.32})$$

For a Bernoulli likelihood distribution, we obtain from eq. (D.10) that  $\partial_{V_{ki}} A_{ki} = \sigma(V_{ki} + V_{0i}) = \pi_{ki}$  and hence:

$$\partial_{V_{ki}} \mathcal{F} = \langle z_k \cdot (y_i - \sigma(V_{ki} + V_{0i})) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \quad (\text{D.33})$$

The update rule eq. (D.33) only depends on local information, namely pre- $(y_i)$  and post- $(z_k)$  synaptic activity and the current synaptic weight  $V_{ki}$ . The same holds true for Poisson distributions with  $\partial_{V_{ki}} A_{ki} = \exp(V_{ki} + V_{0i})$  and Gaussian distributions with  $\partial_{V_{ki}} A_{ki} = \sigma^2 \cdot (V_{ki} + V_{0i})$ . Intuitively, since  $\partial_{V_{ki}} A_{ki} = \langle y_i \rangle_{p(y_i | z_k=1, \boldsymbol{\theta})}$  in any natural exponential family, the plasticity rule (D.32) compares the true input value  $y_i \sim p^*(\mathbf{y})$  with the current expectation of the probabilistic model whenever  $z_k$  is active.

### M-step in $\hat{W}_{kj}^{\text{exc}}$ : Wake-sleep plasticity of recurrent weights

Similarly, we can examine the derivative of  $\mathcal{F}$  with respect to the recurrent weights  $\hat{\mathbf{W}}$  and biases  $\hat{\mathbf{b}}$  in the prior:

$$\partial_{\hat{W}_{kj}} \langle \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} = \langle z_k z_j \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} - \langle z_k z_j \rangle_{p(\mathbf{z}|\boldsymbol{\theta})} \quad (\text{D.34})$$

$$\partial_{\hat{b}_k} \langle \log p(\mathbf{y}, \mathbf{z} | \boldsymbol{\theta}) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} = \langle z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} - \langle z_k \rangle_{p(\mathbf{z}|\boldsymbol{\theta})} \quad (\text{D.35})$$

These update rules compare expected values from the variational posterior with expected values from the model prior, i.e., they are a variant of the wake-sleep algorithm. In the neural sheet model, we will apply recurrent learning only to lateral excitatory connections  $W_{kj}^{\text{exc}}$  and keep lateral inhibitory connections  $W_{kj}^{\text{inh}}$  fixed since lateral inhibition is the foundation for local synaptic plasticity rules of the afferent connections  $V_{ki}$ .

### Plasticity rules for sample-based online learning (Corollary 2)

The above learning scheme revealed local update rules for unsupervised model optimization via Generalized Expectation Maximization. The derived algorithm relies on three ingredients:

1. the variational posterior distribution  $q(\mathbf{z} | \mathbf{y})$  in eq. (D.28),

2. the homeostatic update rule (D.31) to solve the E-step, and
3. the synaptic update rule (D.33) to solve the M-step.

*Ingredient 1:* By applying the sufficient condition (5.5) on the homeostatic posterior  $q(z | \mathbf{y})$ , we find that the network can sample from  $q(z | \mathbf{y})$  when we use the same form of the membrane potential  $u_k$ , but with  $b_k := \hat{b}_k - A_k + \beta_k$  (instead of  $b_k = \hat{b}_k - A_k$  for sampling from the model posterior  $p(z | \mathbf{y}, \boldsymbol{\theta})$ ). This establishes an equivalent to Corollary 1 for the variational posterior.

*Ingredient 2:* A sample-based online approximation of eq. (D.31) is established by

$$\frac{\partial}{\partial t} b_k = \eta_b \cdot (m_k - z_k(t)) \quad , \quad (\text{D.36})$$

with  $\eta_b$  denoting a small learning rate. This homeostatic intrinsic plasticity rule approximates the expected values  $\langle z_k \rangle_{p^*(\mathbf{y})q(z|\mathbf{y})}$  in eq. (D.31) through samples  $z \sim q(z | \mathbf{y})$  in response to the input  $\mathbf{y} \sim p^*(\mathbf{y})$ . The required samples  $z \sim q(z | \mathbf{y}) \cdot p^*(\mathbf{y})$  are naturally provided by the network in response to presented input. The homeostatic plasticity rule (D.36) uses only locally available information. When homeostatic plasticity has converged, i.e., when the average update  $\langle \frac{\partial}{\partial t} b_k \rangle = 0$  vanishes for all neurons, the network implements the variational posterior distribution (D.28) with optimal multipliers  $\beta_k$ , i.e., the network solves the E-step by sampling from  $q^*(z | \mathbf{y})$ . Due to the non-infinitesimal learning rate  $\eta_b$ , this equilibrium is subject to small stochastic fluctuations.

*Ingredient 3:* A sample-based online approximation of eq. (D.33) is established by

$$\frac{\partial}{\partial t} V_{ki} = \eta_V \cdot z_k(t) \cdot (y_i(t) - \sigma(V_{ki} + V_{0i})) \quad , \quad (\text{D.37})$$

with a small learning rate  $\eta_V$ . This synaptic plasticity rule approximates the gradient in eq. (D.33) from samples  $\mathbf{y} \sim p^*(\mathbf{y})$  and the evoked response  $z \sim q(z | \mathbf{y})$ . The required samples  $z \sim q(z | \mathbf{y}) \cdot p^*(\mathbf{y})$  are provided by the network if the homeostatic posterior  $q(z | \mathbf{y})$  is correctly implemented in the E-step. Hence, it must be ensured that homeostatic intrinsic plasticity acts on significantly faster time scales than synaptic plasticity. This can be achieved by separating the time scales of intrinsic and synaptic plasticity via the learning rates  $\eta_b$  and  $\eta_V$ , such that homeostatic intrinsic plasticity can react quickly to changes in the synaptic weights (see also *Interaction of time scales* below).

Since eq. (D.36) approximates the E-step, and eq. (D.37) approximates the M-step, the joint application of these rules approximates the previously described Generalized Expectation Maximization algorithm. This proves Corollary 2.

In addition, theoretically optimal learning of recurrent connections can be realized by sample-based implementations of eq. (D.34) and (D.35). This gives rise to the plasticity rule eq. (5.16). The LTD term  $\phi_{kj}^{\text{opt}}(\hat{\mathbf{W}}^{\text{exc}}, \hat{\mathbf{W}}^{\text{inh}}, \hat{\mathbf{b}}) = \langle z_k z_j \rangle_{p(z|\boldsymbol{\theta})}$  can be determined by using an independent prior sampler, thereby giving rise to a sample-based sleep phase. In computer simulations with wake-sleep learning, we

#### D.4 Approximate plasticity rule for recurrent synapses

used such a prior sampler. This sampler maintained independent parameters  $\hat{\mathbf{b}}$  which were updated according to a sample-based approximation of eq. (D.35), i.e., according to the difference of samples from the variational posterior and the prior. Since it is not known if (or how) the LTD term  $\phi_{kj}^{\text{opt}}(\hat{\mathbf{W}}^{\text{exc}}, \hat{\mathbf{W}}^{\text{inh}}, \hat{\mathbf{b}})$  can be calculated by a single spiking network, we did not include recurrent plasticity in Corollary 2. Nevertheless, algorithmically the theory supports concurrent learning of input synapses and recurrent synapses.

#### D.4 Approximate plasticity rule for recurrent synapses

As shown above, theoretically optimal learning of recurrent weights  $W_{kj}^{\text{exc}}$  is challenging in a spiking network since the model expectations in eq. (D.34) are not directly available from the network response and demand an independent sleep phase. Therefore, we investigated to what extent the simple local plasticity rule (5.17) could entail similar weight configurations as the theoretically exact wake-sleep rule. In the following, we describe how the simple rule was obtained.

For a given learning problem, we first performed wake-sleep learning to obtain optimized weights  $W_{kj}^{\text{opt}}$ , along with covariances

$$c_{kj} := \langle (z_k - \langle z_k \rangle) \cdot (z_j - \langle z_j \rangle) \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} \quad (\text{D.38})$$

that led to these weights. Then, we fit a function  $W(c_{kj})$  to the data. For constructing the local plasticity rule, we use that  $m_k = \langle z_k \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}$  due to homeostatic intrinsic plasticity, i.e.,

$$c_{kj} = \langle z_k \cdot z_j \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} - m_k \cdot m_j \quad (\text{D.39})$$

Importantly, the information  $\langle z_k \cdot z_j \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}$  on the pre-post spike response during inference is locally available to a synapse. The following weight-dependent plasticity rule for  $W_{kj}^{\text{exc}}$  then has a fixed point at  $W_{kj}^{\text{exc}} = W(c_{kj}) \approx W_{kj}^{\text{opt}}(c_{kj})$ :

$$\frac{\partial}{\partial t} W_{kj}^{\text{exc}} = \eta_W \cdot \left[ z_k \cdot z_j - \left( m_k \cdot m_j + W^{-1}(W_{kj}^{\text{exc}}) \right) \right] \quad (\text{D.40})$$

with  $W^{-1}$  denoting the inverse function of  $W(c_{kj})$ . This follows directly from inspecting

$$0 \stackrel{!}{=} \left\langle \frac{\partial}{\partial t} W_{kj}^{\text{exc}} \right\rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} = \langle z_k \cdot z_j \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})} - m_k \cdot m_j - W^{-1}(W_{kj}^{\text{exc}}) \quad (\text{D.41})$$

and solving for  $W_{kj}^{\text{exc}}$ . For this work, we assumed the following functional form of  $W(c_{kj})$ :

$$W(c_{kj}) = \frac{W^{\text{max}}}{\pi/2} \arctan(\gamma \cdot c_{kj}) \quad (\text{D.42})$$

with two free parameters  $W^{\max}$  and  $\gamma$ . This function turned out to match the data points  $(W_{kj}^{\text{opt}}, c_{kj})$ , as obtained from wake-sleep learning, well, and results in the following heuristic plasticity rule:

$$\frac{\partial}{\partial t} W_{kj}^{\text{exc}} = \eta_W \cdot \left[ z_k \cdot z_j - \left( m_k \cdot m_j + \frac{1}{\gamma} \tan \left( \frac{\pi}{2} \frac{W_{kj}^{\text{exc}}}{W^{\max}} \right) \right) \right], \quad (\text{D.43})$$

i.e.,  $\phi_{\vartheta}(W_{kj}^{\text{exc}}) = m_k \cdot m_j + \frac{1}{\gamma} \tan \left( \frac{\pi}{2} \frac{W_{kj}^{\text{exc}}}{W^{\max}} \right)$ . While this plasticity rule preserves the symmetry of recurrent weights and features the desired fixed points  $W_{kj}^{\text{exc}} = W(c_{kj})$ , it must be noted that there exists no theoretical guarantee for convergence under recurrent network dynamics.

## D.5 Interaction of time scales

The learning rates  $\eta_b$ ,  $\eta_V$  and  $\eta_W$  control the typical time scales for significant changes in the parameters  $\mathbf{b}$ ,  $\mathbf{V}$  and  $\mathbf{W}^{\text{exc}}$ . In an online EM learning scenario, these changes are interrelated with the network's spike response and variations in the external input. In total, four different processes are to be distinguished which jointly orchestrate the learning dynamics:

1. On the fastest time scale, the synaptic time constant  $\tau$  sets the typical scale for inference and mixing during sampling from the posterior. For this study, we set  $\tau = 10$  ms.
2. On a slower time scale – let's refer to it as  $\tau_{\text{inp}}$  for this discussion – significant changes in the presented input statistics occur, i.e. the input vector  $\mathbf{y} \in \{0, 1\}^N$  switches to substantially different regions in the space of possible inputs during sampling from  $p^*(\mathbf{y})$ .
3. The variational E-step integrates the average network response  $\langle \mathbf{z} \rangle_{p^*(\mathbf{y})q(\mathbf{z}|\mathbf{y})}$ , and thus, relies on a representative coverage of the input space. Therefore, changes in  $\mathbf{b}$ , which happen on a time scale  $\tau_E \approx \eta_b^{-1}$ , must be slower than the mixing of  $\mathbf{y} \sim p^*(\mathbf{y})$ .
4. Finally, the M-step adapts synaptic weights  $\mathbf{V}$  and  $\mathbf{W}^{\text{exc}}$  based on a reliable E-step, and hence,  $\tau_M \approx \eta_V^{-1}$  (or  $\eta_W^{-1}$ ) need to be large compared to  $\tau_E$ .

In summary, from a strictly theoretical perspective we require  $\tau \ll \tau_{\text{inp}} \ll \tau_E \ll \tau_M$ . In practice, we typically find a factor 5-20 per  $\ll$ -relation to be sufficient. For instance, the network responds on the time scale of few to tens of milliseconds; input statistics vary on the time scale of hundreds of milliseconds to seconds; intrinsic neuronal excitabilities adapt on the time scale of tens of seconds to minutes; and synapses change their weight on the time scale of minutes to hours.

## D.6 Details to the computer simulations

Category	Parameter	Symbol	Fig 2	Fig 3	Fig 4	Fig 5	Comments
Simulation	Sim. time	$T$	2.5	10,000	25,000	10,000	Unit [s]
	Time step	$\delta t$	0.001	0.001	0.001	0.001	Unit [s]
Network	Network nrn's	$K$	6	21	21	144	
	Geometry		$6 \times 1$	$7 \times 3$	$7 \times 3$	$12 \times 12$	
	Neuron distance		3	$3 \times 2$	$6 \times 1$	$2 \times 2$	Rel. to inp. space
	Rec. exc. range		–	–	$\infty$	$\infty$	Max. norm
	Rec. exc. prob.		–	–	1	0.25	
	Inh. conn. dist.		3	3	1	4	Max. norm
	Inh. weight	$W^{\text{inh}}$	-100	-100	-100	-100	
	Inputs	$N$	108	126	252	576	
	Geometry		$18 \times 6$	$21 \times 6$	$42 \times 6$	$24 \times 24$	
	Afferent conn.		$6 \times 6$	$6 \times 6$	$6 \times 6$	$6 \times 6$	
	PSP-/Ref.-time	$\tau$	0.010	0.010	0.010	0.010	Unit [s]
Stimulus	Backgr. act.	$\pi_{0i}$	0.2	0.1	0.1	0.1	
	Activity range		0.2/0.55	0.1/0.6	0.1/0.4	0.1/0.5	Min./Max. $\langle y_i \rangle$
	Pattern duration		–	200ms <sup>1)</sup>	250ms	100ms	During learning
Plasticity	Hom. target	$m_k$	–	0.065	0.95/3	0.025	
	Init bias	$b_k$	$-1 - A_k$	-2	-1	-3	at $t = 0$
	Init inp. weight	$V_{ki}$	eq. (D.10)	0	0	0	at $t = 0$
	Init rec. weight	$W_{kj}^{\text{exc}}$	0 or 1	0	0	0	at $t = 0$
	Learn rate bias	$\eta_b$	0	1	0.1	10	Unit [Hz]
	Learn rate inp.	$\eta_V$	0	0.2	0.1	2	Unit [Hz]
	Learn rate rec.	$\eta_W$	0	–	0.005 <sup>2)</sup>	1	Unit [Hz]

**Table D.1: Parameters of the computer simulations.** <sup>1)</sup> Average duration of locally occurring patterns. <sup>2)</sup> 0.05 for wake-sleep learning.

## D.6 Details to the computer simulations


All computer simulations were performed with custom Python (G. VAN ROSSUM and DRAKE JR, 1995) scripts, using a discrete time version for spiking neurons and synaptic plasticity with simulation time step  $\delta t = 1$  ms and PSP time constant  $\tau = 10$  ms. In order to avoid boundary effects at the edge of the network, a torus-like network topology is used, i.e., neurons at the left-hand edge are adjacent to neurons at the right-hand edge, and neurons at the top edge are adjacent to neurons at the bottom edge of the sheet. Inhibitory connections were non-plastic. Excitatory connections, if plastic, were restricted to positive weight values. The constraint to positive weights was imposed for the purpose of neuroscientific modeling, only. The theory for inference and learning supports positive and negative weights (incl. sign changes). In the following, we first describe the implementation of neurons and synapses as used in all simulations. Then we provide specific simulation details for each figure. An overview of used parameters is provided in Table D.1.

**Stochastic neuron model (common to all figures).** We employed the simplest neuron model from (BUESING et al., 2011) with an absolute refractory period. Neurons are characterized by their membrane potential  $u_k$  and a refractory time  $\tau$  that matches the time constant the associated RV  $z_k$  is active after a spike. In discrete time, the active period lasts for  $\tilde{\tau} := \tau/\delta t = 10$  time steps. The spiking probability in each time step reads  $p(\text{spike}) = \sigma(u_k - \log \tilde{\tau})$  if the neuron is non-refractory. A neuron is non-refractory if  $z_k = 0$  or if it is in its last active time step (to allow an uninterrupted active state, see (BUESING et al., 2011)). Neurons in a network were updated sequentially such that state transitions of one cell are visible in the same time step to subsequently updated neurons. In the computer simulations, the update order was chosen randomly, in every time step. In the limit  $\delta t \rightarrow 0$  and  $\tilde{\tau} \rightarrow \infty$  while keeping  $\tau = \text{const.}$ , we obtain the above continuous-time neuron model and the sequential update policy disappears. In the other extreme, for  $\tilde{\tau} = 1$ , neurons can spike immediately again, and we recover the common Gibbs sampling algorithm. For a discrete time implementation of homeostatic intrinsic plasticity, we updated the excitability  $b$  in every time step according to  $\delta b_k = \delta t \cdot \eta_b \cdot (m_k - z_k)$ . This is a simple Euler integration of the continuous-time plasticity rule.

**Synaptic transmission and plasticity (common to all figures).** A spike of the  $i$ -th neuron elicits a rectangular post-synaptic potential (PSP) at the  $k$ -th neuron with duration  $\tau$  and amplitude  $V_{ki}$  ( $W_{ki}$ ) for afferent (recurrent) connections. Synaptic transmission has zero-delay and is non-additive, and thus, PSPs encode the value of the pre-synaptic random variable times the synaptic weight at any time. In discrete time, PSPs last for  $\tilde{\tau}$  time steps, accordingly. For a discrete time implementation of plasticity, we updated the weights  $V$ ,  $W^{\text{exc}}$  in every time step according to  $\delta V_{ki} = \delta t \cdot \eta_V \cdot z_k \cdot (y_i - \sigma(V_{ki} + V_{0i}))$  and  $\delta W_{kj}^{\text{exc}} = \delta t \cdot \eta_W \cdot (z_k \cdot z_j - \phi)$ . Here  $\phi$  denotes the LTD term of wake-sleep learning and the approximate plasticity rule, respectively. Again, this is a simple Euler integration of the continuous-time plasticity rules.

**Spiking input generation (common to all figures).** In all simulation, spiking input was presented to the network in form of Poisson spike trains with time varying firing rate. In accordance with the synaptic transmission in the network via non-additive, rectangular PSPs, an input RV  $y_i(t)$  has value 1 if a spike had occurred in the  $i$ -th input channel within  $(t - \tau, t]$ , and value 0 otherwise. Firing rates are chosen such that the expected value of  $y_i$  matches a target activity  $x_i$ , i.e.,  $\langle y_i \rangle = x_i$ . The specific target values  $x_i$  used for the simulations are provided below for each figure. In order to achieve  $\langle y_i \rangle = x_i$  in a discrete time simulation, the spiking probability of an input is set to  $p_i := p(\text{spike in } \delta t) = 1 - (1 - x_i)^{1/\tilde{\tau}}$  for each time step. This assignment originates from the following thought. An input is inactive when there was no spike in the last  $\tilde{\tau}$  time steps, i.e.,  $p(y_i = 0) = (1 - p_i)^{\tilde{\tau}}$ . The expected value  $\langle y_i \rangle$  is thus given by  $\langle y_i \rangle = p(y_i = 1) = 1 - p(y_i = 0) \stackrel{!}{=} x_i$ . Solving for  $p_i$





**Fig. D.1: Sampling quality and heuristic learning rule.** (A) Sampling quality of the spiking network. Red: Histogram over the difference between the traces in Fig. 5.2D for every neuron and time point. For visual clarity, the data in Fig. 5.2D had been smoothed with a 20ms box kernel. Gray: Histogram over the non-smoothed, raw data. (B) Recurrent plasticity function in the setup of Fig. 5.4. Weights obtained with wake-sleep learning vs. the covariance of network variables (blue dots), alongside the fitted plasticity function  $W(c_{kj})$  (red line). (C) Same for Fig. 5.5.

yields the above assignment. In the limit  $\delta t \rightarrow 0$  such that  $\tau = \tilde{\tau} \cdot \delta t = const.$ , these spiking dynamics yield a Poisson process with firing rate  $\log[(1 - x_i)^{-1/\tau}]$ .

### Figure 2: Sheets of spiking neurons can perform Bayesian inference on distributed spiking input

The network consists of  $N = 18 \times 6$  input neurons and  $K = 6$  network neurons. Each network neuron receives local input from  $6 \times 6$  inputs, with local connections being shifted by 3 between neighboring network neurons. Neurons with overlapping input inhibit each other, resulting in nearest-neighbor inhibition. In addition, three neuron pairs  $((k, j) = (1, 3), (3, 5)$  and  $(4, 6))$  maintain excitatory recurrent connections of weight  $W_{kj}^{exc} = W_{jk}^{exc} = 1$ . Preferred local activity patterns  $x_{kj} \in (0.2, 0.55)$ ,  $1 \leq k \leq 6$  and  $1 \leq j \leq 36$ , were drawn for each network neuron from a uniform distribution. Background activity was set to  $\pi_{0i} = 0.2$ , and afferent synaptic weights were set to match the patterns according to eq. (D.10). Neuronal excitabilities were set to  $b_k = \hat{b}_k - A_k$  with  $\hat{b}_k = -1$  and  $A_k = \sum_i A_{ki}$  being calculated for each neuron according to eq. (D.10).

Spiking input was generated with the aim that the input distribution  $p^*(\mathbf{y})$  closely resembles the model distribution  $p(\mathbf{y} | \boldsymbol{\theta})$  of the network: At each location and time point, at most one local input pattern  $x_{kj}$  is active. If pattern  $x_{kj}$  is active, it governs the firing rate of all those inputs that are connected to neuron  $k$ , i.e.,  $\langle y_i \rangle = x_{kj}$  for  $i = [(18 \cdot (k - 1) + (j - 1) - 6) \bmod N] + 1$ . The presence of an input pattern is indicated by colored spikes in Fig. 5.2E. If no dedicated pattern is active, inputs fire with the background activity, i.e.,  $\langle y_i \rangle = \pi_{0i}$  (gray spikes). The data shown in Fig. 5.2 and evaluated Fig. D.1A covers 1.5s simulation time. The total simulation time was 2.5s with a short period before and after the shown data being discarded. This serves to provide a burn-in phase for the sampling network, and to prevent boundary artifacts when smoothing the posterior marginals. In panel D, the posterior marginals  $p_{net}$  were estimated from the network response  $z(t)$  during 1000 simulation runs, all with exactly the same input (spike-level identity). The correct posterior  $p_{theo}$  is given by eq. (D.18). For visual clarity, the traces of the posterior marginals have been smoothed with a 20 ms box kernel.

A comparison of non-smoothed data is provided in Fig. D.1A. The figure shows a more systematic analysis of the sampling quality of the spiking network by means of a histogram over the mismatch  $(p_{net}(z_k = 1 | \mathbf{y}(t)) - p_{theo}(z_k = 1 | \mathbf{y}(t)))$ . The red dotted counts were evaluated on the smoothed data of Fig. 5.2D; the gray bars depict the non-smoothed (raw) data, evaluated on a millisecond basis. This quantitative analysis confirms the excellent approximation quality of the sampling network. In conclusion, the data shown in Fig. 5.2D and Fig. D.1A indicate that the sampling network can calculate and represent the general structure as well as quantitative specificities of the time-varying posterior distribution with high accuracy. However, small differences between the traces in Fig. 5.2D are visible. Most notably, rapid and sharp peaks in the posterior, which arise from pronounced but transient jumps in the input, are not fully integrated by the network. The origin of these deviations is of stochastic and systematic nature. Stochastic fluctuations arise from the general sample-based representation. For time-varying input signals  $\mathbf{y}(t)$ , only few independent samples can be drawn from the target posterior distribution  $p(z | \mathbf{y}(t), \theta)$  under (almost) stable conditions, i.e., before the target distribution changes. Any representation based on a limited number of samples can approximate the posterior only with limited precision. Systematic deviations result from incomplete convergence of the Markov chain. The Markov chain, that underlies the network dynamics, is guaranteed to converge to the correct equilibrium distribution  $p(z | \mathbf{y}, \theta)$  only for any constant input  $\mathbf{y}$  in the limit  $t \rightarrow \infty$ . For time-varying input  $\mathbf{y}(t)$ , convergence of the network will typically “lag behind” the “moving target”  $p(z | \mathbf{y}(t), \theta)$ . Theoretical work (HABENSCHUSS et al., 2013a) has shown that the network distribution converges exponentially fast to its equilibrium in almost arbitrary network architectures. Notably, the local WTA architecture is expected to facilitate mixing of the Markov chain since typically only a few competing neurons will attempt to fire in response to the presented input at the same time.

### Figure 3: Emergence of probabilistic local experts through synaptic plasticity

The network consists of  $N = 21 \times 6$  input neurons and  $K = 7 \times 3$  network neurons. Network neurons are organized in 7 local populations. The 3 neurons within a population share the same  $6 \times 6$  field of afferent connections. Afferent fields of neighboring populations are shifted by 3 (measured in the domain of input neurons). Due to the torus-like topology, every network neuron has overlapping inputs with 8 other network neurons (2 in its population and  $2 \times 3$  in the neighboring populations). This is the range of lateral inhibition.

Spatio-temporal spiking input (that determines the samples  $\mathbf{y}$  and thus  $p^*(\mathbf{y})$ ) was generated as follows. For each of the 7 afferent field locations, 3 random activity patterns  $x_{lj}^p$  ( $1 \leq p \leq 3$ ,  $1 \leq l \leq 7$ ,  $1 \leq j \leq 36$ ), were drawn. To facilitate the generation of locally different activity patterns, each input location was drawn from a Dirichlet distribution and scaled to the activity range  $[0.1, 0.6]$ :  $(x_{lj}^1, x_{lj}^2, x_{lj}^3) \sim$

$0.5 \cdot \text{Dir}(0.3, 0.3, 0.3) + 0.1$ . Whenever a local activity pattern  $x_l^p$  was presented, Poisson spike trains were generated such that  $\langle y_i \rangle = x_{l_j}^p$  with  $i = (18 \cdot l + j \bmod N) + 1$ . The presence of local activity patterns  $x_l^p$  was determined as follows. Three chains  $c = 1, 2, 3$  were started at time  $t = 0$ . Each chain can either be active or inactive. If it is active, it appears at a location  $l$  and presents one of the local activity patterns  $x_l^p$ . Initially, all chains were inactive and the initial duration of inactivity (in ms) was drawn for each chain from a Gamma distribution  $\Gamma(k = 10, \theta = 10)$ , leading to an average initial inactivity of 100 ms. Whenever inactivity of a chain ends, it turns active for a duration (in ms) drawn from  $\Gamma(k = 10, \theta = 20)$ , leading on average to 200 ms duration of activity. When a chain turns active, a random pattern  $p$  is drawn uniformly, and a random location  $l$  is drawn such that the invoked activity pattern  $x_l^p$  does not overlap with the local activity pattern of a different currently active chain. Such a valid location always exists in the given architecture. After a chain's active phase, it turns inactive again for a duration (in ms) drawn from  $\Gamma(k = 10, \theta = 10)$ . Inputs, that are not covered by a currently active chain, maintain a background activity with  $\langle y_i \rangle = \pi_{0i} = 0.1$ . This process leads to spatially non-overlapping, but temporally interleaved input spike patterns as shown in Fig. 5.3F.

The network was exposed to this spiking input for 10000 s. Afferent weights  $V_{ki}$  and intrinsic excitabilities  $b_k$  were plastic; recurrent connections  $W_{kj}^{\text{exc}}$  and  $W_{kj}^{\text{inh}}$  were non-plastic, and  $W_{kj}^{\text{exc}} = 0$ .

Details to the plotting: Panel B shows the activity patterns  $x_{l_j}^p$ ,  $p = 1, 2, 3$ , for the input region highlighted in panel A. For estimating the log-likelihood in panel C,  $S = 10000$  input samples  $\mathbf{y}^s$  were randomly drawn from the training data as a proxy for  $p^*(\mathbf{y})$ . Furthermore, the network has 337 possible  $z$ -states that respect the inhibition structure. Thus, the log-likelihood can be calculated via  $\mathcal{L}(\boldsymbol{\theta}) \approx \frac{1}{S} \sum_s \log \sum_z p(\mathbf{z}) \cdot \prod_i p(y_i^s | \mathbf{z}, \boldsymbol{\theta})$ , with  $p(y_i^s | \mathbf{z}, \boldsymbol{\theta})$  given by eq. (D.6), for any afferent weight configuration  $\mathbf{V}$  that emerges over the course of learning. The joint distribution depends on the parameters  $\mathbf{W}^{\text{inh}}$ ,  $\mathbf{V}$  and  $\hat{\mathbf{b}}$ . While the synaptic parameters  $\mathbf{W}^{\text{inh}}$ ,  $\mathbf{V}$  are directly accessible in the spiking network, the biases  $\hat{\mathbf{b}}$  in the prior must be determined differently. We calculated the biases  $\hat{\mathbf{b}}$  offline such that the prior exhibited the homeostatic target activity, i.e.  $\langle z_k \rangle_{p(\mathbf{z} | \boldsymbol{\theta})} = m_k$  for all  $k$ . This is a canonical choice since these are the biases  $\hat{\mathbf{b}}$  a Bayesian observer would determine from observing the network response. In panel D, optimal weights were calculated from the generating input patterns  $x_{l_j}^p$  according to eq. (D.10). This is possible since the data distribution  $p^*(\mathbf{y})$  is structurally similar to the model distribution  $p(\mathbf{y} | \boldsymbol{\theta})$ . For each of the seven input locations, each of the three local network neurons was assigned to the best matching pattern of optimal weights. The assignment was unambiguous, since each network neuron had clearly specialized on one of the local input patterns, and determines the one-to-one mapping between learned weights  $V_{ki}$  and optimal weights plotted in panel D. For panels E and G, four additional simulations were run with network parameters  $(\mathbf{V}(t), \mathbf{b}(t))$  taken from different training time points  $t = 0 \text{ s}, 1000 \text{ s}, 3000 \text{ s}, 10000 \text{ s}$ . Identical

100s-spike patterns were presented to the network in these four simulations. Panel G shows 2.5 s of the input spike pattern alongside the network response for these simulations. Panel F shows the corresponding afferent weights  $V(t)$  for the three highlighted network neurons. For the 2-dimensional linear projection in panel E, the input states  $\mathbf{y}(t)$  of each 100s-simulation were sampled every 10 ms (i.e. 10000 data points per scatter plot) and projected onto the 2d plane. The color of each data point is determined by the network response: red, green, blue if one of the neurons marked in panel A responded; and gray otherwise. The projection plane is spanned by the two leading principle components (PCA) of those input samples the three highlighted network neurons responded to at the end of learning, i.e., the PCA is based on the colored samples in the rightmost panel. This biased selection only concerns the choice of the projection plane with the aim to visually discern the clusters of interest; the plotted data points are unbiased.

#### Figure 4: Plastic recurrent synapses integrate structural knowledge

The network consists of  $N = 7 \times 6 \times 6$  input neurons and  $K = 7 \times 3$  network neurons. Network neurons are organized in 7 local populations. The 3 neurons within a population share the same  $6 \times 6$  field of afferent connections. The afferent fields of different populations are disjoint. Network neurons within the same population share lateral inhibition. Network neurons, which belong to different populations, maintain excitatory recurrent connections  $W_{kj}^{\text{exc}} \geq 0$ .

For the spiking input, a simpler temporal input structure than in Figure 5.3 was used since the focus of this simulation was set at the correlation structure of the input beyond the range of individual input fields. There are three local activity patterns  $x_l^p$ ,  $p = 1, 2, 3$ , referred to as the red, green and blue pattern, each consisting of “vertical stripes” at shifted locations:  $x_{lj}^p = 0.4$  if  $\lfloor (j-1)/3 \rfloor \bmod 3 = (p-1)$ , and 0.1 otherwise, with  $\lfloor \cdot \rfloor$  denoting the floor function. To generate a global activity pattern  $x$ , two (potentially identical) cue patterns  $x_l^p$  were picked for the outermost locations  $l = 1$  and  $l = 7$ . Then the inner locations  $l = 2 - 6$  were filled with a consistent valid pattern, with “validity” referring to the condition to show a pattern that differs from the cues. Thus, in case of two different cue patterns, the choice of the inner pattern was fully determined (e.g. a red and green cue leads to blue inner patterns); in case of two identical cues, the inner pattern could show either of two valid patterns with all inner locations showing the same pattern (e.g. a double-green cue leads to either all-red or all-blue inner patterns). Hence, both cues ( $l = 1$  and  $l = 7$ ) must be taken into consideration to determine the validity of inner patterns ( $2 \leq l \leq 6$ ) during inference, and all inner patterns are supposed to be of equal type. As in previous simulations, input Poisson spike trains were generated such that the average value  $\langle \mathbf{y} \rangle$  matched the activity pattern  $x$ . Input activity patterns  $x$  were presented for a fixed duration before the global activity pattern switched: The patterns iterated over the nine possible cue combinations. In case of differently colored cues the global pattern was presented for 500 ms; in case of identical cues the two valid global patterns were presented for 250 ms each.

Thus all cue combinations were presented for an equal amount of time. After all cues were presented, the presentation was repeated.

For the learning experiments, the network was exposed to this spiking input for 25000 s. In total, three learning experiments were conducted with recurrent plasticity being (a) governed by the theoretically optimal wake-sleep rule, (b) governed by the simple heuristic rule, (c) switched off. In all simulations, afferent weights  $V_{ki}$  and intrinsic excitabilities  $b_k$  were plastic. In (a) and (b) recurrent connections  $W_{kj}^{\text{exc}}$  were plastic and restricted to positive (excitatory) weight values. For (a) “wake-sleep learning”, the theoretically derived learning rule (D.34) was used with  $\eta_W = 0.05$ , and prior samples being drawn from an independent sampling network which shared its recurrent weights  $W_{kj}^{\text{exc}}$  with the learning network but maintained independent (homeostatically regulated) biases and was not exposed to any input. After learning, the covariances  $c_{kj}$  were calculated from the posterior samples  $z(t)$  of the last 1000s of the simulation. To obtain data points  $(W_{kj}^{\text{exc}}, c_{kj})$  for fitting the function  $W(c_{kj})$ , only excitatory synapses were considered that had a weight  $W_{kj}^{\text{exc}} > 0.01$  at the end of learning. This is to prevent distortions in the fit due to synapses between negatively correlated neurons that would have developed negative weights during wake-sleep learning (but were bounded to  $W_{kj}^{\text{exc}} \geq 0$  in the simulation). Fitting the function (D.42) to this data yielded  $W_{\text{max}} = 1.4113$  and  $\gamma = 31.606$  for the free parameters. Data points and fitted function are shown in Fig. D.1. For (b) “heuristic learning”, the fitted learning rule was used with  $\eta_W = 0.005$ . All recurrent weights  $W_{kj}^{\text{exc}}$  converged to stable values. Weights connecting inner neurons ( $2 \leq l \leq 6$ ), that had specialized on equal activity patterns, settled at  $W_{kj}^{\text{exc}} \approx 1.27$ . Weights between cue neurons and compatible inner neurons settled at  $W_{kj}^{\text{exc}} \approx 0.90$ . Weights between cue neurons responsive to the same pattern settled at  $W_{kj}^{\text{exc}} \approx 0.32$ . All other weights settled close to zero ( $W_{kj}^{\text{exc}} < 0.003$ ).

For the demonstration of inference in face of incomplete observations (panels E and F), activity patterns  $\mathbf{x}$  were generated as follows. Two cues were chosen at the outer locations. These cues had increased contrast to ensure that the spike pattern at the cue locations was unambiguous:  $x_{lj}^p = 0.6$  if  $\lfloor (j-1)/3 \rfloor \bmod 3 = (p-1)$ , and 0.1 otherwise (for  $l = 1$  and  $l = 7$ ). All inner locations  $l = 2 - 6$  showed uninformative uniform activity of moderate intensity:  $x_{lj} = (0.6 + 0.1)/2 = 0.35$  for all  $j$ . We refer to the uninformative patterns as “gray” patterns. For each cue combination, Poisson spike trains with  $\langle \mathbf{y} \rangle = \mathbf{x}$  were presented to the network for 100 s. The vertical bars in panel E and F show the mean activity  $\langle z_k \rangle$  of each network neuron during the simulation given an unambiguous cue (panel E) and an ambiguous cue (panel F).

The log-likelihood  $\mathcal{L}(\boldsymbol{\theta}(t))$  in panel D was estimated as follows.  $S = 10000$  input samples  $\mathbf{y}^s$  were randomly drawn from the training data as a proxy for  $p^*(\mathbf{y})$ . Furthermore, the network has  $4^7 = 16384$  possible  $z$ -states (zero or one active neuron in each local population). Thus, the log-likelihood can be calculated via  $\mathcal{L}(\boldsymbol{\theta}) \approx \frac{1}{S} \sum_s \log \sum_z p(z) \cdot \prod_i p(y_i^s | z, \boldsymbol{\theta})$  for each time point  $t$  and each network

$z, z_k$	Network neurons	$k = 1, \dots, K$	Index for $z$
$y, y_i$	Input neurons	$i = 1, \dots, N$	Index for $y$
$x, x_i$	Activity patterns	$\mathcal{I}_k \subseteq \{1, 2, \dots, N\}$	Afferent field (index set)
$\tau$	PSP and ref. time constant	$\mathcal{P}_i \subseteq \{1, 2, \dots, K\}$	Projection field (index set)
$\rho_k$	Network neuron firing rate	$u_k$	Membrane potential
$\theta$	Model params	$Z$	Partition function prior
$\hat{b}, \hat{b}_k$	Prior bias vector	$b, b_k$	Neuronal excitability
$\hat{W}, \hat{W}^{\text{exc}}, \hat{W}^{\text{inh}}$	Prior weight matrices	$\mathbf{W}^{\text{exc}}, \mathbf{W}^{\text{inh}}$	Exc. and inh. weight matrix
$\pi_k, \pi_{ki}$	Cluster center (Bernoulli)	$\mathbf{V}, V_{ki}$	Afferent weights
$\pi_0, \pi_{0i}$	Default hypothesis	$\mathbf{V}_0, V_{0i}$	Exp. family default hypo.
$A_k, A_{ki}$	Log-partition (likelihood)	$\mathbf{m}, m_k$	Target activation
$h(\mathbf{y}), h_i(y_i)$	Base measure (likelihood)	$\beta, \beta_k$	Hom. Lagrange mult.
$t$	Current time	$T$	Total simulation time
$p(\cdot   \theta)$	Model related prob. distr.	$\delta t$	Simulation time step
$p(z   \mathbf{y}, \theta)$	Conditional probability	$\delta W_{kj}, \delta V_{ki}, \delta b_k$	Discrete time plast. updates
$p^*(\cdot)$	Data related prob. distr.	$\eta_b, \eta_V, \eta_W$	Learning rates
$q(\cdot), q^*(\cdot)$	Variational posterior	$\phi_\theta, \phi^{\text{opt}}$	LTD term in rec. learning
$\langle x \rangle_{p(x)}$	Expected value	$W_{kj}^{\text{opt}}$	Opt. weights from WS learning
$\mathcal{Q}$	Homeostatic constraint set	$W(c_{kj})$	Simple rule fit function
$c_{kj}$	Var. posterior covariance	$\gamma, W^{\text{max}}$	Simple rule parameters
$\mathcal{L}(\theta)$	Log-likelihood	$\lambda$	Lagrange mult. for norm. to 1
$\mathcal{F}$	Gen. EM lower bound	$H(p)$	Entropy
$D_{\text{KL}}(q    p)$	Kullback-Leibler divergence	$\Lambda, \Psi$	Lagrange and Dual function

**Table D.2: List of symbols used in Appendix D.**

type (a)-(c). The joint distribution depends on the parameters  $\mathbf{W}^{\text{exc}}, \mathbf{W}^{\text{inh}}, \mathbf{V}$  and  $\hat{b}$ . While the synaptic parameters  $\mathbf{W}^{\text{exc}}, \mathbf{W}^{\text{inh}}, \mathbf{V}$  are directly accessible in the spiking network, the biases  $\hat{b}$  in the prior must be determined differently. We calculated the biases  $\hat{b}$  offline for each weight configuration ( $\mathbf{W}^{\text{exc}}(t), \mathbf{W}^{\text{inh}}(t)$ ) such that the prior matched the homeostatic target activity, i.e.  $\langle z_k \rangle_{p(z|\theta)} = m_k$ . The log-likelihood is shown only for the first 10,000 of the simulation in order to highlight the early stage of learning.

### Figure 5: Emergence of excitatory subnetworks in neural sheets

The network consists of  $N = 24 \times 24$  input neurons and  $K = 12 \times 12$  network neurons. Network neurons are organized in a sparse grid with twice the distance of inputs. Each network neuron maintains afferent connections with  $6 \times 6$  inputs, such that the input field of neighboring neurons is shifted by two. Lateral inhibition has a range of 2 (in the network grid; maximum norm). Additionally, any pair of network neurons (beyond the range of inhibition) maintains a reciprocal excitatory connection  $W_{kj}^{\text{exc}}$  with 25% probability. All excitatory connections are plastic with initial values  $V_{kj}(t=0) = 0$  and  $W_{kj}^{\text{exc}}(t=0) = 0$ .

Spiking input is composed of three prototypic rate patterns  $x_i^p, 1 \leq p \leq 3$  (grid, diagonal stripes, checkerboard), occurring locally at random locations (cp. panel

B). Activity patterns are binary with a high rate of  $x_i^p = 0.5$  and a low rate of  $x_i^p = \pi_{0i} = 0.1$ . The diameter of local patterns is variable, but exceeds the  $6 \times 6$  input field of individual network neurons. Two randomly selected (and possibly equal) patterns are presented simultaneously at non-overlapping locations. New patterns and locations are drawn every 100ms during training. All inputs not covered by a pattern fire with the background activity  $\langle y_i \rangle = 0.1$ . Homeostatic target activations  $m_k = 0.025$  are chosen such that the network explains on average approx. 130 inputs; this is roughly half of the average area covered by rate patterns  $x_i^p$ , i.e., ca. 50% of the input is on average being explained by the network.

The network was first trained with wake-sleep learning and variable learning rates until all parameters had converged to stable values. The simple learning rule was fitted to the resulting  $(W_{kj}^{\text{exc}}, c_{kj})$ -pairs just as for Fig. 5.4, yielding parameters  $W_{\text{max}} = 2.7049$  and  $\gamma = 733.69$ . See Fig. D.1 for data and fitted function  $W(c_{kj})$ . The high value of the sensitivity  $\gamma$ , compared to Fig. 5.4, likely originates from the generally much lower covariance  $c_{kj}$ , which in turn arises from the lower average network activity  $m_k$ .

For the subsequent spiking network simulation with the simple plasticity rule, a new network with different connected pairs of network neurons was generated. Thus, the extracted plasticity rule is only tailored to the general learning setup, but not to a specific network instance. The total simulation time was  $T = 10,000$ s. Learning rates were set to quite high values ( $\eta_b = 10$ ,  $\eta_V = 2$ ,  $\eta_W = 1$ ) since the focus of this simulation was on the general structure of emerging weight configurations, rather than on numerical precision. After training, recurrent weights covered the entire range of positive weight values with  $\max[W_{kj}^{\text{exc}}] = 2.275$ . To verify that the emergent weight configuration was stable even in face of high learning rates, the simulation was continued for another 10,000s, showing no signs of instability.

Details to the plotting: For panels C and D, the trained network was exposed to spiking input with patterns switching every 500ms. This was to obtain more stable estimates for the “expected input”  $\langle y^{\text{gen}} \rangle$  in panel D. For panel E, all neurons could unambiguously be labeled to be responsive to one of the three local patterns, since their afferent weights  $V_{ki}$  showed an evident preference for either one of them. Panel F uses the euclidean distance in the lattice coordinates of the network neurons, i.e., directly neighboring network neurons have distance one.





## Bibliography

- ACKLEY, D., G. HINTON, and T. SEJNOWSKI (1985). "A learning algorithm for Boltzmann machines." In: *Cognitive Science* 9.1, pp. 147–169 (cit. on pp. 4, 16, 20, 35).
- ADESNIK, H. and M. SCANZIANI (2010). "Lateral competition for cortical space by layer-specific horizontal circuits." In: *Nature* 464.7292, pp. 1155–1160 (cit. on p. 108).
- ALAIS, D. and R. BLAKE (2005). *Binocular rivalry*. MIT press (cit. on pp. 30, 34, 35).
- ANDRIEU, C., N. D. FREITAS, A. DOUCET, and M. I. JORDAN (2003). "An Introduction to MCMC for Machine Learning." In: *Mach Learn* 50, pp. 5–43 (cit. on pp. 7, 17, 22).
- ANGELAKI, D. E., Y. GU, and G. C. DEANGELIS (2009). "Multisensory integration: psychophysics, neurophysiology and computation." In: *Current opinion in neurobiology* 19.4, pp. 452–458 (cit. on pp. 3, 38, 76).
- ANTIC, S. D., W. L. ZHOU, A. R. MOORE, S. M. SHORT, and K. D. IKONOMU (2010). "The decade of the dendritic NMDA spike." In: *J Neurosci Res* 88.14, pp. 2991–3001 (cit. on p. 36).
- AVERMANN, M., C. TOMM, C. MATEO, W. GERSTNER, and C. C. PETERSEN (2012). "Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex." In: *Journal of neurophysiology* 107.11, pp. 3116–3134 (cit. on pp. 108, 109).
- AZOUZ, R. and C. M. GRAY (1999). "Cellular mechanisms contributing to response variability of cortical neurons in vivo." In: *J Neurosci* 19, pp. 2209–2223 (cit. on p. 14).
- BARBER, D. (2012). *Bayesian reasoning and machine learning*. Cambridge University Press (cit. on p. 4).
- BARTELS, A. and N. K. LOGOTHETIS (2010). "Binocular rivalry: a time dependence of eye and stimulus contributions." In: *Journal of Vision* 10.12, p. 3 (cit. on pp. 30, 34, 35).
- BARTH, A. L. and J. F. POULET (2012). "Experimental evidence for sparse firing in the neocortex." In: *Trends in neurosciences* 35.6, pp. 345–355 (cit. on p. 109).
- BERKES, P., G. ORBÁN, M. LENGYEL, and J. FISER (2011). "Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment." In: *Science* 331.6013, p. 83 (cit. on pp. 4, 34, 35, 38, 95, 109, 110).
- BERNIKER, M., M. VOSS, and K. KORDING (2010). "Learning priors for bayesian computations in the nervous system." In: *PloS one* 5.9, e12686 (cit. on p. 76).
- BI, G.-Q. and M. POO (1998). "Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type." In: *The Journal of neuroscience* 18.24, pp. 10464–10472 (cit. on pp. 51, 56).
- BILL, J. and R. LEGENSTEIN (2014). "A compound memristive synapse model for statistical learning through STDP in spiking neural networks." In: *Frontiers in neuroscience* 8 (cit. on p. 113).

## Bibliography

- BINZEGGER, T., R. J. DOUGLAS, and K. A. MARTIN (2004). "A quantitative map of the circuit of cat primary visual cortex." In: *J. Neurosci.* 24.39, pp. 8441–8453 (cit. on p. 33).
- BINZEGGER, T., R. DOUGLAS, and K. MARTIN (2009). "Topology and dynamics of the canonical circuit of cat V1." In: *Neural Netw* 22.8. Cortical Microcircuits, pp. 1071–1078 (cit. on p. 33).
- BISHOP, C. M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer (cit. on pp. 4, 5, 7, 8, 34, 41, 46, 76, 79, 90).
- BLAKE, R. and N. K. LOGOTHETIS (2002). "Visual competition." In: *Nat Rev Neurosci* 3, pp. 13–21 (cit. on pp. 30, 34, 35).
- BOSKING, W. H., Y. ZHANG, B. SCHOFIELD, and D. FITZPATRICK (1997). "Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex." In: *The Journal of Neuroscience* 17.6, pp. 2112–2127 (cit. on pp. 80, 82, 109).
- BRASCAMP, J. W., R. van EE, A. J. NOEST, R. H. A. H. JACOBS, and A. V. van den BERG (2006). "The time course of binocular rivalry reveals a fundamental role of noise." In: *J Vis* 6, pp. 1244–1256 (cit. on p. 14).
- BREA, J., W. SENN, and J.-P. PFISTER (2012). "Sequence learning with hidden units in spiking neural networks." In: *Proc. of NIPS 2011*. Vol. 24. MIT Press, pp. 1422–1430 (cit. on pp. 4, 38).
- BRUEDERLE, D., J. BILL, B. KAPLAN, J. KREMKOW, K. MEIER, E. MULLER, and J. SCHEMMELE (2010). "Live demonstration: Simulator-like exploration of cortical network architectures with a mixed-signal VLSI system." In: *Proceedings of the IEEE International Symposium on Circuits and Systems; 30 May - 2 June 2010; Paris, France. ISCAS 2010*. (Cit. on p. 36).
- BUESING, L., J. BILL, B. NESSLER, and W. MAASS (Nov. 2011). "Neural Dynamics as Sampling: A Model for Stochastic Computation in Recurrent Networks of Spiking Neurons." In: *PLoS Comput Biol* 7.11, e1002211 (cit. on pp. 38, 46–48, 77, 78, 80, 82–84, 104, 107, 110, 160).
- CANNON, R., C. O'DONNELL, and M. NOLAN (2010). "Stochastic ion channel gating in dendritic neurons: morphology dependence and probabilistic synaptic activation of dendritic spikes." In: *PLoS Comput Biol* 6.8, e1000886 (cit. on p. 14).
- CAPORALE, N. and Y. DAN (2008). "Spike timing-dependent plasticity: a Hebbian learning rule." In: *Annu. Rev. Neurosci.* 31, pp. 25–46 (cit. on p. 51).
- CARREIRA-PERPINAN, M. A. and G. E. HINTON (2005). "On contrastive divergence learning." In: *Proceedings of the tenth international workshop on artificial intelligence and statistics*. Citeseer, pp. 33–40 (cit. on p. 96).
- CHICCA, E., F. STEFANINI, C. BARTOLOZZI, and G. INDIVERI (2014). "Neuromorphic electronic circuits for building autonomous cognitive systems." In: *Proceedings of the IEEE* 102.9, pp. 1367–1388 (cit. on p. 69).
- CHOI, H., H. JUNG, J. LEE, J. YOON, J. PARK, D.-J. SEONG, W. LEE, M. HASAN, G.-Y. JUNG, and H. HWANG (2009). "An electrically modifiable synapse array of resistive switching memory." In: *Nanotechnology* 20.34, p. 345201 (cit. on p. 50).
- CHUA, L. O. (1971). "Memristor—the missing circuit element." In: *Circuit Theory, IEEE Transactions on* 18.5, pp. 507–519 (cit. on pp. 12, 113).

- CHURCHLAND, M. M., B. M. YU, J. P. CUNNINGHAM, L. P. SUGRUE, M. R. COHEN, G. S. CORRADO, W. T. NEWSOME, A. M. CLARK, P. HOSSEINI, B. B. SCOTT, D. C. BRADLEY, M. A. SMITH, A. KOHN, J. A. MOVSHON, K. M. ARMSTRONG, T. MOORE, S. W. CHANG, L. H. SNYDER, S. G. LISBERGER, N. J. PRIEBE, I. M. FINN, D. FERSTER, S. I. RYU, G. SANTHANA, M. SAHANI, and K. V. SHENOY (2010). "Stimulus onset quenches neural variability: a widespread cortical phenomenon." In: *Nat Neurosci* 13.3, pp. 369–378 (cit. on p. 34).
- DEMPSTER, A. P., N. M. LAIRD, and D. B. RUBIN (1977). "Maximum likelihood from incomplete data via the EM algorithm." In: *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 1–38 (cit. on pp. 62, 143).
- DENEVE, S. (2008a). "Bayesian spiking neurons II: Learning." In: *Neural Computation* 20.1, pp. 118–145 (cit. on pp. 4, 38, 41).
- DENEVE, S. (2008b). "Bayesian spiking neurons I: Inference." In: *Neural Comput* 20.1, pp. 91–117 (cit. on pp. 4, 15, 35, 38).
- DENEVE, S. (2005). "Bayesian inference in spiking neurons." In: *Advances in neural information processing systems* 17, pp. 353–360 (cit. on pp. 4, 77, 111).
- DENISON, S., E. BONAWITZ, A. GOPNIK, and T. GRIFFITHS (2009). "Preschoolers sample from probability distributions." In: *Proceedings of the 32nd Annual Conference of the Cognitive Science Society; 29 July - 1 August 2009; Amsterdam, Netherlands. CogSci 2009*. (Cit. on pp. 3, 34).
- DESAI, N., L. RUTHERFORD, and G. TURRIGIANO (1999). "Plasticity in the intrinsic excitability of cortical pyramidal neurons." In: *Nature Neuroscience* 2.6, p. 515 (cit. on pp. 39, 42).
- DOUGLAS, R. J. and K. A. MARTIN (2004). "Neuronal circuits of the neocortex." In: *Annu. Rev. Neurosci.* 27, pp. 419–451 (cit. on pp. 57, 80, 86, 109).
- DOYA, K., S. ISHII, A. POUGET, and R. P. N. RAO (2007). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT-Press (cit. on pp. 3, 14, 15).
- ECKER, A. S., P. BERENS, G. A. KELIRIS, M. BETHGE, N. K. LOGOTHETIS, and A. S. TOLIAS (2010). "Decorrelated neuronal firing in cortical microcircuits." In: *Science* 327.5965, pp. 584–587 (cit. on p. 109).
- FAISAL, A. A., L. P. SELEN, and D. M. WOLPERT (2008). "Noise in the nervous system." In: *Nature Reviews Neuroscience* 9.4, pp. 292–303 (cit. on p. 3).
- FANG, Z., H. YU, X. LI, N. SINGH, G. LO, and D. KWONG (2011). "HfOx/TiOx/HfOx/TiOx Multilayer-Based Forming-Free RRAM Devices With Excellent Uniformity." In: *IEEE Electron Device Letters* 32 (cit. on pp. 51, 53, 70).
- FETSCH, C., A. POUGET, C. DEANGELIS, and D. ANGELAKI (2012). "Neural correlates of reliability-based cue weighting during multisensory integration." In: *Nat Neurosci* 15, pp. 146–154 (cit. on p. 4).
- FISER, J., P. BERKES, G. ORBAN, and M. LENGYEL (2010). "Statistically optimal perception and learning: from behavior to neural representation." In: *Trends in Cogn. Sciences* 14.3, pp. 119–130 (cit. on pp. 4, 16, 34, 38, 77, 110, 111).
- FISER, J., C. CHIU, and M. WELIKY (2004). "Small modulation of ongoing cortical dynamics by sensory input during natural vision." In: *Nature* 431, pp. 573–583 (cit. on pp. 3, 14, 33).

## Bibliography

- FLIGHT, M. (2010). "Synaptic transmission: On the probability of release." In: *Nat Rev Neurosci* 9, pp. 736–737 (cit. on p. 14).
- FOX, M. D. and M. E. RAICHLER (2007). "Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging." In: *Nat Rev Neurosci* 8, pp. 700–711 (cit. on p. 35).
- FRISTON, K. J., J. DAUNIZEAU, J. KILNER, and S. J. KIEBEL (2010). "Action and behavior: a free-energy formulation." In: *Biol Cybern* 102.3, pp. 227–260 (cit. on p. 15).
- FUSI, S. (2002). "Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates." In: *Biological cybernetics* 87.5-6, pp. 459–470 (cit. on pp. 65, 71).
- GABA, S., P. SHERIDAN, J. ZHOU, S. CHOI, and W. LU (2013). "Stochastic memristive devices for computing and neuromorphic applications." In: *Nanoscale* 5.13, pp. 5872–5878 (cit. on pp. 12, 51, 54, 69, 70).
- GARDINER, C. (2009). *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer Berlin Heidelberg (cit. on p. 29).
- GEMAN, S. and D. GEMAN (1984). "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images." In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, pp. 721–741 (cit. on pp. 14, 16).
- GENTET, L., M. AVERMANN, F. MATYAS, J. STAIGER, and C. PETERSEN (2010). "Membrane Potential Dynamics of GABAergic Neurons in the Barrel Cortex of Behaving Mice." In: *Neuron* (cit. on p. 108).
- GERSHMAN, S. J., E. VUL., and J. TENENBAUM (2009). "Perceptual Multistability as Markov Chain Monte Carlo Inference." In: *Proceedings of the 22nd Conference on Advances in Neural Information Processing Systems; December 2008; Vancouver, Canada. NIPS 2008*. (Cit. on pp. 4, 7, 14, 16, 30, 33, 35).
- GERSTNER, W. and W. M. KISTLER (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press (cit. on pp. 14, 25, 28, 57).
- GOLD, J. I. and M. N. SHADLEN (2007). "The neural basis of decision making." In: *Annu Rev Neurosci* 30, pp. 535–574 (cit. on pp. 14, 23).
- GOODMAN, N. D., J. B. TENENBAUM, J. FELDMAN, and T. L. GRIFFITHS (2008). "A Rational Analysis of Rule-Based Concept Learning." In: *Cognitive Science* 32.1, pp. 108–154 (cit. on p. 76).
- GOPNIK, A. and J. B. TENENBAUM (2007). "Bayesian special section: Introduction; Bayesian networks, Bayesian learning and cognitive development." In: *Dev Sci* 10.3, pp. 281–287 (cit. on p. 14).
- GRACA, J., K. GANCHEV, and B. TASKAR (2008). "Expectation maximization and posterior constraints." In: *Proc. of NIPS 2007*. Vol. 20. MIT Press (cit. on pp. 39, 43, 91, 143, 154).
- GRAVES, A., A.-R. MOHAMED, and G. HINTON (2013). "Speech recognition with deep recurrent neural networks." In: *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, pp. 6645–6649 (cit. on p. 4).
- GRIFFITHS, T. L., C. KEMP, and J. B. TENENBAUM (2008). "Bayesian models of cognition." In: *Handbook of Computational Cognitive Modeling*. Ed. by R. SUN. Cambridge Univ. Press, pp. 59–100 (cit. on p. 14).

- GRIFFITHS, T. L. and J. B. TENENBAUM (2006). "Optimal predictions in everyday cognition." In: *Psychol Sci* 17.9, pp. 767–773 (cit. on pp. 3, 15, 34, 38).
- GRIMMETT, G. R. and D. R. STIRZAKER (2001). *Probability and Random Processes*. 3rd. Oxford University Press (cit. on p. 17).
- HABENSCHUSS, S., J. BILL, and B. NESSLER (2012). "Homeostatic plasticity in Bayesian spiking networks as Expectation Maximization with posterior constraints." In: *Advances in Neural Information Processing Systems*, pp. 773–781 (cit. on pp. 57–59, 62, 70, 90, 91, 104, 143, 153, 154).
- HABENSCHUSS, S., Z. JONKE, and W. MAASS (2013a). "Stochastic Computations in Cortical Microcircuit Models." In: *PLoS Computational Biology* 9.11, e1003311 (cit. on pp. 77, 111, 162).
- HABENSCHUSS, S., H. PUHR, and W. MAASS (2013b). "Emergence of optimal decoding of population codes through STDP." In: *Neural computation* 25.6, pp. 1371–1407 (cit. on pp. 57, 59, 70, 80, 107, 143).
- HAIDER, B., M. HÄUSSER, and M. CARANDINI (2013). "Inhibition dominates sensory responses in the awake cortex." In: *Nature* 493.7430, pp. 97–100 (cit. on p. 108).
- HARVEY, C. D., P. COEN, and D. W. TANK (2012). "Choice-specific sequences in parietal cortex during a virtual-navigation decision task." In: *Nature* 484.7392, pp. 62–68 (cit. on p. 4).
- HINTON, G. E. (2010). "Learning to represent visual input." In: *Philos Trans R Soc Lond B Biol Sci* 365, pp. 177–184 (cit. on pp. 20, 133).
- HINTON, G. and A. BROWN (2000). "Spiking Boltzmann machines." In: *Proceedings of the 13th Conference on Advances in Neural Information Processing Systems; December 1999; Vancouver, Canada. NIPS 1999* (cit. on p. 35).
- HINTON, G. (2002). "Training products of experts by minimizing contrastive divergence." In: *Neural computation* 14.8, pp. 1771–1800 (cit. on pp. 4, 96, 133).
- HINTON, G. E. and T. J. SEJNOWSKI (1986). "Learning and relearning in Boltzmann machines." In: *Cambridge, MA: MIT Press* 1, pp. 282–317 (cit. on p. 96).
- HINTON, G. and R. SALAKHUTDINOV (2006). "Reducing the dimensionality of data with neural networks." In: *Science* 313.5786, pp. 504–507 (cit. on p. 4).
- HINTON, G. E., S. OSINDERO, and Y.-W. TEH (2006). "A Fast Learning Algorithm for Deep Belief Nets." In: *Neural Comput* 18, pp. 1527–1554 (cit. on pp. 16, 20).
- HOPFIELD, J. J. and D. W. TANK (1985). "'Neural' Computation of Decisions in Optimization Problems." In: *Biol Cybern* 52, pp. 141–152 (cit. on p. 34).
- HOYER, P. and A. HYVÄRINEN (2003). "Interpreting neural response variability as Monte Carlo sampling of the posterior." In: *Proceedings of the 16th Conference on Advances in Neural Information Processing Systems; December 2002; Vancouver, Canada. NIPS 2002*. (Cit. on pp. 4, 14, 16, 30, 35, 77, 110).
- HYVÄRINEN, A. (2006). "Consistency of pseudolikelihood estimation of fully visible Boltzmann machines." In: *Neural Computation* 18.10, pp. 2283–2292 (cit. on p. 96).
- INDIVERI, G. (2000). "Modeling selective attention using a neuromorphic analog VLSI device." In: *Neural computation* 12.12, pp. 2857–2880 (cit. on p. 57).
- INDIVERI, G., B. LINARES-BARRANCO, R. LEGENSTEIN, G. DELIGEORGIS, and T. PRODROMAKIS (2013). "Integration of nanoscale memristor synapses in neuromorphic

## Bibliography

- computing architectures." In: *Nanotechnology* 24.38, p. 384010 (cit. on pp. 12, 50, 51, 69, 70, 113).
- JOLIVET, R., A. RAUCH, H.-R. LÜSCHER, and W. GERSTNER (2006). "Predicting spike timing of neocortical pyramidal neurons by simple threshold models." In: *Journal of computational neuroscience* 21.1, pp. 35–49 (cit. on pp. 40, 58, 82).
- JO, S. H., T. CHANG, I. EBONG, B. B. BHADVIYA, P. MAZUMDER, and W. LU (2010). "Nanoscale memristor device as synapse in neuromorphic systems." In: *Nano letters* 10.4, pp. 1297–1301 (cit. on pp. 50, 52, 69).
- JO, S. H., K.-H. KIM, and W. LU (2009a). "High-density crossbar arrays based on a Si memristive system." In: *Nano letters* 9.2, pp. 870–874 (cit. on p. 72).
- JO, S. H., K.-H. KIM, and W. LU (2009b). "Programmable resistance switching in nanoscale two-terminal devices." In: *Nano Letters* 9.1, pp. 496–500 (cit. on pp. 51, 54, 70, 72).
- KAPPEL, D., B. NESSLER, and W. MAASS (2014). "STDP Installs in Winner-Take-All Circuits an Online Approximation to Hidden Markov Model Learning." In: *PLoS computational biology* 10.3, e1003511 (cit. on pp. 4, 57, 77, 110, 112).
- KECK, C., C. SAVIN, and J. LÜCKE (2012). "Feedforward Inhibition and Synaptic Scaling – Two Sides of the Same Coin?" In: *PLoS Computational Biology* 8.3, e1002432 (cit. on pp. 4, 38, 41, 48, 57).
- KENET, T., D. BIBITCHKOV, M. TSODYKS, A. GRINVALD, and A. ARIELI (2003). "Spontaneously emerging cortical representations of visual attributes." In: *Nature* 425.6961, pp. 954–956 (cit. on p. 35).
- KERLIN, A. M., M. L. ANDERMANN, V. K. BEREZOVSKII, and R. C. REID (2010). "Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex." In: *Neuron* 67.5, pp. 858–871 (cit. on p. 108).
- KERSTEN, D., P. MAMASSIAN, and A. YUILLE (2004). "Object perception as Bayesian inference." In: *Annu Rev Psychol* 55.1, pp. 271–304 (cit. on pp. 3, 14).
- KO, H., S. B. HOFER, B. PICHLER, K. A. BUCHANAN, P. J. SJÖSTRÖM, and T. D. MRSIC-FLOGEL (2011). "Functional specificity of local synaptic connections in neocortical networks." In: *Nature* 473.7345, pp. 87–91 (cit. on p. 109).
- KOLLER, D. and N. FRIEDMAN (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press (cit. on p. 34).
- KÖRDING, K. P. and D. M. WOLPERT (2004). "Bayesian integration in sensorimotor learning." In: *Nature* 427.6971, pp. 244–247 (cit. on pp. 3, 14, 38, 76).
- KUZUM, D., R. G. JEYASINGH, B. LEE, and H.-S. P. WONG (2011). "Nanoelectronic programmable synapses based on phase change materials for brain-inspired computing." In: *Nano letters* 12.5, pp. 2179–2186 (cit. on p. 50).
- LANSNER, A. (2009). "Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations." In: *Trends in neurosciences* 32.3, pp. 178–186 (cit. on pp. 57, 86, 109).
- LECUN, Y., L. BOTTOU, Y. BENGIO, and P. HAFFNER (1998). "Gradient-based learning applied to document recognition." In: vol. 86. 11. IEEE, pp. 2278–2324 (cit. on pp. 42, 63, 146).
- LEE, D., D.-J. SEONG, H. jung CHOI, I. JO, R. DONG, W. XIANG, S. OH, M. PYUN, S.-O. SEO, S. HEO, et al. (2006). "Excellent uniformity and reproducible resis-

- tance switching characteristics of doped binary metal oxides for non-volatile resistance memory applications." In: *Electron Devices Meeting, 2006. IEDM'06. International*. IEEE, pp. 1–4 (cit. on p. 70).
- LEE, J. H., X. MA, and K. LIKHAREV (2006). "CMOL Crossnets: Possible neuromorphic nanoelectronic circuits." In: *Advances in Neural Information Processing Systems* 18, p. 755 (cit. on p. 73).
- LEE, T. S. and D. MUMFORD (2003). "Hierarchical Bayesian inference in the visual cortex." In: *J. Opt. Soc. Am. A* 20.7, pp. 1434–1448 (cit. on pp. 3, 14, 15).
- LEGENSTEIN, R. and W. MAASS (2014). "Ensembles of Spiking Neurons with Noise Support Optimal Probabilistic Inference in a Dynamically Changing Environment." In: *PLoS computational biology* 10.10, e1003859 (cit. on pp. 77, 112).
- LEOPOLD, D. A., M. WILKE, A. MAIER, and N. K. LOGOTHETIS (2002). "Stable perception of visually ambiguous patterns." In: *Nat Neurosci* 5, pp. 605–609 (cit. on pp. 30, 34, 35).
- LITVAK, S. and S. ULLMAN (2009). "Cortical Circuitry Implementing Graphical Models." In: *Neural Comput* 21, pp. 1–47 (cit. on p. 15).
- LIU, Z.-W., U. FARAGUNA, C. CIRELLI, G. TONONI, and X.-B. GAO (2010). "Direct evidence for wake-related increases and sleep-related decreases in synaptic strength in rodent cortex." In: *The Journal of Neuroscience* 30.25, pp. 8671–8675 (cit. on p. 96).
- LOCHMANN, T. and S. DENEVE (2011). "Neural processing as causal inference." In: *Current Opinion in Neurobiological aspectgy* 21.5, pp. 774–781 (cit. on p. 79).
- MALENKA, R. C. and M. F. BEAR (2004). "LTP and LTD: an embarrassment of riches." In: *Neuron* 44.1, pp. 5–21 (cit. on p. 50).
- MARKRAM, H., W. GERSTNER, and P. J. SJÖSTRÖM (2012). "Spike-timing-dependent plasticity: a comprehensive overview." In: *Frontiers in Synaptic Neuroscience* 4.2 (cit. on p. 51).
- MARKRAM, H., J. LÜBKE, M. FROTSCHER, and B. SAKMANN (1997). "Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs." In: *Science* 275.5297, pp. 213–215 (cit. on p. 51).
- MARKRAM, H., M. TOLEDO-RODRIGUEZ, Y. WANG, A. GUPTA, G. SILBERBERG, and C. WU (2004). "Interneurons of the neocortical inhibitory system." In: *Nature Reviews Neuroscience* 5.10, pp. 793–807 (cit. on pp. 82, 108).
- MARKRAM, H., Y. WANG, and M. TSODYKS (1998). "Differential signaling via the same axon of neocortical pyramidal neurons." In: *Proceedings of the National Academy of Sciences* 95.9, pp. 5323–5328 (cit. on p. 11).
- MARR, D. and A. VISION (1982). "A computational investigation into the human representation and processing of visual information." In: *WH San Francisco: Freeman and Company* (cit. on p. 5).
- MATEO, C., M. AVERMANN, L. J. GENTET, F. ZHANG, K. DEISSEROTH, and C. C. PETERSEN (2011). "In vivo optogenetic stimulation of neocortical excitatory neurons drives brain-state-dependent inhibition." In: *Current Biology* 21.19, pp. 1593–1602 (cit. on pp. 108, 109).

## Bibliography

- MA, W. J., J. M. BECK, P. E. LATHAM, and A. POUGET (2006). "Bayesian inference with probabilistic population codes." In: *Nature neuroscience* 9.11, pp. 1432–1438 (cit. on pp. 4, 77, 111).
- MAYR, C., P. STÄRKE, J. PARTZSCH, L. CEDERSTROEM, R. SCHÜFFNY, Y. SHUAL, N. DU, and H. SCHMIDT (2012). "Waveform Driven Plasticity in BiFeO<sub>3</sub> Memristive Devices: Model and Implementation." In: *Advances in Neural Information Processing Systems*, pp. 1700–1708 (cit. on pp. 51, 69, 113).
- MEAD, C. and M. ISMAIL (1989). *Analog VLSI implementation of neural systems*. Springer (cit. on pp. 12, 57, 112).
- MEROLLA, P., J. ARTHUR, B. E. SHI, and K. BOAHEN (2007). "Expandable Networks for Neuromorphic Chips." In: *IEEE Transactions on Circuits and Systems I* 54.2, pp. 301–311 (cit. on p. 36).
- MORRISON, A., A. AERTSEN, and M. DIEMANN (2007). "Spike-timing-dependent plasticity in balanced random networks." In: *Neural computation* 19.6, pp. 1437–1467 (cit. on p. 51).
- NEAL, R. and G. HINTON (1998). "A view of the EM algorithm that justifies incremental, sparse, and other variants." In: *Learning in graphical models*. Springer, pp. 355–368 (cit. on pp. 8, 153).
- NESSLER, B., M. PFEIFFER, L. BUESING, and W. MAASS (2013). "Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity." In: *PLoS computational biology* 9.4, e1003037 (cit. on pp. 9, 52, 56, 57, 59, 60, 62, 70, 78–80, 90, 104, 107, 110, 111, 143).
- NESSLER, B., M. PFEIFFER, and W. MAASS (2009a). "STDP enables spiking neurons to detect hidden causes of their inputs." In: *Advances in neural information processing systems*, pp. 1357–1365 (cit. on pp. 4, 9, 38, 40, 41, 45, 46, 48, 59, 70).
- NESSLER, B., M. PFEIFFER, and W. MAASS (2009b). "Hebbian learning of Bayes optimal decisions." In: *Proceedings of the 21th Conference on Advances in Neural Information Processing Systems; December 2008; Vancouver, Canada. NIPS 2008*. (Cit. on p. 20).
- OAKSFORD, M. and N. CHATER (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford University Press (cit. on p. 15).
- ORBAN, G., J. FISER, R. ASLIN, and M. LENGYEL (2008). "Bayesian learning of visual chunks by human observers." In: *Proceedings of the National Academy of Sciences* 105.7, pp. 2745–2750 (cit. on pp. 3, 38, 76).
- PEARL, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann (cit. on p. 14).
- PECEVSKI, D., L. BUESING, and W. MAASS (Dec. 2011). "Probabilistic Inference in General Graphical Models through Sampling in Stochastic Networks of Spiking Neurons." In: *PLoS Computational Biology* 7.12, e1002294 (cit. on pp. 4, 11, 38, 46, 110, 112).
- PETERSEN, C. C. and S. CROCHET (2013). "Synaptic computation and sensory processing in neocortical layer 2/3." In: *Neuron* 78.1, pp. 28–48 (cit. on p. 109).
- PETROVICI, M. A., J. BILL, I. BYTSCHOK, J. SCHEMMELE, and K. MEIER (2013). "Stochastic inference with deterministic spiking neurons." In: *arXiv preprint arXiv:1311.3211* (cit. on pp. 11, 71, 77, 82, 148).



- PEYRACHE, A., M. KHAMASSI, K. BENCHENANE, S. I. WIENER, and F. P. BATTAGLIA (2009). "Replay of rule-learning related neural patterns in the prefrontal cortex during sleep." In: *Nature neuroscience* 12.7, pp. 919–926 (cit. on p. 96).
- PILLOW, J. W., J. SHLENS, L. PANINSKI, A. SHER, A. M. LITKE, E. J. CHICHILNISKY, and E. P. SIMONCELLI (2008). "Spatio-temporal correlations and visual signalling in a complete neuronal population." In: *Nature* 454.7207, pp. 995–999 (cit. on p. 25).
- POSPISCHIL, M., Z. PIWKOWSKA, T. BAL, and A. DESTEXHE (2009). "Characterizing neuronal activity by describing the membrane potential as a stochastic process." In: *J Physiol Paris* 103.1-2, pp. 98–106 (cit. on p. 136).
- POUGET, A., J. M. BECK, W. J. MA, and P. E. LATHAM (2013). "Probabilistic brains: knowns and unknowns." In: *Nature Neuroscience* 16.9, pp. 1170–1178 (cit. on p. 111).
- PROBST, D., M. A. PETROVICI, I. BYTSCHOK, J. BILL, D. PECEVSKI, J. SCHEMMEL, and K. MEIER (2015). "Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons." In: *Frontiers in computational neuroscience* 9 (cit. on p. 11).
- PURVES, D., G. AUGUSTINE, D. FITZPATRICK, W. HALL, A. LAMANTIA, J. MCNAMARA, and L. WHITE (2008). *Neuroscience*. De Boeck, Sinauer, Sunderland, Mass (cit. on p. 6).
- QUERLIOZ, D., O. BICHLER, and C. GAMRAT (2011). "Simulation of a memristor-based spiking neural network immune to device variations." In: *Neural Networks (IJCNN), The 2011 International Joint Conference on*. IEEE, pp. 1775–1781 (cit. on pp. 51, 52, 54, 58, 69, 113).
- RAO, R. P. (2004). "Hierarchical Bayesian inference in networks of spiking neurons." In: *Advances in neural information processing systems*, pp. 1113–1120 (cit. on pp. 4, 77, 111).
- RAO, R. P. (2007). "Neural Models of Bayesian Belief Propagation." In: MIT Press, p. 239 (cit. on pp. 4, 15, 77, 111).
- RAO, R. P. and D. H. BALLARD (1999). "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects." In: *Nature neuroscience* 2.1, pp. 79–87 (cit. on p. 79).
- RAO, R. P. N., B. A. OLSHAUSEN, and M. S. LEWICKI (2002). *Probabilistic Models of the Brain*. MIT Press (cit. on pp. 3, 14, 15).
- REZENDE, D. J., D. WIERSTRA, and W. GERSTNER (2012). "Variational Learning for Recurrent Spiking Networks." In: *Proc. of NIPS 2011*. Vol. 24. MIT Press, pp. 136–144 (cit. on pp. 4, 38).
- RINGACH, D. L. (2009). "Spontaneous and driven cortical activity: implications for computation." In: *Curr Opin Neurobiol* 19, pp. 1–6 (cit. on p. 14).
- ROLLS, E. T. and G. DECO (2010). *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford Univ. Press (cit. on pp. 3, 14).
- ROTH, G. and U. DICKE (2005). "Evolution of the brain and intelligence." In: *Trends in cognitive sciences* 9.5, pp. 250–257 (cit. on p. 1).
- RUNYAN, C. A., J. SCHUMMERS, A. VAN WART, S. J. KUHLMAN, N. R. WILSON, Z. J. HUANG, and M. SUR (2010). "Response features of parvalbumin-expressing

## Bibliography

- interneurons suggest precise roles for subtypes of inhibition in visual cortex." In: *Neuron* 67.5, pp. 847–857 (cit. on p. 108).
- SADAGHIANI, S., G. HESSELMANN, K. J. FRISTON, and A. KLEINSCHMIDT (2010). "The relation of ongoing brain activity, evoked neural responses, and cognition." In: *Front Syst Neurosci* 4, Artikel 20 (cit. on p. 14).
- SAHANI, M. and P. DAYAN (2003). "Doubly Distributional Population Codes: Simultaneous Representation of Uncertainty and Multiplicity." In: *Neural Comput* 15, pp. 2255–2279 (cit. on p. 16).
- SALAKHUTDINOV, R. and G. HINTON (2009). "Deep boltzmann machines." In: *International Conference on Artificial Intelligence and Statistics*, pp. 448–455 (cit. on pp. 4, 96, 112).
- SATO, M. (1999). "Fast learning of on-line EM algorithm." In: *Rapport Technique, ATR Human Information Processing Research Laboratories* (cit. on pp. 8, 41, 81).
- SAVIN, C. and S. DENEVE (2014). "Spatio-temporal Representations of Uncertainty in Spiking Neural Networks." In: *Advances in Neural Information Processing Systems*, pp. 2024–2032 (cit. on pp. 4, 77, 111).
- SCHEMMELE, J., D. BRÜDERLE, A. GRÜBL, M. HOCK, K. MEIER, and S. MILLNER (2010). "A Wafer-Scale Neuromorphic Hardware System for Large-Scale Neural Modeling." In: *Proc. of ISCAS'10*, pp. 1947–1950 (cit. on pp. 12, 39).
- SCHEMMELE, J., J. FIERES, and K. MEIER (2008). "Wafer-scale integration of analog neural networks." In: *Neural Networks, 2008. IJCNN 2008. (IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on.* IEEE, pp. 431–438 (cit. on pp. 50, 112).
- SCHEMMELE, J., A. GRUBL, K. MEIER, and E. MUELLER (2006). "Implementing synaptic plasticity in a VLSI spiking neural network model." In: *Neural Networks, 2006. IJCNN'06. International Joint Conference on.* IEEE, pp. 1–6 (cit. on pp. 12, 113).
- SCHULZ, L. (2012). "The origins of inquiry: Inductive inference and exploration in early childhood." In: *Trends in cognitive sciences* 16.7, pp. 382–389 (cit. on p. 3).
- SERRANO-GOTARREDONA, T., T. MASQUELIER, T. PRODROMAKIS, G. INDIVERI, and B. LINARES-BARRANCO (2013). "STDP and STDP variations with memristors for spiking neuromorphic learning systems." In: *Frontiers in neuroscience* 7 (cit. on pp. 51, 69, 113).
- SHINOMOTO, S., H. KIM, T. SHIMOKAWA, N. MATSUNO, S. FUNAHASHI, K. SHIMA, I. FUJITA, H. TAMURA, T. DOI, K. KAWANO, N. INABA, K. FUKUSHIMA, S. KURKIN, K. KURATA, M. TAIRA, K.-I. TSUTSUI, H. KOMATSU, T. OGAWA, K. KOIDA, J. TANJI, and K. TOYAMA (July 2009). "Relating Neuronal Firing Patterns to Functional Differentiation of Cerebral Cortex." In: *PLoS Comput Biol* 5.7 (cit. on p. 136).
- SHUKOFF, P. and L. AHLQUIST (2012). "ERB: Steve Jobs vs Bill Gates." In: *www.youtube.com/watch?v=njos57Ijf-o#t=93* (cit. on p. 4).
- SIMONCELLI, E. P., L. PANINSKI, J. PILLOW, and O. SCHWARTZ (2004). "Characterization of neural responses with stochastic stimuli." In: *The cognitive neurosciences* 3, pp. 327–338 (cit. on p. 60).
- SIMONCELLI, E. and D. HEEGER (1998). "A model of neuronal responses in visual area MT." In: *Vision Research* 38.5, pp. 743–761 (cit. on p. 40).

- SMITH, S. L. and M. HÄUSSER (2010). "Parallel processing of visual space by neighboring neurons in mouse visual cortex." In: *Nature neuroscience* 13.9, pp. 1144–1149 (cit. on p. 109).
- SNIDER, G. S. (2008). "Spike-timing-dependent learning in memristive nanodevices." In: *Nanoscale Architectures, 2008. NANOARCH 2008. IEEE International Symposium on*. IEEE, pp. 85–92 (cit. on pp. 69, 113).
- SOFTKY, W. and C. KOCH (1993). "The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs." In: *J Neurosci* 13.1, pp. 334–350 (cit. on p. 136).
- SOHYA, K., K. KAMEYAMA, Y. YANAGAWA, K. OBATA, and T. TSUMOTO (2007). "GABAergic Neurons Are Less Selective to Stimulus Orientation than Excitatory Neurons in Layer II/III of Visual Cortex, as Revealed by In Vivo Functional Ca<sup>2+</sup> Imaging in Transgenic Mice." In: *The Journal of Neuroscience* 27.8, pp. 2145–2149 (cit. on p. 108).
- SRIVASTAVA, N. and R. SALAKHUTDINOV (2012). "Multimodal learning with deep boltzmann machines." In: *Advances in neural information processing systems*, pp. 2222–2230 (cit. on p. 4).
- STEIMER, A., W. MAASS, and R. DOUGLAS (2009). "Belief propagation in networks of spiking neurons." In: *Neural Computation* 21.9, pp. 2502–2523 (cit. on pp. 4, 15, 38, 77, 111).
- STEPANYANTS, A., L. M. MARTINEZ, A. S. FERESKÓ, and Z. F. KISVÁRDAY (2009). "The fractions of short- and long-range connections in the visual cortex." In: *Proceedings of the National Academy of Sciences* 106.9, pp. 3555–3560 (cit. on pp. 80, 108, 109).
- STRUKOV, D. B., G. S. SNIDER, D. R. STEWART, and R. S. WILLIAMS (2008). "The missing memristor found." In: *Nature* 453.7191, pp. 80–83 (cit. on pp. 12, 113).
- SUNDARESWARA, R. and P. R. SCHRATER (2008). "Perceptual multistability predicted by search model for Bayesian decisions." In: *Journal of Vision* 8.5, p. 12 (cit. on pp. 14, 16, 30, 33, 35).
- SURI, M., D. QUERLIOZ, O. BICHLER, G. PALMA, E. VIANELLO, D. VUILLAUME, C. GAMRAT, and B. DESALVO (2013). "Bio-inspired stochastic computing using binary CBRAM synapses." In: *Electron Devices, IEEE Transactions on* 60.7, pp. 2402–2409 (cit. on pp. 12, 51, 54, 55, 70, 73, 113).
- SUTTON, R. and A. BARTO (1998). *Reinforcement learning: An introduction*. Cambridge Univ Press (cit. on p. 8).
- TAVOSANIS, G. (2012). "Dendritic structural plasticity." In: *Developmental neurobiology* 72.1, pp. 73–86 (cit. on p. 11).
- TÉGLÁS, E., E. VUL, V. GIROTTO, M. GONZALEZ, J. B. TENENBAUM, and L. L. BONATTI (2011). "Pure reasoning in 12-month-old infants as probabilistic inference." In: *Science* 332.6033, pp. 1054–1059 (cit. on p. 76).
- TENENBAUM, J. B., T. L. GRIFFITHS, and C. KEMP (2006). "Theory-based Bayesian models of inductive learning and reasoning." In: *Trends Cogn Sci* 10.7, pp. 309–318 (cit. on p. 15).

## Bibliography

- TENENBAUM, J. B., C. KEMP, T. L. GRIFFITHS, and N. D. GOODMAN (2011). "How to Grow a Mind: Statistics, Structure, and Abstraction." In: *Science* 331.6022, pp. 1279–1285 (cit. on pp. 3, 39, 76).
- TONONI, G. and C. CIRELLI (2014). "Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration." In: *Neuron* 81.1, pp. 12–34 (cit. on p. 96).
- TOUSSAINT, M. (2009). "Probabilistic inference as a model of planned behavior." In: *Künstliche Intelligenz (German Artificial Intelligence Journal)* 3 (cit. on p. 15).
- TOUSSAINT, M. and C. GOERICK (2010). "A Bayesian view on motor control and planning." In: *From motor to interaction learning in robots. Studies in Computational Intelligence*. Ed. by O. SIGAUD and J. PETERS. Springer, pp. 227–252 (cit. on p. 15).
- VAN ROSSUM, G. and F. L. DRAKE JR (1995). *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam (cit. on pp. 132, 159).
- VAN ROSSUM, M. C., G. Q. BI, and G. G. TURRIGIANO (2000). "Stable Hebbian learning from spike timing-dependent plasticity." In: *The Journal of Neuroscience* 20.23, pp. 8812–8821 (cit. on pp. 52, 56).
- VASILAKI, E., N. FRÉMAUX, R. URBANCIK, W. SENN, and W. GERSTNER (2009). "Spike-based reinforcement learning in continuous state and action space: when policy gradient methods fail." In: *PLoS computational biology* 5.12, e1000586 (cit. on p. 11).
- VINCENT, A., J. LARROQUE, W. ZHAO, N. B. ROMDHANE, O. BICHLER, C. GAMRAT, J.-O. KLEIN, S. GALDIN-RETAILLEAU, and D. QUERLIOZ (2014). "Spin-transfer torque magnetic memory as a stochastic memristive synapse." In: *Circuits and Systems (ISCAS), 2014 IEEE International Symposium on*. IEEE, pp. 1074–1077 (cit. on p. 73).
- VUL, E. and H. PASHLER (2008). "Measuring the Crowd Within: Probabilistic representations Within individuals." In: *Psychol Sci* 19.7, pp. 645–647 (cit. on p. 34).
- WATT, A. and N. DESAI (2010). "Homeostatic plasticity and STDP: keeping a neuron's cool in a fluctuating world." In: *Frontiers in Synaptic Neuroscience* 2 (cit. on pp. 39, 42).
- WESTHEIMER, G. (2008). "Was Helmholtz a Bayesian? a review." In: *Perception* 37, pp. 642–650 (cit. on p. 4).
- WU, S.-Y., J. LIAW, C. LIN, M. CHIANG, C. YANG, J. CHENG, M. TSAI, M. LIU, P. WU, C. CHANG, et al. (2009). "A highly manufacturable 28nm CMOS low power platform technology with fully functional 64Mb SRAM using dual/tripe gate oxide process." In: *VLSI Technology, 2009 Symposium on*. IEEE, pp. 210–211 (cit. on p. 72).
- YANG, J. J., D. B. STRUKOV, and D. R. STEWART (2013). "Memristive devices for computing." In: *Nature nanotechnology* 8.1, pp. 13–24 (cit. on pp. 12, 72, 73, 113).
- YANG, T. and M. N. SHADLEN (2007). "Probabilistic reasoning by neurons." In: *Nature* 447, pp. 1075–1080 (cit. on pp. 3, 4, 14, 23).
- YOSHIMURA, Y. and E. M. CALLAWAY (2005). "Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity." In: *Nature neuroscience* 8.11, pp. 1552–1559 (cit. on p. 108).

- YU, S., B. GAO, Z. FANG, H. YU, J. KANG, and H.-S. P. WONG (2013). "Stochastic learning in oxide binary synaptic device for neuromorphic computing." In: *Frontiers in neuroscience* 7 (cit. on pp. 51, 54, 69, 73).
- ZAMARREÑO-RAMOS, C., L. A. CAMUÑAS-MESA, J. A. PÉREZ-CARRASCO, T. MASQUELIER, T. SERRANO-GOTARREDONA, and B. LINARES-BARRANCO (2011). "On spike-timing-dependent-plasticity, memristive devices, and building a self-learning visual cortex." In: *Frontiers in neuroscience* 5 (cit. on pp. 52, 54).
- ZARIWALA, H., L. MADISEN, K. AHRENS, A. BERNARD, E. LEIN, A. JONES, and H. ZENG (2011). "Visual Tuning Properties of Genetically Identified Layer 2/3 Neuronal Types in the Primary Visual Cortex of Cre-Transgenic Mice." In: *Frontiers in Systems Neuroscience* 4.00162 (cit. on p. 108).
- ZEMEL, R., Q. J. M. HUYS, R. NATARAJAN, and P. DAYAN (2005). "Probabilistic computation in spiking populations." In: *Proceedings of the 17th Conference on Advances in Neural Information Processing Systems; December 2004; Vancouver, Canada. NIPS 2004* (cit. on pp. 4, 35).
- ZEMEL, R. S., P. DAYAN, and A. POUGET (1998). "Probabilistic interpretation of population codes." In: *Neural computation* 10.2, pp. 403–430 (cit. on pp. 4, 77, 111).