DISSERTATION

to obtain the title of

Doctor of Technical Sciences

of

Graz University of Technology

by

Klaus SCHUCH

Data Analysis, Software- and Hardware-Modeling of Computation in Cortical Networks of Neurons

Thesis Advisor: Univ.Prof. DI Dr. Wolfgang MAASS

defended on August 24, 2010

Jury :

Advisor:	Univ. Prof. DI Dr. Wolfgang MAASS	-	TU Graz
Reviewer:	Univ. Prof. Dr. Helmut POCKBERGER	-	Uni Wien
Dean of Studies:	Univ. Prof. DI Dr. Horst BISCHOF	-	TU Graz

Eidesstattliche Erklärung

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

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, 6. August 2010

....

(signature)

Acknowledgments

I would like to thank Wolfgang Maass, my supervisor, for his guidance and support.

Many thanks go to Stefan Häusler, Malte Rasch, Johannes Bill, Dejan Pecevski, Thomas Natschläger, Daniel Brüderle, Johannes Schemmel, Karlheinz Meier, and Nikos Logothetis for enjoyable and fruitful collaborations.

Additionally, I would like to thank Helmut Pockberger for his willingness to be the second referee of this thesis.

Special thanks go to all colleagues at the IGI for the enjoyable working conditions during my PhD studies.

Finally, I would like to thank my family and you, Andrea, for all your patience and your ongoing support.

I was partially supported by the Austrian Science Fund FWF, the FACETS project and the ORGANIC project of the European Union.

Abstract

The recurrent connectivity within in the cortex is highly structured. For example, the connection probability and the mean synaptic strength varies between the six layers of the cortex. The underlying rules of this connectivity structure seem to be crucial for the computational function of the cortex.

This thesis investigates how this structured recurrent connectivity in cortical neural circuits influences the activity of neurons and the computational function of the circuit. This is done by developing both software and hardware models of neural circuits based on experimental data.

First, two data-based templates of connectivity in cortical microcircuits are analyzed and compared regarding their network structure, as well as their dynamical and computational properties. The distribution of network motifs are quite different for the two microcircuit templates. The computational performance of a cortical microcircuit is correlated with specific statistical properties of the circuit dynamics that is induced by a particular distribution of degrees of nodes.

Furthermore, a prototype of a spiking neuromorphic VLSI hardware is used to emulate cortical neural circuits. Such neuromorphic VLSI hardware systems are promising candidates for neuroscientific research tools, especially for tasks, which exhaust the computing power of software simulations. Still, neuromorphic models suffer from a constricted configurability and production-related fluctuations of device characteristics. A cortically inspired self-adjusting network architecture is applied to keep the the activity of emulated generic spiking neural networks within a biologically realistic firing regime. This architecture increases the robustness against transistor-level variations.

Finally, two computer models for a patch of primary visual cortex (V1) are developed. For the first model, a detailed large network for a patch of V1 of macaque monkeys, the statistics of the spike activity in response to a natural movie stimuli is compared with statistics of *in vivo* recordings in response to the same stimuli. Some high level parameters of the recurrent connectivity are varied to investigate their influence on the firing statistics. This knowledge is used to optimize the model to achieve realistic firing statistics. The deviation between the firing regime of the optimized model and the *in vivo* data is on the same level as deviations among monkeys and experimental sessions. A second network model for a patch of V1 of cats is used to reproduce the experimental findings about the temporal integration of information about the visual stimulus in this area. Additionally, the model is used to show how higher order correlations of neurons contribute to the information about the visual stimulus.

Zusammenfassung

Die Verbindungsstruktur zwischen den Neuronen im Kortex ist sehr strukturiert. Beispielsweise variiert die Verbindungswahrscheinlichkeit und die Stärke der Verbindungen zwischen den typischen sechs Schichten des Kortex. Die Regeln hinter dieser Verbindungsstruktur scheinen für die Funktionalität des Gehirns essentiell zu sein. Diese Arbeit untersucht, wie diese Verbindungsstrukturen in kortikalen neuronalen Schaltkreisen deren Funktionalität und die Aktivität der Neuronen beeinflussen. Die Untersuchung wird mittels Simulationen von Modellen in Software und Hardware durchgeführt, die auf experimentellen Daten basieren.

Zunächst werden zwei experimentell ermittelte Muster von Verbindungsstrukturen von kortikalen Netzwerken analysiert und bezüglich der Netzwerkstruktur, der Dynamik und der Berechnungseigenschaften verglichen. Die Verteilungen der Netzwerkmotive der beiden Verbindungsstrukturen unterscheiden sich. Die Berechnungseigenschaften der Netzwerke korrelieren mit der Statistik der Netzwerkdynamik, die von den jeweiligen Verteilungen der Knotengrade abhängt.

Darüber hinaus wird ein Prototyp eines neuromorphen VLSI Chips verwendet um kortikale Netzwerke zu emulieren. Solche VLSI Systeme sind erfolgversprechende Werkzeuge für neurowissenschaftliche Fragestellungen, welche die Leistungsfähigkeit üblicher Softwaresimulationen überschreiten. Die beschränkte Konfigurierbarkeit und die produktionsbedingten Bauteilschwankungen solcher Systeme sind noch problematisch. Eine selbstregulierende Netzwerk-Architektur wird verwendet um das emulierte Netzwerk in einem biologisch realistischen Aktivitätszustand zu stabilisieren. Diese Architektur macht das Netzwerk zusätzlich robuster gegenüber Bauteilschwankungen.

Schließlich werden zwei Computermodelle für einen Ausschnitt des Primären Visuellen Kortex (V1) entwickelt. Für das erste Modell, ein Modell des V1 von Makak Affen, werden die Statistiken der Feueraktivität nach Stimulation mit natürlichen Filmsequenzen mit den Statistiken der *in vivo* aufgezeichneten Feueraktivität verglichen. Einige Parameter der Verbindungsstruktur des Modells werden variiert um deren Einfluss auf die Feuerstatistiken zu untersuchen. Das Modell kann nun optimiert werden um ein möglichst biologisch realistisches Feuerverhalten zu erreichen. Die Abweichung des optimierten Modells ist in etwa im selben Bereich wie die Abweichung zwischen den einzelnen Affen und den Experimenten. Das zweite Modell für einen Ausschnitt des V1 von Katzen reproduziert die experimentellen Ergebnisse über die zeitliche Integration von Information über den visuellen Stimulus in diesem kortikalen Areal. Zusätzlich wird gezeigt, wie Korrelationen höherer Ordnung zwischen den Neuronen zur Information über den visuellen Stimulus beitragen.

Contents

1 Introduction **7**

2 Comparison of two data-based cortical microcircuits **11**
- 2.1 Introduction 12
- 2.2 Methods 14
  - 2.2.1 Microcircuits and computational tasks 14
  - 2.2.2 Control circuits 18
  - 2.2.3 Graph properties 20
- 2.3 Results 25
  - 2.3.1 Graph properties of the two microcircuit templates 25
  - 2.3.2 Relationships between graph theoretical properties 28
  - 2.3.3 Dynamical properties 30
  - 2.3.4 Relationship between the graph structure of neural circuits and their dynamical properties 32
  - 2.3.5 Computational properties 33
  - 2.3.6 Relationship between structure, dynamics, and computational properties of the circuits 36
- 2.4 Discussion 38
- 2.5 Acknowledgments 43

3 Compensating inhomogeneities of neuromorphic VLSI **45**
- 3.1 Introduction 46
- 3.2 Material and Methods 47
  - 3.2.1 The Utilized Hardware System 48
  - 3.2.2 The Parallel neural Circuit SIMulator (PCSIM) 53
  - 3.2.3 Network Configuration 54
  - 3.2.4 Measurement 59
- 3.3 Results 60
  - 3.3.1 Self-Adjustment Ability 60
  - 3.3.2 Comparison to PCSIM 62
  - 3.3.3 Robustness 63
  - 3.3.4 Responsiveness to Input 67
- 3.4 Discussion 71
- 3.5 Acknowledgments 72

4 Spike statistics of area V1 *in vivo* and simulations 73
4.1 Introduction 75
4.2 Materials and Methods 77
  4.2.1 Experimental Methods 77
  4.2.2 Model 79
  4.2.3 Comparing the V1 model to electrophysiological data . . . 89
  4.2.4 Estimating the relative strength of the thalamic input . . . 89
  4.2.5 Characterization of the spiking statistics 91
  4.2.6 Evaluating the deviation between model response and *in vivo* recordings 93
  4.2.7 Simulation techniques 94
4.3 Results 95
  4.3.1 Statistical characterization of the spike response to movies in monkey area V1 95
  4.3.2 Quantification of the discrepancy between the firing regime of a model for a patch of V1 and the firing state exhibited *in vivo* 101
  4.3.3 Improvement of the firing regime when optimizing the model 102
  4.3.4 Deviations of the model response to the *in vivo* firing regime 110
4.4 Discussion 114
  4.4.1 Evidence for a characteristic firing state in V1 114
  4.4.2 Is the model in a realistic state? 115
  4.4.3 Which general parameters should be optimized in a cortical network model? 117
  4.4.4 Conclusion 119
4.5 Acknowledgments 120

5 Temporal dynamics of information in a model of V1 121
5.1 Introduction 122
5.2 Materials and Methods 124
  5.2.1 Model 124
  5.2.2 Setup of the stimulus to the model 132
  5.2.3 Simulation techniques 133
  5.2.4 Classification tasks 133
5.3 Results 134
  5.3.1 Similar neural spike responses 134
  5.3.2 Prolonged responses to single letters can not be explained by local synaptic connectivity 138
  5.3.3 Similar non-linear superposition of information 139
  5.3.4 The variety of firing rate profiles correlates with electrode locations 139
  5.3.5 Similar impact of jitter on the computational performance . 140

5.3.6 Higher-order correlations contain little information about stimulus properties . . . 144
5.4 Discussion . . . 146
5.5 Acknowledgments . . . 148

A List of Publications **149**
A.1 Comments and Contribution to Publications . . . 150

Bibliography **151**

List of Figures

2.1 Cortical microcircuit template estimated from paired intracellular recordings according to (Thomson et al., 2002) 22
2.2 Cortical microcircuit template predicted from bouton and target densities estimated from three-dimensional cell reconstructions according to (Binzegger et al., 2004) 23
2.3 Input distributions for the spike pattern classification/memory and exclusive-or (XOR) tasks 24
2.4 Motif distributions for the cortical microcircuit templates 26
2.5 Distributions of motifs with three nodes for three different control circuits 28
2.6 Distributions of motifs with four nodes 30
2.8 Correlation of statistical properties of data-based circuits and corresponding control circuits 32
2.9 Performance of trained linear readout neurons in layers 2/3 and layer 5 35
2.10 Correlation of the mean computational performance and three statistical properties 38

3.1 STP-mechanism of the FACETS Stage 1 Hardware. 50
3.2 Schematic of the self-adjusting network architecture proposed in (Sussillo et al., 2007). 54
3.4 Results of the emulations on the FACETS Stage 1 Hardware. . . . 59
3.5 Results of the software simulation. 61
3.6 Software simulation: Lower excitatory reversal potential. 63
3.7 Reliable and realistic network activity. 64
3.8 Self-adjusting effect on different platforms. 67
3.9 Network traces of transient input. 68

4.1 Long-range connectivity of the V1 model 90
4.2 Spiking response to movie stimulus in model simulation and *in vivo* 98
4.2 Spiking response to movie stimulus in model simulation and *in vivo* 99
4.4 Improvement in the goodness of fit between *in vivo* and *in silico* firing regimes when varying physiological meaningful general parameters. 104
4.5 Effects of parameter optimization on individual statistical features 105
4.6 Optimization of statistical feature individually 109
4.8 Correlation structure in the model and *in vivo* 113

5.1 Receptive fields and connectivity of recorded neurons 124
5.2 Circuit activity in response the Stimulus presentation 136
5.3 The ability of a readout to classify the identity of a letter 137
5.4 The impact of the NMDA-to-AMPA ratio on the firing rates . . . 138
5.5 The information about different parts of the stimulus sequence is simultaneously available 141
5.6 The temporal profile of the average firing rate for different virtual electrode positions 142
5.7 The classification performance of R_t readouts as a function of the amount of Gaussian jitter to spike times 143
5.8 Performance of readouts using linear and polynomial kernels 145

Chapter 1

Introduction

For decades, response properties of single neurons within the primary sensory areas of the cortex have been studied in great detail. It has been shown that neurons in these areas encode or extract specific features of an input signal such as a movie or a tone sequence. The primary visual cortex (V1) is one of these primary sensory cortical areas for which the typical response properties of single neurons are well known. V1 neurons are tuned on e.g. the orientation or the direction of moving bars in the visual input. Another primary sensory area is the primary auditory cortex (A1). It has been found that neurons of A1 are tuned to respond to distinct frequency bands or intensities of the auditory stimulus. Receptive field properties are typically inferred by studies using artificially generated stimuli. However, much less is known about the tuning properties in response to natural stimuli, for which spatial and temporal statistics are typically more complex than for artificial stimuli.

Neurons in V1 that are tuned for an oriented bar at a given position of the visual field are referred to as simple cells. It is a common belief that the response properties of such simple cells emerge from the connections of neurons of the lateral geniculate nucleus (LGN) of the thalamus, which is thought to function as a relay of the retina. Thus, tuning properties can be attributed to the feed-forward connectivity structure from the LGN to V1. The main recipient layer of V1 is the granular layer (layer 4) and the excitatory neurons of this layer typically show the behavior of simple cells. This fact supports the idea that the properties of simple cells mainly emerge from the feed-forward connectivity. The recurrent connectivity within this cortical area is thought to only modulate the activity of these neurons.

However, in the primary visual cortex only 5-10 % of the synapses targeting a neuron in layer 4 (Peters et al., 1994; Peters and Payne, 1993) originate from the LGN. The remaining synapses originate from neurons within the cortex. Such a connectivity structure does not support the idea of a simple feed-forward process that transforms the input. It rather suggests that the recurrent connectivity plays a very important role in the computational function of the sensory cortex. But how does this recurrent connectivity influence the computations or the general dynamics of such cortical circuits?

Moreover, the connectivity in the cortex seems to be highly structured. The connectivity between specific layers varies with respect to the connection probability and the connection strength. The underlying rules of this connectivity structure seem to be crucial for the computational function of the cortex. The aim of this thesis is to analyze how the recurrent connectivity of a cortical neural circuit influences its spiking activity and its computational function. The functional properties of the cortical neural circuits are studied using the framework of liquid computing. This framework suggests that the circuit acts as a reservoir that supports different simultaneous computations. Readout neurons that have access to the activity of the reservoir (i.e., which receive input from the neurons within the reservoir) can transform this activity in order to produce the desired output for a specific computational task.

Organization of the Thesis

This thesis is comprised of four chapters which are based on publications to which I contributed during my PhD studies.

In chapter 2, two data-based templates of connectivity in cortical microcircuits are analyzed and compared regarding their network structure, as well as their dynamical and computational properties. The distribution of network motifs, i.e. of subcircuits consisting of a small number of neurons, are quite different for the two microcircuit templates, although they both have a characteristic small-world property. The dynamical and computational properties of the two microcircuit templates were studied using computer models consisting of Hodgkin-Huxley point neurons with conductance-based synapses that have a biologically realistic short-term plasticity. The models were tested on seven generic computational tasks that require accumulation and merging of information contained in two afferent spike inputs. Although the two models exhibit a different performance for some of these tasks, their average computational performance is very similar. Changing the connectivity structure of the models, we found that the distribution of degrees of nodes is a common key factor for their computational performance. The computational performance of a cortical microcircuit is also correlated with specific statistical properties of the circuit dynamics that is induced by a particular distribution of degrees of nodes.

The simulation of cortical neural circuits with digital computers is typically very time consuming. In chapter 3, a prototype of a spiking neuromorphic VLSI hardware that was developed within the FACETS research project (FACETS, 2009), is used to emulate a cortical neural circuit. Such neuromorphic VLSI hardware systems are promising candidates for neuroscientific research tools and massively parallel computing devices, especially for tasks, which exhaust the computing power of software simulations. Still, like all analog hardware systems, neuromorphic models suffer from a constricted configurability and production-related fluctuations of device characteristics. A cortically inspired self-adjusting network architecture is applied to keep the the activity of generic spiking neural networks emulated on a neuromorphic hardware system within a biologically realistic firing regime. Additionally, this architecture makes the circuit remarkable robust against transistor-level variations.

In chapter 4, a detailed large network for a 5 mm x 5 mm patch of primary visual cortex (V1) of macaque monkeys that integrates many previously published anatomical and physiological details is developed. The statistics of the spike activity of the model in response to a natural movie stimuli is compared with statistics of *in vivo* recordings in response to the same stimuli. Acquiring a realistic firing regime in cortical network models might be a prerequisite for analyzing their emergent computational functions. This chapter also addresses the question how some high level parameters of the recurrent connectivity influence the firing statistics of the model. This knowledge is used to optimize the model to achieve realistic firing statistics. The deviation between the firing regime of the optimized model and the *in vivo* data is on the same level as deviations among monkeys and sessions.

Finally, in chapter 5 a laminar model of a patch of primary visual cortex of cats is used to reproduce the experimental findings of (Nikolic et al., 2009) about the temporal integration of information about the visual stimulus in this area. The model is similar to the one used in chapter 4, but several parameters are adapted to fit the anatomical and physiological properties of cats. Additionally, the model is used to show how higher order correlations of neurons contribute to the information about the visual stimulus. This analysis had to be conducted on large trial sets produced by the model, because the relatively small number of trials from the experiments did not suffice to obtain accurate results for higher order correlations.

CHAPTER 2

Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates

---

Contents

2.1 Introduction **12**
2.2 Methods **14**
2.3 Results **25**
2.4 Discussion **38**
2.5 Acknowledgments **43**

The neocortex is a continuous sheet composed of rather stereotypical local microcircuits that consist of neurons on several laminae with characteristic synaptic connectivity patterns. An understanding of the structure and computational function of these cortical microcircuits may hold the key for understanding the enormous computational power of the neocortex. Two templates for the structure of laminar cortical microcircuits have recently been published by Thomson et al. and Binzegger et al., both resulting from long-lasting experimental studies (but based on different methods).

We analyze and compare in this article the structure of these two microcircuit templates. In particular, we examine the distribution of network motifs, i.e. of subcircuits consisting of a small number of neurons. The distribution of these building blocks has recently emerged as a method for characterizing similarities and differences among complex networks. We show that the two microcircuit templates have quite different distributions of network motifs, although they both have a characteristic small-world property. In order to understand the dynamical and computational properties of these two microcircuit templates, we have generated computer models of them, consisting of Hodgkin-Huxley point neurons with conductance based synapses that have a biologically realistic short-term plasticity. The performance of these two cortical microcircuit models was studied for 7 generic computational tasks that require accumulation and merging of information contained in two afferent spike inputs. Although the two models exhibit a different performance for some of these tasks, their average computational performance is very similar. When we changed the connectivity structure of these two microcircuit models in order to see which aspects of it are essential for computational performance, we found that the distribution of degrees of nodes is a common key factor for their computational performance. We also show that their computational performance is correlated with specific statistical properties of the circuit dynamics that is induced by a particular distribution of degrees of nodes.

2.1 Introduction

Many complex networks from biochemistry and neurobiology as well as engineering share certain global properties (Strogatz, 2001; Newman, 2003; Watts and Strogatz, 1998), like degree distributions (distribution of the number of edges per node) and small-world properties, i.e. local clustering of edges in a graph while maintaining a short path between nodes. But they often have different local properties, yielding different distributions of stereotypical connectivity patterns for few nodes, called motifs (Shen-Orr et al., 2002; Milo et al., 2002, 2004).

Neurobiological studies have shown that cortical circuits have a distinctive modular and laminar structure, with stereotypical connections between neurons that are repeated throughout many cortical areas (White, 1989; Douglas et al., 1995b; Mountcastle, 1998; Nelson, 2002; Silberberg et al., 2002; Douglas and Martin, 2004; Kalisman et al., 2005). It has been conjectured that these stereotypical canonical microcircuits are not merely an artifact of the specific mapping of afferent and efferent cortical pathways or other anatomical constraints like evolutionary processes or development, but are also advantageous for generic computational operations that are carried out throughout the neocortex.

Over the past years detailed statistical data became available that are based on two different experimental methods became available: dual intracellular recordings *in vitro* and cell morphology. The first dataset assembled by (Thomson et al., 2002) was estimated from 998 paired intracellular recordings with sharp electrodes in slices of somatosensory, motor and visual areas of adult rats and adult cats. It specifies connection probabilities and connection strengths of effectively established synaptic connections between excitatory and inhibitory neocortical neurons, to which we will refer as functional connectivity in this chapter. The second dataset assembled by (Binzegger et al., 2004) was predicted from bouton and target densities in cat primary visual cortex estimated from three-dimensional cell reconstructions. This dataset does not specify the distribution of functional connections, but rather represents potential synaptic connectivity. The probabilities of synaptic connections between excitatory and inhibitory neurons located in different layers, i.e. layer 2/3, 4 and 5, differ significantly for the functional and the potential microcircuit template (see (Thomson and Lamy, 2007)). In addition this dataset also includes neurons on layer 6.

We investigate these two cortical microcircuit templates with regard to structural and functional properties. In order to evaluate the computational properties of microcircuit templates we carried out computer simulations of detailed cortical microcircuit models consisting of 560 Hodgkin-Huxley type point neurons and synaptic connections with stereotypical dynamic properties (such as paired pulse depression and paired pulse facilitation) from (Markram et al., 1998). Similar to (Häusler and Maass, 2007), our analysis is based on the assumption that stereotypical cortical microcircuits have some "universal" computational capabilities, and can support quite different computations in different cortical areas. Consequently we concentrate on generic information processing capabilities that are likely to be needed for many concrete computational tasks: to accumulate, hold and fuse information contained in Poisson input spike trains from two different sources (modeling thalamic or cortical feedforward input that arrives primarily in layer 4, and lateral or top-down input that arrives primarily in layers 2/3). In addition we examine the capability of such circuit models to carry out linear and nonlinear computations on time-varying firing rates of these two afferent input streams. In order to avoid rather arbitrary assumptions about the specific type of neuronal encoding of the results of such computations, we analyzed how much information is available about the results of such computations to the generic "neural users", i.e., to pyramidal neurons in layers 2/3 (which typically project to higher cortical areas) and to pyramidal neurons in layer 5 (which typically project to lower cortical areas or to subcortical structures, but also project for example from V1 back to nonspecific thalamus, i.e. to the intralaminar and midline nuclei that do not receive direct primary sensory input, and through this relay to higher cortical areas, see (Callaway, 2004)).

In (Häusler and Maass, 2007) it was shown that the cortical microcircuit model based on the template from (Thomson et al., 2002) exhibits specific computational advantages over various types of control circuits that have the same components and the same global statistics of neurons and synaptic connections, but are missing the lamina-specific structure of real cortical microcircuits. Furthermore it was demonstrated that the connectivity graphs defined by this cortical microcircuit template has a small-world property. However we had shown that the degree distribution of neurons is more salient for their computational performance than the small world property.

Here we extend this study by showing that the two cortical microcircuit templates of (Thomson et al., 2002) and (Binzegger et al., 2004) share some global structural properties, like degree distributions and small-world properties, but have significantly different local structural properties, i.e. network motif distributions. A comparison of the information processing capabilities of both microcircuit templates reveals that they have a similar average computational performance but significantly different computational properties for specific tasks. We also address the question which aspect of the microcircuit template of (Binzegger et al., 2004) is essential for its computational performance, by scrambling specific aspects of their connectivity pattern in a variety of control circuits. We find that, like for the template of (Thomson et al., 2002), the degree distribution of nodes is essential for its computational performance. This result is, besides their similar average computational performance, a second common property of these otherwise quite dissimilar microcircuit templates. We also identify specific properties of the dynamics of the two networks that correlate with their superior computational performance.

2.2 Methods

2.2.1 Microcircuits and computational tasks

We analyzed cortical microcircuit models based on the laminae-specific connectivity pattern specified by two different cortical microcircuit templates. The first cortical microcircuit template assembled by (Thomson et al., 2002) was estimated from paired intracellular recordings with sharp electrodes from 998 pairs of identified neurons from somatosensory, motor and visual areas of adult rats, and visual areas of adult cats. The sampling was made randomly within a lateral spread of $50 - 100\,\mu m$ (Thomson, 2005). This cortical microcircuit template specifies functional synaptic connectivity, i.e. connection probabilities and efficacies of synaptic connections between neurons located in 6 different populations (excitatory and inhibitory neurons in layers 2/3, 4, 5). For those pairs where both data from rat and from cat are given in (Thomson et al., 2002), we have taken the data from rat (see Fig. 2.1). Only for pairs of neurons within layer 4 no data from rat are given in (Thomson et al., 2002), hence the corresponding data are from cat.[1] We analyzed a model of this microcircuit template that consisted of 560 neurons, with 30%, 20%, and 50% of the neurons assigned to layers 2/3, layer 4, and layer 5, respectively (the number 560 was chosen somewhat arbitrarily, based on required simulation speed and programming details). We will refer to the microcircuit model based on this cortical microcircuit template as Thomson et al. circuit[2].

---

[1] Some of the pairings were rarely observed and the corresponding entries suffer from small sample size (see Thomson et al. (2002) for details). Also very small neurons in rat may have been missed in this study (Thomson, 2005). In addition it is likely that in some cortical microcircuits connections exist between pairs of neurons for which no connections were reported in (Thomson et al., 2002) (see for example (Dantzker and Callaway, 2000) for the case of connections to inhibitory neurons in layers 2/3).

[2] One should note that the layers defined by this and other microcircuit templates are not induced by their graph structure (like the layers in a multi layer perceptron). Rather the layer

The second cortical microcircuit template assembled by (Binzegger et al., 2004) was predicted from bouton and target densities in cat primary visual cortex estimated from three-dimensional reconstructions of cells *in vivo*. This cortical microcircuit template specifies potential synaptic connectivity between neurons located in 10 specific populations (excitatory and inhibitory neurons in layers 1, 2/3, 4, 5 and 6). We omitted layer 1, because it only receives input and provides no synaptic connections back to other layers, and is therefore irrelevant for the analysis in this study. The connectivity graph of the 8 modeled populations is shown in Fig. 2.2. A number of additional unassigned connections were reported in (Binzegger et al., 2004). These were not considered for the calculation of connection probabilities in this article. The synaptic connection probabilities were further rescaled to achieve on average the same number of 42540 synaptic connections as obtained for the Thomson et al. circuits. This resulted in a mean connection probability of 13% or, equivalently, an average number of 76 recurrent synaptic connections per neuron. The microcircuit model consisted of 560 neurons, with 36%, 36%, 7% and 21% of the neurons assigned to layers 2/3, layer 4, layer 5 and layer 6, respectively (see (Binzegger et al., 2004) for a discussion of data which justify these estimates).[3] Because the cortical microcircuit template of (Binzegger et al., 2004) provides no strengths of synaptic connections (i.e., synaptic weights), we modeled the distribution of synaptic weights according to the cortical microcircuit template of (Thomson et al., 2002). The strengths of synaptic connections from and to layer 6, which do not occur in Thomson et al. circuits, were set to the average values of all other connections of the corresponding synapse type. [4] These weights are labeled with "?" in Fig. 2.2. We will refer to microcircuit models based on the second cortical microcircuit template as Binzegger et al. circuits. Each layer of both cortical microcircuit templates consisted of a population of excitatory neurons and a population of inhibitory neurons with a ratio of 4:1.

The short term dynamics of cortical synapses (i.e., their specific mixture of paired pulse depression and paired pulse facilitation) is known to depend on the type of the presynaptic and postsynaptic neuron. We modeled the short term synaptic dynamics according to the model proposed in (Markram et al., 1998), with synaptic parameters chosen as in (Maass et al., 2002) to fit data from microcircuits in rat somatosensory cortex (based on (Gupta et al., 2000) and (Markram et al., 1998)). The maximum conductances of synapses were chosen from a Gaussian distribution with a SD of 70% of its mean (the negative values were replaced

---

to which a neuron belongs should formally be viewed as a label of the corresponding node in the graph.

[3] We verified that the reported results are qualitatively the same for circuits consisting of up to 1000 neurons. In general, the performance scales with the network size (see Häusler and Maass, 2007).

[4] Synapse types are defined according to the post-and presynaptic neuron type

by values chosen from an uniform distribution between zero and two times the mean) [5]. The mean maximum conductances of synapses were chosen to reproduce the mean amplitude of PSPs given in Fig. 2.1 and Fig. 2.2 at the resting membrane potential (in the presence of synaptic background activity). Synaptic transmission delays between neurons were chosen from Gaussian distributions with mean 1.5 ms (0.8 ms) for connections between excitatory neurons (all other connections) and a standard deviation of 0.1 times the mean.

Excitatory and inhibitory neurons were modeled as conductance based single compartment Hodgkin-Huxley neuron models with passive and active properties modeled according to (Destexhe et al., 2001). A cortical neuron receives synaptic inputs not only from immediately adjacent neurons, but also smaller background input currents from a large number of more distal neurons, causing in awake animals a depolarization of the membrane potential commonly referred to as 'high conductance state'. This was reflected in our computer model by background input currents that were injected into each neuron. The conductances of these background currents were modeled as a one-variable stochastic process similar to an Ornstein-Uhlenbeck process with parameters obtained from a biophysical model matched to intracellular recordings from a L5 neuron from cat cerebral cortex (Destexhe et al., 2001).

Two afferent input streams, each consisting of 4 or 40 spike trains (i.e., 4 or 40 input channels), were injected into the circuit. Each of the channels of the first input stream (representing thalamic, or feedforward cortical input) was injected mainly into layer 4, i.e. to 50% of its inhibitory neurons and 80% of its excitatory neurons, but also into 20% of the excitatory neurons in layers 2/3 and 10% of the excitatory neurons in layer 5 (all randomly chosen). Each of the channels of the second afferent input stream was injected into 20% of the excitatory neurons on layers 2/3. The mean maximum conductances of all input synapses were chosen to generate a PSP with a mean amplitude of 1.9 mV at the resting membrane potential in the presence of background input. Input synapses were chosen to be static and their maximum conductances were chosen from a Gaussian distribution with a SD of 70% of its mean (with negative values replaced by values chosen from an uniform distribution between zero and two times the mean).

In addition to these data there remain 3 parameters S_{I1}, S_{I2} and S_{RW} that scale (in the form of multiplicative factors) the amplitudes of PSPs for all recurrent synapses, and EPSPs from the two input streams (1 and 2). These parameters were adjusted for inputs consisting of 40 Poisson spike trains at 20 Hz to induce biologically plausible average firing rates. The parameter S_{I1} (S_{I2}) was chosen so that its respective afferent input stream caused in the absence of the other input stream S_{I2} (S_{I1}) and without recurrent connections ($S_{RW} = 0$) but with

---

[5] No observable differences occurred in case of replacing this distribution with a gamma distribution with a scale parameter that equals the specified mean and a shape parameter of 1.

background input currents injected into each neuron an average firing rate of 15 Hz in layer 4 (10 Hz in layers 2/3). The parameter S_{RW} was set to a value that produced for the previously fixed parameter S_{I1} and S_{I2} in the presence of synaptic background noise a realistic low but significant firing activity of 8.5 Hz in layer 5 and an average firing rate of 24 Hz in layer 2/3, layer 4 and layer 5. To achieve these firing rates for the Binzegger et al. circuits, the parameter S_{RW} had to be scaled down by a factor of 3.3 compared with its value for the Thomson et al. circuits. For the analysis reported in Fig. 2.8 and Fig. 2.10 we performed simulations with 40 randomly chosen scaling parameters (S_{I1}, S_{I2} and S_{RW}) that were drawn uniformly from the interval $[0.2, 2]$ times the previously described standard values. The resulting average firing rates ranged from 4 to 55 Hz.

We further modeled hypothetical projection neuron in layers 2/3 and layer 5. The set of presynaptic neurons for such hypothetical readout neuron was chosen according to Fig. 2.1 for Thomson et al. circuits and Fig. 2.2 for Binzegger et al. circuits, but no synaptic connections from a readout neuron back into the circuit were included. This amounted for the Thomson et al. circuits to an average of 84 presynaptic neurons for a readout neuron on layers 2/3, and 109 presynaptic neurons for a readout neuron on layer 5. To allow a fair comparison between the two cortical microcircuit templates the connections probabilities for the hypothetical readout neurons of Binzegger et al. circuits were rescaled by a multiplicative factor to obtain the same number of presynaptic neurons for each of the two readout neurons.

The projection or readout neurons themselves were modeled as linear neurons, i.e., their output was a weighted sum of low pass filtered spikes (exponential decay with a time constant of 15 ms, modeling the time constants of synaptic receptors and membrane of a readout neuron). The weights of synaptic connections from the presynaptic neurons to the readout neuron were optimized for specific tasks. Care was taken to make sure that weights from excitatory (inhibitory) presynaptic neurons could not become negative (positive) by using a standard method of linear regression with sign-constraints.

The information processing tasks comprised spike pattern classification tasks, i.e. classification of spike patterns in either of the two afferent input streams, memory tasks (classification of earlier spike patterns in one of the two input streams), and non-linear fusion of information from spike patterns in both input streams, as well as real-time computations on the firing rates of both input streams. For information processing tasks with spike patterns we randomly generated spike pattern templates consisting of 30 ms segments of 40 Poisson spike trains at 20 Hz (see Fig. 2.3). More precisely, the spike trains of each of the two input streams were of length 450 ms and consisted of 15 consecutive time segments of length 30 ms. For each time segment 2 spike pattern templates were generated randomly. For the actual input one of the two templates of each time segment was chosen randomly (with equal probability) and a noisy variation of it was injected into the circuit, where each spike was shifted by an amount drawn from a Gaussian distribution with mean 0 and SD 1 ms. Readout neurons were trained to classify which of the two spike templates fixed for input 1 (input 2) was injected during the last time interval $[t-30\ ms, t\ ms]$, denoted as task $\mathrm{tcl}_1(t)$ ($\mathrm{tcl}_2(t)$), or during the preceding time interval $[t-60\ ms, t-30\ ms]$, denoted as task $\mathrm{tcl}_1(t-\Delta t)$ ($\mathrm{tcl}_2(t-\Delta t)$). Note that the latter task is more difficult because the relevant spike input during time $[t-60\ ms, t-30\ ms]$ is overwritten by new input before the readout takes place. It may be viewed as a memory task (with distractors).

Nonlinear fusion of information from both input streams was tested by training readouts to output the exclusive-or (XOR) of the two bits that represent the labels of the two templates from which the most recent spike patterns in the two input streams had been generated. Note that this computation involves a nonlinear "binding" operation on spike patterns, since it has to give a low output value if and only if the labels of the noisy spike templates injected in input stream 1 and 2 are identical. The XOR task has been used in the neural network literature as a standard example for a nonlinear computational task.

In addition we analyzed nonlinear computations on time-varying firing rates of the two input streams. The spike trains of each of the two input streams were of length 450 ms and consisted of 15 time segments of length 30 ms. For each input stream and each time segment 4 Poisson spike trains were generated with a randomly chosen time-varying frequency between 15 Hz and 25 Hz. The actual firing rates of both input streams, i.e. $r1$ and $r2$, used for the computations on these input firing rates, i.e. $r_1(t)/r_2(t)$ and $(r_1(t)-r_2(t))^2$, were calculated from these spike trains with a sliding window of 15 ms width. The error bars in Fig. 2.9 denote standard errors. All performance results are for test inputs, and freshly generated random initial conditions and random background noise for all neurons in the circuit. All simulations were carried out with the CSIM software (Natschläger et al., 2003) in combination with MATLAB. For further details see (Häusler and Maass, 2007).

2.2.2 Control circuits

Control circuits have the same components and the same global statistics of neurons and synaptic connections, but are missing the lamina-specific connectivity structure of data-based cortical microcircuits. Random control circuits were generated from data-based circuits by randomly rewiring recurrent synaptic connections whereas no synaptic connection was allowed to occur more than once. In order to maintain the stereotypical neuron to synapse type alignment observed for short-term synaptic plasticity (Gupta et al., 2000) the rewiring was carried out under the constraint that the pre- and postsynaptic neuron type (i.e. excitatory or inhibitory) of each synaptic connection stays the same. This constraint introduces a difference in the randomized networks generated from Thomson et al. circuits and Binzegger et al. circuits, because the numbers of synaptic connections between excitatory neurons (N_{EE}) are different in these two circuit templates. The same holds for the number of synaptic connections between inhibitory neurons (N_{II}), from inhibitory to excitatory (N_{EI}), and from excitatory to inhibitory neurons (N_{IE}), although the total number of synapses is identical for both circuits[6]. The difference is largest for N_{II}, which is for the Thomson et al. circuits more than 3 times larger. We will refer to the randomized networks generated from the two microcircuit templates as amorphous Thomson et al. (amorphous Binzegger et al.) circuits[7].

Degree-controlled circuits (Kannan et al., 1999; Maslov and Sneppen, 2002) preserve the degree distributions of neurons in all layers but otherwise lack a laminae-specific connectivity pattern. The degree of a neuron is defined as the total number of incoming and outgoing synaptic connections. Degree-controlled circuits were constructed from data-based circuits by randomly exchanging the target neurons of pairs of recurrent synaptic connections that emerge from the same neuron, and have neurons of the same type (excitatory/inhibitory) as target.

Small-world control circuits have identical cluster coefficient and average shortest path lengths as Thomson et al. circuits (but without laminae-specific synaptic connectivity pattern). They were constructed with two different algorithms, i.e. the spatial growth algorithm described in (Kaiser and Hilgetag, 2004) (with parameters $\alpha = 4$, $\beta = 1.32$) and the algorithm proposed by (Watts and Strogatz, 1998) (with parameter $\beta = 0.319$). Both algorithms were applied to generate undirected graphs that have the same size (560 nodes), clustering coefficient, and average shortest path length as Thomson et al. circuits. Each node in these graphs was replaced by a randomly drawn excitatory or inhibitory neuron (without replacement) located in one of the 6 populations of Thomson et al. circuits. Subsequently the undirected graphs were converted to directed graphs by randomly replacing each edge with a synapse (that is randomly oriented) or a reciprocal synaptic connection, with a probability so that the total number of synaptic connections and reciprocal synaptic connections is identical to the corresponding number for Thomson et al. circuits [8].

For all control circuits the assignment of neurons to layers, the target neurons of input synapses, and the set of presynaptic neurons for hypothetical readouts was the same as for the corresponding data-based circuits.

---

[6] For Thomson et. al circuits (Binzegger et al. circuits) the average values of N_{EE}, N_{II}, N_{EI}, and N_{IE} are 31011 (32243), 2533 (785), 5370 (4762), and 3086 (4210), respectively.

[7] Note that the connection probabilities for amorphous circuits are not uniform, but differ for each of the four synaptic connection types.

[8] It should be noted that this procedure does not reproduce the same fraction of synapse types as for data-based circuits and amorphous circuits.

2.2.3 Graph properties

A cortical microcircuit can be described by a directed graph with nodes (neurons) and edges (synaptic connections). A connected sub-graph with M nodes is called a motif of size M and has at least $M-1$ and at most $M(M-1)$ edges (ignoring self-edges). The motif count COUNT is defined as the total number of motifs of a certain type occurring in a directed graph that corresponds to a cortical microcircuit generated from a specific microcircuit template. The probability of a motif is defined as the probability that M randomly drawn neurons within a microcircuit (that was randomly drawn from a specific microcircuit template) are connected according to the sub-graph defined by the motif (with no additional edges except self-edges). We analyzed sets of motifs of size $M = 2$, $M = 3$ and $M = 4$ consisting of 2, 13 and 199 sub-graphs, respectively. According to (Milo et al., 2002, 2004) we didn't consider motifs that contain self-edges.

The motif probabilities for data-based microcircuits, amorphous circuits and degree-controlled circuits can be calculated analytically from the corresponding microcircuit templates. The templates for the connection probabilities of amorphous circuits can be obtained from data-based circuits (defined in Fig. 2.1 and Fig. 2.2) by setting the probability of each synaptic connection to the weighted average probability of all synaptic connections of the same type (defined by the type of the pre- and postsynaptic population of neurons, i.e. excitatory or inhibitory). The contribution of each synaptic connection probability to the average is weighted proportional to the product of the pre- and postsynaptic population size. The templates for the connection probabilities of degree-controlled circuits can be obtained from data-based templates by carrying out the following procedure for each type of synaptic connection. First the probability $P^s(i)$ that a randomly selected synapse of a specified type targets a specific population i is calculated from the data-based template. Subsequently the probability of a synaptic connection between presynaptic population j and postsynaptic population k is set to the value $P^s(k)$ resulting in a connectivity pattern that is independent of the presynaptic population. Finally, the probabilities of all connections with identical presynaptic population are multiplied by a common scaling factor to obtain the same average total number of outgoing synapses for each population as for the data-based microcircuit template.

The probability of a certain motif P^M is defined as the product of the probabilities for the existence or absence of each edge of a sub-graph (excluding self-edges). Care has to be taken to account for permutations of neurons and different populations sizes. For a distribution of circuits consisting of N neurons that were generated from a specific circuit template the mean motif count is given by

$$\overline{\mathrm{COUNT}} = \binom{N}{M} \cdot P^M$$

and the standard deviation of the motif count is defined by

$$std(\text{COUNT}) = \sqrt{\binom{N}{M} \cdot P^M (P^M - 1)}$$

The motif distributions for small-world circuits were sampled using 200 circuits.

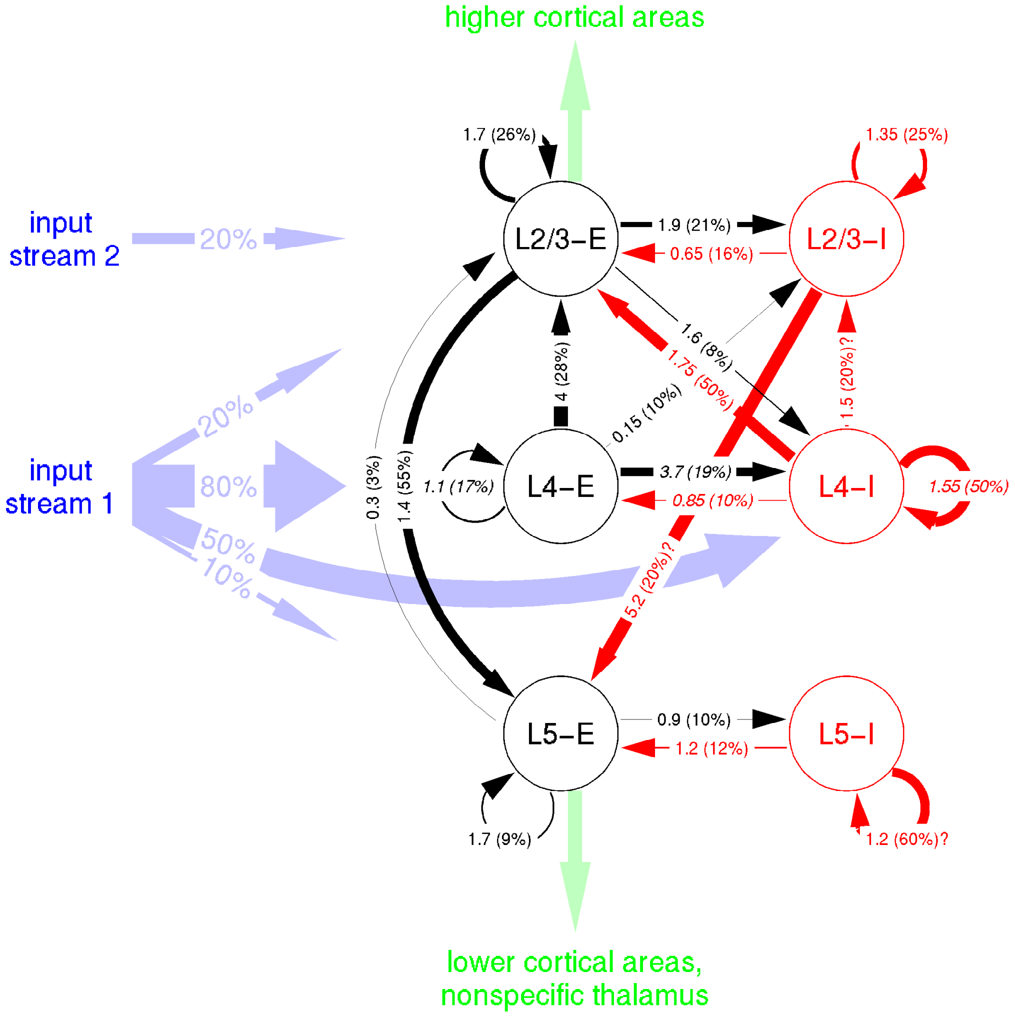

Figure 2.1: Cortical microcircuit template estimated from paired intracellular recordings according to (Thomson et al., 2002)
Numbers at arrows denote connection strengths (mean amplitude of postsynaptic potentials, PSPs, measured at soma in mV) and connection probabilities (in parentheses), for connections between cortical neurons in 3 different layers, each consisting of an excitatory (E) and an inhibitory (I) population, with an estimated maximal horizontal distance of up to 100 μm. The width of arrows is proportional to the product of these two numbers. Most of the data are from rat cortex, except for interconnections in layer 4 (italic), which are from cat. Input stream 1 models feedforward inputs, and input stream 2 models top-down or lateral input to the cortical microcircuit. Percentages at input streams denote connection probabilities for input neurons (that produce these input streams) in our simulations. Figure reproduced with permission of (Häusler and Maass, 2007).

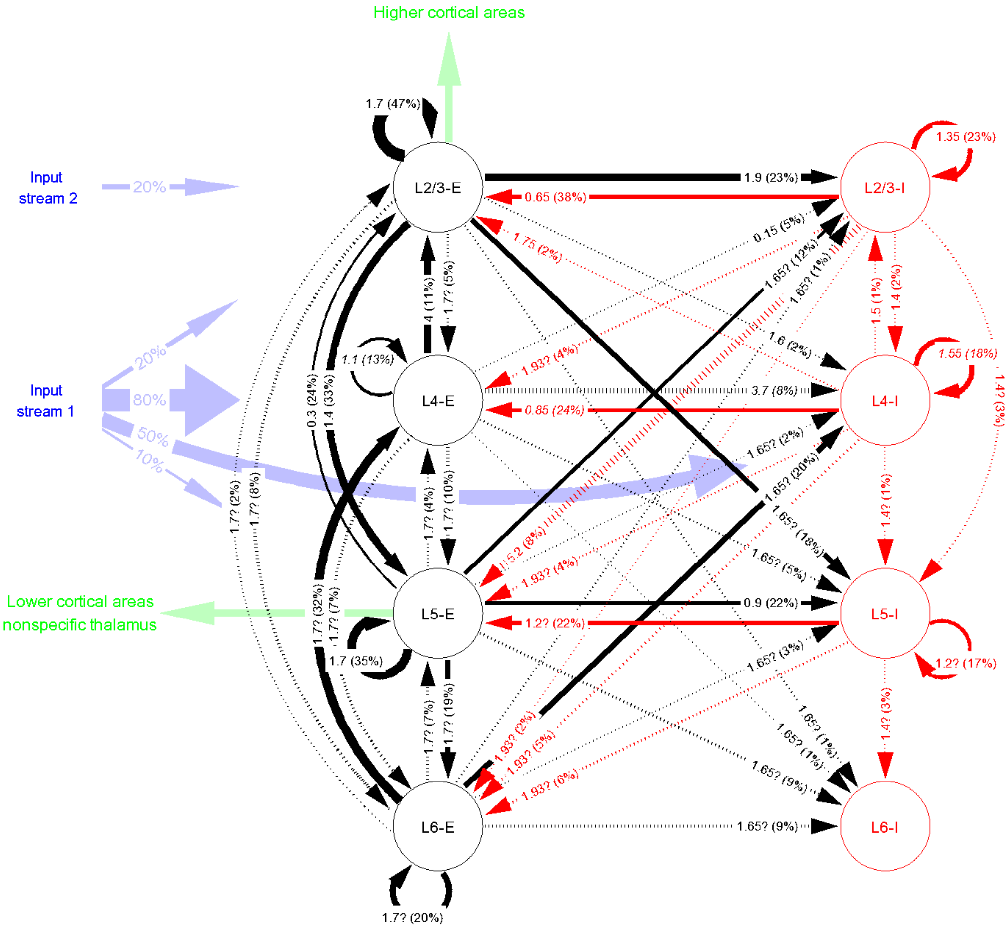

Figure 2.2: Cortical microcircuit template predicted from bouton and target densities estimated from three-dimensional cell reconstructions according to (Binzegger et al., 2004)

Numbers at arrows denote connection strengths (mean amplitude of postsynaptic potentials, PSPs, measured at soma in mV) and connection probabilities (in parentheses), for connections between cortical neurons in 4 different layers, each consisting of an excitatory (E) and an inhibitory (I) population. Connection probabilities were estimated from 39 three-dimensional single neuron reconstructions from cat primary visual cortex. Connection strengths were taken from the cortical microcircuit template described in (Thomson et al., 2002) (see Fig. 2.1). Weights with question marks denote unspecified connections (see text).

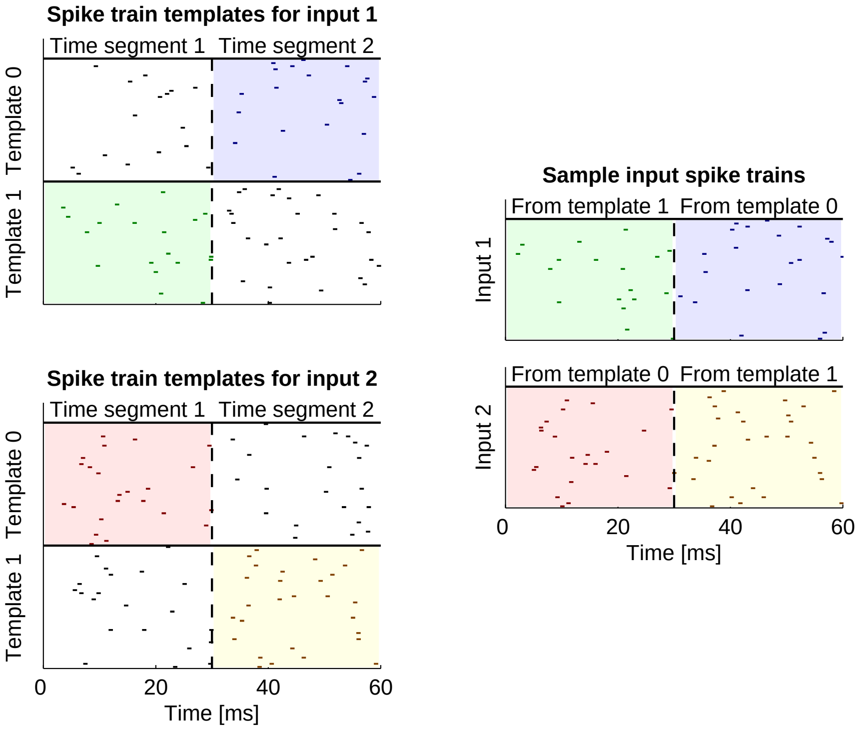

Figure 2.3: Input distributions for the spike pattern classification/memory and exclusive-or (XOR) tasks
The spike trains of each of the two input streams were of length 450 ms and consisted of 15 time segments of length 30 ms. For each segment 2 templates were generated randomly (40 Poisson spike trains at 20 Hz). The actual spike trains of each input of length 450 ms used for training or testing were generated by choosing for each segment one of the two previously chosen associated templates, and then generating a jittered version by moving each spike by an amount drawn from a Gaussian distribution with mean 0 and a SD 1 ms (a sample for two time segments is shown in the panel on the right hand side). Figure reproduced with permission of (Häusler and Maass, 2007)

2.3 Results

2.3.1 Graph properties of the two microcircuit templates

We analyzed the two data-based cortical microcircuit templates[9] shown in Fig. 2.1 and Fig. 2.2 for their differences and similarities in connectivity structure. In order to evaluate the significance of specific structural features of the two data-based microcircuit templates we compared them with random control circuits which consist of the same number of components, i.e. neurons and synapses, but lack a laminae-specific connectivity pattern. These random control circuits were generated from data-based circuits by randomly rewiring synaptic connections[10]. This rewiring was carried out under the constraint that the pre- and postsynaptic neuron type (i.e. excitatory or inhibitory) of each synaptic connection stays the same, in order to maintain the stereotypical neuron to synapse type alignment observed for short-term synaptic plasticity (Gupta et al., 2000) that we have implemented in our models. We will refer to the randomized networks generated from the two microcircuit templates as amorphous Thomson et al. (amorphous Binzegger et al.) circuits.

In general the connectivity graph of Binzegger et al. circuits is more similar to the connectivity graph of the corresponding amorphous circuits than for Thomson et al. circuits. The connection probabilities defined by the laminae-specific connectivity templates (see Fig. 2.1 and Fig. 2.2) correlate with the connection probabilities of the corresponding amorphous templates with a correlation coefficient of 0.3 for Binzegger et al. circuits and 0.2 for Thomson et al. circuits.

Three graph properties have primarily been used for the characterization of naturally occurring directed graphs: clustering, degree distribution and motif distribution. A quantity often studied in relation with clustering is the small-world property defined by (Watts and Strogatz, 1998). In small-world networks neighbors of a node are more likely to be neighbors themselves when compared to random graphs, thereby causing a so-called small-world effect. Nevertheless any two nodes in a small-world network are connected by a relative small number of edges, providing fast communication between any two nodes[11]. Thomson et al. circuits as well as the Binzegger et al. circuits exhibit a significant small-world property. Their cluster coefficient, that is defined as the fraction of existing edges between direct neighbors of a node, is 0.36 and 0.33, respectively, which is 37% and 30% larger than in the corresponding amorphous circuits. Both data-based microcircuit templates imply an average shortest path length of 1.77 edges, which is comparable to the average shortest path length of 1.74 edges for amorphous

---

[9] To be precise, each of these two templates is actually a probability distribution over graphs, rather than a specific graph.

[10] But no synaptic connection was allowed to occur more than once.

[11] Note that both properties refer to the structure of the underlying undirected graph, where directed edges are replaced by undirected edges.

control circuits.

The higher cluster coefficient of Thomson et al. circuits is mainly due to excitatory and inhibitory neurons located in layer 2/3, which form highly connected hubs. The amount of convergence and divergence of a node in a graph can be specified by its degree, defined as the total number of its incoming and outgoing connections. The average degree of a layer 2/3 neuron in Thomson et al. circuits is 251, which is 20.3% larger than the average degree of 208.7 for a layer 2/3 neuron in Binzegger et al. circuits. The average degree in the remaining layers was somewhat smaller and ranged from 77.6 to 137.3 for Thomson et al. circuits and from 14.1 to 188.3 for Binzegger et al. circuits (without layer 1). Overall the average degrees of neurons in layer 2/3, 4 and 5 are correlated for both data-based circuits with a correlation coefficient of 0.74.

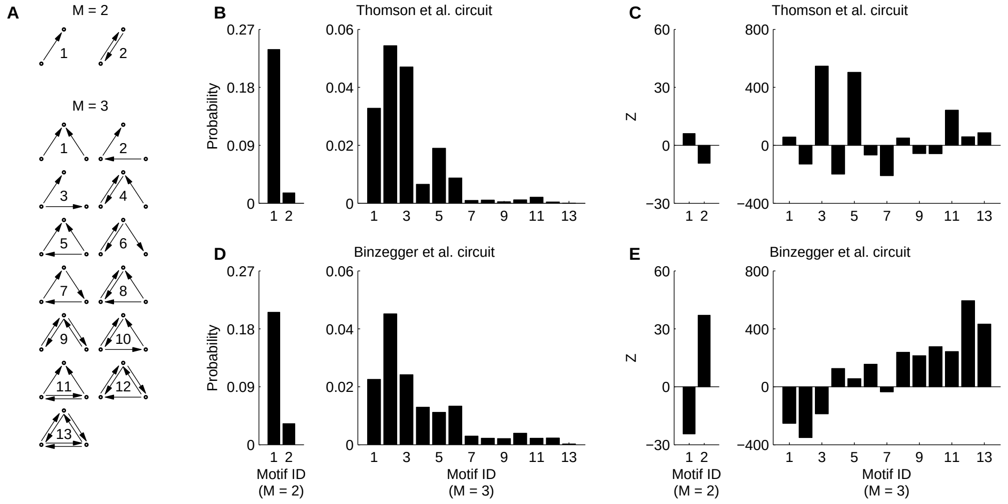

Figure 2.4: Motif distributions for the cortical microcircuit templates Motif distributions of Thomson et al. (2002) (see Fig. 2.1) and Binzegger et al. (2004) (see Fig. 2.2). **A:** Definition of motifs. **B:** Motif probabilities for pairs (left panel) and triplets (right panel) of neurons for Thomson et al. circuits. **C:** Z score (defined as the difference in the average motif count for the specified circuits and corresponding amorphous circuits measured in units of the standard deviation of the motif count for amorphous circuits) for pairs and triplets of neurons for Thomson et al. circuits consisting of 560 neurons and on average 42540 synapses. **D:** Motif probabilities for pairs and triplets of neurons for Binzegger et al. circuits. **E:** Z score for Binzegger et al. circuits consisting of 560 neurons and on average 42540 synapses. The Motif distribution for the two cortical microcircuit templates differs significantly.

A third approach to characterize graphs, which also takes the direction of edges into account, is to analyze which subgraphs (motifs) occur with a frequency significantly higher or lower than in corresponding amorphous networks (see section 2.2.3). The motif distributions of Thomson et al. circuits and Binzegger et al.circuits are shown in Fig. 2.4 for motifs consisting of two and three nodes.

The deviation of the motif distribution of the data-based microcircuits from the motif distribution of the corresponding amorphous circuits was evaluated for each motif by means of its Z score (Milo et al., 2002, 2004) defined as

$$Z = \frac{\overline{\mathrm{COUNT}}_{data-based} - \overline{\mathrm{COUNT}}_{amorphous}}{std\,(\mathrm{COUNT}_{amorphous})},$$

where std denotes the standard deviation and $\overline{\mathrm{COUNT}}$ denotes the mean motif count obtained for a distribution of circuits that were generated from a specific circuit template. Remarkably the procedure of generating amorphous control circuits generates more reciprocal connections for the Thomson et al. circuits but fewer reciprocal connections for the Binzegger et al. circuits, as indicated by the Z scores for motif number 2 for the motif class consisting of two nodes ($M = 2$), see Fig. 2.4C, E. Also the Z scores for motifs consisting of three nodes differ significantly for the two templates. In particular, the motifs 3, 5 and 11 (which represent converging or diverging sub-graphs) are over-represented, whereas the motifs 2, 4 and 7 (which represent feed-forward or circular sub-graphs) are under-represented in Thomson et al. circuits. The more frequent appearance of motif 5 in Thomson et al. circuits can be attributed to the typical structure of connections from layer 4 to layer 2/3 and from layer 2/3 to layer 5. In both cases excitatory neurons within the target and the source layer are often directly synaptically connected (corresponding to the edge from the right to the top node of motif 5). Additionally, excitatory neurons in the target layer receive input from inhibitory neurons in the source layer (edge from the left to the top node of motif 5) that receive input from excitatory neurons within the source layer (edge from the right to the left node of motif 5). Therefore the excitation spreading from layer 4 (layer 2/3) to layer 2/3 (layer 5) is balanced by an inhibitory pathway passing through inhibitory neurons in layer 4 (layer 2/3).

In contrast, motifs with 3 nodes in Binzegger et al. circuits that consist only of two edges are significantly under-represented, whereas motifs with many edges appear more frequently than in corresponding amorphous circuits (see Fig. 2.4E. This may be partially attributed to the more frequent appearance of reciprocal connections compared to amorphous circuits.

The characteristic shape of the motif distribution of Binzegger et al. circuits is not due to the additional connections with layer 6, which are not specified in Thomson et al. circuits. The motif distributions of Binzegger et al. circuits with and without layer 6 have a correlation coefficient of 0.96. Moreover the characteristic shape of their motif distribution is also not caused by the comparatively high number of connections with low probability ($< 10\%$). The correlation coefficient of the motif distributions of Binzegger et al. circuits and pruned Binzegger et al. circuits, for which all connections with a connection probability lower than 10% were removed (dashed lines in Fig. 2.2), is 0.97.

2.3.2 Relationships between graph theoretical properties

In order to identify the structural aspects of data-based microcircuit templates that are responsible for the characteristic deviation in their motif distributions, we generated additional random control circuits which preserve the degree distributions of neurons in all layers, but otherwise lack a layer specificity of synaptic connections. These degree-controlled circuits (Kannan et al., 1999; Maslov and Sneppen, 2002) were constructed from data-based circuits by randomly exchanging the target neurons of pairs of synapses that emerge from the same neuron, and have neurons of the same type (excitatory/inhibitory) as target. In contrast to the scrambling procedure that generates amorphous circuits, this procedure preserves not only the type but also the identity of the presynaptic neuron.

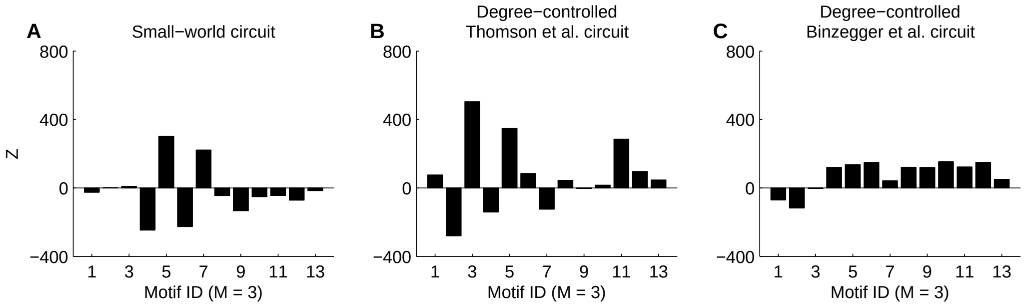

Figure 2.5: Distributions of motifs with three nodes for three different control circuits
A: Z score for small-world control circuits generated with the spatial growth algorithm described in (Kaiser and Hilgetag, 2004). The result for a second type of small-world control circuits generated with the algorithm proposed by (Watts and Strogatz, 1998) is nearly identical (correlation coefficient 0.99). **B:** Z score for degree-controlled control circuits generated from Thomson et al. circuits. **C:** Z score for degree-controlled control circuits generated from Binzegger et al. circuits. The degree-controlled circuits generated from Thomson et al. circuits have a Z score that resembles the Z score of Thomson et al. circuits (the average absolute difference in Z score represents 40% of the average absolute Z score for data-based circuits). In contrast, the motif distribution of degree-controlled circuits generated from Binzegger et al. circuits is more similar to the one for corresponding amorphous circuits (the average absolute difference in Z score is 65% of the average absolute Z score for data-based circuits).

The motif distribution of Thomson et al. circuits changes very little through this scrambling procedure (compare Fig. 2.5B with Fig. 2.4C). The motifs 2 and 5 that represent feed-forward pathways from the right to the top node of the sub-graph (see Fig. 2.4A) appear less frequent, but nevertheless the Z score for data-based circuits and degree-controlled circuits correlates with a correlation coefficient of 0.94. Therefore the characteristic motif distribution of Thomson et al. circuits is mainly induced by the specific distribution of degrees of neurons over layers.

In contrast, the motif distribution of degree-controlled circuits generated from Binzegger et al. circuits resembles the one for amorphous circuits (see the Z-score plotted in Fig. 2.5C). In particular, motifs consisting of many edges appear less frequently in degree-controlled circuits compared to data-based circuits. Although the Z score is smaller for the majority of motifs in degree-controlled circuits, it nevertheless correlates with the Z score of the Binzegger et al. circuits with a correlation coefficient of 0.79. Thus the motif distribution of Binzegger et al. circuits is induced by the laminae-specific connectivity pattern between layers and less by the specific degree distributions of neurons. Similar results are obtained for the distributions of motifs consisting of 4 nodes (see Fig. 2.6). [12]

In order to demonstrate that the motif distribution of Thomson et al. circuits can not only be attributed to its specific clustering properties (small world property), we additionally generated control circuits with identical cluster coefficient and average shortest path length as Thomson et al. circuits but without laminae-specific synaptic connectivity pattern. In order to show that the motif distribution does not depend on a specific construction algorithm we generated two types of small-world control circuits. The first one was constructed with the spatial growth algorithm described in (Kaiser and Hilgetag, 2004) and the second type of control circuits was generated with the algorithm proposed by (Watts and Strogatz, 1998). The spatial growth algorithm of (Kaiser and Hilgetag, 2004) is capable of constructing networks with multiple, interconnected clusters, whereas the algorithm described by (Watts and Strogatz, 1998) doesn't preferentially connect highly connected hubs with each other. Note that both algorithms generate undirected graphs that were subsequently converted to directed graphs by randomly replacing each edge with a synapse (that is randomly oriented) or a reciprocal synaptic connection, with a probability so that the total number of synaptic connections and reciprocal synaptic connections is identical to the corresponding number for data-based circuits. [13] The $M = 3$ motif distribution for both types of small-world circuits is similar but differs significantly from the motif distribution of data-based microcircuits (see Fig. 2.5A). The correlation coefficient for the

---

[12] For $M = 4$ motifs the Z score for Thomson et al. circuits (Binzegger et al. circuits) and the corresponding degree-controlled circuits correlates with a correlation coefficient of 0.91 (0.41).

[13] It should be noted that this procedure does not reproduce the same fraction of synapse types as for data-based circuits and amorphous circuits.

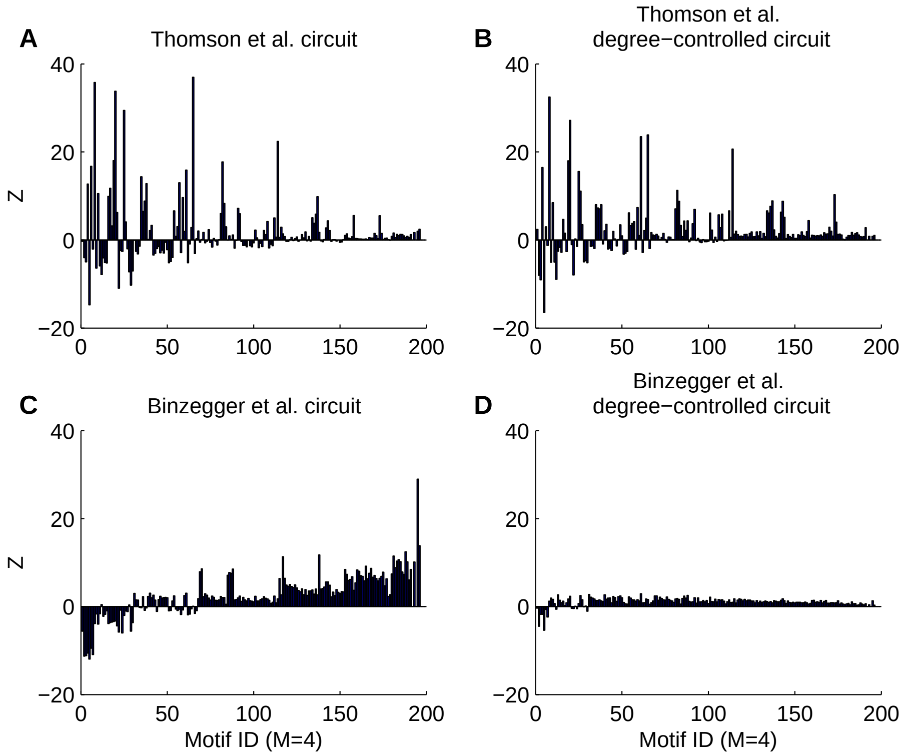

Figure 2.6: Distributions of motifs with four nodes
A: Z score for Thomson et al. circuits. **B:** Z score for degree-controlled circuits generated from Thomson et al. circuits. **C:** Z score for Binzegger et al. circuits. **D:** Z score for degree-controlled circuits generated from Binzegger et al. circuits. These plots show that the degree-controlled scrambling of the connectivity structure mainly preserves the Z score for Thomson et al. circuits (correlation coefficient of 0.91) but not for Binzegger et al. circuits (correlation coefficient of 0.41).

motif distribution of small-world circuits, and Thomson et al. circuits and the Binzegger et al. circuits is 0.37 and -0.29, respectively.

2.3.3 Dynamical properties

We investigated the dynamical properties of Thomson et al. circuits and Binzegger et al. circuits by analyzing computer simulations of detailed cortical microcircuit models (see Methods). Fig. 2.7 shows a comparison of various statistical properties of the circuit dynamics in response to generic spike inputs consisting of 40 Poisson spike trains at 20 Hz into layer 4 and layers 2/3 for both data-based circuit templates (black lines).

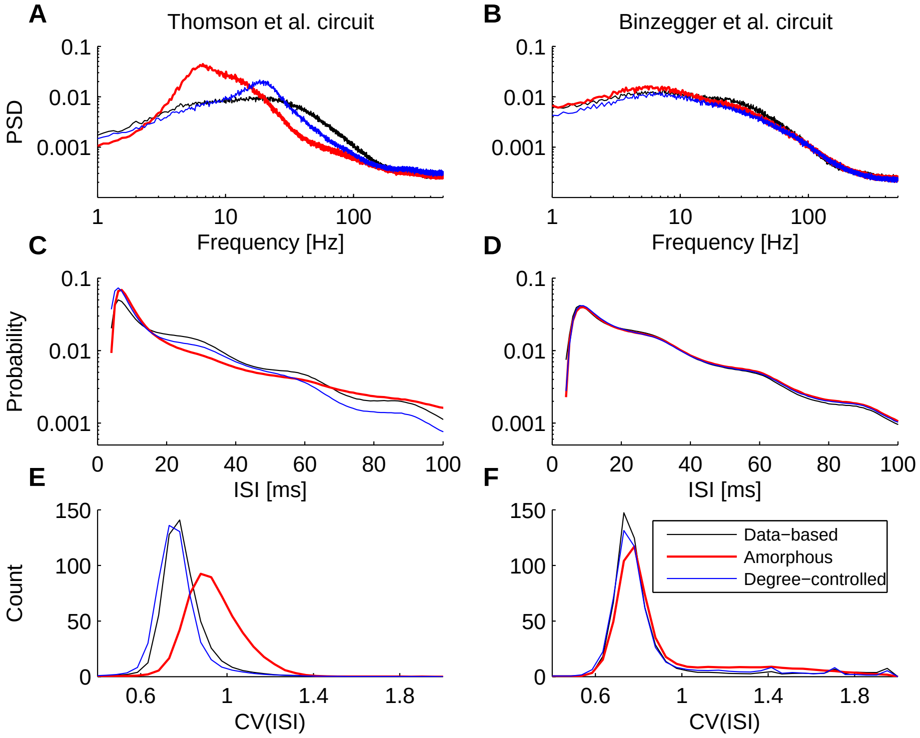

Figure 2.7: Statistical properties of the circuit dynamics of Thomson et al. circuits (left panels) and Binzegger et al. circuits (right panels). The statistical properties of each of the two data-based microcircuits (black lines) are compared to the properties of corresponding control circuits, i.e. amorphous circuits (red lines) and degree-controlled circuits (blue lines). **A,B:** Power spectral densities (PSD) of the mean firing rate (bin size 1 ms). The majority of neurons has excess power at low frequencies. **C,D:** Interspike interval (ISI) distributions for different neurons. All ISI distributions are well described by exponential distributions (corresponding to straight lines in semi-log plots) predicted by Poisson processes. **E,F:** Coefficient of variation of interspike intervals (CV(ISI)) for different neurons. Only for Thomson et al. circuits the corresponding amorphous circuits show different statistical properties. Error bars are negligible and omitted for clarity.

We found that Thomson et al. circuits and Binzegger et al. circuits have very similar dynamical properties. The mean firing rate of neurons is about 24 Hz and the power spectral densities (PSD) of the mean firing rates show an excess power at low frequencies between 5 and 50 Hz. Moreover for both circuit templates the interspike interval (ISI) distributions have an exponential tail as predicted by Poisson processes. A peak in the ISI distribution around 7 ms indicates the occurrence of bursts, that is a higher probability of observing two or more spikes within a short time window of a few ms than predicted by Poisson processes

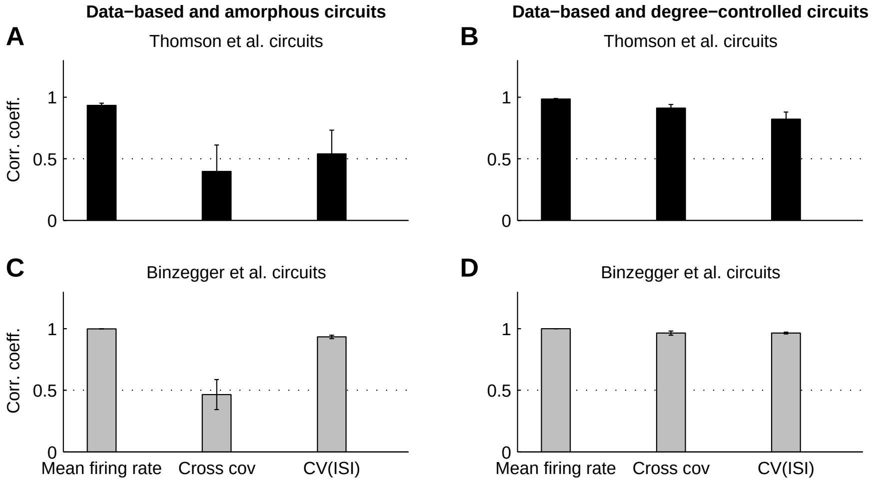

Figure 2.8: Correlation of statistical properties of data-based circuits and corresponding control circuits for 40 randomly chosen scaling parameters for the synaptic weights of the input and the recurrent connections. Shown are correlation coefficients of the mean firing rates, the mean cross covariances of the firing rates of single neurons and the mean coefficients of variation of interspike intervals (CV(ISI)) for **A:** Thomson et al. circuits and corresponding amorphous circuits, **B:** Thomson et al. circuits and corresponding degree-controlled circuits, **C:** Binzegger et al. circuits and corresponding amorphous circuits, and **D:** Binzegger et al. circuits and corresponding degree-controlled circuits. Only degree-controlled circuits preserve the mean cross covariance and the CV(ISI) of data-based circuits. Results shown are bias corrected bootstrap estimates.

with identical mean firing rates. Furthermore the distribution of the coefficient of variation of interspike intervals (CV(ISI)) of single neurons peaks at 0.72 for both circuit templates. However, the mean CV(ISI) differs significantly with a value of 0.78 (SD of 0.008) and 0.99 (SD of 0.009) for Thomson et al. circuits and Binzegger et al. circuits, respectively. As a measure of synchronous spiking activity we calculated the average cross covariance of firing rates (CC) within time bins of 1 ms for pairs of neurons (cross correlogram at time lag 0). The CC for Thomson et al. circuits and Binzegger et al. circuits is 0.035 and 0.066, respectively. Therefore, Binzegger et al. circuits operate in a more synchronous firing regime than Thomson et al. circuits.

2.3.4 Relationship between the graph structure of neural circuits and their dynamical properties

We related the dynamical properties of the data-based microcircuits to their connectivity structure by comparing the statistical properties of their circuit dynamics with the properties of corresponding control circuits (see Sec. 2.2.2).

Fig. 2.7 shows that degree-controlled circuits (blue lines) have similar dynamical properties when compared to corresponding data-based circuits (black lines). Only the power spectral density of degree-controlled circuits generated from Thomson et al. circuits shows a small deviation with a peak around 20 Hz. Additionally, for degree-controlled circuits generated from Binzegger et al. circuits the cross covariance of the firing rates of single neurons drops to a value of 0.05 (not shown).

In contrast, amorphous circuits have different statistical properties than data-based circuits. In particular, amorphous circuits generated from Thomson et al. circuits have a different power spectral density with a shift in power from high to low frequencies (see Fig. 2.7A). Furthermore the distribution of CV(ISI) differs significantly for amorphous circuits generated from Thomson et al. circuits with a higher mean CV(ISI) of 0.93 (Fig. 2.7E). For amorphous circuits generated from Binzegger et al. circuits the change in the CV(ISI) distribution is less pronounced with a significant drop in peak count by 7%. The mean cross covariance of the firing rates of single neurons in amorphous circuits drops significantly to a value of 0.033 and 0.045 for Thomson et al. circuits and Binzegger et al. circuits, respectively.

In order to verify that these results are in fact related to the connectivity structure and not an artifact of the specific set of scaling parameters that we had chosen for the synaptic weights of the input and the recurrent connections (see Methods) we repeated the analysis for 40 randomly chosen scaling parameters (S_{RW}, S_{I1}, and S_{I2}) that were drawn uniformly from the interval $[0.2, 2]$ times the standard values. Subsequently we correlated the 40 values for each of three statistical properties for pairs of circuit types to quantify their difference in circuit dynamics. Fig. 2.8A illustrates the results for Thomson et al. circuits and corresponding amorphous circuits. Amorphous circuits preserve the mean firing rate of Thomson et al. circuits but change the cross covariance of the firing rates of single neurons and the coefficient of variation of ISI. In contrast, degree-controlled circuits largely preserve all three statistical properties of Thomson et al. circuits as indicated in Fig. 2.8B. Similar results are obtained for Binzegger et al. circuits (Fig. 2.8C and D).

2.3.5 Computational properties

We compared the computational properties of Thomson et al. circuits (as investigated in detail in (Häusler and Maass, 2007)) with the corresponding results for Binzegger et al. circuits. For this purpose we analyzed to what extent these cortical microcircuit templates support computations on information contained in generic spike inputs into layer 4 and layers 2/3, and how well they make results of these computations accessible to (hypothetical) projection neurons in layers 2/3 and layer 5. These projection neurons were modeled as linear readout neurons that were trained in a supervised manner to perform a variety of information processing tasks that are likely to be related to actual computations of cortical microcircuits.

A comparison of the performance of Thomson et al. circuits (black bars) and Binzegger et al. circuits (gray bars) for these information processing tasks is shown in Fig. 2.9. The performance of trained readout for test inputs (which are generated from the same distribution as the training examples, but not shown during training) was measured for all binary classification tasks by the kappa coefficient, which ranges over [-1,1], and assumes a value ≥ 0 if the resulting classification of test examples makes fewer errors than random guessing. [14] For tasks that require an analog output value the performance of the trained readouts was measured on test examples by its correlation coefficient with the analog target output.

It turns out that the mean performance (averaged over all tasks and readouts) for both data-based circuit models is the same, i.e. 0.50, but the performance for individual tasks differs significantly [15]. In particular for tasks involving memory or non-linear computations the performance changes heterogeneously and does not depend on the readout type or task in an obvious manner. Binzegger et al. circuits perform significantly better for the non-linear spike template classification task that involves information from both input streams during the last 30 ms (XOR task), although they have only a similar performance as Thomson et al. circuits for the classification of individual spike templates from this time segment (on which the non-linear classification is based). This suggests that Binzegger et al. circuits better support non-linear fusion of recent information from both input streams. Also the classification of earlier spike patterns in input stream 2 ($\mathrm{tcl}_2(t - \Delta t)$) is significantly better supported by these circuits. A similar effect can be observed for layer 5 readout neurons performing computations based on the firing rates of both input streams. The performance for the total task including the linear component of the computation doesn't differ significantly, whereas the performance on the purely non-linear component is significantly better for Binzegger et al. circuits. [16]

One of the main structural features of Binzegger et al. circuits is the additional feedback loop from excitatory neurons in layer 6 to excitatory neurons

---

[14] The kappa coefficient measures the percentage of agreement between two classes expected beyond that of chance and is defined as (P_o - P_c)/(1 - P_c), where P_o is the observed agreement and P_c is the chance agreement. Thus for classification into 2 equally often occurring classes one has $P_C = 0.5$.

[15] The performance difference for each classification tasks stays qualitatively the same for Poisson inputs at different rates (i.e. 10 Hz and 30 Hz), although the precise value of the difference and its dependence on the input firing rates differs for individual tasks and readouts.

[16] This non-linear component of the target functions $r1/r2$ and $(r1-r2)^2$ resulted by subtracting from these functions an (for the considered distribution of input firing rates $r1$ and $r2$) optimally fitted linear function.

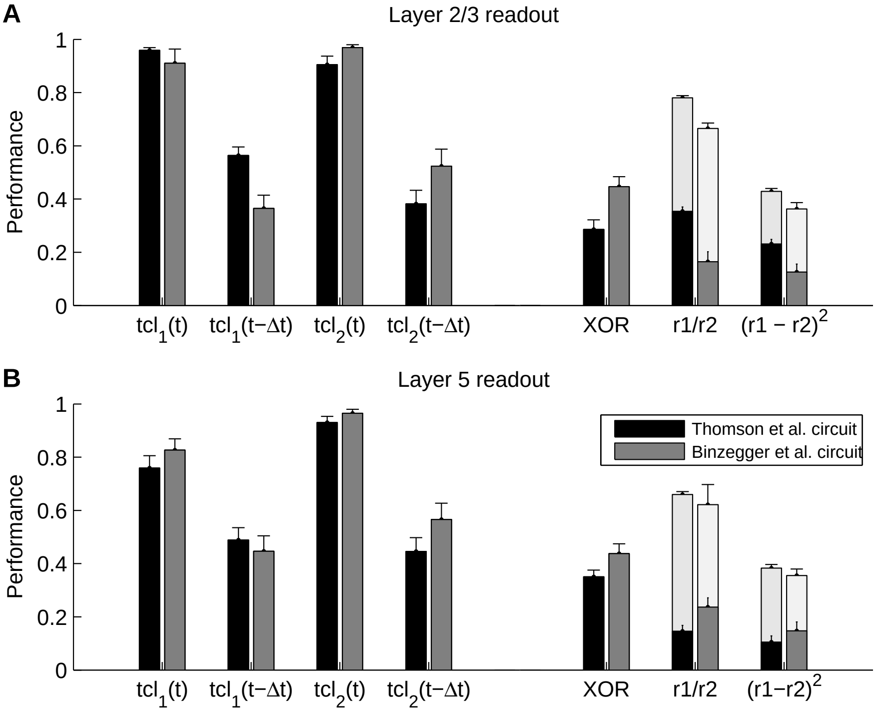

Figure 2.9: Performance of trained linear readout neurons in layers 2/3 and layer 5 (see Methods) for various classification tasks on spike patterns and computations performed on the rates of the two input streams for Thomson et al. circuits (black bars) and for Binzegger et al. circuits (gray bars). $\mathrm{tcl}_{1/2}$(t) denotes retroactive classification of noisy spike patterns in input streams 1 or 2 that were injected during the preceding time interval [t-30ms,t] into two classes according to the template from which each spike pattern had been generated (see Methods for details). $\mathrm{tcl}_{1/2}(t-\Delta t)$ refers to the more difficult task to classify at time t the spike pattern before the last one that had been injected during the time interval $[t-60\ ms, t-30\ ms]$. For XOR classification the task is to compute at time $t = 450$ ms the XOR of the template labels (0 or 1) of both input streams injected during the preceding time segment [420 ms, 450 ms]. On the right hand side the performance results for real-time computations on the time-varying firing rates $r_1(t)$ of input stream 1 and $r_2(t)$ of input stream 2 (both consisting of 4 Poisson spike trains with independently varying firing rates in the two input streams). The light bars show performance results for the two target functions $r_1(t)/r_2(t)$ and $(r_1(t) - r_2(t))^2$, and the bold bars for the performance on the nonlinear components of these real-time computations at any time t (on the actual firing rates in both input streams during the last 30 ms). The average computational performance for both types of cortical microcircuits is similar (with a value of 0.50) but the performance for specific tasks differs significantly.

in layer 4. This feedback loop turns the cortical feed-forward pathway into an intracortical closed-loop system. Remarkably, removal of this synaptic feedback doesn't significantly change the average readout performance (averaged over tasks and readouts), but changes the performance for specific tasks. Binzegger et al. circuits without feedback perform significantly better for the two template classification tasks, i.e. classification of the previously injected spike template in input stream 1 ($\mathrm{tcl}_1(t)$) for layer 2/3 readout neurons and classification of the spike template injected before the last one in input stream 1 ($\mathrm{tcl}_1(t - \Delta t)$) for layer 5 readout neurons. On the other hand the feedback loop improves the average performance for rate tasks significantly by 25% (not shown).

2.3.6 Relationship between structure, dynamics, and computational properties of the circuits

In (Häusler and Maass, 2007) it was shown that the data-based laminae-specific cortical microcircuit model introduced by (Thomson et al., 2002) exhibits specific computational advantages compared to various control circuits that have the same components and the same global statistics of neurons and synaptic connections, but are missing the laminae-specific connectivity pattern. In particular it was shown that degree-controlled circuits have similar average computational performance (averaged over all tasks and readouts) when compared to Thomson et al. circuits.

Here we show that the procedure of generating degree-controlled circuits, which just leaves the distribution of degrees of excitatory and inhibitory neurons in the circuit intact (as well as their roles as input-receiving node or readout-node), but randomizes their interconnectivity, does not change the computational performance of both data-based microcircuit templates in any significant manner: For Thomson et al. circuits it causes a drop by 1.6% in the average performance for the 7 computational tasks, and for Binzegger et al. circuits a drop by 1.1%. Therefore the degree distribution not only preserves statistical properties of the circuit dynamics, but also computational properties for both microcircuit templates.

On the other hand it was shown in (Häusler and Maass, 2007) that the average performance of amorphous circuits (generated from Thomson et al. circuits), which no longer have the motif distribution of Thomson et al. circuits, drops by 25%. In contrast, for Binzegger et al. circuits the procedure of generating amorphous control circuits, which changes also for these circuits the distribution of motifs, causes a small but significant average performance improvement by 5.0%.

In order to verify that these results are general properties of the connectivity structure, and do not depend on the specific choice of the three scaling parameters S_{I1}, S_{I2} and S_{RW} for the synaptic weights of the input and the recurrent connections we repeated the analysis for 40 randomly chosen values of these scaling parameters. We obtained similar results as for our standard values of these scaling parameters: The average performance of degree-controlled circuits doesn't change significantly compared to data-based circuits with an average performance improvement by 2.6% and 0.8% for degree-controlled circuits generated from Thomson et al. circuits and Binzegger et al. circuits, respectively. In contrast, the average performance of amorphous circuits generated from Thomson et al. circuits drops significantly by 26.7% and the performance of amorphous circuits generated from Binzegger et al. circuits increases significantly by 6.0%.

That scrambling of the connectivity structure has a different effect on the average computational performance of both data-based cortical microcircuits models may be attributed to three difference in the connectivity structure of their amorphous circuits, i.e. 1) different fractions of excitatory and inhibitory synapses for both microcircuit templates, 2) a different distribution of synaptic weights for Binzegger et al. circuits (e.g. differences in synaptic strengths due to additional synaptic connections from and to layer 6 and a rescaled average synaptic weight to achieve similar firing rates for both data-based microcircuit templates) and 3) a different set of neurons within the microcircuit that is connected to the readout neurons. Note that the strengths of synaptic connection to and from layer 6 are unknown and were set to average values. Presumably a change in the strength of these connections has an effect on the computational performance of Binzegger et al. circuits.

The results for the two different microcircuit templates leave open the possibility, that specific motif distribution enhance the computational performance. But they certainly do not verify such conjecture. To do that, one would need to be able to construct circuits with many different motif distributions. The procedure of generating degree-controlled circuits destroys the motif distribution of Binzegger et al. circuits, see Fig. 2.4E, 2.5C, but leaves the motif distribution of Thomson et al. circuits largely unchanged (Fig. 2.4C, 2.5B). It also changes the number of reciprocal connections, because the magnitude of the Z score of the second $M = 2$ motif drops consistently by 56% and 38% for Thomson et al. circuits and Binzegger et al. circuits, respectively. These results show that there are cases where the motif distribution changes significantly, but the computational performance remains largely the same.

In order to investigate whether the computational properties can be in fact related to specific dynamical properties that are preserved by degree-controlled circuits we correlated three statistical measures for circuit dynamics with the average computational performance (averaged over all tasks and readouts) for 40 randomly chosen values of the scaling parameters for the synaptic weights of the input and the recurrent connections. Fig. 2.10 shows that the mean coefficient of variation of ISIs and the mean cross covariance of the firing rate of single neurons correlate with the average computational performance for Thomson et al. circuits

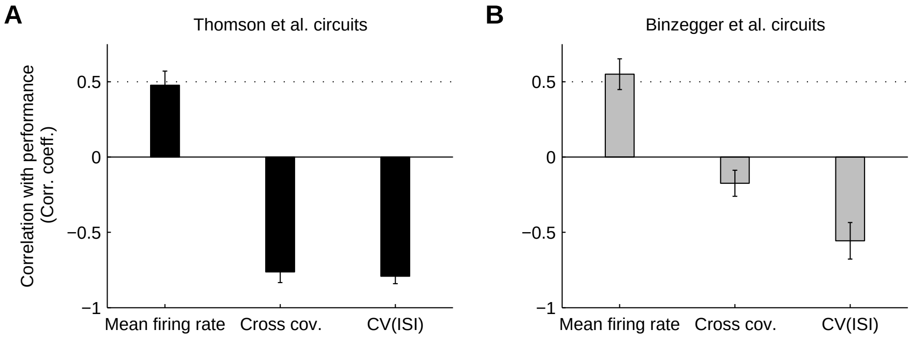

Figure 2.10: Correlation of the mean computational performance (averaged over all tasks and readouts in Fig. 2.9) and three statistical properties of the circuit dynamics for 40 randomly chosen scaling parameters for the synaptic weights of the input and the recurrent connections. **A:** Results for Thomson et al. circuits. **B:** Results for Binzegger et al. circuits. These results suggest that the best performance is achieved in the regular asynchronous firing regime. Results shown are bias corrected bootstrap estimates.

and Binzegger et al. circuits (although with a somewhat smaller value for the correlation coefficient between performance and mean cross covariance of the firing rates). Therefore, for the set of computational tasks considered in this article the best performance is achieved in the regular asynchronous firing regime with small values for the mean cross covariance and the mean coefficient of variation of ISIs. For both data-based microcircuit models the mean firing is positively correlated with the average performance.

2.4 Discussion

We found that (i) the microcircuit template by (Binzegger et al., 2004) and the cortical microcircuit template by (Thomson et al., 2002) have significant small-world properties, but quite different motif distributions ; (ii) both data-based microcircuits have similar circuit dynamics and average computational capabilities but different computational properties for individual tasks ; and (iii) the degree distribution is the aspect of the connectivity structure of both data-based microcircuit templates that is responsible for their dynamical properties and their computational properties.

(i) Similar global structural properties but different motif distributions

Thomson et al. circuits and Binzegger et al. circuits have similar small-world properties and their degree-distributions share characteristic main features. For both circuit templates the average degree was highest for excitatory neurons in layer 2/3, suggesting the importance of hubs in a layer where both cortical input streams merge. Furthermore the average degree of neurons in layer 2/3, layer 4 and layer 5 correlates for both cortical microcircuit templates (correlation coefficient of 0.74).

The somewhat more amorphous-like connectivity pattern of Binzegger et al. circuits may be attributed to a few diffuse synaptic connections that are emerging at the boundaries between two adjacent layers. These might be the results of slightly shifted layer boundaries for several 3-dimensional cell reconstructions obtained at different locations in the primary visual cortex. Dendritic and axonal arbors of two cells, which are non-overlapping in case of precise layer boundaries, could in case of shifted layer-boundaries account for additional diffuse synaptic connectivity.

The most distinct structural difference between the two microcircuit templates is expressed in their motif distribution (Fig. 2.4 and 2.6).[17]

The number of reciprocal connections for Thomson et al. circuits is smaller than for the corresponding amorphous circuits (Fig. 2.4C), whereas for the Binzegger et al. circuits it is twice as large as for amorphous circuits (Fig. 2.4D). Furthermore only for Binzegger et al. circuits highly connected motifs consisting of 3 nodes are over-represented (Fig. 2.4E), whereas the motif distribution for Thomson et al. circuits (Fig. 2.4C) can be primarily attributed to their degree-distribution. Scrambling the connectivity structure while leaving the degrees of nodes invariant does not change the motif distribution for Thomson et al. circuits (see Fig. 2.5 and Fig. 2.6) but for Binzegger et al. circuits.

Remarkably, the motif distribution of Thomson et al. circuits matches to some degree the motif distribution of a superfamily of biological information-processing networks reported in (Milo et al., 2004). Their connectivity patterns show highly over-represented three node motifs 5, 8, and 11 and are lacking motifs 1, 3, 4, and 6 when compared to random networks. Analogous, in Thomson et al. circuits motifs 5 and 11 occur with a frequency significantly higher than in random networks, whereas motif 4 is under-represented. The feedforward motif 5 has been theoretically and experimentally linked to signal-processing tasks such as persistence detection and pulse generation (see Alon, 2006). However, the frequent occurrence of motif 3 is atypical for this superfamily.

It should be noted that the different results for Thomson et al. circuits and Binzegger et al. circuits are not contradictory but rather point out structural differences between potential synaptic connectivity specified by Binzegger et al. circuits and functional synaptic connectivity specified by Thomson et al. circuits.

---

[17] The motif distributions for both data-based cortical microcircuits differ significantly from the motif distributions of corresponding amorphous circuits that were generated from the two data-based cortical microcircuits by scrambling their connectivity structure, but these differences are not consistent for both data-based cortical microcircuits.

In principle Thomson et al. circuits can be considered as sub-graphs of Binzegger et al. circuits which were shaped by synaptic plasticity according to their specific functional role. Furthermore the number of synapses in a microcircuit has an effect on the local and global structural properties. We introduced for each cortical microcircuit template to some extent arbitrarily a global constant factor that scales all synaptic connection probabilities to match the average total number of synaptic connections for Thomson et al. circuits and Binzegger et al. circuits.

(ii) Similar circuit dynamics and average computational capabilities but different computational properties

Both microcircuit templates have similar circuit dynamics (Fig. 2.7). The power spectral densities show an excess power at low frequencies between 5 and 50 Hz, which is consistent with experimental results showing peaks in the 20-60 Hz range for cells in area MT in behaving Monkey (Bair et al., 1994). The interspike interval distributions show no evidence for power-law behavior but have exponential tails as predicted by Poisson processes. Neurons slightly tend to burst with longer periods of low firing activity resembling experimental results obtained in cat cerebral cortex during slow wave sleep (Destexhe et al., 1999; Steriade, 2003; Bedard et al., 2006). The most frequently occurring value of the coefficient of variation of interspike intervals for neurons in the circuits that we considered was for both microcircuit templates 0.72. This value is lower than the reported value of about 1 for MT cells and V1 cells in behaving monkey (Softky and Koch, 1992, 1993) in response to bars and textured stimuli. This might be attributed to a missing diversity in neuron types (e.g. bursting and stuttering cells), synapse types (e.g. NMDA receptors), and a missing calcium dynamics that introduces processes on many different time scales.

It turned out that the average information processing capabilities of Thomson et al. circuits and Binzegger et al. circuits are similar but differ for specific tasks (Fig. 2.9). In particular for the chosen set of seven computational tasks Binzegger et al. circuits support slightly superior non-linear fusion of information contained in both input streams.

(iii) The degree distribution determines the dynamical properties and the computational properties

The degree distribution of neurons largely determines the circuit dynamics for both microcircuit templates. Degree-controlled circuits largely preserve the mean cross covariances of the firing rates of neurons and the mean coefficients of variation of interspike intervals of data-based circuits independent of the dynamic regime controlled by the scaling parameters for the synaptic weights of the input and the recurrent connections (Fig. 2.8).

When we changed the connectivity structure of these two microcircuit models in order to see which aspects of it are essential for the computational performance, we found in agreement with (Häusler and Maass, 2007) that the distribution of degrees of nodes is also a key factor for their computational performance. Scrambling the connectivity structure but leaving the degree distribution of neurons invariant did not result in a change in the average computational performance (see section 2.3.6).

These results suggest that the degree-controlled scrambling of the connectivity structure preserves specific statistical properties of the circuit dynamics that are crucial for information processing and correlate with the average computational performance (Fig. 2.10). For the chosen set of information processing tasks the best performance is achieved in the regular asynchronous firing regime.

It is interesting to relate these results to theoretical results obtained for much simpler models. (Ganguli et al., 2008) showed for networks consisting of linear neurons that the graph structure has an impact on the computational properties, more precisely the memory capacity of the network.

Furthermore (Schrauwen et al., 2009) have shown by means of mean-field analysis that the performance of randomly connected recurrent networks built from neurons whose output can only take one of two values (binary output) depends strongly on the network connectivity structure and is related to the in-degree of neurons in addition to the distribution of synaptic weights.

That scrambling of the connectivity structure can have different effects on the average computational performance can be verified by amorphous control circuits. The average performance of amorphous circuits generated from Binzegger et al. circuits is improved, whereas the performance for amorphous circuits generated from Thomson et. al circuits decreases. This suggests that only Thomson et al. circuits are optimized for a specific set of computations, whereas Binzegger et al. circuits represent potential synaptic connections that provide the possibility of implementing various different sets of computations.

Suggestions for further experimental and simulation based research

The results of this article point to three directions for further research. First, further work is needed to provide more reliable microcircuit models. The two data-based cortical microcircuit templates were estimated with different experimental techniques, i.e. dual intracellular recordings *in vitro* and three-dimensional cell reconstructions, and it is up to now unclear how to relate them to each other or if possible merge them to one unified cortical microcircuit template. Both microcircuit templates do not account for the lateral connectivity patterns of the neocortex. The Thomson et al. microcircuit template was estimated from neurons in slices within a maximum horizontal distance of 100 μm, whereas the Binzegger et al. microcircuit template was obtained by averaging statistical data of whole cell reconstructions discarding information about horizontal locations of neurons. Current work in progress (Potjans and Diesmann, 2008) suggests that the assumption of Gaussian lateral connectivity patterns presumably explains to some extent the discrepancies between both microcircuit templates. Merging both datasets could potentially result in a unified microcircuit template with computational properties and circuit dynamics in between those obtained for the two microcircuit templates analyzed in this study.

Secondly, a thorough empirical analysis of the distribution of motifs (especially for more than two nodes) in cortical microcircuits is needed. So far we can only analyze the distribution of motifs that is induced by data connection probabilities for any two neurons A and B. Preliminary data suggest however, that the probability of a synaptic connection from B to A depends on the presence of a synaptic connection from A to B. For instance whole-cell recordings of layer 5 pyramidal neurons of somatosensory, visual and prefrontal areas have shown that reciprocal connections are ≥ 3 times more likely than in random networks (Markram, 1997; Holmgren et al., 2003; Song et al., 2005; Wang et al., 2006). Similarly for 3 neurons A, B, C (of specific types) the probability of a synaptic connection from C to A is likely to depend on the presence or absence of other synaptic connections between A, B, C. This was confirmed by triple and quadruple whole-cell recordings of layer 5 pyramidal neurons in the rat visual cortex that showed that motifs consisting of many edges are over-represented when compared to random networks (Song et al., 2005). Likewise the probability that a layer 2/3 pyramidal neuron in rat somatosensory cortex makes a synaptic connection with two layer 5 neurons is fourfold higher compared with random connectivity if the layer 5 neurons are synaptically connected (motif 5 and 8) (Kampa et al., 2006). In contrast the probability that a layer 5 pyramidal neuron receives input from two layer 2/3 pyramidal neurons is threefold higher compared with random networks if the layer 2/3 pyramidal neurons are not connected (motif 1). Reliable data for such conditioned connection probabilities are needed not only for all possible types of neurons A, B, C, but also for all possible laminar locations of these three neurons.

Finally, the microcircuit templates should be tested on a larger variety of computational tasks. Our results on the computational performance of the two microcircuit templates and several variations of them depend on a somewhat arbitrary choice of 7 concrete computational tasks, and on the decision to only specialize the synaptic weights of readout neurons for a specific computational task, while having the weights within the circuit chosen from data-based probability distributions(hence not specialized for a particular computational task). Furthermore the results presumably depend also on the specific choice of distributions for input spike trains from external sources. In particular, geniculate relay cells have a firing pattern that can vary between tonic and bursting (Sherman, 1996). Burst firing has been shown to make a strong contribution to the initial phasic part of visual responses to gratings (Guido et al., 1992) and flashed spots (Guido and Sherman, 1998).

Conclusions

This article has shown that the two available templates for cortical microcircuits have quite interesting structural, dynamical, and computational features. In particular we have shown that it is possible to relate the structure and the (conjectured) computational function of these two microcircuit templates. This positive result will hopefully stimulate further systematic experimental work on the anatomy and physiology of cortical microcircuits, that is needed in order to arrive at a definite understanding of the computational function of cortical microcircuits and their genetically encoded structural basis.

2.5 Acknowledgments

This chapter is based on the paper *Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates*, which was written by Stefan Haeusler (SH), myself (KS) and Wolfgang Maass (WM). The experiments and analysis were performed by SH and KS, and the paper was written by SH and KS with additional input from WM.

Chapter 3

Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity

Contents

3.1 Introduction 46
3.2 Material and Methods 47
3.3 Results 60
3.4 Discussion 71
3.5 Acknowledgments 72

Recent developments in neuromorphic hardware engineering make mixed-signal VLSI neural network models promising candidates for neuroscientific research tools and massively parallel computing devices, especially for tasks which exhaust the computing power of software simulations. Still, like all analog hardware systems, neuromorphic models suffer from a constricted configurability and production-related fluctuations of device characteristics. Since also future systems, involving ever-smaller structures, will inevitably exhibit such inhomogeities on the unit level, self-regulation properties become a crucial requirement for their successful operation. By applying a cortically inspired self-adjusting network architecture, we show that the activity of generic spiking neural networks emulated on a neuromorphic hardware system can be kept within a biologically realistic firing regime and gain a remarkable robustness against transistor-level variations. As a first approach of this kind in engineering practice, the short-term synaptic depression and facilitation mechanisms implemented within an analog VLSI model of I&F neurons are functionally utilized for the purpose of network level stabilization. We present experimental data acquired both from the hardware model and from comparative software simulations which prove the applicability of the employed paradigm to neuromorphic VLSI devices.

3.1 Introduction

Software simulators have become an indispensable tool for investigating the dynamics of spiking neural networks (Brette et al., 2007). But when it comes to studying large-scale networks or long-time learning, their usage easily results in lengthy computing times (Morrison et al., 2005). A common solution, the distribution of a task to multiple CPUs, raises both required space and power consumption. Thus, the usage of neural networks in embedded systems remains complicated.

An alternative approach implements neuron and synapse models as physical entities in electronic circuitry (Mead, 1989). This technique provides a fast emulation at a maintainable wattage (Douglas et al., 1995a). Furthermore, as all units inherently evolve in parallel, the speed of computation is widely independent of the network size. Several groups have made significant progress in this field during the last years (see for example Indiveri et al., 2006; Merolla and Boahen, 2006; Vogelstein et al., 2007; Schemmel et al., 2007, 2008; Mitra et al., 2009). The successful application of such neuromorphic hardware in neuroscientific modeling, robotics and novel data processing systems will essentially depend on the achievement of a high spatial integration density of neurons and synapses. As a consequence of ever-smaller integrated circuits, analog neuromorphic VLSI devices inevitably suffer from imperfections of their components due to variations in the productions process (Dally and Poulton, 1998). The impact of such imperfections can reach from parameter inaccuracies up to serious malfunctioning of individual units. In conclusion, the particular, selected emulation device might distort the network behavior.

For that reason, designers of neuromorphic hardware often include auxiliary parameters which allow to readjust the characteristics of many components. But since such calibration abilities require additional circuitry, their possible extent of use usually has to be limited to parameters that are crucial for the operation. Hence, further concepts are needed in order to compensate the influence of hardware variations on network dynamics. Besides increasing the accuracy of unit parameters like threshold voltages or synaptic time constants, a possible solution is to take advantage of self-regulating effects in the dynamics of neural networks. While individual units might lack adequate precision, populations of properly interconnected neurons can still feature a faultless performance.

Long-term synaptic potentiation and depression (Morrison et al., 2008) might be effective mechanisms to tailor neural dynamics to the properties of the respective hardware substrate. Still, such persistent changes of synaptic efficacy can drastically reshape the connectivity of a network. In contrast, short-term synaptic plasticity (Zucker and Regehr, 2002) alters synaptic strength transiently. As the effect fades after some hundred milliseconds, the network topology is preserved.

We show that short-term synaptic plasticity enables neural networks, that are emulated on a neuromorphic hardware system, to reliably adjust their activity to a moderate level. The achievement of such a *substrate on a network level* is an important step towards the establishment of neuromorphic hardware as a valuable scientific modeling tool as well as its application as a novel type of adaptive and highly parallel computing device.

For this purpose, we examine a generic network architecture as proposed and studied by (Sussillo et al., 2007), which was proven to feature self-adjustment capabilities. As such networks only consist of randomly connected excitatory and inhibitory neurons and exhibit little specialized structures, they can be found in various cortical network models. In other words, properties of this architecture are likely to be valid in a variety of experiments.

Still, the results of (Sussillo et al., 2007) not necessarily hold for neuromorphic hardware devices: The referred work addressed networks of 5000 neurons. As the employed prototype hardware system (Schemmel et al., 2006, 2007) only supports some hundred neurons, it remained unclear whether the architecture is suitable for smaller networks, too. Furthermore, the applicability to the specific inhomogeneities of the hardware substrate have not been investigated before. We proof that even small networks are capable of leveling their activity. This suggests that the studied architecture can enhance the usability of upcoming neuromorphic hardware systems, which will comprise millions of synapses.

The successful implementation of short-term synaptic plasticity into neuromorphic hardware has been achieved by several work groups, see e.g. (Boegershausen et al., 2003) or (Bartolozzi and Indiveri, 2007). Nevertheless, this work presents the first functional application of this feature within emulated networks. It is noteworthy, that the biological interpretation of the used hardware parameters is in accord with physiological data as measured by (Markram et al., 1998) and (Gupta et al., 2000).

Since the utilized system is in a prototype state of development, the emulations have been prepared and counter-checked using the well-established software simulator PCSIM (Pecevski et al., 2009). In addition, this tool allowed a decent analysis of network dynamics because the internal states of all neurons and synapses can be accessed and monitored continuously.

3.2 Material and Methods

The applied setup and workflow involve an iterative process using two complementary simulation back-ends: Within the FACETS research project (FACETS, 2009), the FACETS Stage 1 Hardware system (Schemmel et al., 2006, 2007) and the software simulator PCSIM (Pecevski et al., 2009) are being developed.

First, it had to be investigated whether the employed network architecture exhibits its self-adjustment ability in small networks fitting onto the current prototype hardware system. For this purpose, simulations have been set up on PCSIM which only roughly respected details of the hardware characteristics, but comprised a sufficiently small number of neurons and synapses. Since the trial yielded promising results, the simulations were transferred to the FACETS Hardware. At this stage the setup had to be readjusted in order to meet all properties and limitations of the hardware substrate. Finally, the parameters used during the hardware emulations were transferred back to PCSIM in order to verify the results.

In Section 3.2.1 and Section 3.2.2 both back-ends are briefly described. Section 3.2.3 addresses the examined network architecture and the parameters applied. In Section 3.2.4, the experimental setup for both back-ends is presented.

3.2.1 The Utilized Hardware System

The present prototype FACETS Stage 1 Hardware system physically implements neuron and synapse models using analog circuitry (Schemmel et al., 2006, 2007). Beside the *analog neural network core* (the so-called *Spikey* chip) it consists of different (mostly digital) components that provide communication and power supply as well as a multi-layer software framework for configuration and readout (Grübl, 2007; Brüderle et al., 2009).

The Spikey chip is built using a standard 180 nm CMOS process on a 25 mm^2 die. Each chip holds 384 conductance-based leaky integrate-and-fire point neurons, which can be interconnected or externally stimulated via approximately 100,000 synapses whose conductance courses rise and decay exponentially in time. As all physical units inherently evolve both in parallel and time-continuously, experiments performed on the hardware are commonly referred to as *emulations*. The dimensioning of the utilized electronic components allows a highly accelerated operation compared to the biological archetype. Throughout this work, emulations were executed with a speedup factor of 10^5.

In order to identify voltages, currents and the time flow in the chip as parameters of the neuron model, all values need to be translated between the hardware domain and the biological domain. The configuration and readout of the system has been designed for an intuitive, biological description of experimental setups: The Python-based (Rossum, 2000) meta-language *PyNN* (Davison et al., 2008) provides a back-end independent modeling tool, for which a hardware-specific implementation is available (Brüderle et al., 2009). All hardware-specific configuration and data structures (including calibration and parameter mapping), which are encapsulated within low-level machine-oriented software structures, are addressed automatically via a *Python Hardware Abstraction Layer* (PyHAL).

Using this translation of biological values into hardware dimensions and vice versa which is performed by the PyHAL, all values given throughout this work reflect the biological interpretation domain.

3.2.1.1 Short-Term Synaptic Plasticity

All synapses of the FACETS Stage 1 Hardware support two types of synaptic plasticity (Schemmel et al., 2007). While a spike-timing dependent plasticity (STDP) mechanism (Bi and Poo, 1997; Song et al., 2000) is implemented in every synapse, short-term plasticity (STP) only depends on the spiking behavior of the pre-synaptic neuron. The corresponding circuitry is part of the so-called *synapse drivers* and, thus, STP-parameters are shared by all synaptic connections operated by the same driver. Each pre-synaptic neuron can project its action potentials (APs) to two different synapse drivers. Hence, two freely programmable STP-configurations are available per pre-synaptic neuron. The short-term plasticity mechanism implemented in the FACETS Stage 1 Hardware is inspired by (Markram et al., 1998). But while the latter model combines synaptic *facilitation* and *depression*, the hardware provides the two modes separately. Each synapse driver can either be run in facilitation or in depression mode or simply emulate static synapses without short-term dynamics. Despite this restriction, these short-term synapse dynamics support dynamic gain-control mechanisms as, e.g., reported in (Abbott et al., 1997).

In the Spikey chip, the conductance $g(t)$ of a synapse is composed of a discrete synaptic weight multiplier w_n, the base efficacy $w_0(t)$ of a synapse driver and the conductance course of the rising and falling edge $p(t)$:

$$g(t) = w_n \cdot w_0(t) \cdot p(t) =: w(t) \cdot p(t)$$

with $w_n \in \{0, 1, 2, \ldots, 15\}$. In this framework, STP alters the base efficacy $w_0(t)$ while the double-exponential conductance course of a single post-synaptic potential is modeled via $p(t) \in [0, 1]$. Whenever an action potential is provoked by the pre-synaptic neuron, $p(t)$ is triggered to run the conductance course. To simplify matters, the product $w_n \cdot w_0(t)$ often is combined to the synaptic weight $w(t)$ or just w in case of static synapses.

Both STP-modes, facilitation and depression, alter the synaptic weight in a similar manner using an *active partition* $I(t) \in [0, 1]$. The strength w_{stat} of a static synapse is changed to

$$\begin{aligned} w_{\mathrm{fac}}(t) &= w_{\mathrm{stat}} \cdot [1 + \lambda \cdot (I(t) - \beta)] \ , \\ w_{\mathrm{dep}}(t) &= w_{\mathrm{stat}} \cdot [1 - \lambda \cdot I(t)] \end{aligned} \tag{3.1}$$

in case of facilitation and depression, respectively. The parameters λ and β are freely configurable. For technical reasons, the change of synaptic weights by STP cannot be larger than the underlying static weight. Stronger modifications are truncated. Hence, $0 \leq w_{\mathrm{fac}\,/\,\mathrm{dep}} \leq 2 \cdot w_{\mathrm{stat}}$.

The active partition I obeys the following dynamics: Without any activity I decays exponentially with time constant τ_{STP}, while every AP processed increases

I by a fixed fraction C towards the maximum,

$$\frac{\mathrm{d}I}{\mathrm{d}t} = -\frac{I}{\tau_{\mathrm{STP}}} + C \cdot (1 - I) \cdot \delta\,(t - t_{\mathrm{AP}}) \quad .$$

For $C \in [0, 1]$, I is restricted to the interval mentioned above. Since the active partition affects the analog value $w_0(t)$, the STP-mechanism is not subject to the weight-discretization w_n of the synapse arrays but alters weights continuously.

Fig. 3.1 shows examples of the dynamics of the three STP-modes as measured on the FACETS Stage 1 Hardware. The applied parameters agree with those of the emulations presented throughout this work.

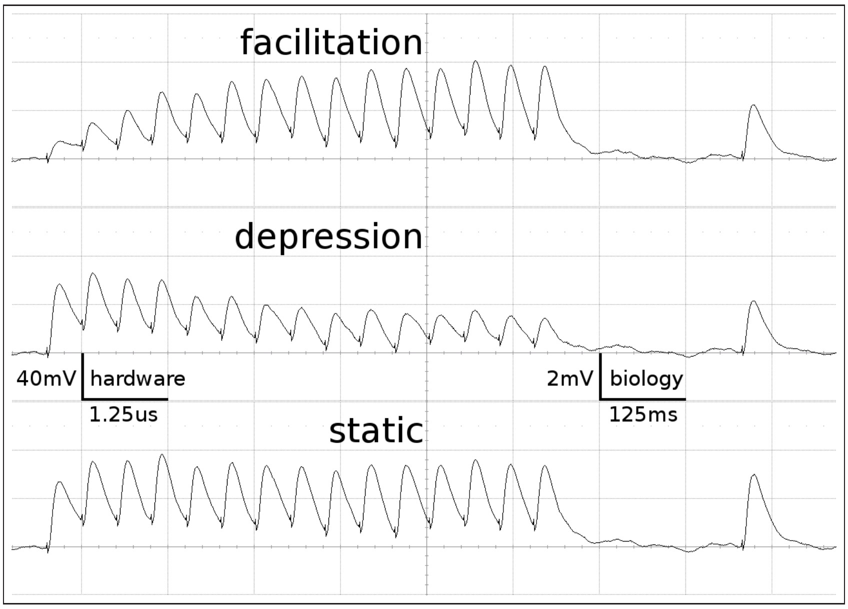

Figure 3.1: STP-mechanism of the FACETS Stage 1 Hardware.
A neuron is excited by an input neuron that spikes regularly at 20 Hz. 300 ms after the last regular spike a single spike is appended. Additionally, the neuron is stimulated with Poisson spike trains from further input neurons. The figure shows the membrane potential of the post-synaptic neuron, averaged over 500 experiment runs. As the Poisson background cancels out, the EPSPs provoked by the observed synapse are revealed. Time and voltage are given in both hardware values and their biological interpretation. The three traces represent different modes of the involved synapse driver. *Facilitation:* The plastic synapse grows in strength with every AP processed. After 300 ms without activity the active partition has partly decayed. *Depression:* High activity weakens the synapse. *Static:* The synapse keeps its weight fixed.

3.2.1.2 Hardware Constraints

Neurons and synapses are represented by physical entities in the chip. As similar units reveal slightly different properties due to the production process, each unit exhibits an individual discrepancy between the desired configuration and its actual behavior. Since all parameters are controlled by voltages and currents, which require additional circuitry within the limited die, many parameters and subcircuits are shared by multiple units. This results in narrowed parameter ranges and limitations on the network topology.

Beyond these intentional design-inherent fluctuations and restrictions, the current prototype system suffers from some malfunctions of different severity. These errors are mostly understood and will be fixed in future systems. In the following, the constraints which are relevant for the applied setup will be outlined. For detailed information the reader may refer to the respective literature given below.

Design-inherent constraints

- As described above, synaptic weights are discrete values $w = w_n \cdot w_0$ with $w_n \in \{0, 1, 2, \ldots, 15\}$ (Schemmel et al., 2006). Since biological weights are continuous values, they are mapped probabilistically to the two closest discrete hardware weights. Therefore, this constraint is assumed to have little impact on large, randomly connected networks.

- Each pre-synaptic neuron allocates two synapse drivers to provide both facilitating and depressing synapses. Since only 384 synapse drivers are available for the operation of recurrent connections, this restricts the maximum network size to $384/2 = 192$ neurons. After establishing the recurrent connections, only 64 independent input channels remain for excitatory and inhibitory external stimulation via Poisson spike trains. (see Bill, 2008, Chapter VI.3)

- Bottlenecks of the communication interface limit the maximum input bandwidth for external stimulation to approximately 12 Hz per channel when 64 channels are used for external stimulation with Poisson spike trains. Future revisions are planned to run at a speedup factor of 10^4 instead of 10^5, effectively increasing the input bandwidth by a factor of 10 from the biological point of view (see Grübl, 2007, Chapter 3.2.1 and Brüderle, 2009, Chapter 4.3.7).

Malfunctions

- The excitatory reversal potential was found to be unstable under high load. A frequent global activity of excitatory synapses has been shown to decrease EPSP amplitudes up to a factor of two. This load-dependent error cannot be counterbalanced by calibration or tuning the configuration and is considered crucial for the presented experimental setup (see Brüderle, 2009, Chapter 4.3.4).

- The current system suffers from a disproportionality between the falling-edge synaptic time constant $\tau_{\mathrm{syn}} \approx 30\,\mathrm{ms}$ and the membrane time constant $\tau_{\mathrm{mem}} \approx 5\,\mathrm{ms}$, i.e. a fast membrane and slow synapses. This was taken into consideration when applying external stimulation, as presented in Section 3.2.3.2 (see Kaplan et al., 2009 and Brüderle, 2009, Chapter 4.3.5).

- Insufficient precision of the neuron threshold comparator along with a limited reset conductance result in a rather wide spread of the neuron threshold and reset voltages V_{thresh} and V_{reset}. As both values are shared by multiple neurons, this effect can only be partially counterbalanced by calibration. The used calibration algorithms lead to $\sigma_{V_{\mathrm{thresh}}} \approx 3\,\mathrm{mV}$ and $\sigma_{V_{\mathrm{reset}}} \approx 8\,\mathrm{mV}$ (see Brüderle, 2009, Chapter 4.3.2 and Bill, 2008, Chapter IV.4).

- Insufficient dynamic ranges of control currents impede a reasonable configuration of the short-term plasticity parameters λ and β in Equation 3.1 without additional technical effort. The presented emulations make use of a workaround which allows a biologically realistic setup of the STP-parameters at the expense of further adjustability. The achieved configuration has been measured and is used throughout the software simulations, as well (see Bill, 2008, Chapter IV.5.4).

- An error in the spike event readout circuitry prevents a simultaneous recording of the entire network. Since only three neurons of the studied network architecture can be recorded per emulation cycle, every configuration was rerun $192/3 = 64$ times with different neurons recorded. Thus, all neurons have been taken into consideration in order to determine average firing rates. But since the data is obtained in different cycles, it is unclear to what extent network correlation and firing dynamics on a level of precise spike timing can be determined (see Müller, 2008, Chapter 4.2.2).

A remark on parameter precision. The majority of the parameter values used in the implemented neuron model are generated by complex interactions of hardware units, as transistors and capacitors. Each type of circuitry suffers from different variations due to the production process, and these fluctuations sum up to intricate discrepancies of the final parameters. For that reason, both shape and extent of the variances often cannot be calculated in advance. On the other hand, only few parameters of the neuron and synapse model can be observed directly. Exceptions are all kind of voltages, e.g. the membrane voltage or reversal potentials. The knowledge of all other parameters was obtained from indirect measurements by evaluating spike events and membrane voltage traces. The configuration given in Section 3.2.3.2 reflects the current state of knowledge. This means that some specifications – especially standard deviations of parameters – reflect estimations which are based on long-term experience with the device. But, compared to the above-described malfunctions of the prototype system, distortions arising from uncertainties in the configuration can be expected to be of minor importance.

3.2.2 The Parallel neural Circuit SIMulator (PCSIM)

All simulations were performed using the PCSIM simulation environment and were set up and controlled via the associated Python interface (Pecevski et al., 2009).

The neurons were modeled as leaky integrate and fire cells (LIF) with conductance based synapses. The dynamics of the membrane voltage $V(t)$ is defined by

$$\begin{aligned}
C_\mathrm{m}\frac{\mathrm{d}V(t)}{\mathrm{d}t} = & - & g_\mathrm{leak}\cdot(V(t)-V_\mathrm{rest}) \\
& - & \sum_{k=1}^{N_\mathrm{e}} g_{\mathrm{e},k}(t)\cdot(V(t)-E_\mathrm{e}) \\
& - & \sum_{k=1}^{N_\mathrm{i}} g_{\mathrm{i},k}(t)\cdot(V(t)-E_\mathrm{i}) \\
& + & I_\mathrm{noise}(t)\ ,
\end{aligned}$$

where C_m is the membrane capacity, g_leak is the leakage conductance, V_rest is the leakage reversal potential, and $g_{\mathrm{e},k}(t)$ and $g_{\mathrm{i},k}(t)$ are the synaptic conductances of the N_e excitatory and N_i inhibitory synapses with reversal potentials E_e and E_i, respectively. The white noise current $I_\mathrm{noise}(t)$ has zero mean and a standard deviation $\sigma_\mathrm{noise} = 5\,\mathrm{pA}$. It models analog noise of the hardware circuits.

The dynamics of the conductance $g(t)$ of a synapse is defined by

$$\frac{\mathrm{d}g(t)}{\mathrm{d}t} = -\frac{g(t)}{\tau_\mathrm{syn}} + w\cdot\delta(t-t_\mathrm{AP})\,,$$

where $g(t)$ is the synaptic conductance and w is the synaptic weight. The conductances decrease exponentially with time constant τ_syn and increase instantaneously by adding w to the running value of $g(t)$ whenever an action potential occurs in the pre-synaptic neuron at time t_AP. Modeling the exponentially rising edge of the conductance course of the FACETS Stage 1 Hardware synapses was considered negligible, as the respective time constant was set to an extremely small value for the hardware emulation.

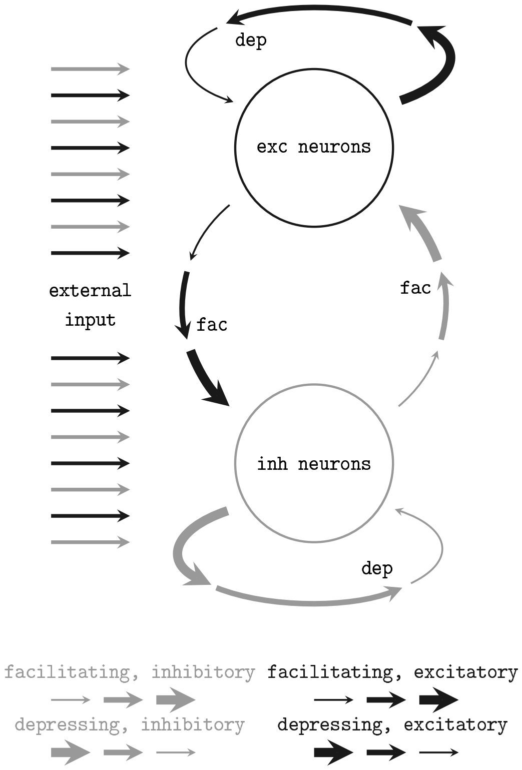

Figure 3.2: Schematic of the self-adjusting network architecture proposed in (Sussillo et al., 2007).
Depressing (`dep`) and facilitating (`fac`) recurrent synaptic connections level the network activity.

If we used static synapses the weight w of a synapse was constant over time. Whereas for simulations with dynamic synapses, the weight $w(t)$ of each synapse was modified according to the short-term synaptic plasticity rules described in Section 3.2.1.1.

The values of all parameters were drawn from random distributions with parameters as listed in **Table 3.1**.

3.2.3 Network Configuration

In the following, the examined network architecture is presented. Rather than customizing the configuration to the employed device, we aimed for a generic, back-end agnostic choice of parameters. Due to hardware limitations in the input bandwidth, a dedicated concept for external stimulation had to be developed.

Description	Name	Unit	Mean μ	σ/μ	π/μ	Comment
Network Architecture						
Number of exc neurons	N_{e}		144			
Number of inh neurons	N_{i}		48			
Conn prob from exc to exc neurons	p_{ee}		0.1			
Conn prob from exc to inh neurons	p_{ie}		0.2			
Conn prob from inh to exc neurons	p_{ei}		0.3			
Conn prob from inh to inh neurons	p_{ii}		0.6			
Neurons (Excitatory and Inhibitory)						
Membrane capacitance	C_{m}	nF	0.2	0	0	by definition
Leakage reversal potential	V_{rest}	mV	-63...-55			variable parameter
Firing threshold voltage	V_{thresh}	mV	-55.0	0.05	0.1	
Reset potential	V_{reset}	mV	-80.0	0.1	0.2	
Excitatory reversal potential	E_{e}	mV	0.0	0	0	unstable
Inhibitory reversal potential	E_{i}	mV	-80.0	0	0	
Leakage conductance	g_{leak}	nS	40.0	0.5	0.5	*)
Refractory period	τ_{ref}	ms	1.0	0.5	0.5	
Recurrent Synapses						
Weight of exc to exc synapses	w_{ee}	nS	1.03	0.6	0.7	*) values refer to
Weight of exc to inh synapses	w_{ie}	nS	0.52	0.6	0.7	*) static synapses
Weight of inh to exc synapses	w_{ei}	nS	3.10	0.6	0.7	*)
Weight of inh to inh synapses	w_{ii}	nS	1.55	0.6	0.7	*)
Cond time constant for all synapses	τ_{syn}	ms	30.0	0.25	0.5	
Conversion factor for facilitation			1.10			to match with static syns
Conversion factor for depression			1.65			at regular firing of 20 Hz
Strength of STP	λ		0.78	0.1	0.2	
Bias for facilitation	β		0.83	0.1	0.2	
STP decay time constant	τ_{STP}	ms	480	0.2	0.4	
Step per spike for facilitation	C_{fac}		0.27	0.1	0.2	
Step per spike for depression	C_{dep}		0.11	0.1	0.2	
External Stimulus: Poisson Spike Trains						
Number of exc external spike sources	$N_{\mathrm{ext,\,e}}$		32			
Number of inh external spike sources	$N_{\mathrm{ext,\,i}}$		32			
Number of exc inputs per neuron			4 - 6			uniform distribution
Number of inh inputs per neuron			4 - 6			uniform distribution
Firing rate per input spike train	ν_{inp}	Hz	11.8	0.2	0.2	*)
Weight of exc input synapses	$w_{\mathrm{inp,\,e}}$	nS	0.26...1.29	0.6	0.7	*) varied via W_{input} and
Weight of inh input synapses	$w_{\mathrm{inp,\,i}}$	nS	0.77...3.87	0.6	0.7	*) refer to $V_{\mathrm{rest}} = -60\,\mathrm{mV}$
Cond time constant for all synapses	τ_{syn}	ms	30.0	0.25	0.5	
Experiment						
Simulated time per exp run	T_{exp}	ms	4500			only $t \geq 1000\,\mathrm{ms}$ evaluated
Number of exp runs per param set	n_{run}		20			×64 in hardware with same network

Table 3.1: Full set of parameters. All values given in biological units. If not stated otherwise, values are drawn from a bound normal distribution with mean μ, standard deviation σ and bound π. Parameters marked by a *) have been spread for the hardware emulations by configuration.

3.2.3.1 Network Architecture

We applied a network architecture similar to the setup proposed and studied by (Sussillo et al., 2007) which was proven to feature self-adjustment capabilities. A schematic of the architecture is shown in Fig. 3.2. It employs the short-term plasticity mechanism presented above. Two populations of neurons – both similarly stimulated externally with Poisson spike trains – are randomly connected obeying simple probability distributions (see below). Connections within the populations are depressing, while bridging connections are facilitating. Thus, if excitatory network activity rises, further excitation is reduced while inhibitory activity is facilitated. Inversely, in case of a low average firing rate, the network sustains excitatory activity.

(Sussillo et al., 2007) studied the dynamics of this architecture for sparsely connected networks of 5000 neurons through extensive computer simulations of leaky integrate-and-fire neurons and mean field models. In particular, they examined how the network response depends on the mean value and the variance of a Gaussian distributed current injection. It was shown that such networks are capable of adjusting their activity to a moderate level of approximately 5 to 20 Hz over a wide range of stimulus parameters while preserving the ability to respond to changes in the external input.

3.2.3.2 Applied Parameters

With respect to the constraints described in Section 3.2.1.2, we set up recurrent networks comprising 192 conductance-based leaky integrate-and-fire point neurons, 144 (75%) of which were chosen to be excitatory, 48 (25%) to be inhibitory. Besides feedback from recurrent connections, each neuron was externally stimulated via excitatory and inhibitory Poisson spike sources. The setup of recurrent connections and external stimulation is described in detail below.

All parameters specifying the networks are listed in **Table 3.1**. Most values are modeled by a *bound normal distribution* which is defined by its mean μ, its standard deviation σ and a bound π: The random value x is drawn from a normal distribution $\mathcal{N}(\mu, \sigma^2)$. If x exceeds the bounds, it is redrawn from a uniform distribution within the bounds.

In case of hardware emulations, some of the deviations σ only reflect chip-inherent variations, i.e. fluctuations that remain when all units are intended to provide equal values. For other parameters – namely for all synaptic efficacies w, the leakage conductance g_{leak} and the input firing rate ν_{inp} – the major fraction of the deviations σ was intentionally applied by the experimenter. If present, the variations of hardware parameters are based on (Brüderle et al., 2009).

In case of software simulations, all inhomogeneities are treated as independent statistical variations. Especially, systematic effects, like the instability of the excitatory reversal potential E_e or the unbalanced sensitivity between the neuron populations (see Section 3.3.1), have not been modeled during the first simulation series.

Recurrent connections. Any two neurons are synaptically connected with probability $p_{\mathrm{post,\,pre}}$ and weight $w_{\mathrm{post,\,pre}}$. These values depend only on the populations the pre- and post-synaptic neurons are part of.

Synaptic weights always refer to the strength of static synapses. When a synapse features short-term plasticity, its weight is multiplicatively adjusted such that the strengths of static and dynamic synapses match at a constant regular pre-synaptic firing of 20 Hz for $t \to \infty$. This adjustment is necessary in order to enable dynamic synapses to be both stronger or weaker than static synapses according to their current activity.

Although the connection probabilities and synaptic weights used for the experiments do not rely on biological measurements or profound theoretical studies, they follow some handy rules. The mean values of the probability distributions are determined by three principles:

1. Every neuron has as many excitatory as inhibitory recurrent input synapses: $p_{\mathrm{post,\,e}} \cdot N_\mathrm{e} = p_{\mathrm{post,\,i}} \cdot N_\mathrm{i}$.
2. Inhibitory neurons receive twice as many recurrent synaptic inputs as excitatory neurons. This enables them to sense the state of the network on a more global scale: $p_{\mathrm{i,\,pre}} \cdot N_{\mathrm{pre}} = 2 \cdot p_{\mathrm{e,\,pre}} \cdot N_{\mathrm{pre}}$
3. Assuming a uniform global firing rate of 20 Hz and an average membrane potential of $V = -60\,\mathrm{mV}$, synaptic currents are well-balanced in the following terms:
    (a) For each neuron the excitatory and inhibitory currents have equal strength,
    (b) each excitatory neuron is exposed to as much synaptic current as each inhibitory neuron.

Formally, we examine the average current induced by a population *pre* to a single neuron of the population *post*:

$$I_{\mathrm{post,\,pre}} \propto p_{\mathrm{post,\,pre}} \cdot N_{\mathrm{pre}} \cdot w_{\mathrm{post,\,pre}} \cdot |E_{\mathrm{pre}} - V| \ .$$

Principle 3 demands that $I_{\mathrm{post,\,pre}}$ is equal for all tuples (post, pre) under the mentioned conditions. Given the sizes of the populations and the reversal potentials, the principles 1 and 2 determine all recurrent connection probabilities $p_{\mathrm{post,\,pre}}$ and weights $w_{\mathrm{post,\,pre}}$ except for two global multiplicative parameters: one scaling all recurrent connection probabilities, the other one all recurrent weights. While the ratios of all $p_{\mathrm{post,\,pre}}$ as well as the ratios of the $w_{\mathrm{post,\,pre}}$ are fixed, the scaling factors have been chosen such that the currents induced by recurrent synapses exceed those induced by external inputs in order to highlight the functioning of the applied architecture.

External stimulation. In order to investigate the modulation of activity by the network, external stimulation of different strength should be applied. One could think of varying the total incoming spike rate or the synaptic weights of excitation and inhibition. In order to achieve a biologically realistic setup, one should choose the parameters such that the stimulated neurons will reach a *high-conductance state* (see Destexhe et al., 2003 and Kaplan et al., 2009). Neglecting the influence of recurrent connections and membrane resets after spiking, the membrane would tune in to an average potential μ_V superposed by temporal fluctuations σ_V.

As mentioned above, the FACETS Stage 1 Hardware suffers from a small number of input channels if 2×192 synapse drivers are reserved for recurrent connections. At the same time, even resting neurons exhibit a very short membrane time constant of $\tau_{\mathrm{mem}} \approx 5\,\mathrm{ms}$. Due to these limitations, we needed to apply an alternative type of stimulation to approximate appropriate neuronal states:

Regarding the dynamics of a conductance-based leaky integrate-and-fire neuron, the conductance course towards any reversal potential can be split up into a time-independent average value and time-dependent fluctuations with vanishing mean. Then, the average conductances towards all reversal potentials can be combined to an effective resting potential and an effective membrane time constant (Shelley et al., 2002). In this framework, only the fluctuations remain to be modeled via external stimuli.

From this point of view, the hardware neurons appear to be in a high-conductance state with an average membrane potential $\mu_V = V_{\mathrm{rest}}$ without stimulation due to the short membrane time constant τ_{mem}. Ex post, the available input channels can be used to add fluctuations. The magnitude σ_V of the fluctuations is adjusted via the synaptic weights of the inputs.

Throughout all simulations and emulations, 32 of the 64 input channels were used for excitatory stimulation, the remaining 32 input channels for inhibitory stimulation. Each neuron was connected to 4 to 6 excitatory and 4 to 6 inhibitory inputs using static synapses. The number of inputs was randomly drawn from a uniform distribution for each neuron and reversal potential. The synaptic weights of the connections were drawn from bound normal distributions. The mean value of these distributions was chosen such that the average traction $w \cdot (E_{\mathrm{rev}} - \mu_V)$ was equal for excitatory and inhibitory synapses. The values listed in **Table 3.1** refer to $\mu_V = V_{\mathrm{rest}} = -60\,\mathrm{mV}$. In case of other resting potentials, the synaptic weights were properly adjusted to achieve an equal average current towards the reversal potentials: In case of excitatory inputs the weight was set to $w_{\mathrm{inp,\,e}} \cdot \left| \frac{E_{\mathrm{e}} - (-60\,\mathrm{mV})}{E_{\mathrm{e}} - V_{\mathrm{rest}}} \right|$.

Similarly, inhibitory input weights were adjusted to $w_{\mathrm{inp,\,i}} \cdot \left|\frac{E_\mathrm{i}-(-60\,\mathrm{mV})}{E_\mathrm{i}-V_\mathrm{rest}}\right|$.

Thus, neglecting the influence of recurrent connections and resets of the membrane after APs, the average input-induced membrane potential μ_V always equals V_rest. The magnitude of the fluctuations was controlled via a multiplicative *weight factor* W_input affecting all input synapses.

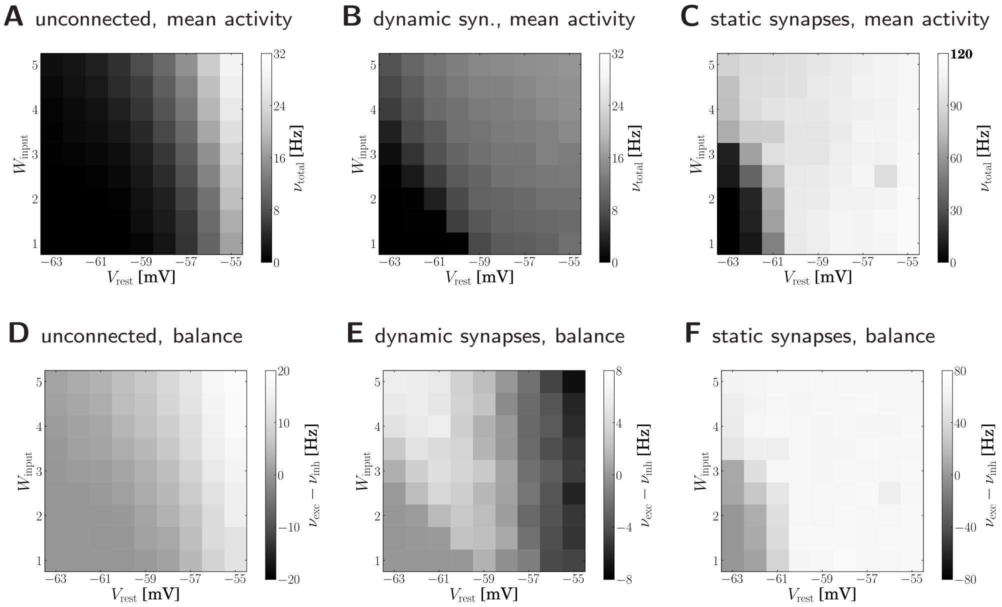

A unconnected, mean activity

B dynamic syn., mean activity

C static synapses, mean activity

D unconnected, balance

E dynamic synapses, balance

F static synapses, balance

Figure 3.4: Results of the emulations on the FACETS Stage 1 Hardware.
External stimulation of diverse strength is controlled via V_rest and W_input. For every tile, 20 randomly connected networks with new external stimulation were generated. The resulting average firing rates are illustrated by different shades of gray. Inevitably, differing saturation ranges had to be used for the panels. *HORIZONTAL:* different types of recurrent synapses. *Left:* Solely input driven networks without recurrent connections. *Center:* Recurrent networks with dynamic synapses using short-term plasticity. *Right:* Recurrent networks with static synapses. *VERTICAL:* Mean activity of the entire network (*top*) and the balance of the populations, measured by the difference between the mean excitatory and inhibitory firing rates (*bottom*).

3.2.4 Measurement

In order to study the self-adjustment capabilities of the setup, three types of networks were investigated:

- **unconnected** All recurrent synapses were discarded ($w = 0$) in order to determine the sole impact of external stimulation.

- **dynamic** All recurrent synapses featured short-term plasticity. The mode (facilitating, depressing) depended on the type of the connection as shown in Figure 3.2.

- **static** The STP-mechanism was switched off in order to study the relevance of short-term plasticity for the self-adjustment ability.

Rather than on the analysis of the dynamics of a specific network, we aimed at the investigation of the universality of application of the examined network architecture.

Therefore, random networks were generated obeying the above described probability distributions. Besides the three fundamentally different network types (unconnected, dynamic and static), external stimulation of different strength was applied by sweeping both the average membrane potential V_{rest} and the magnitude of fluctuations W_{input}.

For every set of network and input parameters, $n_{\text{run}} = 20$ networks and input patterns were generated and run for $T_{\text{exp}} = 4.5$ seconds. The average firing rates of both populations of neurons were recorded. To exclude transient initialization effects, only the time span $1\,\text{s} \leq t \leq T_{\text{exp}}$ was evaluated. Networks featuring the self-adjustment property are expected to modulate their activity to a medium level of about 5 Hz to 20 Hz over a wide range of external stimulation.

This setup was both emulated on the FACETS Stage 1 Hardware system and simulated using PCSIM in order to verify the results.

3.3 Results

First we present the results of the hardware emulation and compare them with the properties of simulated networks. Beside the capability of adjusting network activity in principle, we examine to what extent the observed mechanisms are insusceptible to changes in the hardware substrate. Finally we take a look at the ability of such networks to process input streams.

3.3.1 Self-Adjustment Ability

The results of the hardware emulation performed according to the setup description given in Sections 3.2.3 and 3.2.4 are shown in Fig. 3.4. The axes display different input strengths, controlled by the average membrane potential V_{rest} and the magnitude of fluctuation W_{input}. Average firing rates are indicated by the shade of gray of the respective tile.

The average response of networks without recurrent connections is shown in Fig. 3.4**A**. Over a wide range of weak stimulation (lower left corner) almost no spikes occur within the network. For stronger input, the response steadily rises up to $\nu \approx 29\,\mathrm{Hz}$. In Fig. 3.4**D** the activity of the excitatory and the inhibitory population are compared. Since external stimulation was configured equally for either population, one expects a similar response $\nu_{\mathrm{exc}} - \nu_{\mathrm{inh}} \approx 0$, except for slight stochastic variations. Obviously, the used hardware device exhibits a strong and systematic discrepancy of the sensitivity between the populations, which were located on different halves of the chip. The mean firing rate of excitatory neurons is about three times as high as the response of inhibitory neurons.

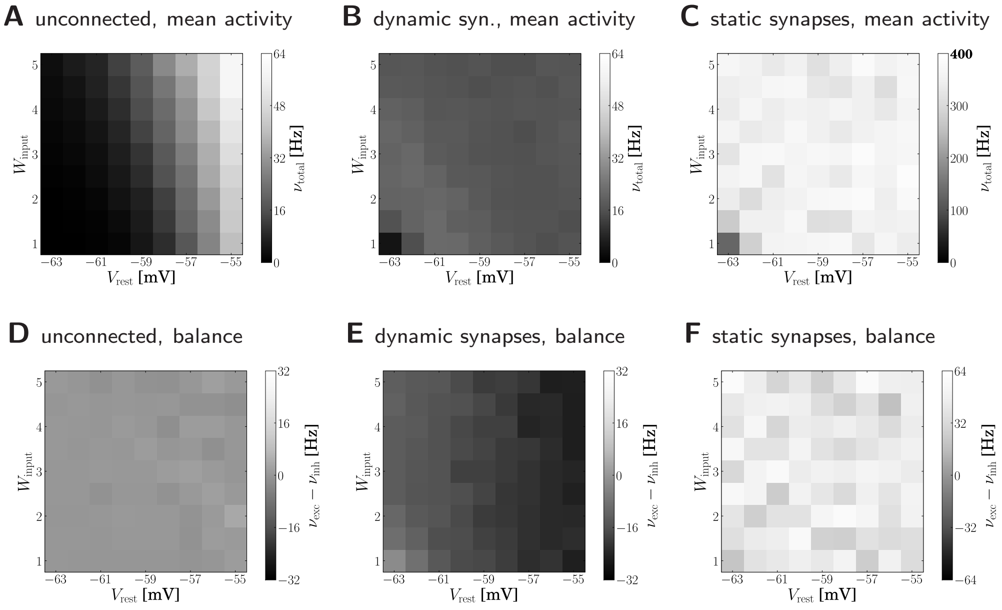

Figure 3.5: Results of the software simulation.
The experimental setup and the arrangement of the panels are equal to Figure 3.4. Also, the general behavior is consistent with the hardware emulation, though the average network response is more stable against different strengths of stimulation and all firing rates are higher. Accordingly, in case of dynamic recurrent synapses, the plateau is located at $\nu_{\mathrm{total}} \approx 17\,\mathrm{Hz}$.

The mid-column – Fig. 3.4**B** and **(E)** – shows the response of recurrent networks featuring dynamic synapses with the presented short-term plasticity mechanism. Over a wide range of stimulation, the mean activity is adjusted to a level of 9 Hz to 15 Hz. A comparison to the solely input driven setup proves that recurrent networks with dynamic synapses are capable of both raising and lowering their activity towards a smooth plateau. A closer look at the firing rates of the populations reveals the underlying mechanism: In case of weak external stimulation, excitatory network activity exceeds inhibition, while the effect of strong stimuli is attenuated by intense firing of inhibitory neurons. This functionality agrees with the concept of depressing interior and facilitating bridging connections, as described in Section 3.2.3.1.

In spite of the disparity of excitability between the populations, the applied setup is capable of properly adjusting network activity. It is noteworthy that the used connection probabilities and synaptic weights completely ignored this characteristic of the underlying substrate.

To ensure that the self-adjustment ability originates from short-term synaptic plasticity, the STP-mechanism was switched off during a repetition of the experiment. The respective results for recurrent networks using static synapses are shown in Fig. 3.4**C** and **(F)**. The networks clearly lack the previously observed self-adjustment capability, but rather tend to extreme excitatory firing. It must be mentioned that such high firing rates exceed the readout bandwidth of the current FACETS Stage 1 Hardware system. Thus, an unknown amount of spike events was discarded within the readout circuitry of the chip. The actual activity of the networks is expected to be even higher than the measured response.

3.3.2 Comparison to PCSIM

While the results of the hardware emulation draw a self-consistent picture, it ought to be excluded that the observed self-adjustment arises from hardware specific properties. Therefore, the same setup was applied to the software simulator PCSIM. The results of the software simulation are shown in Fig. 3.5. The six panels are arranged like those of the hardware results in Figure 3.4.

In agreement with the hardware emulation, the average response of networks without recurrent connections rises with stronger stimulation, see Fig. 3.5**A**. But as the disparity in the population excitability was not modeled in the simulation, their balance is only subject to statistical variations, see Fig. 3.5**D**.

Generally, the software simulation yields significantly higher firing rates than the hardware emulation. Two possible causes are:

- The instability of the excitatory reversal potential E_e certainly entails reduced network activity in case of the hardware emulation.
- The response curve of hardware neurons slightly differs from the behavior of an ideal conductance-based LIF model (Brüderle, 2009, Figure 6.4).

Consistently, an increased activity is also observed in the simulations of recurrent networks. Fig. 3.5**B** and **(E)** show the results for networks with synapses featuring short-term synaptic plasticity. Obviously, the networks exhibit the expected self-adjustment ability. But the plateau is found at approximately 17 Hz compared to 12 Hz in the hardware emulation. Finally, in case of static recurrent synapses – see Fig. 3.5**C** and **(F)** – the average network activity rises up to 400 Hz and lacks any visible moderation.

In conclusion, the hardware emulation and the software simulation yield similar results regarding the basic dynamics. Quantitatively, the results differ approximately by a factor of two.

In order to estimate the influence of an unstable reversal potential, which is suspected to be the leading cause for the inequality, the excitatory reversal potential was globally set to $E_\mathrm{e} = -20\,\mathrm{mV}$ during a repetition of the software simulation.

Indeed, the results of the different back-ends become more similar. The average activity of networks with dynamic synapses (corresponding to Figure 3.4**B** and 3.5**B**) is shown in Fig. 3.6.

Due to the obviously improved agreement, all further software simulations have been performed with a lower excitatory reversal potential $E_\mathrm{e} = -20\,\mathrm{mV}$.

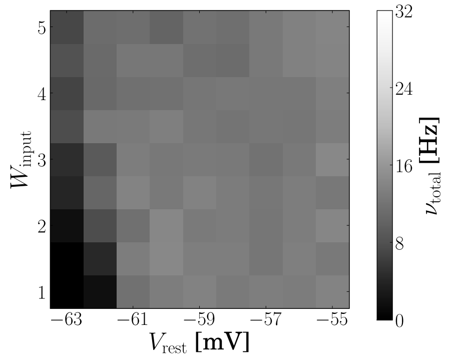

Figure 3.6: Software simulation: Lower excitatory reversal potential. Average network response of recurrent networks with dynamic synapses. In order to approximate the instability of the excitatory reversal potential in the chip, E_e was set to $-20\,\mathrm{mV}$ for subsequent software simulations. Compare with Figure 3.4**B**.

3.3.3 Robustness

We show that the observed self-adjustment property of the network architecture provides certain types of activity robustness that are beneficial for the operation of neuromorphic hardware systems.

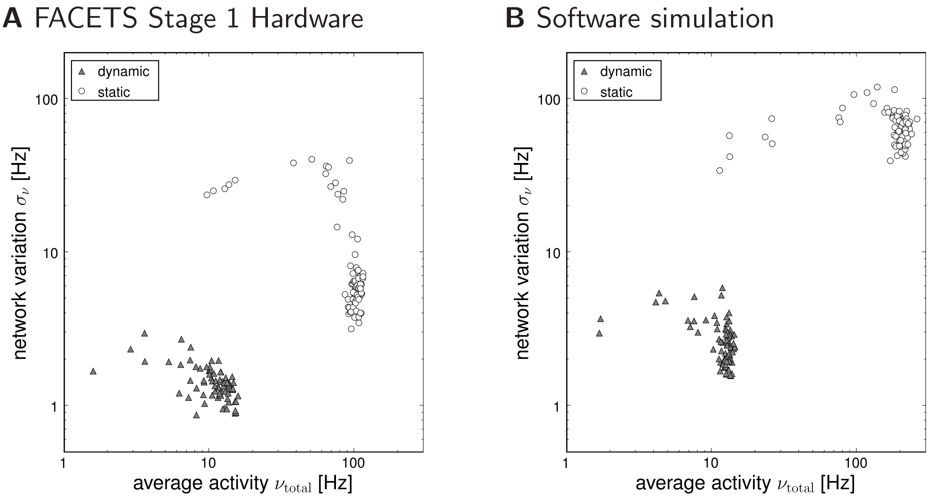

Figure 3.7: Reliable and realistic network activity.
Each point is determined by 20 random networks generated from equal probability distributions. The average firing rate of all networks is plotted on the x-axis, the standard deviation between the networks on the y-axis. Recurrent networks featuring short-term plasticity (triangles) can reliably be found within a close range. Setups with static synapses (circles) exhibit both larger average firing rates and larger standard deviations. (**A**) Emulation on the FACETS Stage 1 Hardware. (**B**) Software simulation with lowered excitatory reversal potential E_e.

Reliable and relevant activity regimes. By applying the network architecture presented in Section 3.2.3.1, we aim at the following two kinds of robustness of network dynamics:

- A high reliability of the average network activity, independent of the precise individual network connectivity or stimulation pattern. All networks with dynamic synapses that are generated and stimulated randomly, but obeying equal probability distributions, shall yield a similar average firing rate ν_{total}.
- The average firing rate ν_{total} shall be kept within a biologically relevant range for a wide spectrum of stimulation strength and variability. For awake mammalian cortices, rates in the order of 5 to 20 Hz are typical, see e.g. (Steriade, 2001; Steriade et al., 2001; Baddeley et al., 1997).

The emergence of both types of robustness in the applied network architecture is first tested by evaluating the PCSIM data. Still, it is not a priori clear that the robustness is preserved when transferring the self-adjusting paradigm to the hardware back-end. The transistor-level variations discussed in Section 3.2.1.2 might impede the reliability of the moderating effects, e.g. by causing an increased excitability for some of the neurons, or by too heterogeneous characteristics of the synaptic plasticity itself. Therefore, the robustness is also tested directly on a hardware device and the results are compared with those of the software simulation.

While each tile in Figure 3.4 represents the averaged overall firing rate ν_{total} of 20 randomly generated networks and input patterns, Fig. 3.7 shows the standard deviation σ_ν of the activity of networks obeying equal probability distributions as a function of ν_{total}. Networks using dynamic synapses are marked by triangles, those with static synapses by circles. Only setups with $\nu_{\text{total}} > 1\,\text{Hz}$ are shown.

For both the hardware device and the software simulation, the data clearly show that the required robustness effects are achieved by enabling the self-adjusting mechanism with dynamic synapses. The fluctuation σ_ν from network to network is significantly lower for networks that employ dynamic recurrent connections. Moreover, only for dynamic synapses the average firing rate ν_{total} is reliably kept within the proposed regime, while in case of static synapses most of the observed rates are well beyond its upper limit.

This observation qualitatively holds both for the hardware and for the software data. In case of networks with static synapses emulated on the hardware system, the upper limit of observed firing rates at about 100 Hz is determined technically by bandwidth limitations of the spike recording circuitry. This also explains the dropping variation σ_ν for firing rates close to that limit. If many neurons fire at rates that exceed the readout bandwidth, the diversity in network activity will seemingly shrink.

While the software simulation data prove that the self-adjusting principle provides the robustness features already for networks as small as those tested, the hardware emulation results show that the robustness is preserved despite of the transistor-level variations. Even though the different biological network descriptions are mapped randomly onto the inhomogeneous hardware resources, the standard deviation of firing rates is similar in hardware and in software.

Independence of the emulation device. Besides the ambiguous mapping of given biological network descriptions to an inhomogeneous neuromorphic hardware system as discussed above, the choice of the particular emulation device itself imposes another source of possible unreliability of results. Often, multiple instances of the same system are available to an experimenter. Ideally, such chips of equal design should yield identical network dynamics. But due to process-related inhomogeneities and due to the imperfections as discussed in Section 3.2.1.2, this objective is unachievable in terms of precise spike timing whenever analog circuitry is involved. Nevertheless, one can aim for a similar behavior on a more global scale, i.e. for alike results regarding statistical properties of populations of neurons.

All previous emulations have been performed on a system which was exclusively assigned to the purpose of this work. In order to investigate the influence of the particular hardware substrate, a different randomly chosen chip was set up with the same *biological configuration*. In this context, biological configuration denotes that both systems had been calibrated for general purpose. The high-level pyNN-description of the experiment remained unchanged – only the translation of biological values to hardware parameters involved different calibration data. This customization is performed automatically by low-level software structures. Therefore, the setup is identical from the experimenter's point of view.

In the following, the two devices will be referred to as *primary* and *comparative*, respectively. Just as on the primary device, networks emulated on the comparative system featured the self-adjustment ability if dynamic synapses were used for recurrent connections. But network activity was moderated to rather low firing rates of 2 Hz to 6 Hz. The response of networks without recurrent connections revealed that the used chip suffered from a similar disparity of excitability as the primary device. But in this case, it was the inhibitory population which showed a significantly heightened responsiveness.

Apparently, the small networks were not capable of completely compensating for the systematic unbalance of the populations. Nevertheless, they still were able to both raise and lower their firing rate compared to input induced response. Fig. 3.8 shows the difference of the activity between recurrent networks with short-term synaptic plasticity and solely input driven networks without recurrent connections,

$$\Delta\nu := \nu_{\text{total, dyn}} - \nu_{\text{total, input}} \quad .$$

For this chart, the V_{rest} - W_{input} diagonal of Figure 3.4 has been mapped to the x-axis, representing an increasing input strength. $\Delta\nu$ is plotted on the y-axis. Independent of the used back-end, recurrent networks raise activity in case of weak external excitation, while the effect of strong stimulation is reduced.

To allow for the inverse disparity of excitability of the comparative device, the mapping of the excitatory and the inhibitory population, which were located on different halves of the chip, was mirrored during a repetition of the emulation. Thus, the excitatory population exhibited an increased responsiveness resembling the disparity of the primary device. The $\Delta\nu$-curve of the mirrored repetition on the comparative system can also be found in Figure 3.8. As expected, with this choice of population placing, the moderating effect of the applied self-adjusting paradigm matches better the characteristics of the primary device.

These observations suggest that differing emulation results rather arise from large-scaled systematic inhomogeneities of the hardware substrate than from statistically distributed fixed pattern noise of individual units.

Therefore, it can be stated that the applied architecture is capable of reliably compensating statistical fluctuations of hardware unit properties, unless varia-

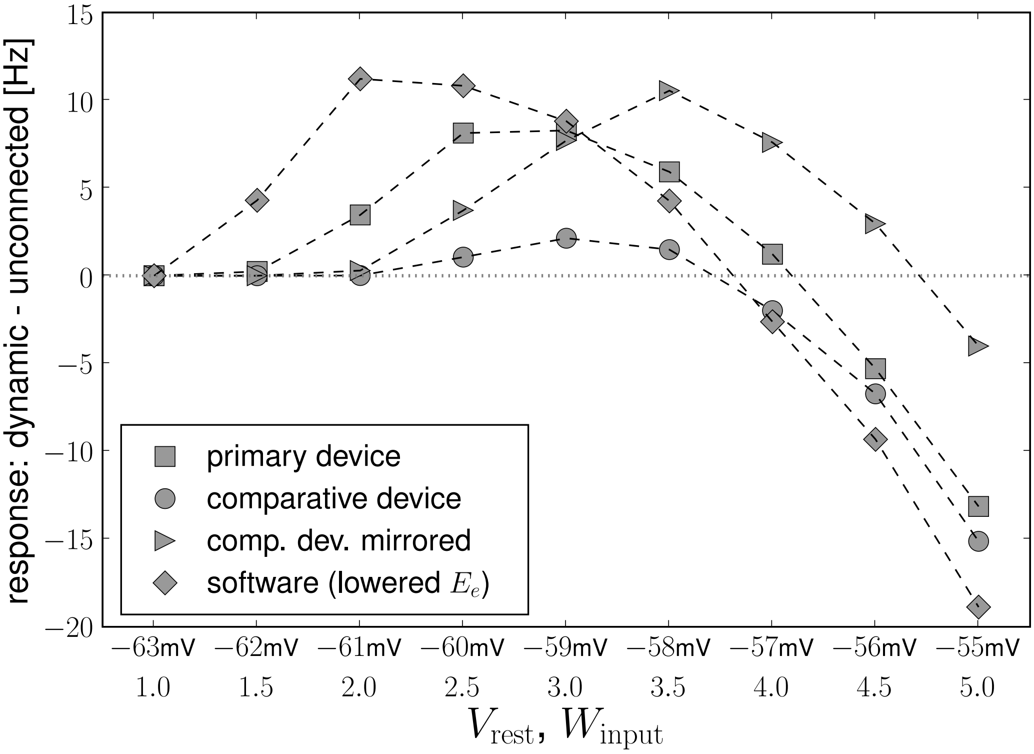

Figure 3.8: Self-adjusting effect on different platforms.
The difference $\Delta\nu := \nu_{\mathrm{total,\,dyn}} - \nu_{\mathrm{total,\,input}}$ is plotted against an increasing strength and variability of the external network stimulation. The diamond symbols represent the data acquired with PCSIM. The square (circle) symbols represent data measured with the primary (comparative) hardware device. Measurements with the comparative device, but with a mirrored placing of the two network populations, are plotted with triangle symbols. See main text for details.

tions extend to a global scale. But even in case of large-scale deviations, the applied construction principle preserves its self-adjustment ability and provides reproducible network properties, albeit at a shifted working point.

3.3.4 Responsiveness to Input

While it was shown that the applied configuration provides a well-defined network state in terms of average firing rates, it remains unclear whether the probed architecture is still able to process information induced by external input. It can be suspected that the strong recurrent connectivity "overwrites" any temporal structure of the input spike trains. Yet, the usability of the architecture regarding a variety of computational tasks depends on its responsiveness to changes in the

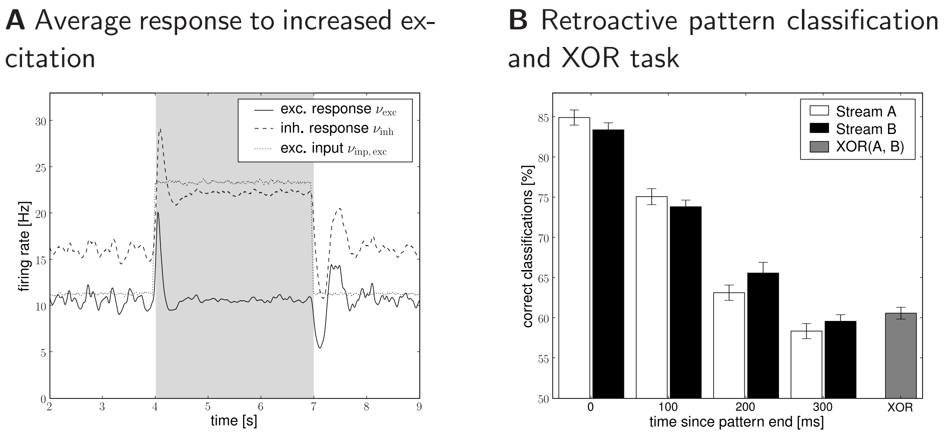
A Average response to increased excitation

B Retroactive pattern classification and XOR task

Figure 3.9: Network traces of transient input.
Results of software simulations testing the response of recurrent networks with dynamic synapses to transient input. (**A**) Firing rate of the excitatory input channels and average response of either population to an excitatory input pulse lasting for 3 seconds. The steep differential change in excitation is answered by a distinct peak. After some hundred milliseconds the networks attune to a new level of equilibrium. (**B**) Average performance of the architecture in a retroactive pattern classification task. The network states contain information on input spike patterns which were presented some hundred milliseconds ago. The latest patterns presented are to be processed in a non-linear XOR task.

input. A systematic approach to settle this question exceeds the scope of this work. Therefore, we address the issue only in brief.

First, we determine the temporal response of the architecture to sudden changes in external excitation. Then, we look for traces of previously presented input patterns in the current network state and test whether the networks are capable of performing a non-linear computation on the meaning assigned to these patterns.

For all subsequent simulations the input parameters are set to $V_{\mathrm{rest}} = -59\,\mathrm{mV}$ and $W_{\mathrm{input}} = 4.0$ (cf. Figure 3.6). Only networks featuring dynamic recurrent connections are investigated. Due to technical limitations of the current hardware system as discussed in Section 3.2.1.2, the results of this section are based on software simulations, only. For example, the additional external stimulation, as applied in the following, exceeds the current input bandwidth of the prototype hardware device. Furthermore, the evaluation of network states requires access to (at least) the spike output of all neurons, simultaneously. The current hardware system only supports the recording of a small subset of neurons at a time.

In Fig. 3.9**A** the average response of the excitatory and inhibitory populations to increased external excitation are shown. For this purpose, the firing rate of all excitatory Poisson input channels was doubled from 11.8 Hz to 23.6 Hz at $t = 4\,\mathrm{s}$. It was reset to 11.8 Hz at $t = 7\,\mathrm{s}$, i.e. the applied stimulation rate was shaped as a rectangular pulse. In order to examine the average response of the recurrent networks to this steep differential change in the input, $n_{\mathrm{run}} = 1000$ networks and input patterns have been generated. While the network response obtained from a single simulation run is subject to statistical fluctuations, the influence of the input pulse is revealed precisely by averaging over the activity of many different networks. For analysis, the network response was convolved with a box filter (50 ms window size). In conclusion, the temporal response of the recurrent networks is characterized by two obvious features:

1. Immediately after the additional input is switched on or off, the response curves show distinct peaks which decay at a time scale of $\tau \approx 100\,\mathrm{ms}$.
2. After some hundred milliseconds, the networks level off at a new equilibrium. Due to the self-adjustment mechanism, the activity of the inhibitory population clearly increases.

These findings confirm that the investigated networks show a significant response to changes in the input. This suggests that such neural circuits might be capable of performing classification tasks or continuous-time calculations if a readout is attached and trained.

We tested this conjecture by carrying out a computational test proposed in (Häusler and Maass, 2007). The 64 external input channels were assigned to two disjunct *streams* A and B. Each stream consisted of 16 excitatory and 16 inhibitory channels. For each stream two Poisson spike train templates (referred to as $+_S$ and $-_S$, $S \in \{A, B\}$) lasting for 2400 ms were drawn and partitioned to 24 segments $\pm_{S,i}$ of 100 ms duration. In every simulation run the input was randomly composed of the segments of these templates, e.g.

$$\begin{aligned} &Stream\,A: \; +_{A23} \; -_{A22} \; \cdots \; -_{A1} \; +_{A0} \\ &Stream\,B: \; -_{B23} \; +_{B22} \; \cdots \; -_{B1} \; -_{B0} \end{aligned}$$

leading to 2^{24} possible input patterns for either stream. Before the input was presented to the network, all spikes were jittered using a Gaussian distribution with zero mean and standard deviation 1 ms. The task was to identify the last 4 segments presented ($0 \leq i \leq 3$) at the end of the experiment. For that purpose, the spike response of the network was filtered with an exponential decay kernel ($\tau_{\mathrm{decay}} = \tau_{\mathrm{syn}} = 30\,\mathrm{ms}$). The resulting network state at $t = 2400\,\mathrm{ms}$ was presented to linear readout neurons which were trained via linear regression as in (Maass et al., 2002). The training was based on 1500 simulation runs. Another 300 runs were used for evaluation. In order to determine the performance of the architecture for this retroactive pattern classification task, the above setup was repeated 30 times with newly generated networks and input templates.

The average performance of networks with recurrent dynamic synapses is shown in Fig. 3.9**B**. The error bars denote the standard error of the mean. Obviously, the network state at $t = 2400\,\mathrm{ms}$ contains significant information on the latest patterns presented and preserves traces of preceding patterns for some hundred milliseconds. For comparison, recurrent networks using static synapses performed only slightly over chance level (not shown). In addition to the pattern classification task, another linear readout neuron was trained to compute the non-linear expression $\mathrm{XOR}(\pm_{A0}, \pm_{B0})$ from the network output. Note that this task cannot be solved by a linear readout operating directly on the input spike trains.

Summing up, the self-adjusting recurrent networks are able of perform multiple computational tasks in parallel. Since the main objective of this work was to verify the self-adjustment ability of small networks on a neuromorphic hardware device, both connection probabilities and synaptic weights of recurrent connections had been chosen high compared to the strength of external stimulation. Still, the networks significantly respond to changes in the input and provide manifold information on present and previous structure of the stimulus.

Recent theoretical work (Buesing et al., 2010) stressed that the computational power of recurrent networks of spiking neurons strongly depends on their connectivity structure. As a general rule, it has been shown to be beneficial to operate a recurrent neural network in the edge-of-chaos regime (Bertschinger and Natschläger, 2004). Nevertheless, as addressed in (Legenstein and Maass, 2007), the optimal configuration for a specific task can differ from this estimate. Accordingly, task-dependent recurrent connectivity parameters might be preferable to achieve good experimental results (see, e.g., (Haeusler et al., 2009)). While networks of randomly connected neurons feature favorable kernel qualities, i.e. they perform rich non-linear operations on the input, theoretical studies of (Ganguli et al., 2008) prove that networks with hidden feedforward structures provide superior memory storage capabilities. Future research might identify such connectivity patterns in seemingly random cortical circuits and improve our understanding of working memory.

While the examined recurrent network architecture was not optimized for computation, neither regarding its kernel quality nor its memory traces, the cited studies suggest that the performance will increase if network parameters are attuned to particular tasks. Further research is needed to explore under which conditions the examined architecture provides a stable operating point, a high responsiveness to stimuli, and appropriate memory traces.

3.4 Discussion

We showed that recurrent neural networks featuring short-term synaptic plasticity are applicable to present neuromorphic mixed-signal VLSI devices. For the first time dynamic synapses play a functional role in network dynamics during a hardware emulation. Since neuromorphic hardware devices model neural information processing with analog circuitry, they generally suffer from process-related fluctuations which affect the dynamics of their components. In order to minimize the influence of unit variations on emulation results, we applied a self-adjustment principle on a network level as proposed by (Sussillo et al., 2007).

Even though the employed prototype system only supports a limited network size, the expected self-adjustment property was observed on all used back-ends. The biological description of the experimental setup was equal for all utilized chips, i.e. the configuration was not customized to characteristics of the specific hardware system. Beyond the validation of the basic functioning of the self-adjusting mechanism, we addressed the robustness of the construction principle against both statistical variations of network entities and systematic disparities between different chips. We showed that the examined architecture reliably adjusts the average network response to a moderate firing regime. While congeneric networks emulated on the same chip yielded a widely similar behavior, the operating point achieved on different systems still was affected by large-scale characteristics of the utilized back-end.

All outcomes of the hardware emulation were qualitatively confirmed by software simulations. Furthermore, the major imperfection of the current revision of the FACETS Stage 1 Hardware, the instability of the excitatory reversal potential, was identified by the accompanying application of the simulator PCSIM.

Presumably, the performance of the applied architecture will improve with increasing network size. Upcoming neuromorphic emulators like the FACETS Stage 2 Wafer-scale Integration system (see Schemmel et al., 2008; Fieres et al., 2008) will comprise more than 100,000 neurons and millions of synapses. Even earlier, the present chip-based system will sustain the interconnection of multiple chips and thus provide a substrate of some thousand neurons. As such large-scale mixed-signal VLSI devices will inevitably exhibit variations in unit properties, detailed knowledge of circuitry design is required by the user to reduce distortions of experimental results on the level of single units. On the other hand, the beneficial application of neuromorphic VLSI devices as both neuro-scientific modeling and novel computing tools will require that it does not demand an expert in electronic engineering to run the system. We showed that self-regulation properties of neural networks can help to overcome disadvantageous effects of unit level variations of neuromorphic VLSI devices. The employed network architecture might ensure a highly similar network behavior independent of the utilized system. Therefore this work displays an important step towards a reliable and practicable operation of neuromorphic hardware.

The applied configuration required strong recurrent synapses at a high connectivity. The results of (Sussillo et al., 2007) show that even sparsely connected networks can manage to efficiently adjust their activity, provided they comprise a sufficiently large number of neurons which will be sustained by future hardware systems. Thereby, the examined construction principle will become applicable to a variety of experimental setups and network designs. As touched upon in Section 3.3.4, the presented self-adjusting networks still are sensitive and responsive to changes in external excitation. Furthermore, we verified that even networks with disproportionately strong recurrent synapses can perform simple non-linear operations on transient input streams. By applying biologically more realistic connectivity parameters, it has been shown that randomly connected networks of spiking neurons are able to accomplish ambitious computational tasks (Maass et al., 2004) and that short-term synaptic plasticity can improve the performance of such networks in neural information processing (Maass et al., 2002). Thus, this architecture provides a promising application for neuromorphic hardware devices while the high configurability of novel systems as well supports the emulation of circuits tailored to specific tasks.

3.5 Acknowledgments

This chapter is based on the paper *Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity*, which was written by Johannes Bill (JB), myself (KS), Daniel Brüderle (DB), Johannes Schemmel (JS), Karlheinz Meier (KM) and Wolfgang Maass (WM). The work for this chapter was done within the FACETS project as a colaboration with the University of Heidelberg. The experiments were designed and planned by JB and KS during the visit of JB in our lab for his master thesis. The hardware experiments were performed by JB and DB. The software experiments were perfomed by KS and JB, and the paper was written by JB, KS, and DB with additional input from JS, KM, and WM.

CHAPTER 4

Statistical comparison of spike responses to natural stimuli in monkey area V1 with simulated responses of a detailed laminar network model for a patch of V1

Contents

4.1 Introduction 75

4.2 Materials and Methods 77

4.3 Results 95

4.4 Discussion 114

4.5 Acknowledgments 120

A major goal of computational neuroscience is the creation of computer models for cortical areas whose response to sensory stimuli resembles that of cortical areas in-vivo in important aspects. Although often based on detailed anatomical data, it is seldom considered whether the simulated spiking activity is realistic (in a statistical sense) in response to relevant natural stimuli. Because certain statistical properties of spike responses were suggested to facilitate ongoing computations in the cortex, acquiring a realistic firing regime in cortical network models might be a prerequisite for analyzing their emergent computational functions. We here present a characterization and comparison of the statistical response properties of the primary visual cortex (V1) *in vivo* and *in silico* in response to natural movie stimuli. We recorded from multiple electrodes in area V1 of 4 macaque monkeys, and developed a large state-of-the-art network model for a 5 mm x 5 mm patch of V1 composed of 35000 neurons and 3.9 million synapses that integrates many

previously published anatomical and physiological details. By quantitative comparison of the model response to the “statistical fingerprint” of responses *in vivo*, we find that our model for a patch of V1 responds to the same movie (using a simple model for retina and LGN as pre-processors) in a way which matches the statistical structure of the recorded data surprisingly well. In particular, the deviation between the firing regime of the model and the *in vivo* data is on the same level as deviations among monkeys and sessions. To reach a realistic firing state, it was not only necessary to include both NMDA and GABAB synaptic conductances in our model, but also to dramatically increase the strength of excitatory synapses onto inhibitory neurons (more than two-fold) in comparison to literature values, suggesting that the strength of inhibition is underestimated in current network models.

4.1 Introduction

Numerical simulations of detailed biophysical models of cortical microcircuits or even whole brain regions provide powerful tools to approach complex questions in neuroscience, and are commonly regarded as a promising tool to understand the mechanistic link from anatomical structure and physiological properties to computational functions of cortical circuits. In general, approaches along this line incorporate selected aspects of the known anatomy and physiology to replicate experimental data on emergent functional properties, such as for instance the structure of preferred orientation maps of the primary visual cortex (Adorjan et al., 1999; Bartsch and van Hemmen, 2001; Blumenfeld et al., 2006), direction selectivity maps (Ernst et al., 2001; Wenisch et al., 2005), and simple/complex cells (Chance et al., 1999; Wielaard et al., 2001; Tao et al., 2004), or successfully exemplifying theoretical ideas about information processing in the brain (e.g. (Diesmann et al., 1999; Maass et al., 2002; Vogels and Abbott, 2005)). However, these increasingly complex recurrent network models are often still a strong abstraction from reality and it is not clear whether the responses of such network models exhibit at least a general likeness to its counterpart in reality.

Constraining the firing regime of *in silico* models with that observed *in vivo* is important for at least two reasons: First, it will benchmark current models to achieve a realistic firing response, and thus will further help to open new research directions because it will hint at current short-comings of existing models. Second, it has been suggested theoretically, that there might be a firing regime or state that is favorable for ongoing computation within recurrent neural networks (Brunel, 2000; Vogels and Abbott, 2005; Legenstein and Maass, 2007). One might thus postulate that during evolution the brain has shaped a particular useful firing regime which is in some way supporting the computational function of the neural tissue. Therefore achieving a realistic firing activity in cortical circuit models might be an important but rarely considered prerequisite to employing these models for analyzing aspects of cortical computational functions. If a realistic firing regime cannot be achieved easily, the validity of conclusions drawn from these model circuits might be corrupted, or efforts have to be made to tune these models towards a realistic regime. To investigate this issue, we ask in this study if a state-of-the-art network model of a cortical circuit is able to reproduce the characteristic firing regime of the cortex.

We focus on the primary visual system (V1). Its anatomical and neurophysiologic details are relatively well known, and its position in visual sensory processing is relatively well established. Moreover, V1 already serves as a reference cortical area to investigate large scale network models (Johansson and Lansner, 2007; Kremkow et al., 2007), although still many aspects of its computational organization and the underlying mechanisms remain poorly understood (Olshausen and Field, 2005). To compare the firing state of V1 *in vivo* with that of simulated responses of a cortical network model *in silico*, we first recorded spike responses with multi-electrode arrays in V1 of 4 anesthetized monkeys while presenting several minutes lasting semi-natural movies. Given the complex naturalistic stimuli, we thus expect that V1 will likely be in an operating regime, where its computations are usually performed. We characterized this firing regime by its "statistical fingerprint" using a number of salient statistical features, measuring the spike variability, the burst behavior, and the correlation structure. We then compared this "statistical fingerprint" to that obtained from the response of a state-of-the-art cortical circuit model of a 5 mm x 5 mm patch of V1, comprising about 35000 neurons and 3.9 million synapses situated in several hypercolumns. The developed spiking neuron network model is based on the cortical microcircuit model described in (Häusler and Maass, 2007), which implements experimental data from (Thomson et al., 2002) on lamina-specific connection probabilities, and data from (Markram et al., 1998) and (Gupta et al., 2000) regarding stereotypical dynamic properties (such as paired pulse depression and paired pulse facilitation) of synaptic connections. We extended this cortical microcircuit model laterally and incorporated many anatomical properties of V1 in macaques to ensure the comparability to our *in vivo* recordings.

Our combinative approach, using both electrophysiological recordings and model circuit simulation, provided us with the unique possibility to use the same movie stimuli for the model simulations and during *in vivo* recordings. Given this comparability, we were able to investigate whether the firing regime of a model achieves a realistic state, and if not, whether a set of global parameters were sufficient to tune the models firing regime to become more realistic. We found that the response of the detailed model circuit adopts a firing regime that is remarkably similar to the *in vivo* response and is on average close to the deviations across different sessions and different monkeys. This close match was achieved by tuning only a few parameters: an overall synaptic weight scaling factor compensating for the reduced number of modeled neurons, the relative synaptic weight from excitatory to inhibitory neurons, and the relative strength of patchy lateral long-range excitatory weights. We found that the firing response statistics was not simply induced by the statistics of the complex input stimuli but instead depended significantly on the internal dynamics. This good fit suggests that current network models comprising realistic neuron dynamics, as well as realistic time courses of synaptic activation, which included short-term depression and facilitation, are capable of generating a similar diverse network response behavior as can be observed in *in vivo* recordings. We are convinced that the analysis of computational function of realistic circuits is closely linked to its firing regime. We therefore expect that this characterization of the firing regime provided here and the possibility to use a few parameters to calibrate a complex model will greatly ease the analysis of the computational properties of realistic, detailed circuit models in future.

4.2 Materials and Methods

4.2.1 Experimental Methods

4.2.1.1 Electrophysiological recording

The electrophysiological recordings were previously described (Montemurro et al., 2008), where the same data had been analyzed from a different perspective. However for completeness we include a detailed description here. Four adult rhesus monkeys (Macaca mulatta) participated in these experiments. All procedures were approved by the local authorities (Regierungspräsidium) and were in full compliance with the guidelines of the European Community (EUVD 86/609/EEC) for the care and use of laboratory animals. Prior to the experiments, form-fitted head posts and recording chambers were implanted during an aseptic and sterile surgical procedure (see e.g. (Logothetis et al., 2002)). To perform the neurophysiological recordings, the animals were anesthetized (remifentanil (0.5-2 μg/kg/min)), intubated, and ventilated. Muscle relaxation was achieved with a fast acting paralytic, mivacurium chloride (5 mg/kg/h). Body temperature was kept constant and lactated Ringer's solution was given at a rate of 10 ml/kg/h. During the entire experiment, the vital signs of the monkey and the depth of anesthesia were continuously monitored (as described in (Logothetis et al., 1999)). For the protocol used in these experiments, we had previously examined the concentration of all stress hormones (catecholamines) (Logothetis et al., 1999) and found them to be within the normal limits. Drops of 1% ophthalmic solution of anticholinergic cyclopentolate hydrochloride were instilled into each eye to achieve cycloplegia and mydriasis. Refractive errors were measured and contact lenses (hard PMMA lenses by Wöhlk GmbH, Germany) with the appropriate dioptric power were used to bring the animal's eye into focus on the stimulus plane. The electrophysiological recordings were performed with electrodes that were arranged in a 4 x 4 square matrix (inter-electrode spacing varied from 1 mm to 2.5 mm) and introduced each experimental session into the cortex through the overlying dura mater by a microdrive array system (Thomas Recording, Inc., Giessen, Germany). Electrode tips were typically (but not always) positioned in the upper or middle cortical layers. The impedance of the electrode varied from 300 to 800 kOhm. Both spontaneous and stimulus-induced neural activity were collected and recorded for periods up to 6 minutes. Signals were amplified using an Alpha Omega amplifier system (Alpha Omega GmbH, Ubstadt-Weiher, Germany). The amplifying system filtered out the frequencies below 1 Hz. Recordings were performed in a darkened booth (Illtec, Illbruck acoustic GmbH, Germany). The site receptive fields were plotted manually and the position and size of each field were stored together with the stimulus parameters and the acquisition data. The visual stimulator was a dual processor Pentium II workstation running Windows NT (Intergraph Corp., Huntsville, Alabama) and equipped with OpenGL graphics cards (3Dlabs Wildcat series). The resolution was set to 640 by 480 pixels. The refresh rate was 60 Hz and the movie frame rate was 30 Hz. All image generation was in 24 bit true-color, using hardware double buffering to provide smooth animation. The 640x480 VGA output drove the video interface of a fiber-optic stimulus presentation system (Avotec, Silent Vision, Florida) and also drove the experimenter's monitor. The field of view (FOV) of the system was 30Hx23V degrees of visual angle and the focus was fixed at 2 diopters. The system's effective resolution, determined by the fiber-optic projection system was 800Hx225V pixels (the fiber-optic bundle is 530x400 fibers). Binocular presentations were accomplished through two independently positioned plastic, fiber-optic glasses. Positioning was aided by a modified fundus camera (Zeiss RC250) that permitted simultaneous observation of the eye-fundus. The fundus camera has a holder for avotec projector so that the center of camera lens and avotec projector is aligned in the same axis. This process ensured the alignment of the stimulus-center with the fovea of each eye. To ensure accurate control of stimulus presentation a photo-diode was attached to the experimenter's monitor permitting the recording of the exact presentation time of every single frame. The visual stimuli were binocularly presented 3.5 to 6 minute long natural color movies (segments of the commercial movie "Star Wars"). During each of 10 recording sessions the movie was repeated 12–40 times.

4.2.1.2 Spike detection

In order to extract spike times from the electrophysiological recordings, the 20.83 kHz neural signal was filtered in the high-frequency range of 500–3500 Hz. The threshold for spike detection was set to 3.5 standard deviations. A spike was recognized as such only if previous spikes have occurred more than one millisecond earlier. Spikes detected with this simple threshold method represented the spiking activity of a small population of cells rather than well separated spikes from a single neuron. For spike sorting we used the method described by (Quian Quiroga et al., 2004). The spike waveforms were extracted around the detection times as described above (in a region of 0.25 ms before to 0.5 ms after the detected spike). These spike forms were interpolated and 10 wavelet features (with 4 scales) were extracted (Quian Quiroga et al., 2004). From this feature pool the 10 features (KS-test) were used as input for the clustering algorithm. We then sorted the spikes using the paramagnetic algorithm of (Quian Quiroga et al., 2004). For each electrode a few reasonable clusters were selected by visual inspection of the spike waveforms ensuring a reasonable distinguishable average waveform among clusters. After this initial selection, spikes which initially were not classified in a particular cluster (or belonging to not selected clusters) were forced to belong to the nearest selected cluster (Mahalanobis distance (Quian Quiroga et al., 2004)). A cluster that maintained very similar waveforms after this step was deemed to be a well-isolated cluster and was considered for further analysis. Otherwise the cluster was not considered further for spike sorting.

4.2.2 Model

In this section we describe a data based model, developed to compare its spiking activity with the electrophysiological recordings from macaque. It consists of an input model (representing the retina and lateral geniculate nucleus (LGN) of the thalamus) and a model of a patch of V1, receiving and processing the spikes of the input model. In the following we will first describe the V1 model and subsequently the input model.

4.2.2.1 V1 Model

Our model for a 5 mm x 5 mm patch of area V1 consisted of 34596 neurons and 3.9 million synapses. Various anatomical and physiological details were included in our model. The connectivity structure of the V1 model was similar to that of the generic cortical microcircuit model discussed in (Häusler and Maass, 2007). The neurons of that model were equally distributed on three layers, corresponding to the cortical layers 2/3, 4 and 5. Each layer contained a population of excitatory neurons and a population of inhibitory neurons with a ratio of 4:1 (Beaulieu et al., 1992; Markram et al., 2004). The inter- and intra-layer connectivity (probability and strength) was chosen according to experimental data from rat and cat cortex assembled in (Thomson et al., 2002). Although there are differences, the connectivity structure in macaque is similar to that of the cat (Callaway, 1998). In particular, if one identifies layer 2/3 and 4 in cat with 2-4B and 4C in macaque, respectively. The major geniculate input reaches in both species first layer 4C. Layer 4C projects to layer 2-4B, which in turn projects further to layer 5 (and layer 6 via layer 5), where feedback connections are made to layers 2-4B (see (Callaway, 1998) for a review). Further, the sub-laminar organization, e.g., the structure built by cytochrome-oxidase blobs in layer 2/3 (Callaway 1998), was neglected for simplicity, and for the lack of precise data. However, as described below, the V1 model contained in addition to the microcircuit model of (Häusler and Maass, 2007) a realistic thalamic input, a smooth orientation map, and patchy long-range connections in the superficial layer.

In contrast to (Häusler and Maass, 2007), we set the relative amount of neurons per layer to 33%. This partitioning corresponds to experimental data from macaques (O'Kusky and Colonnier, 1982; Beaulieu et al., 1992; Tyler et al., 1998), although we slightly adjusted the relative amount of neurons compared to the experimental values (where layer 4 has about 33% more neurons), compensating for the fact that our model neglects the magnocellular and koniocellular pathways in favor of the parvocellular pathway (Callaway, 1998). The three layers of the model can be identified with layers 2-4B, 4Cβ and 5-6 in macaque V1. To avoid confusion with the terminology of (Häusler and Maass, 2007) we will, nevertheless, call them layer 2/3, 4, and 5 in the following.

In macaque each of these three layers contains approximately 50000 neurons under a surface area of 1 mm^2 (Beaulieu et al., 1992). In our model we neglected that neuron density varies up to 1.5 fold between the layers (Beaulieu et al., 1992) and assumed instead that the neurons are uniformly distributed throughout the cortex. Thus, for simplicity, we positioned all neurons on a cuboid grid with a constant grid spaces. Using the experimentally measured neuron density, e.g. for layer 2/3, the grid spacing would be $20 \mu m$ for all directions. Because the simulation of such a dense network would take too much computation time, we diluted the neuron density by increasing the lateral grid spacing to $80 \mu m$ and the vertical spacing to about $200 \mu m$.

We used a conductance-based single-compartment neuron model. Due to a considerable gain in computational speed we employed a neuron model suggested by (Izhikevich, 2003), which can be adjusted to express different firing dynamics (Izhikevich, 2006). We randomly drew the parameters for each neuron in the network according to the bounds provided by (Izhikevich et al., 2004). On the basis of these parameter distributions the excitatory pools consisted of regular spiking cells, intrinsically bursting cells, and chattering cells, with a bias towards regular spiking cells. The inhibitory pools consisted of fast spiking neurons and low-threshold spiking neurons.

In addition to the synaptic input from other neurons in the model, each neuron received as additional input synaptic background input, modeling the bombardment of each neuron with synaptic inputs from a large number of neurons that are not represented in our model. This synaptic background input causes a depolarization of the membrane potential and a lower membrane resistance, commonly referred to as the “high conductance state” (Destexhe et al., 2001). The conductances of the background input was modeled according to (Destexhe et al., 2001) by Ornstein-Uhlenbeck processes with means $g_{exc} = 0.012\ \mu S$ and $g_{inh} = 0.047\ \mu S$, variances $\sigma_{exc} = 0.003\ \mu S$ and $\sigma_{inh} = 0.0066\ \mu S$, and time constants $\tau_{exc} = 2.7\ ms$ and $\tau_{inh} = 10.5\ ms$, where the indices exc/inh refer to excitatory and inhibitory background conductances, respectively. During the parameter optimization we scaled the variances of both processes. The scaling factor of both variances affects the amount of noise added to the conductance of a neuron.

Short term synaptic dynamics was implemented according to (Markram et al., 1998), with synaptic parameters chosen as in (Maass et al., 2002) to fit data from microcircuits in rat somatosensory cortex (based on (Gupta et al., 2000) and (Markram et al., 1998)). For further details we refer to (Häusler and Maass, 2007).

Lateral connectivity structure The generic microcircuit model of (Häusler and Maass, 2007) was based on data for a column of about 100 μm diameter with uniform connectivity per layer and neuron type. Here we extended the model laterally to several millimeters. Thus connection probabilities in our model depend on the lateral distance. For inter- and intra-cortical connections we generally used a bell-shaped (Gaussian) probability distribution for determining the lateral extent. The standard deviation of the Gaussian was set to 200 μm for excitatory neurons (Lund et al., 2003; Blasdel et al., 1985; Buzas et al., 2006) and to 150 μm for inhibitory neurons to incorporate the observed occurrence of extremely narrow inhibitory dendritic and axonal spreads (70 μm, (Lund et al., 2003)). The arborization of excitatory neurons in layer 5 seems to be wider, more diffuse and has a spread of more than 500 μm laterally from the soma (Blasdel et al., 1985). Thus, for these connections we set the standard deviation to 300μm. Note that the value for the standard deviation is about half the expected maximal extent of 95% of the arborizations.

To ensure consistency with the connectivity data of (Thomson et al., 2002), we scaled the Gaussian profiles such that the peak probabilities correspond to their experimentally measured connection probabilities. Therefore, their connectivity pattern was locally preserved.

According to (Song et al., 2005) the number of bidirectional connections between excitatory neurons in layer 5 is four times higher than the expected number under the assumption that the conditional probabilities, whether an unidirectional connection exists or not, are the same. We incorporated this probability increase into our model.

Patchy lateral long-range connections In cat and macaque, many pyramidal cells in layer 2/3 of the striate cortex (and also elsewhere in the cortex (Lund et al., 2003)) have characteristic long-range projections targeting laterally 80% excitatory and 20% inhibitory cells (McGuire et al., 1991) which are up to 6 mm and more away (Gilbert et al., 1996; Lund et al., 2003; Buzas et al., 2006). Moreover, the targeted neurons tend to have similar feature preference as its origin, resulting in patchy connections linking similar preferred orientations (Gilbert et al., 1996; Buzas et al., 2006). Combining anatomical reconstructions of neurons and optical imaging of orientation maps, (Buzas et al., 2006) proposed a formula to calculate the bouton density ρ of a typical layer 2/3 pyramidal cell:

$$\rho(r, \Delta\phi) = Z \left(e^{-\frac{r^2}{2\sigma_1^2}} + m\, e^{-\frac{r^2}{2\sigma_2^2}}\, e^{\kappa \cos(2(\Delta\phi - \mu))} \right) \tag{4.1}$$

r is the lateral (Euclidean) distance between the pre- and the post-synaptic neuron, and $\Delta\phi$ is the difference of preferred orientations of the two neurons. Parameter μ is an offset in the orientation preference and parameter m is a scaling factor which accounts for the importance of the long-range orientation dependent term against the local orientation independent term. Standard deviations σ_1 and σ_2 regulate the spatial width of the non-oriented and oriented term, respectively. Parameter κ signifies the "peakiness" of the density on the orientation axis. Z is a normalization constant.

Since we defined preferred orientations in a hard-wired manner via "oriented" input connections (see Section 4.2.2.1), we could apply Eq. 4.1 for the lateral connections in layer 2/3, more specifically,for projections from excitatory cells targeting excitatory and inhibitory cells (McGuire et al., 1991).

Analogous to connections between other layers, we set $\sigma_1 = 200\mu$m for the local non-oriented term. We set $\mu = 0°$ and $\sigma_2 = 1000\mu$m (estimated from the measurements of (Buzas et al., 2006)). We chose a higher $\kappa = 20$ than reported because of the following reasoning. As described above, the neuron density of our circuit model is much smaller than in reality. We compensated this neuron dilution by a noise process fed into each modeled neuron, which implicitly models activation of omitted neurons. Because any (implicit) input from omitted neurons is independent of orientation preference, neurons in the circuit should have a strong bias towards orientation preference dependent connections. To account for this bias, we therefore increased κ.

Finally, the parameter m was set so that 58 % of the excitatory synapses onto an excitatory neuron in layer 2/3 were long-range connections. As before, the connection probability was scaled, according to (Thomson et al., 2002), by setting Z to appropriate values. Thus, locally, i.e., for a neuron at the same lateral position (and orientation preference), such as a neuron located in the same layer beneath or above the pre-synaptic neuron, the connection probabilities were preserved. However, the weight distribution of the long-range connection was not constrained by (Thomson et al., 2002). Hence, we scaled the recurrent weight reported by (Thomson et al., 2002). We used a value of 1.0 for this parameter that is the same average weight as in (Thomson et al., 2002).

Distance dependent synaptic delay Synaptic delays differ for inhibitory and excitatory neurons. They were set according to measurements by (Gupta et al., 2000) (for details see (Häusler and Maass, 2007)). These delays stem from molecular processes of synaptic transmission. In addition, a second delay originating from finite spike propagation velocity of the fibers was included. This delay depends on the (Euclidean) distance between the pre- and the post-synaptic neuron. (Girard et al., 2001) found a median conduction velocity of 0.3 m/s for the upper layers, and 1 m/s for the lower layers of V1 in macaque monkeys. Thus, we sampled the velocity for each excitatory synapse in layer 2/3 from a Gaussian distribution with mean 0.3 m/s and standard deviation 0.5 m/s (with enforced lower and upper bounds of 0.05 and 5 m/s, respectively). For the other layers, the conduction velocities were drawn from a Gaussian distribution with mean 1 m/s and standard deviation 0.9 m/s (with same bounds as before). Due to myelination, conduction velocities of inhibitory fibers are generally higher than for excitatory cells (Thomson et al., 2002). Lacking exact measurements in the literature for all inhibitory cells, we sampled the velocities from a distribution with mean and standard deviation twice as high as for excitatory neurons in the deep layers (the enforced upper bound was set to 10 m/s).

Synaptic conductances A spike, arriving at a synapse, causes a change in the synaptic conductance in the post-synaptic neuron. The dynamic of the conductance depends on the receptor kinetics. Each excitatory synapse in our model contains α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptors having relatively fast kinetics (modeled as exponential decay with time constant $\tau_{\mathrm{AMPA}} = 5$ ms, reversal potential 0 mV). A fraction f_{NMDA} of all excitatory synapses contain additionally relatively slow, post-synaptic voltage dependent N-methyl-D-aspartate (NMDA) receptors ($\tau_{\mathrm{NMDA}} = 150$ ms, reversal potential 0 mV,(Gerstner and Kistler, 2002; Dayan and Abbott, 2001)), and therefore exhibit a superposition of conductance kinetics. The ratio of NMDA to AMPA receptors in a synapse was drawn from a Gaussian distribution with mean $\mu_{\mathrm{NMDA/AMPA}} = 0.47$ and standard deviation $\sigma = 0.2\mu_{\mathrm{NMDA/AMPA}}$ (Myme et al., 2003).

Analogously, the inhibitory synaptic synapses were modeled as a mixture of $\mathrm{GABA_A}$ and $\mathrm{GABA_B}$ receptors. Whereas the $\mathrm{GABA_A}$ kinetic was again modeled as arelatively fast exponential decay ($\tau_{\mathrm{GABA}_A} = 5$ ms, reversal potential -70 mV), the conductance kinetic of the $\mathrm{GABA_B}$ receptors was implemented according to a model proposed by (Destexhe et al., 1994) with parameter values taken from (Thomson and Destexhe, 1999) (reversal potential -90 mV). The $\mathrm{GABA_B}$-to-$\mathrm{GABA_A}$-ratio of an individual inhibitory synapse was drawn from an uniform distribution between zero and a maximum ratio $m_{\mathrm{inh}} = 0.3$.

Orientation map generation It is well established that orientation preference and other features (such as visual field position, ocular dominance, or direction preference) form intertwined maps, where neighboring neurons tend to respond to similar features (Hubel and Wiesel, 1977; Obermayer and Blasdel, 1993).

We employed Kohonen's Self-Organizing Map algorithm (Kohonen, 1982) to create orientation maps across the cortical surface. An orientation attribute was necessary for each neuron for defining thalamic inputs, as well as for preferred orientation dependent patchy lateral long-range connections. The algorithm has been used to generate feature maps, which resembled cortical measured feature maps in their overall appearance, as well as e.g. the occurrence of pinwheels (Obermayer et al., 1990, 1992; Obermayer and Blasdel, 1993; Erwin et al., 1995;

Brockmann et al., 1997). Basically, the Kohonen's Self-Organizing Map algorithm tries to map a low-dimensional manifold (a horizontal sheet of neurons) to a high-dimensional feature space, while ensuring that neighboring points on the manifold exhibit similar feature preference. Let $\mathbf{z} = (x, y, q\cos(2\phi), q\sin(2\phi))^T$ define a feature vector, where $0 \leq x, y < k$ are the positions in visual space, $0 \leq q < 1$ is the orientation preference (or tuning strength) and $0 \leq \phi < \pi$ is the preferred orientation. We did not model ocular dominance because our V1 model received input only from one retina. If one uses the low-dimensional variant of the learning rule (Erwin et al., 1995; Obermayer and Blasdel, 1993), one attributes to each point on the manifold, i.e. each neuron having cortical 2D surface coordinates $u = (u_1, u_2)^T$, its current "optimal" feature vector $\mathbf{w}(u)$. Relations between neurons u and v are enforced by the neighborhood function $h(u,v) = \exp\left(-\frac{|u-v|^2}{2\delta^2}\right)$. The update of the feature vector of a neuron v can be written as

$$\Delta\mathbf{w}(v) = \alpha h(u_*, v)\left(\mathbf{z} - \mathbf{w}(u_*)\right). \tag{4.2}$$

Note that in each learning step the neuron u_*, showing maximal response to the current input $\mathbf{z}$, is updated in the direction of the input, weighted by a learning rate α. Depending on the *cortical* distance to the maximally activated neuron, the preferred features of the remaining neurons will be updated to a lesser extent in the same direction (mediated by the neighborhood function). In this rule we took the maximally activated neuron to be the nearest in feature space to the current input, $u_* = \text{argmin}|\mathbf{w}(u) - \mathbf{z}|$. We sampled the input features from uniform distribution (within the above bounds). k regulates the hierarchy between different features (Obermayer et al., 1992) and was set to $k = 5$. If one starts from a retinotopic initial condition, a high value for k ensures that cortical position corresponds to visual space in an approximate one-to-one map. The characteristic length scale δ was set to match the experimental observed correlation length in cortical orientation maps (corresponding to the distance of neighboring pinwheel center) of $d_{\text{pin}} = 660\mu\text{m}$ (Obermayer and Blasdel, 1993). We used the approximate formula $\delta = \sqrt{k}\, d_{\text{pin}}/D/8$, where D denotes the lateral extent of our V1 model.

4.2.2.2 Input model

The electrophysiological recordings were done during presentation of natural movies. Although our modeling effort was concentrated on the V1 model, we needed a sufficiently realistic transformation of movie stimulus to (V1 input) spike trains. Therefore, the retina and the lateral geniculate nucleus (LGN) were modeled, according to (Dong and Atick, 1995), as a spatio-temporal filter bank with nonlinearities, which seems to be a good compromise between simplicity and realism (Gazeres et al., 1998). The filter bank converted time varying input signals on the retina, such as movies, into firing rates of LGN neurons. Note that this feed-forward, rate-based model neglects any feedback connections from V1 to

LGN (Callaway, 1998). Moreover, we neglected that the ganglion cells typically react to color opponency rather than to pure luminance differences (Perry et al., 1984). Thus, the color movie was converted to a gray scale movie.

Retina model The 2-dimensional retinal inputs (movie frames) were filtered by "Mexican hat" (difference of Gaussians) spatial filters (Rodieck, 1965; Enroth-Cugell and Robson, 1966; Dong and Atick, 1995). Filter sizes (describing the receptive fields of ganglion cells) were adapted to the geometry of parvocellular cells of macaque, where the standard deviations of the Gaussian for center and surround were estimated to be $\sigma_{\text{center}} = (0.0177^\circ + 0.00196\epsilon)$ and $\sigma_{\text{surround}} \approx 6.67\sigma_{\text{center}}$ at eccentricity ϵ, respectively (in visual degrees; estimated from Figure 4 a and b in (Croner and Kaplan, 1995)). After the convolution of the stimulus luminance portrait with these kernels (yielding S_{center} and S_{surround}), the response of a retinal ON-cell at visual field position r can be described by

$$R_{\text{ON}}(r) = C(r)\left[S_{\text{center}}(r) - \omega S_{\text{surround}}(r)\right]_{+} \tag{4.3}$$

Following (Croner and Kaplan, 1995), we set the ratio of center to surround $\omega = 0.642$. The positive part of the center and surround interaction (indicated by the brackets $[\ldots]_{+}$) was assigned to the response of an ON-cell and, analogously, the absolute value of the negative part to the response of an OFF-cell (Dong and Atick, 1995). For simplicity we assumed that the origins of the center and surround summation fields are identical, although a recent study suggests that there might be an offset between them (Conway and Livingstone, 2006).

Applying the Difference-of-Gaussians model to the luminance of a stimulus results in a quantity called "contrast gain" (Croner and Kaplan, 1995; Enroth-Cugell and Robson, 1966; Rodieck, 1965). To calculate the firing rate of ganglion cells, one has to multiply the "contrast gain" with the *local* contrast $C(r)$ (as done in Eq. 4.3), if one neglects non-linear saturation in the high contrast regime, which is typically not reached for the natural stimuli we used here. Locality is important because the concept of a global contrast, easily defined for full-field grating stimuli commonly used in experiments, is not applicable for real world images and movies (Tadmor and Tolhurst, 2000). Following (Tadmor and Tolhurst, 2000) we estimated the local contrast using the same kernels as

$$C(r) = \frac{|S_{\text{center}}(r) - S_{\text{surround}}(r)|}{S_{\text{center}}(r) + S_{\text{surround}}(r)} \tag{4.4}$$

where we additionally set the contrast to be zero in the case of darkness. Note that applying Eq. 4.4 results in a response $R_{\text{ON}}(r)$ that is sparser than for a constant global contrast, since the response is now quadratic in the center and surround luminance difference (see Eq. 4.3).

LGN model The retinal output was filtered by the LGN model using a temporal kernel. The temporal kernel combines a phasic (taken from (Dong and Atick, 1995)) and a tonic component (as in (Gazeres et al., 1998)), i.e. $k_{\mathrm{LGN}} = k_{\mathrm{phasic}} + k_{\mathrm{tonic}}$. It is for non-negative times

$$k_{\mathrm{phasic}}(t) = t\,(1 - \pi w_c t)\,\exp(-2\pi w_c t) \tag{4.5}$$

and

$$k_{\mathrm{tonic}}(t) = A\,\exp(-t/\tau)/\tau. \tag{4.6}$$

Parameter $A = 0.3$ is the fraction of tonic activation (with respect to the peak firing rate) for a given stimulus, integrated over a time window of $\tau = 15$ ms. Parameter $w_c = 5.5\ \mathrm{s}^{-1}$ defines the shape of the phasic kernel (Dong and Atick, 1995).

The positive parts and the absolute values of the negative parts of the temporal convolutions were assigned to non-lagged and lagged cells, respectively. Altogether, there are four different time-varying rate outputs, i.e. that of any combination of non-lagged or lagged cells in the LGN with either ON- or OFF-cells from the retina (Dong and Atick, 1995). Following (Gazeres et al., 1998) a so-called "switching Gamma renewal process" was used to convert these time-varying rates to spike trains. This process, which was suggested to fit experimental data from cat LGN X-cells (Gazeres et al., 1998), adopts a higher spike time regularity for high input rates (≥ 30 Hz; regularity parameter $r = 5$) and switches to a Poisson process for low rates (< 30 Hz). The spontaneous background activity of each LGN neuron was set to a low value of 0.15 Hz. The peak LGN spike rate $f_{\max}$ was adjusted to achieve a mean firing rate of about 7 Hz under movie stimulation, when the four input channels were combined. The 7 Hz mean rate was estimated from our electrophysiological data from macaque monkey. Applying a typical 50 s movie section, we found that a mean rate of 7 Hz was achieved for $f_{\max} = 250$ Hz. The peak response would be evoked by a dot of highest contrast filling the center region of a ganglion cell with optimal duration. This value is in good agreement with (Gazeres et al., 1998), they reported peak rates in the range of 50 to 400 Hz.

Input connectivity to V1 The visual field is retinotopically arranged on the cortical surface. However, while there exists only one retinal ganglion cell per LGN cell corresponding to the same visual field position at all eccentricities in macaque, there is a considerable magnification in density of cortical neurons in V1 per degree of visual field (Schein and de Monasterio, 1987; Tootell et al., 1982). Comparing several earlier studies, (Schein and de Monasterio, 1987) estimated the cortical magnification factor (CMF) at eccentricity ϵ to be (in mm cortex per degree of visual field)

$$\mathrm{CMF} = \frac{12.2}{\epsilon + 0.94}. \tag{4.7}$$

This definition of the cortical magnification factor (Eq. 4.7) is very convenient: for a fixed eccentricity and distance between adjacent neurons (grid spacing), one can calculate the lateral extent of the network needed to cover a given visual field size. Note, however, that this estimate is only useful when the lateral exten of the network model can be regarded as small compared to the variation in lateral cell density.

LGN neurons, belonging to the parvocellular pathway, typically project to layer 4Cβ of V1. There is still an ongoing debate, to which extent oriented input shapes the orientation selectivity of neurons in the primary visual cortex or to what extent local cortical processing is involved (see (Teich and Qian, 2006) for a review). It seems that in macaques, orientation selectivity is thought to arise from the interaction of cells with gradually shifted input characteristics across the sublamina of the layer 4C (Lund et al., 2003; Callaway, 1998), whereas the inputs to a single cell might not be oriented in macaques as suggested for the cat (Hubel and Wiesel, 1977). However, since we did not model sublamina, we simplified the circuitry by, nevertheless, assuming that input connections to each neuron generate orientation tuning. This allows the definition of orientation maps in a straightforward "hard-wired" manner in our model.

We employed Kohonen's Self-Organizing Map algorithm (Kohonen, 1982) to create orientation maps across the cortical surface, which is known to generate orientation maps with good correspondence to V1 orientation maps (Obermayer et al., 1990, 1992; Obermayer and Blasdel, 1993; Erwin et al., 1995; Brockmann et al., 1997). See Fig. 4.1 C for a typical orientation map generated by this algorithm (see Section 4.2.2.1 for details on the implementation of the algorithm).

Based on the generated orientation preferences for each cortical position, the thalamic input connection probability to a cell in the circuit could thus be modeled as an oriented Gabor function, i.e. a 2-dimensional Gaussian multiplied by a cosine function. The absolute value of the Gabor function corresponds to the connection probability of LGN neurons with a cortical cell positioned at the cortical equivalent position of the origin of the Gabor patch in the visual field. Positive and negative regions correspond to the connection probabilities of LGN ON- and OFF-response cells, respectively. Lagged and non-lagged cells connected equally likely to cortical cells. Following (Troyer et al., 1998), we expressed the Gabor function in parameters defining the number of sub-regions n_s, the aspect ratio of the width and the height of the Gaussian envelope a, the orientation ϕ, the offset of the cosine ψ, and the frequency of the cosine f. Given these parameters, one calculates the standard deviation of the Gaussian envelope as (see (Troyer et al., 1998))

$$\Sigma^{\frac{1}{2}} = \frac{1}{4 \cdot 2.448\, f} \begin{pmatrix} n_s & 0 \\ 0 & a \end{pmatrix} \tag{4.8}$$

while using coordinates rotated by ϕ. The advantage of using these parameters is that the frequency defines implicitly the size of the Gabor patch, while the number of sub-regions is kept constant. Therefore, the receptive fields of macaque V1, which are much smaller than those of the cat, can be easily included in this framework. We used data from (Bredfeldt and Ringach, 2002) and chose the frequency f from a Gaussian distribution with a mean of 3.7deg^{-1} and a standard deviation of 2.1deg^{-1} (with an enforced minimum of 0.7deg^{-1} and maximum of 8.0deg^{-1}). The number of sub-regions n_s and phase shifts ψ were drawn from uniform distributions with ranges of $(1.85, 2.65)$ and $(0, 2\pi)$, respectively (experimental values from cat as in (Troyer et al., 1998)).

To incorporate the smooth maps of preferred orientation ϕ and orientation preference q depending on cortical position u (see Section 4.2.2.1), we set $\phi = \phi(u)$ and the aspect ratio to $a(u) = (a_{\text{max}} - a_{\text{min}})\, q(u) + a_{\text{min}}$. We used values reported by (Troyer et al., 1998) for the bounds $a_{\text{min}} = 3.8$ and $a_{\text{max}} = 4.54$ for excitatory neurons, and for the generally less well tuned inhibitory neurons $a_{\text{min}} = 1.4$ and $a_{\text{max}} = 2.0$.

The overall connection probability, defined by the Gabor functions, was scaled to achieve an average number of 24 input synapses for both excitatory and inhibitory neurons, which is the estimated number of parvocellular afferent connection per cortical neuron in layer 4C of macaques (Peters et al., 1994). There is evidence that layer 6 receives occasional collaterals of the LGN input to layer 4 (Callaway, 1998). Thus, we set the connection probability to excitatory neurons in layer 5 (comprising layer 5 and layer 6 in our model) to 20% of that of the input to layer 4. These values are in good agreement with the data from (Binzegger et al., 2004) estimated from cat. In macaques, layer 2/3 receives only koniocellular input (Callaway, 1998). As we omitted the koniocellular pathway in our model, layer 2/3 did not receive any thalamic input.

Due to finite conduction velocities of the fibers, signals from the retina reach V1 with a characteristic delay of about 30 ms (Maunsell et al., 1999). We sampled the delay of the LGN input synapses from a Gaussian distribution with mean 31 ms and standard deviation 5 ms (and additionally enforced delays below 24 ms and above 50 ms to a value uniformly in the latter range). These values were taken from Figure 3 of (Maunsell et al., 1999).

Top-down connections In addition to the thalamic input, V1 neurons receive multiple feedback connections from extra-striate cortical areas (Felleman and Essen, 1991), especially from V2, where the feedback connections are almost as numerous as the feed-forward connections (see (Sincich and Horton, 2005) for a review). Feedback projections predominantly project to targets in the upper layers but also to layer 5 (Rockland and Virga, 1989; Sincich and Horton, 2005).

Because our model is restricted to V1 and we do not have any recordings from V2 available, we decided to not include any top-down input stream explicitly.

However, implicitly additional input to V1 neurons is included by modeling the "high conductance state" of each neuron, which reflects the synaptic background input arriving from distal neurons (Destexhe et al., 2001) (see Section 4.2.2.1).

4.2.3 Comparing the V1 model to electrophysiological data

4.2.3.1 Setup of the stimulus to the model

The stimulus, presented to the V1 model during simulation, resembled the one presented to the monkeys. We used a 10 s fragment of one of the movie segments (sw21) shown during the electrophysiological recordings as input movie for the model. However, modeling the whole $10° \times 7°$ visual field was not feasible because of computational costs. Therefore, we trimmed the movie frames to a smaller size, covering $3 \times 3°$ of the visual field. The center of the extracted region was aligned to the center of a receptive field of one of the electrodes (channel 7) of a particular session ("d04nm1"). Since the diameter of the receptive field of that electrode was experimentally determined to be 1.2°, the reduced stimulus should at least contain all direct input information available for neurons recorded by that electrode. On the retina, this receptive field was centered at $(0.69°, -2.39°)$ eccentricity relative to the fovea. In the model, we set the eccentricity, nevertheless, to 5°, since otherwise the lateral extent (and, therefore, the amount of neurons in the model) per visual degree would be prohibitively large (compare to Eq. 4.7). At 5° eccentricity a V1 model covering $2.4° \times 2.4°$ has a lateral extent of $5 \times 5\,\mathrm{mm}^2$ cortical surface and neurons are positioned on a virtual grid of size $62 \times 62 \times 9$ if one assumes a lateral grid spacing of 80 μm. Vertically, the grid spacing corresponds to 200 μm (see Method Section 4.2.2.1). The visual field covered by the V1 model is somewhat smaller than the stimulus to avoid boundary effects in the input connectivity. For analogous reasons, the LGN neurons were set to cover an intermediate area of $2.8° \times 2.8°$ (77×77 grid).

4.2.4 Estimating the relative strength of the thalamic input

In the recorded spike trains, the mean firing rate of multiple trials (5 min duration) across monkeys and V1 electrode channels was on average 5.1 (4.8) Hz (standard deviation) during movie stimulation and 1.9 (3.3) Hz during spontaneous activity (blank screen). Thus, one could state that due to the thalamic input the mean firing rate of the circuit increases by about 3 Hz. From simultaneous extracellular recordings in LGN, we analogously find a mean firing rate of 7.1 (2.9) Hz during visual stimulation and 4.4 (2.1) Hz during absence of visual stimulation. Hence, in the LGN the movie stimulus increases the mean firing rate by about 60 % of the spontaneous activity.

We used these values to determine the synaptic input weight scale ($W_{\mathrm{In,scale}}$), i.e. the scaling factor of the peak conductances originating from LGN neurons, in

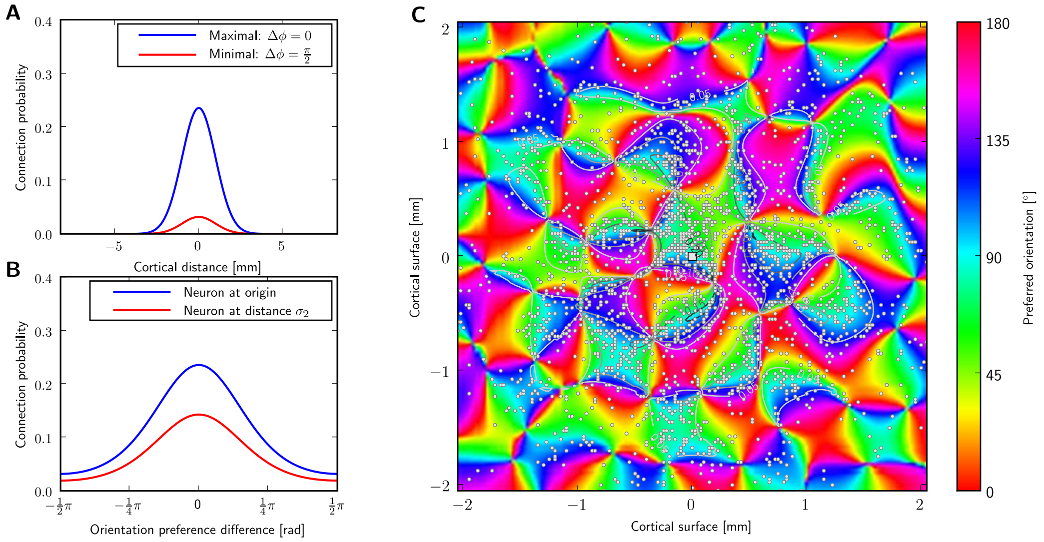

Figure 4.1: Long-range connectivity of the V1 model.
Long-range patchy connectivity of an example neuron implemented in a model circuit having $165 \times 165 \times 3$ neurons in layer 2/3 positioned on a cuboid grid with a spacing of 25 μm. (Note, these dimensions are different from that used in the simulations of the Results section, but they are used here for better visualization) **A** and **B**: Conditional probability that the neuron (marked with a white square in the center of **C**) is connected to a neuron having lateral distance r or orientation selectivity ϕ, respectively. The connection probability to a post-synaptic neuron at zero lateral distance and same orientation preference was scaled to experimental data ($\approx 0.24\%$; (Thomson et al., 2002)). Blue and red curves show the connection probabilities for neurons which have aligned or orthogonal preferred orientation to the pre-synaptic neuron, respectively. **C**: Connections established according to the probability distributions for a pre-synaptic neuron in the origin of the circuit (white square). Small white dots represent lateral positions of post-synaptic neurons. Colors code for orientation tuning of a neuron (generated by a self-organizing map). The conditional connection probabilities are indicated by contour lines. One notes that the connection probability rises for regions with similar orientation as the pre-synaptic neuron (about 90°) thereby generating a patchy appearance. Only the orientated (long-range) part of Eq. 4.1 (second term) is used for establishing connections in this example plot. The orientation map additionally determines the orientation of thalamic input connections (see Methods section).

Table 4.1: Parameters investigated in their optimization potentials together with their standard value $\mathbf{p}_{st}$. Standard values for 5 of the parameters could be extracted from the literature. If no reference is given, the standard value was chosen heuristically.

	Parameter	std. value	Reference
p_1	noise level scale	1.0	(Destexhe et al., 2001)
p_2	fraction of synapses with NMDA	0.9	
p_3	NMDA-to-AMPA ratio	0.47	(Myme et al., 2003)
p_4	width of inh. connections (μm)	150	(Lund et al., 2003))
p_5	max. fraction of $\mathrm{GABA_B}$ conductance	0.3	
p_6	inh. to exc. connections weight scale	1.0	(Thomson et al., 2002)
p_7	exc. to inh. connections weight scale	1.0	(Thomson et al., 2002)
p_8	long-range weight scale	1.0	

the following manner: In the absence of *all* inter-cortical connections, the weight scaling factor of the input stream was set to a value achieving closest match to a given target mean firing rate r_{target} in each neuron population (minimal Euclidean distance). Assuming that the main input drive to V1 (during visual stimulation) is from the thalamus, we set the target mean rate for the circuit to 2 Hz, which roughly corresponds to the activity increase seen during visual stimulation in our experimental data.

4.2.5 Characterization of the spiking statistics

We defined 10 statistical measures to characterize the spike train statistics. Each statistical measure was calculated on all available spike trains of the experimental session and the simulation, respectively. If not mentioned otherwise, we calculated the spike statistics using time windows with a length of 2 s. The time windows were overlapping with a step size of 0.2 s. The 10 chosen statistical measures are the following:

Firing rate distribution The distribution of mean firing rates in windows of 2 s duration.

Fano factor over different time scales To characterize spike count variability over a range of time scales, we estimated Fano factors of the spike count-distribution. These Fano factors are defined as the ratio of the variance to the mean of the spike-count distribution estimated over time windows of fixed length. We compared the average Fano factors for a range of time window durations (from $10^{-2.5}$ s to $10^{0.4}$ s).

Population Fano factors For different time windows (from $10^{-2.5}$ s to $10^{0.4}$ s), we calculated the temporal mean spike count of the whole population of neurons and its variance. From these values we calculated the Fano factors. These Fano factors describe the variability of the population firing rate.

Burst rate distribution Burst events can be defined as having at least 2 spikes with an average ISI of at most 5 ms (Lisman, 1997). The burst rate is the frequency of burst occurrence estimated within a period of time. The burst rate distribution describes the probability of different burst frequencies and was again estimated on time windows of 2 s duration.

Burst rates for different burst sizes To analyze the occurrence of larger bursts, we calculated the burst rate distribution for burst events having at least n spikes with an average ISI of at most 5 ms. We then took the average burst rates for a set of n minimal spikes ($n = 2 \ldots 25$) as another statistical measure.

Inter-spike interval (ISI) distribution This is the distribution of the intervals between two consecutive occurring spikes of one neuron. The distribution was estimated on the full length of the spike response to include longer intervals in the analysis. As before, we took all available trials and neurons to calculate a population statistic.

2-ISI distributions For neurons, placed in a recurrent network, the generation of a spike might depend in a systematic way on the relative timing of the previous spikes. To compare such dependencies, we estimated the 2-ISI distributions. We define the 2-dimensional 2-ISI distribution $p(\tau_1, \tau_2)$ as the probability of occurrence of two sequential spike intervals of lengths τ_1 and τ_2. These distributions were estimated on the full length of the spike response.

ISI coefficient of variation distribution Another estimation of the variability of spike trains is the coefficient of variation of the inter-spike interval distribution (CV(ISI). The coefficient is defined as the ratio of the standard deviation to the mean of the ISI-distribution. The CV(ISI) was estimated on the available ISIs of each window of 2 s duration, and the resulting population statistics were taken for comparison. If there were less then 3 spikes in a given time window, we set the value of the CV(ISI) to zero.

Neuron synchronization We defined synchronization as the mean cross-correlogram of spike activity between two neurons, that is the cross-covariance of the binned spike trains divided by the square root of the product of the variances (with a temporal bin size of 25 ms). We averaged over all spike clusters available in the sorted electrophysiological recorded spike trains and over 20000 randomly drawn neuron pairs of the model circuit.

Spike time correlation as a function of distance Lateral decorrelation of neural activity might be another important prerequisite for computational function of a neural circuit. Thus, we calculated the cross-correlogram as a function of the Euclidean distance between two model neurons or recording sites. The temporal bin size was 50 ms and the spatial bin size was 500 μm. The correlation was averaged over time lags from -0.5 s to 0.5 s.

4.2.6 Evaluating the deviation between model response and *in vivo* recordings

To compare the firing regime of the model with that of the *in vivo* recordings, we evaluated the discrepancy between a set of 10 statistical features calculated from the model response and the recorded spike trains (see Section 4.2.5). After estimating a statistical feature on the experimental data and the model response, their deviation was calculated using Kulback-Leibler divergence, or by calculated the mean squared error, depending whether the features resulted in an estimated probability distribution or not, respectively. This deviation was normalized by the average deviation seen in this features if tested between any two experimental sessions (different monkeys or different movie stimulus). We call this experimental data weighted deviation the “normalized deviation” (ND) for each statistical features. We report the normalized deviation averaged across all statistical features as a measure for the goodness of fit, and abbreviate it in the following with “mean normalized deviation” (MND). Note that by construction a MND value of 1 indicates that the deviation between the model response and the *in vivo* data (average over all sessions) equals (on average over the 10 statistical features) the average deviation between individual experimental sessions. We used only one model random seed for the evaluation of the fitting error for each parameter setting to reduce computational costs. To compensate for a lack of synaptic drive due to a much smaller neuron density in the model compared to reality, we introduce two scaling parameters $W_{\mathrm{In,scale}}$ and W_{scale}. The $W_{\mathrm{In,scale}}$ parameter, a multiplicative factor applied to all weights of the input connections, was set by a heuristic approach to approximately match the input strength observed in the experiments (described in Methods Section 4.2.4). The second scaling parameter, the weight scale parameter W_{scale}, accounts for the recurrent synaptic drive adjustments and is a multiplicative factor applied to all recurrent weights. As this parameter is inherent to the model design, it cannot be constrained by literature values. Therefore, to estimate the weight scale parameter W_{scale}, we used the value that minimizes the deviation of the model firing response statistics to the “statistical fingerprint” of the firing regime of the *in vivo* recordings. To measure its deviation, we used the MND as described above. We restricted the analysis on the response of excitatory neurons only, because we expect that due to the generally larger size of excitatory neurons, the experimental recordings were strongly biased to record spikes originating from excitatory cells.

4.2.7 Simulation techniques

All simulations were performed using the PCSIM simulation environment (Pecevski et al., 2009). It takes about 5 hours on a quad core machine (2664 MHz) to simulate the described model for 10 seconds of biological time (depending on the mean firing rate). All simulations were performed in a distributed fashion on a cluster of 30 such quad core machines.

4.3 Results

We first established the “statistical fingerprint” of the spiking activity of the primary visual cortex (V1) under naturalistic stimulus conditions *in vivo*. The extracted statistical features provided the grounds for comparison with the simulated firing response of a detailed circuit model. Since we hypothesize that V1 works in a characteristic firing regime favorable for its ongoing computations, we were particular interested in features possibly characterizing a computational advantageous regime. For instance, such a regime might consist of highly irregular firing and low correlations between neurons (Brunel, 2000; Legenstein and Maass, 2007). We therefore extracted 10 salient statistical features, which are sensitive to various aspects of the spiking response, such as response strength, response variability, spike correlations, bursting behavior, and the possible usage of spiking codes with non-linear dependencies on consecutive spike intervals (see Section 4.2.5, for exact definitions).

4.3.1 Statistical characterization of the spike response to movies in monkey area V1

We first analyzed electrophysiological recordings from V1 of anesthetized macaque monkeys during stimulation with natural movies. The data comprised spike responses measured in 10 sessions (from 4 anesthetized macaque monkeys), each with 12-40 repeated representations of a movie of 3.5 to 6 minute length. In Fig. 4.2 D and E, typical spiking responses of selected neurons are depicted. We characterized the firing statistics of this experimental data using a set of 10 statistical features (see Fig. 4.3).

We found that spike responses of V1 under naturalistic stimuli conditions were typically highly variable over time and moderately low correlated between different neurons having a smooth fall-off for long time lags. Firing rate and burst rate distributions followed exponential distributions, burst size frequencies and ISI distributions exhibited a power-law structure. This described general picture is consistent with previous published values. In detail, the exponential distributed firing rates (Fig. 4.3 A) exhibited exponents varying between monkeys and experimental trials in the range of -2.35 s to -0.23 s (mean -0.81 s, standard deviation (SD) 0.62 s). The overall mean firing rate of the experimental data averaged over the different sessions was 5.06 ± 0.75 Hz (mean±SD). The exponential distribution of firing rates is consistent with results from the V1 of cats (Baddeley et al., 1997). Spike train variability is generally high. We tested for the variability in the spiking response using the distribution of Fano factors of individual neurons for multiple time scales (see e.g. (Teich et al., 1997)), and the Fano factor of the network population spike response, to measure the response variability of the population code. For individual neurons (Fig. 4.3 B) the Fano factor approached 1 on average for small window sizes in the order of 10 milliseconds, indicating a Poisson process with stationary rates. On larger time scales, however, the Fano factor increased. This increase in variability could reflect the internal dynamics, but might be partly induced by the movie stimulus, which mean brightness often varies on a time scale in the order of seconds. The population Fano factor (Fig. 4.3 C), measuring the response variability of the neuron population, showed a similar time window dependence as the Fano factor of individual neurons. However, the absolute value of the population Fano factor was markedly smaller, indicating that the population response was less variable over time on short time scales. On a longer time scale, however, the Fano factor of the population response increased, indicating that the neurons in the recorded population tend be active or silent together. This might hint at population burst-like activity, also evident when examining the concrete spike trains in the recordings (see (Montemurro et al., 2008) for a discussion of how these clusters of spikes relate to the local field potential fluctuations in the same data). Another evidence for the high variability in the *in vivo* data is given by the coefficient of variation of the inter-spike intervals (CV(ISI)), plotted in Fig. 4.3 H. High probabilities were typically found for CV(ISI) values above one, indicating a high variability in the spike response. Such a high peak value is consistent with previously published data (Softky and Koch, 1993; Holt et al., 1996; Shadlen and Newsome, 1998; Stevens and Zador, 1998). Spike bursts, i.e. abrupt events of high spiking activity, have been suggested to be an important aspect of neuronal coding of information. For instance, bursts might convey additional and independent information about the sensory inputs (Cattaneo et al., 1981; Lisman, 1997). Thus we included two statistics to measure the occurrence of bursts in neuron spike trains, a feature rarely examined in the literature. Fig. 4.3 D shows the burst rate distribution, measuring the frequency of spiking events having at least two spikes within an (average) inter-spike interval of 5 ms. Fig. 4.3 E shows the average burst rate for different sizes of bursts (see Section 4.2.5 for exact definitions). Qualitatively, the burst rate distributions of different sessions looked similar, having exponential distributions. The exponent varied in the range of -5.07 s to -3.18 s (mean -4.54 s, SD 0.56 s). The average bursts rate of all experimental data was 0.51 ± 0.17 Hz (mean ± SD). However, in some sessions there was a deviation from the exponential distribution and higher burst rates occurred more often than expected. Fig. 4.3 E shows that the burst rate as a function of the burst size can be described by a power law (a straight line in a log-log plot). We found exponents in the range of -3.52 to -2.29 (mean -2.88, SD 0.43). Individual spike precision might be important for information coding. It is conceivable that a certain inter-spike interval (ISI) distribution might be characteristic for the firing regime of the cortex. Inter-spike interval (ISI) distributions (Fig. 4.3 F) were very similar for different monkeys and different sessions. There was a high probability for the occurrence of long inter-spike intervals. The distribution of inter-spike intervals similarly followed a power law with an average exponent of -1.20 ± 0.19 (range from -1.47 to -0.98). Additionally, to account for any local temporal correlations in the spike timings, we also estimated the 2-ISI distribution, which is a 2-dimensional distribution of the joint event of one inter-spike interval and the immediately following inter-spike interval (Fig. 4.7 K). The ISI distribution for the following ISI when conditioned on a very short current ISI had a similar power-law shape as the marginal ISI distribution (Fig. 4.3 G), although the occurrence of a short ISI following a short ISI was more likely. Similar to the full ISI distribution, we found a relatively low variability across sessions and monkeys. In general, we found that two neurons in V1 were on average correlated for lags up to about 250 ms having moderately low peak correlations. Other studies also reported low (signal) correlations between pairs of neurons for naturalistic stimuli (Yen et al., 2007; Reich et al., 2001) and even lower correlations in awake animals (Vinje and Gallant, 2000). To be able to better compare our data to the literature, we calculated the shift-corrected cross-correlogram (Bair et al., 2001; Kohn and Smith, 2005; Smith and Kohn, 2008) and the noise correlations (r_{sc}, e.g. (Ecker et al., 2010)) and found that the correlation structure in our data agreed very well with that of Smith and Kohn (see Fig. 4.8). The strength of correlations however depended on the monkey and movie stimuli (Fig. 4.8 A and C). The mean value was $r_{sc} = 0.26 \pm 0.03$ (mean ± SEM).

We further analyzed the cross-correlation for pairs of neurons as a function of their distance (Fig. 4.3 J, and Fig. 4.8 A). In agreement with others (Smith and Kohn, 2008), the cross-correlation was higher for neurons (clusters) recorded by the same electrode and decreased for longer electrode distances, where the correlation remained on a low level. In summary, we have computed a set of statistical features characterizing the “statistical fingerprint” of the spiking activity under semi-natural movie stimulus condition *in vivo*. Certain features of the obtained “fingerprint”, namely the high variability of inter-spike intervals, low cross-correlation, and the power-law distributions of burst events suggest that the V1 during movie stimulation might indeed reach an operating state, which is favorable for recurrent neural networks for performing computational tasks. The results presented here agree in general with published literature. However, since we characterized the firing regime not only by a small set of mean values but instead by 10 different functions (or estimated probability distributions), we were able to quantify the deviation of the firing regime of a simulated model from that exhibited *in vivo* in great detail. Moreover, the dataset provided us with the unique possibility to test the importance of physiological meaningful parameters to optimize the model response behavior to closely reach a realistic state.

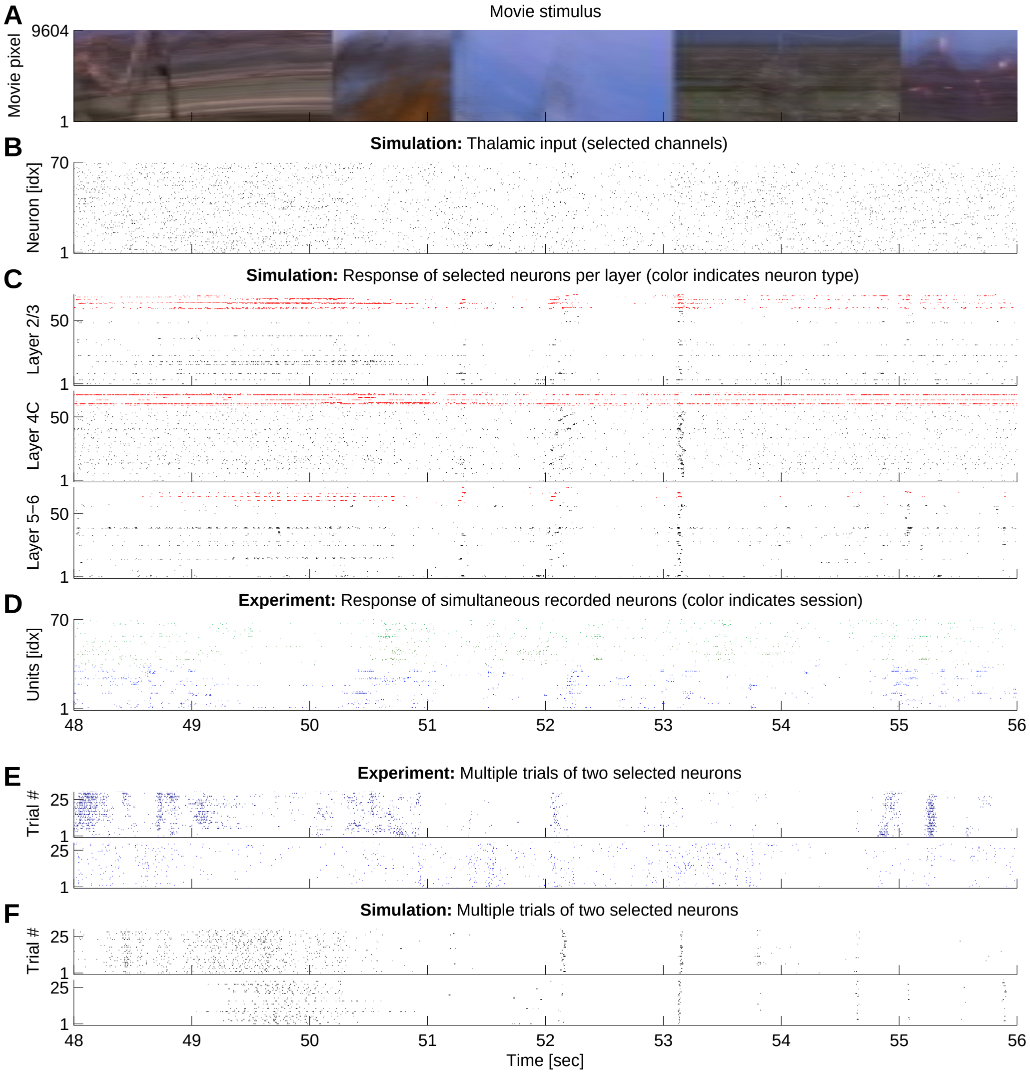

Figure 4.2: Spiking response to movie stimulus in model simulation and *in vivo*
see next page for figure caption

Figure 4.2: Spiking response to movie stimulus in model simulation and *in vivo*
A: Movie input to the model circuit in true colors (in the model we used a grayscale version of this movie). Pixels of the movie frames are lined up vertically. **B:** LGN model response to the movie in A. 70 input channels were randomly selected for plotting (in total there are 4900 LGN inputs) **C:** Spike trains elicited by neurons in the V1 model in response to the LGN output (B) are plotted in separate panels for each of the layers 2/3, 4C, and 5-6. For better visualization, 70 neurons (of 11532) are randomly chosen from each layer. Inhibitory and excitatory neurons are colored red and black, respectively. One notes a high variability in the statistical structure across neurons. **D:** Spike trains of the spike sorted experimental data in response to the same movie segment are shown. Different colors represent different sessions of the same monkey - green (blue): two trials of session d04nm1 (d04nm2). We show two trials to allow for an easier comparison of the statistical structure of the spike trains *in vivo* with the model response (C). Note that the receptive field of some electrode channels lie outside of the depicted movie region (A) (see Method section for details). **E-F:** Multiple trials of two selected neurons in experiment (E) and model (F). Note that trial-to-trial variability is comparable *in silico* and *in vivo*.

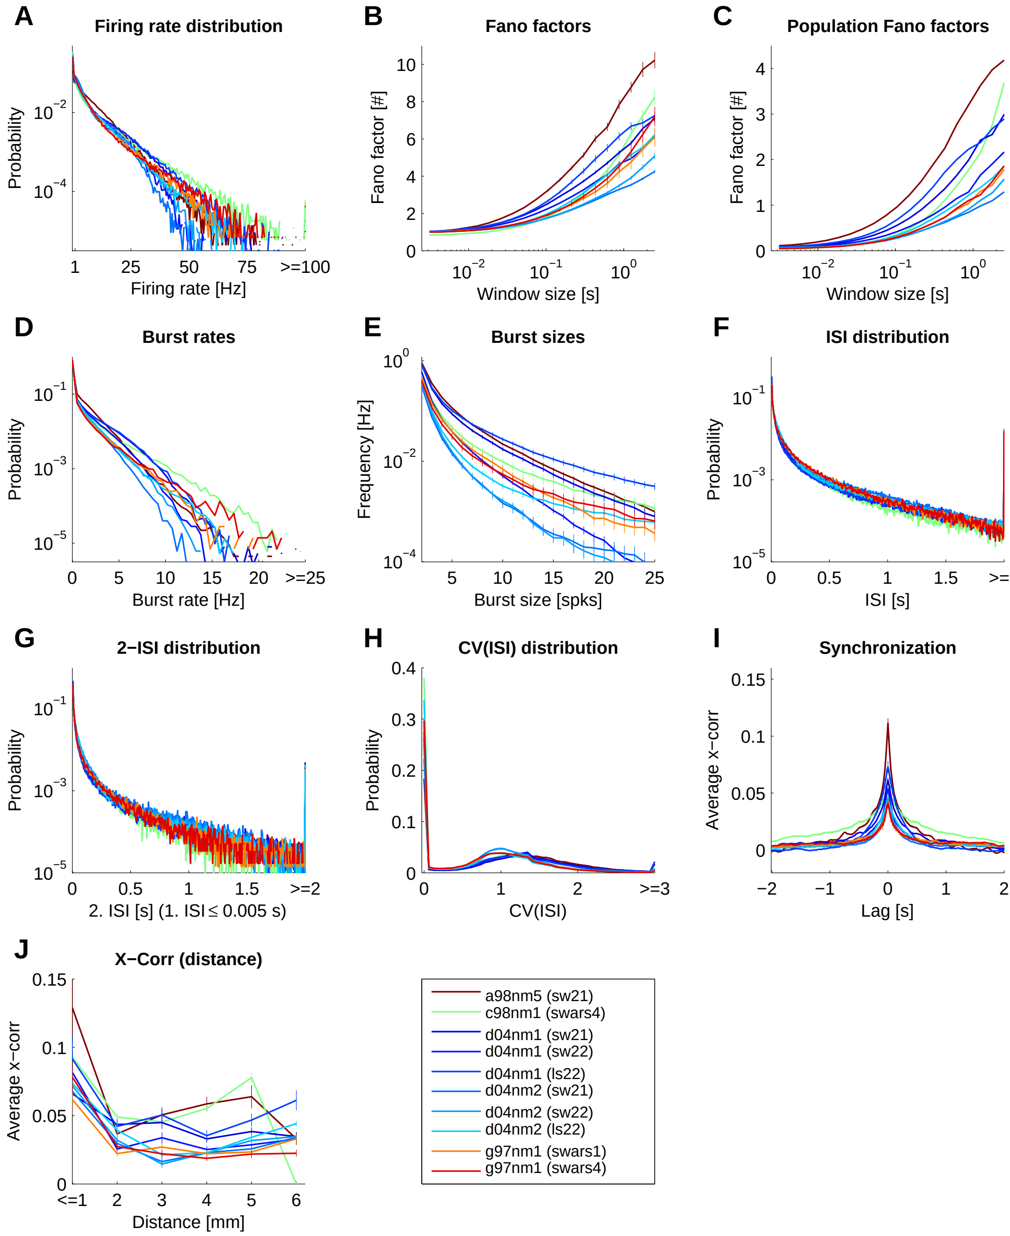

Figure 4.3: Spike statistics of the experimental data.
Each panel corresponds to a particular statistical feature (see Section 4.2.5 for exact definitions). Data from 10 experimental sessions (from 4 monkeys) during movie stimulation are plotted separately in color code (first monkey - brown, second monkey - green, third monkey - shades of blue, fourth monkey - shades of red). In the legend the first three letters of a session code indicate the animal, the second three letters the recording session. The shown segment of the presented movie (Star Wars) is indicated in brackets. All available trials are included. Values greater than the plot limits were included in the last bin (where applicable), resulting in a disproportional large probability in the last bin. Note that the 2-ISI distribution is a 2-dimensional distribution. G shows the conditional ISI distribution given that the first ISI is shorter than 5 ms.

4.3.2 Quantification of the discrepancy between the firing regime of a model for a patch of V1 and the firing state exhibited *in vivo*

Having characterized the V1 *in vivo* recordings, we proceeded with characterizing the simulated responses of the circuit model of V1 *in silico*. The V1 model was based on anatomical and physiological details of macaque monkeys and was built to model the neural activation in a 5 mm x 5 mm cortical patch of V1 (see Methods Section for a detailed description of the V1 model). We simulated the model and recorded the spiking activity in response to 10 s of a typical movie segment (sw21) that had also been used for *in vivo* recordings.

Differences in the firing regime *in silico* and *in vivo* were quantified by estimating the deviations in all statistical features. We calculated the mean normalized deviation (MND) between the model and the *in vivo* response (see Methods for definitions). Note that MND $= 1$ indicates that the deviation of the model response to the mean response over all sessions equals the mean deviation between all pairs of sessions. Our measure thus relates to the deviation among individual experimental sessions. Moreover, the MND weights the importance of each statistical feature in a manner that features showing a high variability between experimental sessions are deemed less important and those features conserved across sessions are emphasized. By setting parameters of the model to values derived from the literature (see Table 4.1) and minimizing the fitting error in respect to the overall recurrent connections weight scale, which is inherent to the model design (W_{scale}; see Methods), we found a mean normalized deviation of MND $= 1.97$, indicating that the deviation is on average about twice as high as between experimental sessions and monkeys. Since we presented complex movie stimuli, it is not clear whether the firing regime of the model was indeed generated by internal dynamics or was instead solely induced by the statistics of the input. To test the possibility of induced dynamics, we calculated the MND on the input spike trains generated by the LGN model (omitting the now meaningless lateral cross-correlation feature), and found a value of MND $= 2.5$. This value is considerably higher than for the model response. We repeated the statistical analysis for the model network after abolishing all recurrent connections, leaving only the input connections intact. By varying the strength of the synaptic input connections, we found a minimal value of MND $= 4.37$. Thus the fit of the firing statistics was much worse than with intact recurrent connections, implicating that the recurrent dynamics of the network indeed shaped the firing response. We concluded that by simply optimizing an overall scale parameter (W_{scale}), the model dynamics shaped its statistical response properties in direction of that of the *in vivo* response. However, deviations from the realistic firing regime *in vivo* were still considerable.

4.3.3 Improvement of the firing regime when optimizing the model

Can the firing regime of the model be adjusted by physiological meaningful parameters to improve the fit to the *in vivo* data? Finding such parameters would shed light on parameters that exert control over certain statistics. Therefore we chose 8 physiological meaningful parameters (see Table 4.1 for an overview), which we believed to influence the firing dynamics. We then optimized the model in respect to each parameter and evaluated each parameter's ability to improve the discrepancy between model and *in vivo* recordings. Unfortunately, simultaneous optimization of multiple parameters was computational prohibitive. Therefore, we varied each parameter individually around the "standard" parameter values taken from the literature ($\mathbf{p}_{st}$), which we used to establish the initial "fingerprint" of the model's response (see above). Since the optimal W_{scale} might change during the variation of a parameter, we additionally varied W_{scale} resulting in two-dimensional landscapes (see Fig. 4.4, see Table 4.2 for a summary of the quantitative results of the optimization). The effects of parameter optimizations on the improvement of each statistical feature are analyzed in Fig. 4.5. We first chose a parameter varying the background noise strength (parameter p_1). The background noise strength implicitly regulates the strength of how neurons not modeled in the circuit affect the modeled neurons (see Methods section). When varying this parameter, we did not find a strong dependence on the quality of the fit (Fig. 4.4 A), suggesting that this background noise strength was of minor importance. Although varying the noise strength improved some individual statistical features in respect to the literature values (such as Fano factors, burst sizes and ISI distributions, see Fig. 4.5 A), the effect was typically below 10 % ND improvement. Indeed, even if we disabled the background noise, the fit to all statistical features simultaneously was only compromised by a negligible decrease of the MND of 3 % (Fig. 4.5 B). This suggests that our network was already big enough to explicitly provide realistic synaptic background inputs to any neuron. In our model, the lateral width of inhibitory neurons was relatively small (standard deviation 150 μm, see Methods). We tested whether the fit could be improved by varying the lateral spread of the inhibitory connections (p_4). However, this was not the case: a good range for this parameter lied between 150 μm and 250 μm, depending on the overall strength of the synapses (W_{scale}, see Fig. 4.4 D). Although, the ND of burst sizes and Fano factors could be markedly improved (see Fig. 4.5 A, and Fig. 4.6), these features had only a small influence on the MND because their variance between experimental sessions was high and, moreover, they were already well fitted by a model with parameters set to standard values (compare to Fig. 4.6 right marginal plot). In consequence, the MND could only be improved by about 5 % by optimizing the lateral connection width of inhibitory neurons, suggesting that our original value was adequate.

Table 4.2: Optimized parameter values.
The best value (and the corresponding W_{scale}) for each parameter were inferred by grid search (see Fig. 4.4). The "5 % range" indicates the range where the MND changed by at most 5 % (in respect to its best value) and was estimated using a quadratic fit around the best value (with fixed W_{scale}).

Parameter	std. value	best value	5 % range	W_{scale}	MND
p_1 noise level scale	1.00	0.76	0.42 - 1.10	160.4	1.92
p_2 fraction of synapses with NMDA	0.90	0.95	0.78 - 1.13	156.5	1.88
p_3 NMDA-to-AMPA-ratio	0.47	0.86	0.73 - 0.98	107.4	1.74
p_4 width of inh. connections (μm)	150	208	182 - 233	223.8	1.79
p_5 max. fraction of $\mathrm{GABA_B}$ cond.	0.30	0.40	0.25 - 0.56	150.8	1.89
p_6 inh. to exc. weight scale	1.00	1.11	0.96 - 1.27	151.5	1.90
p_7 exc. to inh. weight scale	1.00	2.25	2.14 - 2.36	146.1	1.57
p_8 long-range weight scale	1.00	0.11	0.0 - 0.46	139.3	1.65

In general, we expected the synaptic receptor composition to be critical for achieving a realistic regime. Since NMDA conductances are activated on a slow time scale and thus might affect the variability of the model especially on a longer time scale, we tested two parameters varying the amount of NMDA receptors in different ways: the fraction of synapses having NMDA receptors (p_2) and the average NMDA-to-AMPA ratio of a synapse (having NMDA receptors) (p_3). Knowing that the latter ratio shows a relatively high fluctuation in experimental literature (Myme et al., 2003) and that NMDA receptor function might be influenced by anesthesia (Guntz et al., 2005), these parameters might need to be adjusted in the model. Remarkably, when NMDA conductances were not included in the model at all, the fit degraded significantly (about 20 % decrease in MND), compromising mostly the fit to the ISI structure and the Fano Factors, but also the fit to the lateral cross-correlation (see Fig. 4.5 B). This suggests that the NMDA conductances were a necessary component of the network model to achieve a realistic firing regime especially for the variability on a longer time scale. However, we also noticed that varying these parameters led to only minor improvements (within 10 % change of MND in respect to the standard parameters, see Fig. 4.4 B and C). Thus we concluded that the standard literature values for the NMDA-to-AMPA ratio and the fraction of synapses having NMDA receptors were already adequately chosen. $\mathrm{GABA_B}$ conductances are activated non-linearly only in case of high pre-synaptic activity events (Thomson and Destexhe, 1999), and furthermore exhibit relatively slow dynamics. We thus expected that the adjustment of the maximal fraction of $\mathrm{GABA_B}$ conductances (p_5) would affect the population spike structure. Indeed, we found that the $\mathrm{GABA_B}$ conductances were critical

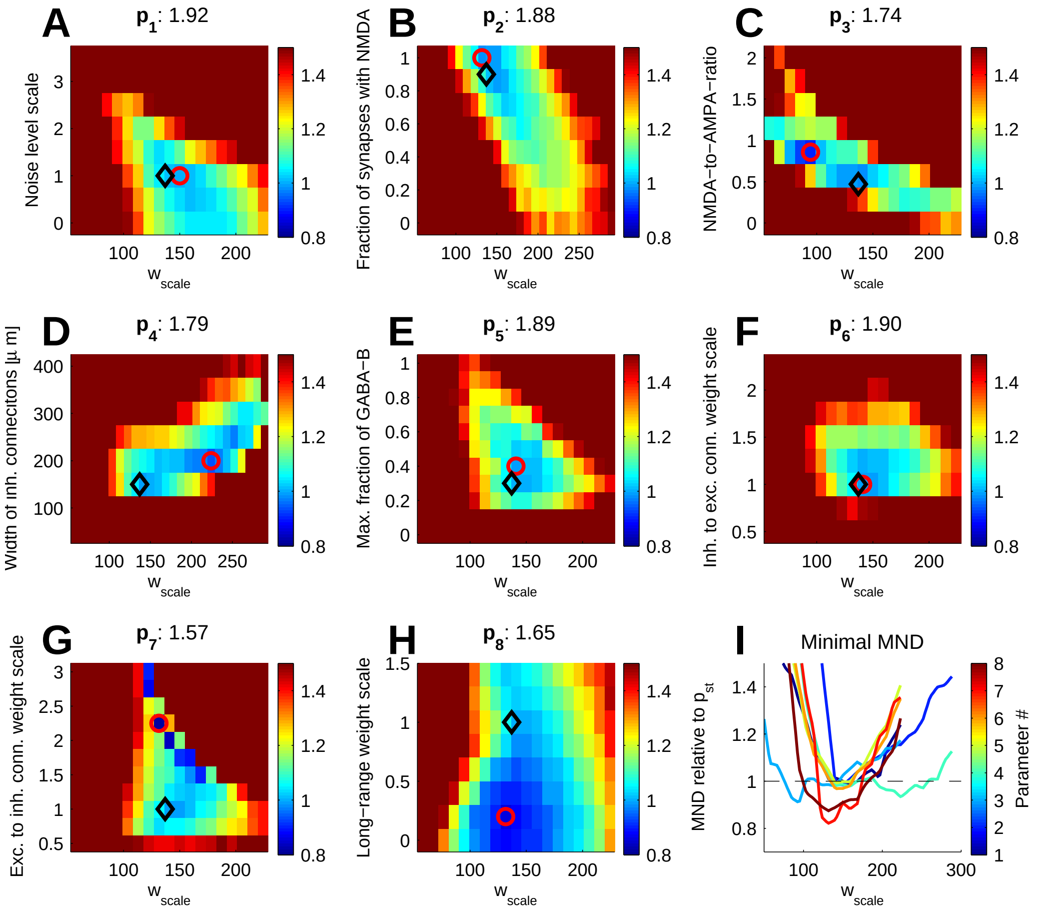

Figure 4.4: Improvement in the goodness of fit between *in vivo* and *in silico* firing regimes when varying physiological meaningful general parameters.
The improvement in MND in respect to the standard parameters is plotted in color code when varying 8 general parameters individually (see Table 4.1 for a description of the parameters). Each parameter was varied together with an overall scaling factor applied to all synaptic weights (W_{scale}), while other parameters were held constant. Adjusting some of the parameters considerably improved the fit to *in vivo* data. For instance, the relative synaptic weights of excitatory to inhibitory neurons needed to be increased dramatically (G). The standard parameter values and the settings showing the best fit in the statistical properties (minimal MND) are indicated with black diamonds and red circles, respectively. The minimal MND values are indicated in the titles. I shows the minimal MND (relative to standard parameters) versus W_{scale} for the 8 parameters.

in our model: including these conductances in the model dramatically improved the fit (about 80 % improvement, see Fig. 4.5 B). One possible reason for this dependence on GABA$_B$ conductances could be the crucial lack of long lasting or the non-linear activation of inhibitory neurons when GABA$_B$ conductances were absent. The strong effect suggested that sufficient activation of inhibitory neurons was necessary for achieving a realistic firing state.

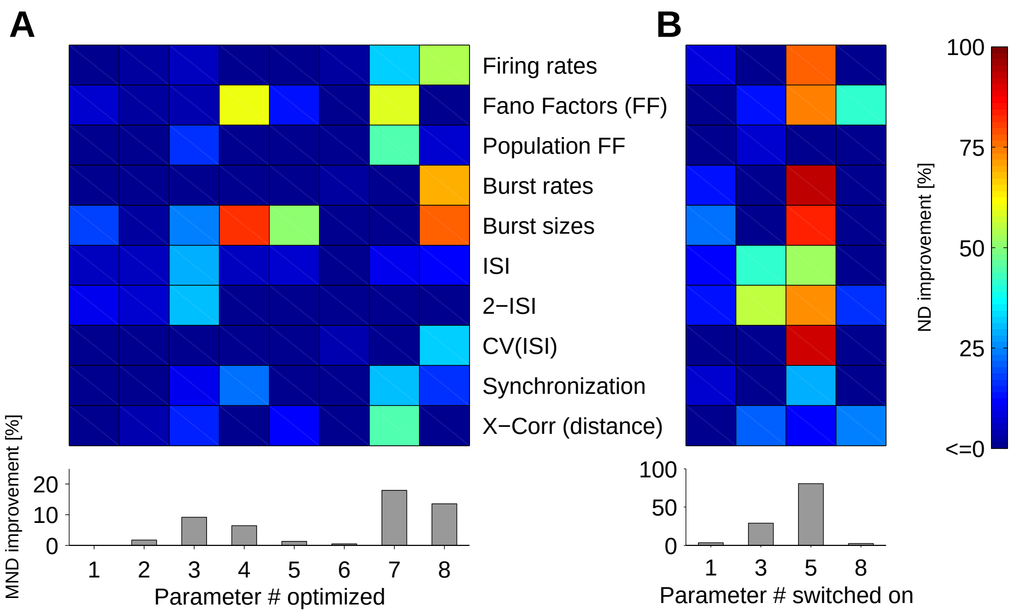

Figure 4.5: Effects of parameter optimization on individual statistical features A: The percentage change in ND in respect to standard parameter settings when optimizing parameters $p_1, \ldots, p_8$ individually is plotted in color code (see Table 4.1 for a description of the parameters). One notes that individual parameters have different influences on statistical features. The bottom margin displays the improvement in MND (averaged over all statistical features). Same simulation data as in Fig. 4.4. B: Impact of the inclusion of different components in the model. Selected components of the model: background noise (p_1), NMDA conductances (p_3), GABA$_B$ conductances (p_5), or patchy long-range connections (p_8). The improvements of the fit when including a component are plotted in color code (relative to the standard parameter settings, having all components included). One notes that including GABA$_B$ conductances had the most pronounced effects, improving the fit to multiple statistics profoundly. When components were switched on and off W_{scale} was again optimized in respect to MND.

However, similar to the NMDA conductances varying the maximal fraction of GABA$_B$ conductances did not considerably improve the MND value in respect to standard parameters (Fig. 4.4 E). One might hypothesize that the balance of excitation and inhibition was not established appropriately in the network model. To vary the overall connection strength between neuron pools, we chose relative synaptic weight scaling factors from inhibitory to excitatory neurons (p_6) and from excitatory to inhibitory neurons (p_7) as parameters. Whereas varying the inhibitory to excitatory connection strengths did not yield any overall improvement (Fig. 4.4 F), varying the reverse, the excitatory to inhibitory connection strengths had a strong effect. We noticed that increasing p_7 2.25-fold resulted in an 18 % improvement of the fit to *in vivo* data (Fig. 4.4 G), indicating the importance of correctly balancing inhibition and excitation for acquiring a realistic firing regime. Judging from the discontinuity of the error landscape (Fig. 4.4 G), an almost 3-fold increase in p_7 seemed to switch the firing regime into a new state, which was much more similar to the firing regime in nature. This strong overall improvement in the MND was mainly mediated by the ND improvement in the correlation structure (lateral cross-correlation and synchronizations), which could be improved by about 40 % in comparison to the simulation using standard literature values (see Fig. 4.5 A). Additionally, deviations in firing rate distribution and both Fano factors were also decreased by high percentages (see Fig. 4.5 A). Finally, we chose the relative synaptic weights scaling factor of the patchy lateral long-range connections (p_8) because it is not well constrained by the literature (see Methods Section for details). We found that the initial weight scale was somewhat too high: decreasing the weight of the long-range connections improved the variability of the network response. Indeed, the removal of long-range connections decreased the MND only by 3 % (Fig. 4.5 B). The decrease of the MND by 14 %, when optimizing for the relative strength of the long-range connection (Fig. 4.4 H), was mainly mediated by improving the burst structure (more than 50 % improvement in the burst sizes and the burst rate distribution), as well as the CV(ISI) distribution (about 35 %, see Fig. 4.5 A). When inspecting the spike responses visually, we noticed a slow rhythmic bursting for high p_8 values (near 1). These periodic population bursts were not seen after decreasing p_8. The relative weight of the lateral long-range connections therefore had to be reduced to avoid the induction of population bursts resulting in a much better fit to responses *in vivo*, in particular reducing the deviation in the statistical features sensitive to the burst structure.

In summary, for the majority of the selected parameters its literature value could not be markedly improved. The improved MND deviated less than 5 % from the MND values in case of standard parameters. An intermediate effect could be seen when varying the NMDA-to-AMPA ratio (p_3). Here the improvement with respect to the standard parameters reached 9 %. The most striking improvement, however, could be gained by varying the relative weight scaling factors of the long-range connections (p_8) and of the excitatory to inhibitory connections (p_7). Here the MND improved by 14 % and 18 %, respectively.

Next, we tested whether the fit could be further improved by varying the combination of the two most promising parameters together, i.e. the relative weight factors of excitatory to inhibitory connections and of patchy long-range connections, respectively p_7 and p_8. By setting p_7 to its best value (2.25) and again varying p_8 (as in Fig. 4.4 H) the goodness of fit improved further to MND $= 1.19$ (for $p_8 = 0.3$). We simulated this optimized model for multiple trials (changing the random seed of the simulation) and found a mean MND value of 1.30 ± 0.01 (mean ± standard error of mean, SEM). This is a 32 % improvement over the model using standard parameters. Finally, if a longer, non-intersecting section of the movie (25 s) was tested with this optimized parameters, the MND value decreased to a value of 1.10. Varying other parameters while using the best value for p_7 did not further improve the fit (not shown).

Are these improvements robust towards changes in the network structure? Since we used only one network to optimize the firing regime to reduce computational costs, we have to test the robustness of our findings in respect to a change in the network instance. Exemplarily, we tested whether ten networks with different construction random seeds achieved similar fits for the case of adjusting parameter p_7 and p_8. The goodness of fit generally depended somewhat on the particular network instance. We found an average MND value of 1.32 ± 0.09 (mean±SEM for different trials and different networks on the 25 s movie segment), when variables were initiated by the optimized parameters described above (and connection weights, time delays etc. of individual neurons and synapses were drawn randomly from distributions specified by general parameters). In general, the MND values therefore remained very stable when changing the network instance. However, in one exceptional network instance we observed an outlier value of MND $= 2.81$ for the same movie section. To test whether the optimal parameters are different in case of this outlier network and whether it could be re-adjusted to a realistic firing regime, we recomputed the optimization exemplarily for the parameter p_7. We found that the overall shape of the fitness landscape was almost identical to the original network (Fig. 4.4 G), except that we had to reduce the overall weight scale (from 146 to 131) to reach a good MND value of 1.41. Remarkably, the optimal value for parameter p_7 remained the same (2.25).

We conclude that the particular instance of the random network structure will commonly not have major influences as long as not too extreme weight configurations are drawn by accident. Therefore the amount of neurons in the network is large enough to sample over random instances of connection weights for individual neurons.

In conclusion, by comparing the model response statistics to *in vivo* data, the contribution of physiological meaningful parameters for achieving a realistic firing state could be revealed and the effect on statistical features quantified. Individual statistical features as well as the overall fit could be robustly improved by varying selected parameters. It was most effective to adjust the synaptic weights of the lateral long-range connection and to balance inhibition and excitation by strengthening the connections from excitatory to inhibitory neurons. The optimal parameter values were generally robust across network instances. The optimized network achieved a mean normalized deviation of MND = 1.10 (calculated on a long movie section), which is remarkably close to the average deviation between experimental sessions. In addition, for achieving a realistic regime, both NMDA and $\mathrm{GABA_B}$ conductances were crucially important components of the network model.

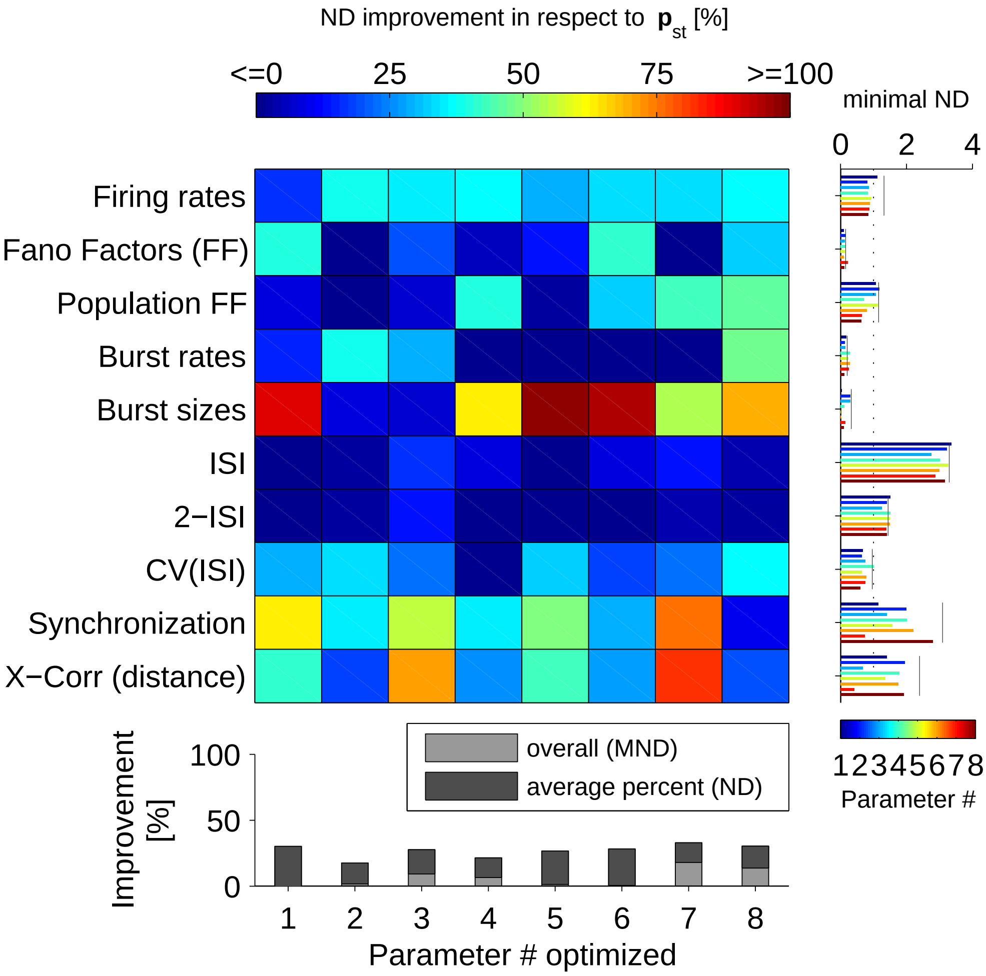

Figure 4.6: Optimization of statistical feature individually (as compared to optimize the mean over all).
The improvement in normalized deviation (ND) relative to the standard parameters is plotted in color code. The ND of individual statistical features was always optimized in respect to W_{scale}. Parameters (listed in Table 4.1) have different impact on statistical features. The bottom margin indicates the cumulative ND improvements and the improvements in MND (where W_{scale} is optimized on MND instead of ND). Left margin displays the actual ND values for each parameter (color coded bars). Note that an ND of 1 means a deviation equal to the average deviation between experimental sessions (and monkeys). Black lines are plotted in case of standard parameters. One notes that some statistical features were more difficult to fit whereas others were less problematic (reaching a value well below 1).

4.3.4 Deviations of the model response to the *in vivo* firing regime

After having optimized the firing regime of the model, how does the model response still deviate from the *in vivo* data? To illustrate the spike responses of the model (after improving parameters), we plotted its response to a section of the movie together with the *in vivo* responses (Fig. 4.2). The general appearances of *in vivo* and *in silico* responses were very similar: high activity periods were followed by low activity periods, bursts were induced by salient features in the movie, and trial-to-trial variability was comparable. Fig. 4.7 plots the statistical features for the optimized model (25 s movie presentation, MND = 1.10) together with the average over *in vivo* data. As the MND value already indicated, the overall correspondences were good, but deviations were still noticeable. In particular, there was a tendency that high-activity periods were overrepresented across neurons of the network, as can be seen in the tails of the firing rate and burst rate distributions (Fig. 4.7 A and D). A lack of long ISIs (above 500 ms) was evident in the ISI distribution (Fig. 4.7 F). This lack of long intervals was consistent across all performed parameter variations. In Fig. 4.6 we examined the improvement of the ND of individual statistical features when varying a parameter (in contrast to the improvement in MND, see Fig. 4.4 and Fig. 4.5). In fact, the ISI distribution is most difficult to fit to *in vivo* responses, as the best ND is only around 3, i.e. three times worse than the deviation between sessions on average (see Fig. 4.6 right margin plot). It seems that neurons *in vivo* can exhibit dynamics on multiple time scales in contrast to the model, which tended to be strongly active for certain times and just silent for others. This stronger dependence of one ISI on the following ISI in the model response can be observed in Fig. 4.7 K and L. If the current ISI was very short, *in vivo* and model responses matched very well, having a relatively high probability that the following ISI was also very short and an exponential fall-off for the probability of longer ISIs (Fig. 4.7 G) that was described well by a power-law behavior (i.e. straight line in log-log plot, not shown). However, for longer ISIs *in vivo* and model responses qualitatively differed. We found that *in vivo* the shape of the distribution for the following ISI did not change qualitatively for longer ISI if conditioned on the current ISI (despite a small increase of the probability for longer ISIs). In particular, the conditional probability had still a power-law shape implicating that the probability for an ISI below and up to 50 ms was relatively high regardless of the current ISI (Fig. 4.7 K). In contrast, in the model responses the shape of the distribution changed if conditioned on longer ISIs: if the current ISI was long ($>$ 100 ms) either the next ISI was very short ($<$ 10 ms), possibly belonging to the onset of a population burst, or the length of the next ISI had nearly uniform probability up to about 120 ms (Fig. 4.7 L). Varying the parameters had little effect on the deviation of the 2-ISI distributions, the strongest effect was exerted by the NMDA-to-AMPA ratio (p_3), reaching 13 % improvement in respect to the standard parameters. This lack of structured variability on multiple time scales in the model response was corroborated by the systematic underrepresentation of periods of high CV(ISI) (see Fig. 4.7 H).

Finally, the synchronization between two neurons and the lateral cross-correlation were generally too low in comparison with our experimental data (Fig. 4.7 J and K). In particular, synchronization between neurons on lags longer than 50 ms was much weaker (Fig. 4.7 J, see also Fig. 4.8 D), suggesting that dynamics on slow times scales were still lacking in the model. The correlation structure was most effectively influenced by four parameters (see Fig. 4.6): The strength of the background noise (p_1), the NMDA-to-AMPA ratio (p_3), the $\mathrm{GABA_B}$ fraction (p_5), and the relative weight of the inhibitory to excitatory connections (p_7). As expected, introducing synaptic dynamics on a longer time scale (p_3 and p_5) improved the synchronization structure. Background noise likely helps to smooth the sharp peak in the synchronizations. It is not immediately clear why the strongest improvement in the synchronization (76 %) and the lateral cross-correlation (83 %) was mediated by the increase of the strength of excitatory to inhibitory connections. We think that this strengthening of excitatory synapses onto inhibitory neurons might have recruited local negative feedback loops and thus initiated dynamics on intermediate and longer time scales.

In summary, analyzing the deviations of the statistical features in detail suggested that the response of the model was limited in reproducing the broad temporal range of the dynamics *in vivo*. Especially, abruptly switching from activity to silence, low probability of bursting, and the temporal correlation of neurons on a longer time scale were difficult to achieve.

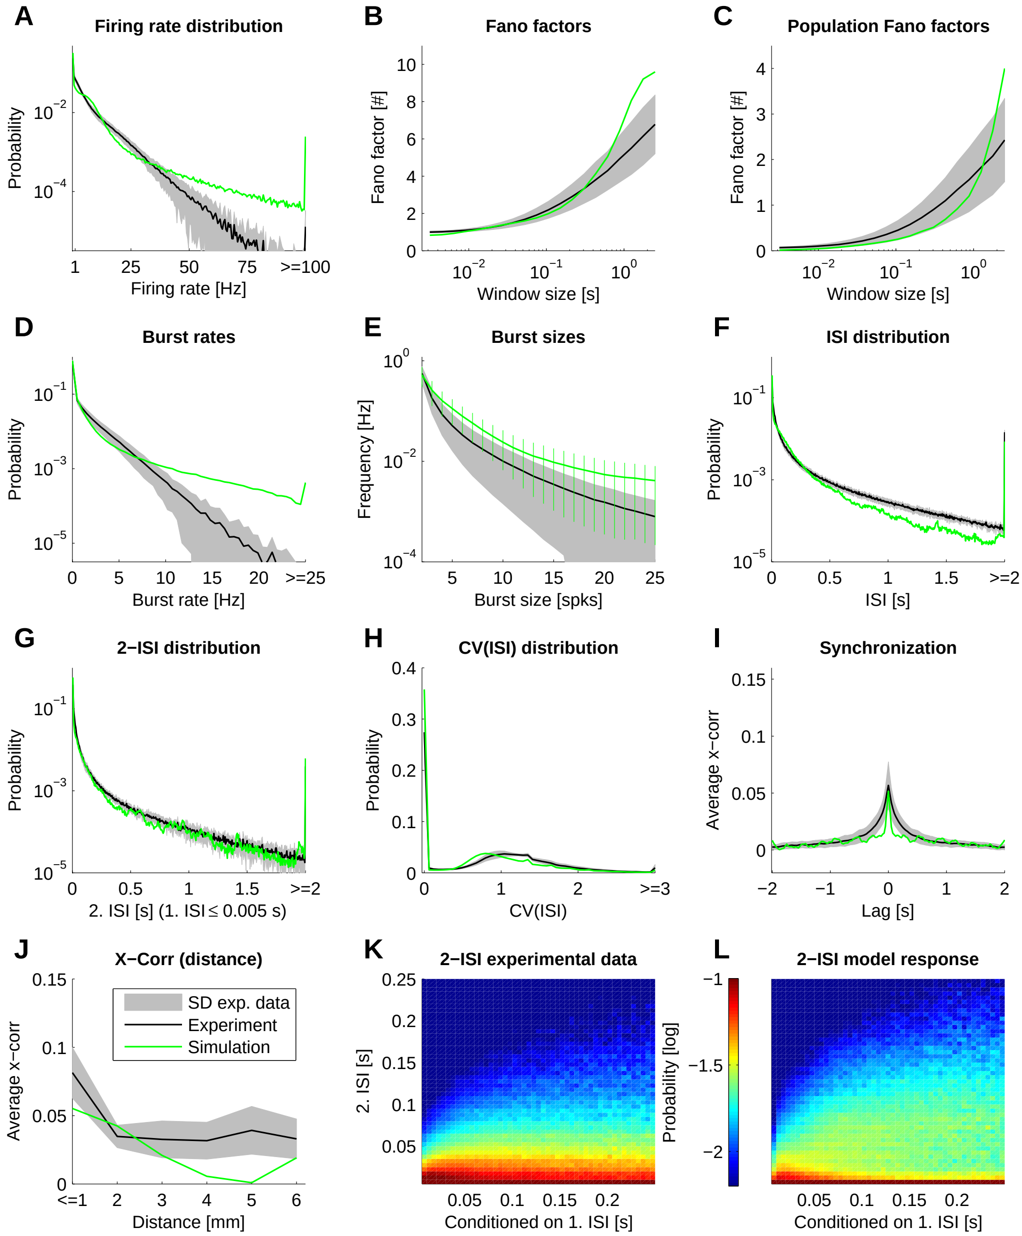

Figure 4.7: The spike statistics of the model in comparison to the statistics of the in-vivo recordings.
The mean of the statistical feature estimated on the *in vivo* data is plotted as a black line. The gray area indicates the standard deviation between different sessions. We used a 25 s part of the movie “sw21” as stimulus for the model. The spike statistics of the simulated model with optimized parameters (adjusted relative weight scaling factors of the patchy lateral long-range connections and of the connections from excitatory to inhibitory neurons) are plotted in green. The goodness of fit is MND=1.10. Each panel corresponds to a particular statistical feature analogous to Fig. 4.3. Note that the 2-ISI distributions (K and L) are plotted conditioned on the first ISI.

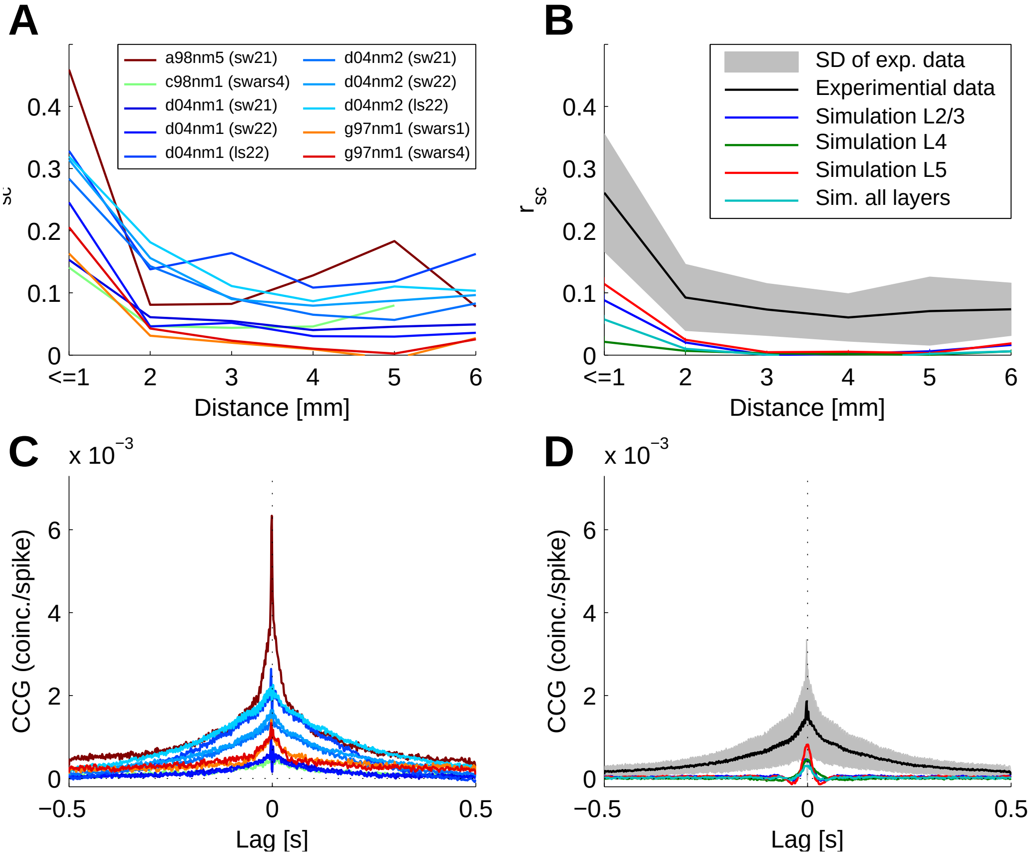

Figure 4.8: Correlation structure in the model and *in vivo*.
For better comparability with literature values correlations were plotted in terms of noise correlations (A and B) and shift-corrected cross-correlograms (CCG, C and D). Sessions of experimental data were plotted as in Fig. 4.3 (A and C). B and D depict the corresponding statistics calculated on the model response (for optimized parameter settings). Correlation structure in the model response were calculated on each layer separately showing a systematic difference in the strength of correlation in each layer (no layer information was available for the *in vivo* data). Correlations in the model were generally lower than in the *in vivo* data.

4.4 Discussion

In this article we investigated network spiking activity from the primary visual cortex under naturalistic stimulus presentation in vivo and in silico. We asked the question whether a state-of-the-art connectionists' model is capable of reproducing the firing statistics observed in vivo, without any arbitrary (unconstrained) adjustments of millions of variables, such as synaptic weights or temporal parameters of synapse dynamics. We concentrated on characterizing the firing regime of the primary visual cortex (V1), since it is an anatomically and physiologically well studied cortical area. We characterized its firing regime by extracting ten statistical features from in vivo spike recordings from V1 of macaques under naturalistic stimulus conditions. We then built and simulated a state-of-the-art circuit model reproducing a 5 mm x 5 mm patch of V1 cortical circuitry, and investigated the discrepancy between the firing regimes of the simulated model and in vivo recordings. The comparison revealed that the firing regime of the detailed laminar circuit model was comparable to in vivo recordings if parameters were set to values constrained by the literature: Responses were rich, showing high coefficient of variation of inter-spike intervals of individual neurons. We obtained a deviation averaged over all extracted statistics of about 2, meaning that the deviation was about twice as large as the average deviation between any two recording sessions. We found that the deviation can be decreased dramatically (by 32 %) if the relative synaptic weight of excitatory neurons to inhibitory neurons was increased more than 2–fold (as compared to the literature values) and relative synaptic weights of the patchy long-range connections were adjusted to avoid slow rhythmic population activity. This indicates the importance of the balance of excitation and inhibition to achieve a realistic firing state in a network model. We further found that selected physiological meaningful parameters affected statistical features of the response in a selective manner, and that some ingredients of our network, such as the NMDA and $\mathrm{GABA_B}$ conductances, were crucially important for achieving a realistic firing state.

4.4.1 Evidence for a characteristic firing state in V1

This study does not answer whether there is indeed a computational advantageous firing regime in neural tissue but instead tries to generally characterize a "statistical fingerprint" of spiking activity in vivo in order to compare the spiking activity to simulated responses in silico. Our characterization of the firing statistics of in vivo responses generally agrees very well with findings of previous studies, which also used natural stimuli. We have found some evidence that V1 might be in a characteristic state. Our finding that the ISI distribution follows a power law is consistent with the ISI distribution of V1 cells of anesthetized cats and of IT cells of awake monkeys (Baddeley et al., 1997; Yen et al., 2008), as well as of neurons in motor cortex of awake rats (Tsubo et al., 2009). A power law distribution has partially been reported for in-vivo spike data (Bedard et al., 2006) and the (spatial) size of neural avalanches recorded via LFP (Beggs and Plenz, 2003, 2004). Power-law distributions have been conjectured on theoretical grounds as characteristic features of self-organized critical states in large complex systems (Jensen, 1998). On the other hand, (Bedard et al., 2006) report that the ISI distributions derived from cat parietal association cortex during wakefulness and slow wave sleep show no evidence for a power-law behavior. This might indicate that the primary visual cortex and the parietal association cortex are working in different firing regimes.

There is still an ongoing debate on how strongly neighboring neurons are correlated, especially under natural conditions. Adding to this debate, Ecker and coauthors recently reported that in awake and behaving monkeys the spike count correlations r_{sc} are surprisingly low even between nearby neurons (in the range from 0.001 to 0.01). They claim that earlier studies have overestimated this value considerably by generally reporting values that are at least 10 times higher (Ecker et al., 2010). In our dataset we found an average value of $r_{sc} = 0.26 \pm 0.03$ (Mean±SEM). This finding is in agreement with the earlier studies and therefore stands similarly in contrast to the very low spike count correlations reported by (Ecker et al., 2010). The discrepancy is likely to arise from superior recording techniques used by Ecker and colleagues. As they convincingly show, the use of chronically implanted tetrodes considerably improves the quality of the recordings from single neurons. In contrast, when using single tip electrodes, spikes are far less distinguishable. Thus our spike trains likely contained a mixture of spikes originating from different neurons rather than from single neurons. These erroneous classifications may result in an artificial increase of spike count correlations between spike sorted neuron clusters. Indeed, it was very difficult to reach a high spike count correlation between neurons of the model under our stimulus paradigm. In fact, the average spike count correlation in our model is $r_{sc} = 0.06 \pm 0.01$ (mean±SEM) for nearby neurons (distances below 1 mm) and thus is very similar to the very low values reported by (Ecker et al., 2010). Moreover, the intra-layer spike count correlation in the model response (plotted in Fig. 4.8 B and D) was very similar to recently reported layer-wise recorded V1 data (Smith and Kohn, 2009) in that it showed a very low correlation for layer 4, an intermediate correlation in deep layers and the highest correlation within neurons from superficial layers. This low correlation in our model match well the picture presented by (Ecker et al., 2010) and suggests that the correlation of our data is artificially enhanced by the recording techniques used.

4.4.2 Is the model in a realistic state?

We used recordings done in anesthetized macaques during the presentation of a commercial movie stimulus. In previous complementary studies, with participation of some of the present authors, this data has already been proven insightful in the analyses of information coding of the movie stimulus in the neural responses (Belitski et al., 2008; Montemurro et al., 2008) and in investigating the relation of spikes to local field potentials (Rasch et al., 2007, 2009). We thus expect this dataset useful for describing a firing regime under naturalistic stimulus conditions. The firing regime will likely depend on the current behavioral state of the animal. Because we used recordings from anesthetized animals, we actually characterized the firing regime during anesthesia. Since most of the parameters constraining the model were not measured in behaving animals, the anesthetized state might actually be "more natural" for our network model. We expect that the firing regime of awake and behaving animals will put stronger constraints on a current connectionist's model to achieve. Our model showed a good correspondence to our in vivo recordings. However, this does not necessarily implicate that the model would be able to perform any computational function of V1, nor does this mean that the fit of the model to any kind of in vivo data could not in principle be improved further. We expect if more neurons could simultaneously be recorded, the deviations between model and experiment will increase because more constraints were set on the model design. Due to the small number of recorded neurons our data was limited in the sense that we could only test whether our model does reproduce the general likeness of randomly selected neurons from V1. However, we argue here that benchmarking of network models of cortical areas to reproduce the general likeness is an important prerequisite to analyze realistic computational functions in network models.

Given the diversity of cortical functions even within V1 (Olshausen and Field, 2005), we assume that a computational advantageous firing regime has to be quite general. We follow earlier setups of neural networks used for analyzing computational functions (Maass et al., 2002; Häusler and Maass, 2007) in that we draw synaptic weights and neuron to neuron connectivity from random distributions, without any specific learning of weights. Although the specific weight structure will influence the network response, we suppose that the general firing statistics will nevertheless be similar in a statistical sense if the network is not too specialized. Supporting this view, we found that different random seed in the network generation in general only slightly changed the goodness of fit. It is therefore a promising research direction to investigate self-organization of synaptic weights for instance by using an intrinsic plasticity rule to achieve a computationally advantageous firing regime of recurrent networks (Triesch, 2007; Schrauwen et al., 2008).

Since we extract a number of different statistical features, measuring variability, correlation structure, and spike time dependency, it is likely that a hypothetical computational advantageous firing state of the cortex leaves its traces in the pool of extracted statistical features. Using a plethora of features has the advantage to minimize the bias of the investigator towards a particular aspect of the response. We also tested different error measures (such as a normalized mean squared error on the logarithm of the distributions instead of the Kullback-Leibler divergence), which slightly differed in the obtained optimal parameter values, because different aspects are deemed more important by other error measures (such as a good correspondence of the tail of the firing rate distribution). If there would be prior knowledge about the importance of statistical features, one could include this knowledge by changing the weighting for important features in the error measure (such as the deviation of the firing rate distribution from an exponential shape).

4.4.3 Which general parameters should be optimized in a cortical network model?

We developed a cortical network model for a 5 mm x 5 mm patch of area V1. This model, consisting of about 35000 neurons and 3.9 million synapses, expands a previously published model for generic cortical microcircuits (Häusler and Maass, 2007) that was based on data from (Thomson et al., 2002) regarding layer-specific connection probabilities and data from (Markram et al., 1998) and (Gupta et al., 2000) regarding short-term depression and facilitation of synaptic connections. Additional, our model comprised data-based patchy long-range connections, two types of excitatory receptors (AMPA and NMDA), and two types of inhibitory receptors ($\mathrm{GABA_A}$ and $\mathrm{GABA_B}$). The last years have seen several attempts to model large areas of the brain with similar components, such as inter- and intralaminar connectivity, laterally structured connectivity, synaptic depression and facilitation, and neurons having one or a few compartments (e.g. (Izhikevich and Edelman, 2008; Johansson and Lansner, 2007; Tao et al., 2004; Kremkow et al., 2007)) These models, as well as ours, incorporate many anatomical and physiological details, but they are, of course, still a strong abstraction of reality.

Given the complexity of these models, it would be desirable to pin-point a few general parameters which are sufficient to tune for achieving a realistic firing regime. Then models could easily be adjusted and investigations of its computational functions could start from a realistic basis. Our analyses of optimizing a state-of-the-art cortical network model resulted in the following observations. At the outset, the overall synaptic connection weight scale has to be adapted to account for the specific synaptic drive of each neuron, which is commonly lower in models because of the much smaller number of synapses. Furthermore, we had to adjust the relative weight of the patchy long-range connections, to dampen the excitability of the network to avoid the tendency to produce periodic population bursts. This parameter was not well defined in the literature because synaptic contacts and weights of the long-range connections are unknown. We suspect that lateral interaction is tightly linked to computation in V1, and it is therefore likely that synaptic targets and weights are carefully selected by experience dependent learning mechanisms. Although we incorporated a higher probability for long-range connections towards similar orientated hypercolumns, our approach of drawing random weights is likely too unspecific in the target neuron selection. Thus long-range connections in our network might form a too generic source of excitation, which has to be damped to avoid rhythmic population bursts. On the other hand, the easiness of inducing periodic patterns on a very slow time scale (below 10 Hz and down to less than 1 Hz) by varying the overall strength of the lateral long-range connections in our network model indicates a possible mechanism for the generation of slow-wave activity commonly seen in the visual cortex and suggested to be important for information coding (Montemurro et al., 2008). In the present study, the most effective parameter for tuning the model behavior was the relative synaptic weight scaling factor for the connections from excitatory to inhibitory neurons. To achieve a realistic regime, we had to dramatically increase this parameter (more than two fold compared to literature values). Additionally, the fact that the $\mathrm{GABA_B}$ inhibition was crucially important further supports that inhibition has to be powerful enough over a long periods to cope with the excitatory drives. The necessity to increase the strength of inhibitory action in relation to literature values suggests that our implementation of the interaction between excitatory and inhibitory neurons underestimates the drive of inhibitory neurons in nature. Since our point neuron model ignores any spatial extent, dendritic and axonal tree architectures were not part of our model. However, it is known that dendrites are capable of non-linearly integrate their synaptic inputs (Koch et al., 1983; Borg-Graham et al., 1998). For instance, some types of inhibitory neurons tend to target more soma proximal regions, whereas excitatory synapses are usually located in more distal parts of the dendritic tree (Markram et al., 2004). This arrangement gives rise to shunting inhibition (with appropriate reversal potentials, (Koch et al., 1983), where inhibitory inputs non-linearly overrule excitatory input. This mechanism effectively increases the strength of inhibitory action (as compared to our point neuron model). It is likely that by increasing the weight of excitatory synapses on inhibitory targets in our model, the lack of non-linear inhibition was partly compensated because inhibitory neurons where made more sensitive, which resulted in a more realistic state. We conclude that incorporating dendritic morphology (or other means to render the effect of inhibitory neurons more realistically) is a promising research direction and will likely improve firing states of network models.

Additionally, to keep the modeling effort tractable, many details of the primary visual system were neglected. For instance, the color processing pathways (Sincich and Horton, 2005), motion processing pathways (dorsal stream), and a more accurate model of the LGN including feedback from V1 to the thalamus were not modeled. Whereas color and motion processing will only influence the movie features that are processed in V1, we expect the LGN to have more profound influences on the V1 firing state. For instance, it has been shown that irregular but correlated inputs to a neuron increase its CV(ISI) (Softky and Koch, 1993; Stevens and Zador, 1998; Salinas and Sejnowski, 2000) and amplification of synchronous inputs due to nonlinear interactions within dendrites further enhances the variability. Because the LGN is likely to deliver highly synchronous inputs to V1 (Wang et al., 2010) and furthermore might code information in bursts of activity (Reinagel et al., 1999), as well as because cortical feedback are known to induce thalamocortical spindle oscillations (Steriade et al., 1993), a more accurate LGN model may further improve the fit to in vivo data. Although we used a switching Gamma process as a model for the LGN input, and thus adopted special means to more realistically model episodes of high firing rates, the input still has independent Poissonian character for low firing rates. Our LGN model might not be realistic enough, in that generated spikes are less synchronous than in vivo (Wang et al., 2010).

In general, we noticed that variability in the dynamics on a multiple time scale was still underrepresented in our model. This was observable in the lack of high variability regions (CV(ISI) distribution), the underrepresentation of long ISIs, and the too weak correlations between neurons spiking for lags longer than 10 ms. It is therefore possible that our network is still not complex enough to generate a non-stereotyped long-lasting dynamics seen in vivo. Since we used point neurons models, neither compartimentization of dendrites or axons nor any diffusion of neurotransmitter, nor dynamics of second messengers (such as Calcium ions), nor any other cellular process were integrated in our model. The only processes reaching a time scale of a few hundreds of milliseconds in our model were slow synaptic conductances and synaptic short term facilitation and depression. We suspect that dynamics of second messengers, which generally happen on a slower times scale than spike interactions, might be necessary ingredients for achieving a more realistic firing regime (in particular on longer time scales).

4.4.4 Conclusion

The characterization of the in-vivo response of neurons in monkey area V1 that we have presented provides useful information for the investigation of large scale models for cortical areas. It is remarkable that a model for a patch of V1 that is based on previously published anatomical and physiological data produces (after adjusting a few general parameters) a spike response that matches the statistical properties of our in vivo data quite well. However, although similar in general statistical measures, we are still a long way of understanding the detailed neural coding properties of the cortex, which are manifested in the fine-structure of the interaction between neurons. The advent of techniques with the possibility to record from hundreds of neurons simultaneously will put forward new challenges for cortical network models. We expect that benchmarking models with in vivo data, as exemplified in this study, will foster the development of new and more realistic models for cortical areas, which will be important tools to ultimately understand neural functions.

4.5 Acknowledgments

This chapter is based on the paper *Statistical comparison of spike responses to natural stimuli in monkey area V1 with simulated responses of a detailed laminar network model for a patch of V1*, which was written by Malte Rasch (MR), myself (KS), Nikos Logothetis (NL), and Wolfgang Maass (WM). MR and KS equally contributed to this work. The in-vivo data was recorded in the lab of NL. The data analysis was performed by MR and KS. The circuit model was developed and implemented by MR and KS. The paper was written by MR, KS and WM with addtional input from NL.

Chapter 5

Temporal dynamics of information in a model for cat primary visual cortex

---

Contents

5.1 **Introduction** 122
5.2 **Materials and Methods** 124
5.3 **Results** 134
5.4 **Discussion** 146
5.5 **Acknowledgments** 148

We show that the temporal dynamics of information about sequences of visual stimuli contained in neural responses in cat primary visual cortex can be replicated by a detailed model of a 4.5 mm×4.5 mm patch of cat primary visual cortex that implements stereotypical lamina-specific cortical synaptic connectivity. In particular we verify that the generic model accounts for experimental results on fading information about previously shown visual stimuli, the non-linear superposition of information about visual input presented at different moments in time, and the information content in the fine temporal structure of neural responses.

Additionally, we make the following three predictions about properties of cat primary visual cortex that could not be deduced from experimental data. First, the firing rate profiles of neurons in response to sequences of visual stimuli depend on the depth of the neurons recorded from within the neural tissue. Secondly, all information that could be extracted from the neural responses is contained in first and second order correlations between discharge patterns of cortical neurons. Thirdly, the long duration of neural responses to single letter presentations of up to 1 s can not only be attributed to local generic synaptic connectivity but is presumably induced by feedback loops with higher cortical areas.

5.1 Introduction

The fundamental aspects of the organization of computations in visual cortex are still unknown (Olshausen and Field, 2006). Traditionally visual processing is considered to be carried out by specialized neural circuits that are organized through precisely structured maps of columns. Complex vision tasks, such as visual object recognition, are then implemented by a subsequent hierarchy of other processing layers that carry out additional computational operations (Serre et al., 2005). In this postulated precisely organized sequence of processing steps, each layer works at any moment in time on a single frame of current or past sensory input.

But recently, (Nikolic et al., 2009) reported that the response properties of neurons in cat primary visual cortex (V1) are in conflict with these classical models for visual processing, which emphasize sequential step-by-step processing of single frames of visual input. In contrast, they conjectured that the reported evidence of slowly fading information about a previously shown visual stimuli and the non-linear superposition of information about visual input presented at different moments in time supports alternative computational models that are based on generic recurrent microcircuits (Maass et al., 2002). Such superposition of information about subsequent frames of online inputs was shown in (Natschläger and Maass, 2005) to be a direct consequence of the recurrent network connectivity. The exact computational operations that are carried out by neurons in the recurrent network are less relevant for this computational model, as long as a number of large scale properties for the simultaneous response of large ensembles of neurons are met. Such desirable properties are diversity of responses of individual neurons and fading information about previously shown stimuli.

Here we show that indeed most of the results reported in (Nikolic et al., 2009) can be replicated by a detailed generic recurrent microcircuit model of a 4.5 mm×4.5 mm patch of cat V1. The model is based on the cortical microcircuit model described in (Häusler and Maass, 2007) that implements experimental data on stereotypical lamina-specific connectivity that is repeated throughout many cortical areas (Thomson et al., 2002). Additionally, the cortical model complemented by a model for the visual pathway located upstream to V1, i.e. the retina and the lateral geniculate nucleus (LGN) of the thalamus. In particular we verify that the generic model accounts for the experimental results on fading information about previously shown visual stimuli, the non-linear superposition of information about visual input presented at different moments in time, and the information content in the fine temporal structure of neural responses.

In addition, we address three questions about the dynamical and computational properties of V1 that could not be answered based on the experimental results reported in (Nikolic et al., 2009). This was in part because of a lack of data and in part because of a lack of means to manipulate specific properties of the synaptic connectivity. First, a variety of different firing rate profiles of the neural responses was observed for different animals and recording sessions. Can this variety of rate profiles be explained in terms of the location of the recording electrodes within the neural tissue? Secondly, all information about stimulus properties that could be extracted from the neural responses was contained in first and second order correlations between the discharge patterns of neurons. Can the lack of additional information contained in higher-order correlations be attributed to the limited size of the datasets? Thirdly, neural responses to presentations of single letter stimuli showed increased firing rates within durations of more than one second. Can these prolonged neural responses be explained in terms of persistant activity maintained by local synaptic connectivity mediated by NMDA receptors? NMDA receptor activity is known to be altered due to the application of anesthetics. This raises the question if the corresponding results for awake animals might differ from the results for anesthetized animals reported in (Nikolic et al., 2009).

Simulation studies represent an appropriate tool to answer these questions because of the possibility to carry out a large number of trial repetitions and the capability to analyze the effect of each model component on the performance properties of the model. For each of the three questions the cortical microcircuit model provides a prediction that remains to be verified experimentally.

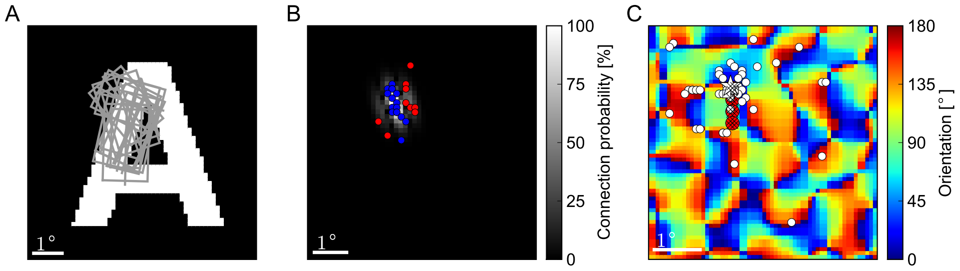

Figure 5.1: Receptive fields and connectivity of recorded neurons **A**: Outlines of the receptive fields of 16 recorded neurons located in layer 4 of the model. **B**: The connection probabilities of LGN neurons to a given neuron in layer 4 of the cortex. The actual selected connections are marked as red and blue circles originating from ON- and OFF-cells of the LGN, respectively. **C**: Orientation map of the simulated model. The electrodes are denoted as red circles. The positions of the recorded neurons are marked as black crosses. For one of the recorded neurons (white star) in layer 2/3 the incoming connections are plotted (white circles and white squares denote excitatory and inhibitory connections, respectively).

5.2 Materials and Methods

5.2.1 Model

The model consists of an input model (representing the retina and the lateral geniculate nucleus (LGN) of the thalamus) and a model of a patch of V1, receiving and processing the spikes of the input model. We will first describe the V1 model and subsequently the input model.

5.2.1.1 V1 Model

The V1 model is based on the cortical microcircuit model described in (Häusler and Maass, 2007), which implements experimental data from (Thomson et al., 2002) on lamina-specific connection probabilities and connection strengths between excitatory and inhibitory neurons of three cortical layers, and data from (Markram et al., 1998) and (Gupta et al., 2000) regarding stereotypical dynamic properties (such as paired pulse depression and paired pulse facilitation) of synaptic connections between excitatory and inhibitory cortical neurons. We extended this microcircuit model laterally and incorporated specific anatomical properties of V1 in cat. Due to the lateral extent of the model, it was possible to include lateral long-range connections, which are thought to be essential for the computational functions of V1, such as spatial integration of extra-classical receptive field context (Gilbert et al., 1996).

The model for a 4.5 mm×4.5 mm patch of area V1 consisted of 32400 neurons and almost 2.8 million synapses. Since the number of neurons in the model was substantially smaller than in a 4.5 mm×4.5 mm patch of area V1, the scale of all synaptic weights was adjusted in order to achieve biologically realistic firing rates. This was carried out by choosing a suitable value for the scaling factor, with which the weights of the synaptic connections between neurons in the V1 model were multiplied. Since the LGN model also consisted of much fewer neurons (namely 10000 neurons) than the LGN in cat, another scaling factor was used, which scaled the strength of all synaptic connections from the LGN model to the V1 model.

The neurons of the V1 model were equally distributed on three layers, corresponding to the cortical layers 2/3, 4 and 5. Each layer contained a population of excitatory neurons and a population of inhibitory neurons with a ratio of 4:1 (Beaulieu et al., 1992; Markram et al., 2004). The inter- and intra-layer connectivity (probability and strength) was chosen according to experimental data from rat and cat cortex assembled in (Thomson et al., 2002).

In our model, we assumed that the neurons are uniformly distributed throughout the cortex. Thus we positioned all neurons on a cuboid grid with a constant grid spacings. Using the experimentally measured neuron density, e.g. for layer 2/3, the grid spacing would be $20\mu m$ for all directions. Because the simulation of such a dense network would be computationally to expensive, we diluted the neuron density by increasing the lateral grid spacing to $75\mu m$ and the vertical spacing to about $200\mu m$.

Due to a considerable gain in computational speed, we used the neuron model suggested by (Izhikevich, 2003), which can be adjusted to express different firing dynamics (Izhikevich, 2006). We randomly drew the parameters for each neuron in the network according to the bounds provided by (Izhikevich et al., 2004). On the basis of these parameter distributions the excitatory pools consisted of regular spiking cells, intrinsically bursting cells, and chattering cells, with a bias towards regular spiking cells. The inhibitory pools consisted of fast spiking neurons and low-threshold spiking neurons.

In addition to the synaptic input from other neurons in the model, each neuron received synaptic background input, modeling the bombardment of each neuron with synaptic inputs from a large number of neurons that are not represented in our model. This synaptic background input causes a depolarization of the membrane potential and a lower membrane resistance, commonly referred to as the “high conductance state” (Destexhe et al., 2001). The conductances of the background input was modeled according to (Destexhe et al., 2001) by Ornstein-Uhlenbeck processes with means $g_{exc} = 0.012\ \mu S$ and $g_{inh} = 0.047\ \mu S$, variances $\sigma_{exc} = 0.003\ \mu S$ and $\sigma_{inh} = 0.0066\ \mu S$, and time constants $\tau_{exc} = 2.7\ ms$ and

$\tau_{inh} = 10.5\ ms$, where the indices exc/inh refer to excitatory and inhibitory background conductances, respectively. Furthermore, these conductances account for the low neuron density in the model in comparison to the density in the cat cortex.

Short term synaptic dynamics were implemented according to (Markram et al., 1998), with synaptic parameters chosen as in (Maass et al., 2002) to fit data from microcircuits in rat somatosensory cortex (based on (Gupta et al., 2000) and (Markram et al., 1998)). For further details we refer to (Häusler and Maass, 2007).

Lateral connectivity structure The generic microcircuit model of (Häusler and Maass, 2007) was based on data for a column of about 100 μm diameter with uniform connectivity per layer and neuron type. Here we extended the model laterally to several millimeters. Thus connection probabilities in our model depend on the lateral distance. For inter- and intra-cortical connections we used a bell-shaped (Gaussian) probability distribution for determining the lateral extent. The standard deviation of the Gaussian was set to 200 μm for excitatory neurons (Lund et al., 2003; Blasdel et al., 1985; Buzas et al., 2006) and to 150 μm for inhibitory neurons to incorporate the observed occurrence of extremely narrow inhibitory dendritic and axonal spreads (70 μm, (Lund et al., 2003)). To ensure consistency with the connectivity data of (Thomson et al., 2002), we scaled the Gaussian profiles such that the peak probabilities correspond to their experimentally measured connection probabilities. Therefore, their connectivity pattern was locally preserved.

Patchy lateral long-range connections In cat, many pyramidal cells in layer 2/3 of the striate cortex (and also elsewhere in the cortex (Lund et al., 2003)) have characteristic long-range projections targeting laterally 80% excitatory and 20% inhibitory cells (McGuire et al., 1991) which are up to 6 mm and more away (Gilbert et al., 1996; Lund et al., 2003; Buzas et al., 2006). Targeted neurons tend to have similar feature preference as its origin, resulting in a patchy connectivity, which connects neurons with similar preferred orientations (Gilbert et al., 1996; Buzas et al., 2006). (Buzas et al., 2006) combined the anatomical reconstructions of neurons and the optical imaging of orientation maps and proposed a formula to calculate the bouton density ρ of a typical layer 2/3 pyramidal cell:

$$\rho(r, \Delta\phi) = Z \left(e^{-\frac{r^2}{2\sigma_1^2}} + m\, e^{-\frac{r^2}{2\sigma_2^2}}\, e^{\kappa \cos(2(\Delta\phi - \mu))} \right) \tag{5.1}$$

r is the lateral (Euclidean) distance between the pre- and the post-synaptic neuron, and $\Delta\phi$ is the difference of preferred orientations of the two neurons. The scaling factor m accounts for the importance of the long-range orientation dependent term against the local orientation independent term. The orientation dependent term is a product of a Gaussian and a von Mises distribution and accounts for the higher likelihood of connecting neurons with preferred orientations, which differ in the range of μ degree. Standard deviations σ_1 and σ_2 regulate the spatial width of the non-oriented and oriented term, respectively. Parameter κ regulates the "peakiness" of the density on the orientation axis. Z is a normalization constant.

The bouton density defined by Eq. 5.1 can be seen as an estimate for the connection probability of a neuron to neurons at distance r having preferred orientation difference $\Delta\phi$ if one assumes that in first approximation dendrites and axons have no spatial extent (i.e. they collapse to a point at the position of the soma). Since we defined the preferred orientations of neurons in a hard-wired manner via "oriented" input connections (see Section 5.2.1.3), we could apply Eq. 5.1 for the lateral connections in layer 2/3. Thus, we applied Eq. 5.1 for the connections of excitatory cells targeting excitatory and inhibitory cells (McGuire et al., 1991).

As for the connections between other layers, we set $\sigma_1 = 200\mu$m for the local non-oriented term. In (Buzas et al., 2006) the values of μ jitter around zero degree for individual cells. Therefore, we set $\mu = 0°$, implicating that the connection probability was highest for iso-oriented cells. We set other parameters to values in between the two populations described by (Buzas et al., 2006), but in the vicinity of the population 1, where the experimentally mapped cortical surface region was bigger. Thus we set $\sigma_2 = 1000\,\mu$m, $\kappa = 1$, and $m = 10$. As before, the connection probability was scaled, according to (Thomson et al., 2002), by setting Z to appropriate values. Thus, locally, i.e., for neurons at the same lateral position (and orientation preference), the connection probabilities were preserved. The weights of the the long-range connections were drawn from the same distribution as the short range connections in layer 2/3.

Distance dependent synaptic delay Synaptic delays differ for inhibitory and excitatory neurons. They were set according to measurements by (Gupta et al., 2000) (for details see (Häusler and Maass, 2007)). These delays stem from molecular processes of synaptic transmission. In addition, a second delay originating from finite spike propagation velocity of the fibers was included. This delay depends on the (Euclidean) distance between the pre- and the post-synaptic neuron. (Girard et al., 2001) measured spike propagation orthodromically as well as antidromically. They found a median conduction velocity of 0.3 m/s for the upper layers, and 1 m/s for the lower layers of V1 in macaque monkeys. Thus, we sampled the velocity for each excitatory synapse in layer 2/3 from a Gaussian distribution with mean 0.3 m/s and standard deviation 0.5 m/s (with enforced lower and upper bounds of 0.05 and 5 m/s, respectively). For the other layers, the conduction velocities were drawn from a Gaussian distribution with mean 1 m/s and standard deviation 0.9 m/s (with same bounds as before). Due to myelination, conduction velocities of inhibitory fibers are generally higher than for excitatory cells (Thomson et al., 2002). Lacking exact measurements in the literature for all inhibitory cells, we sampled the velocities from a distribution with mean and standard deviation twice as high as for excitatory neurons in the deep layers (the enforced upper bound was set to 10 m/s).

Synaptic conductances A spike, arriving at a synapse, causes a change in the synaptic conductance in the post-synaptic neuron. The dynamic of the conductance depends on the receptor kinetics. Each excitatory synapse in our model contains α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) receptors having relatively fast kinetics (modeled as exponential decay with time constant $\tau_{\text{AMPA}} = 5$ ms, reversal potential 0 mV). A fraction f_{NMDA} of all excitatory synapses contain additionally relatively slow, post-synaptic voltage dependent N-methyl-D-aspartate (NMDA) receptors ($\tau_{\text{NMDA}} = 150$ ms, reversal potential 0 mV,(Gerstner and Kistler, 2002; Dayan and Abbott, 2001)), and therefore exhibit a superposition of conductance kinetics. The ratio of NMDA to AMPA receptors in a synapse was drawn from a Gaussian distribution with mean $\mu_{\text{NMDA/AMPA}} = 0.47$ and standard deviation $\sigma = 0.2\mu_{\text{NMDA/AMPA}}$ (Myme et al., 2003).

5.2.1.2 Input model

We needed a sufficiently realistic transformation of stimulus sequences to spike trains. Therefore, the retina and the lateral geniculate nucleus (LGN) were modeled, according to (Dong and Atick, 1995), as a spatio-temporal filter bank with nonlinearities, which seems to be a good compromise between simplicity and realism (Gazeres et al., 1998). The filter bank converted time varying input signals on the retina into firing rates of LGN neurons.

Retina model The 2-dimensional retinal inputs (movie frames) were filtered by difference of Gaussians (DoG) spatial filters (Rodieck, 1965; Enroth-Cugell and Robson, 1966; Dong and Atick, 1995). Filter sizes (corresponding to the receptive fields of ganglion cells) were adapted to the geometry of cats, where the standard deviations of the Gaussian for center and surround were estimated to be $\sigma_{\text{center}} = (0.0177^\circ + 0.00196\epsilon)$ and $\sigma_{\text{surround}} \approx 6.67\sigma_{\text{center}}$ at eccentricity ϵ, respectively (in visual degrees; estimated from (add correct ref.). After the convolution of the stimulus luminance portrait with these kernels (yielding S_{center} and S_{surround}), the response of a retinal ON-cell at visual field position r can be described by

$$R_{\text{ON}}(r) = C(r)\,[S_{\text{center}}(r) - \omega S_{\text{surround}}(r)]_+ \tag{5.2}$$

We set the ratio of center to surround $\omega = 1.0$. The positive part of the center and surround interaction (indicated by the brackets $[\ldots]_+$) was assigned to the response of an ON-cell and, analogously, the absolute value of the negative part to the response of an OFF-cell (Dong and Atick, 1995). For simplicity we assumed that the origins of the center and surround summation fields are identical, although a recent study suggests that there might be an offset between them (Conway and Livingstone, 2006).

Applying the Difference-of-Gaussians model to the luminance of a stimulus results in a quantity called “contrast gain” (Croner and Kaplan, 1995; Enroth-Cugell and Robson, 1966; Rodieck, 1965). To calculate the firing rate of ganglion cells, one has to multiply the “contrast gain” with the *local* contrast $C(r)$ (as done in Eq. 5.2), if one neglects non-linear saturation in the high contrast regime. Locality is important because the concept of a global contrast, easily defined for full-field grating stimuli commonly used in experiments, is not applicable for real world images and movies (Tadmor and Tolhurst, 2000). Following (Tadmor and Tolhurst, 2000) we estimated the local contrast using the same kernels as

$$C(r) = \frac{|S_{\mathrm{center}}(r) - S_{\mathrm{surround}}(r)|}{S_{\mathrm{center}}(r) + S_{\mathrm{surround}}(r)} \tag{5.3}$$

where we additionally set the contrast to be zero in the case of darkness. Note that applying Eq. 5.3 results in a response $R_{\mathrm{ON}}(r)$, which is sparser than for a constant global contrast, since the response is now quadratic in the center and surround luminance difference (see Eq. 5.2).

LGN model The output of the retina was filtered by the LGN model using a temporal kernel. The temporal kernel consists of a phasic (taken from (Dong and Atick, 1995)) and a tonic component (as in (Gazeres et al., 1998)), i.e. $k_{\mathrm{LGN}} = k_{\mathrm{phasic}} + k_{\mathrm{tonic}}$. It is for non-negative times

$$k_{\mathrm{phasic}}(t) = t\,(1 - \pi w_c t)\,\exp(-2\pi w_c t) \tag{5.4}$$

and

$$k_{\mathrm{tonic}}(t) = A\,\exp(-t/\tau)/\tau. \tag{5.5}$$

Parameter $A = 0.3$ is the fraction of tonic activation (with respect to the peak firing rate) for a given stimulus, integrated over a time window of $\tau = 15$ ms. Parameter $w_c = 5.5\ \mathrm{s}^{-1}$ defines the shape of the phasic kernel (Dong and Atick, 1995).

The positive parts and the absolute values of the negative parts of the temporal convolutions were assigned to non-lagged and lagged cells, respectively. Altogether, there are four different time-varying rate outputs, i.e. that of any combination of non-lagged or lagged cells in the LGN with either ON- or OFF-cells from the retina (Dong and Atick, 1995). Following (Gazeres et al., 1998) a so-called “switching Gamma renewal process” was used to convert these time-varying rates to spike trains. This process, which was suggested to fit experimental data from cat LGN X-cells (Gazeres et al., 1998), adopts a higher spike time regularity for high input rates ($\geq$ 30 Hz; regularity parameter $r = 5$) and switches to a Poisson process for low rates ($<$ 30 Hz). The spontaneous background activity of each LGN neuron was set to a value of 10 Hz and the peak LGN spike rate $f_{\max} = 600$ Hz The peak response would be evoked by a dot of highest contrast filling the center region of a ganglion cell with optimal duration. This value is in good agreement with (Gazeres et al., 1998).

Input connectivity to V1 The visual field is retinotopically arranged on the cortical surface. However, there is a considerable magnification in density of cortical neurons in V1 per degree of visual field. A shift of 1° in the visual field corresponds to a shift of 0.9 mm in the cortex (Troyer et al., 1998).

LGN neurons typically project to layer 4 of V1. In our model, we defined the orientation maps in a "hard-wired" manner (see Section 5.2.1.3). Therefore, the thalamic input connection probability to a cell in the circuit was modeled as an oriented Gabor function, i.e. a 2-dimensional Gaussian multiplied by a cosine function. The absolute value of the Gabor function corresponds to the connection probability of LGN neurons with a cortical cell positioned at the cortical equivalent position of the origin of the Gabor patch in the visual field. Positive and negative regions correspond to the connection probabilities of LGN ON- and OFF-response cells, respectively. Lagged and non-lagged cells connected equally likely to cortical cells. Following (Troyer et al., 1998), we expressed the Gabor function in parameters defining the number of sub-regions n_s, the aspect ratio of the width and the height of the Gaussian envelope a, the orientation ϕ, the offset of the cosine ψ, and the frequency of the cosine f. Given these parameters, one calculates the standard deviation of the Gaussian envelope as (see (Troyer et al., 1998))

$$\Sigma^{\frac{1}{2}} = \frac{1}{4 \cdot 2.448\, f} \begin{pmatrix} n_s & 0 \\ 0 & a \end{pmatrix} \tag{5.6}$$

while using coordinates rotated by ϕ. The advantage of using these parameters is that the frequency defines implicitly the size of the Gabor patch, while the number of sub-regions is kept constant. Therefore, the receptive fields of can be easily included in this framework. In correspondence with the data of (Movshon et al., 1978), we chose the frequency f from a Gaussian distribution with a mean of 0.8deg^{-1} and a standard deviation of 0.1deg^{-1} (with an enforced minimum of 0.6deg^{-1} and maximum of 1.0deg^{-1}). The number of sub-regions n_s and phase shifts ψ were drawn from uniform distributions with ranges of $(1.85, 2.65)$ and $(0, 2\pi)$, respectively (experimental values from cat as in (Troyer et al., 1998)).

To incorporate the smooth maps of preferred orientation ϕ and orientation preference q depending on cortical position u (see Section 5.2.1.3), we set $\phi = \phi(u)$ and the aspect ratio to $a(u) = (a_{\max} - a_{\min})\, q(u) + a_{\min}$. We used values reported by (Troyer et al., 1998) for the bounds $a_{\mathrm{min}} = 3.8$ and $a_{\mathrm{max}} = 4.54$ for excitatory neurons, and for the typically weakly tuned inhibitory neurons $a_{\mathrm{min}} = 1.4$ and $a_{\mathrm{max}} = 2.0$.

Lastly, the overall connection probability defined by the Gabor functions, was scaled to achieve an average number of 32 input synapses for both excitatory and inhibitory neurons. Thus, we set the connection probability to excitatory neurons in layer 5 (comprising layer 5 and layer 6) to 20% of that of the input to layer 4. These values are in good agreement with the data from (Binzegger et al., 2004) estimated from cat. An example of such connection probability from LGN neurons to a given neuron in layer 4 with the actually chosen LGN connections is shown in Fig. 5.1 B.

We sampled the delay of the LGN input synapses from a Gaussian distribution with mean 1.5 ms and standard deviation 0.15 ms (and additionally enforced delays below 0.5 ms and above 18 ms to a value uniformly in the latter range).

Top-down connections In addition to the thalamic input, V1 neurons receive multiple feedback connections from extra-striate cortical areas (Felleman and Essen, 1991), especially from V2, where the feedback connections are almost as numerous as the feed-forward connections (see (Sincich and Horton, 2005) for a review). Feedback projections predominantly project to targets in the upper layers. In the model, this top-down input consisted of 50×50 Poisson spiking neurons projecting to layer 2/3 with firing rates randomly drawn every 30 ms independently for each neuron uniformly between 10 and 15 Hz. The connection probability for such a connections was set to be 10 %.

5.2.1.3 Orientation map

It is well established that orientation preference and other features (such as visual field position, ocular dominance, or direction preference) form intertwined maps, where neighboring neurons tend to respond to similar features (Hubel and Wiesel, 1977; Obermayer and Blasdel, 1993).

We employed Kohonen's Self-Organizing Map algorithm (Kohonen, 1982) to create orientation maps across the cortical surface. An orientation attribute was necessary for each neuron for defining thalamic inputs, as well as for preferred orientation dependent patchy lateral long-range connections. The algorithm has been used to generate feature maps, which resembled cortical measured feature maps in their overall appearance, as well as e.g. the occurrence of pinwheels (Obermayer et al., 1990, 1992; Obermayer and Blasdel, 1993; Erwin et al., 1995; Brockmann et al., 1997). Basically, the algorithm tries to map a low-dimensional manifold (a horizontal sheet of neurons) to a high-dimensional feature space, while ensuring that neighboring points on the manifold exhibit similar feature preference. Let $\mathbf{z} = (x, y, q\cos(2\phi), q\sin(2\phi))^T$ define a feature vector, where $0 \leq x, y < k$ are the positions in visual space, $0 \leq q < 1$ is the orientation preference (or tuning strength) and $0 \leq \phi < \pi$ is the preferred orientation. We did not model ocular dominance because our V1 model received input only from one retina. If one uses the low-dimensional variant of the learning rule (Erwin et al., 1995; Obermayer and Blasdel, 1993), one attributes to each point on the manifold, i.e. each neuron having cortical 2D surface coordinates $u = (u_1, u_2)^T$, its current "optimal" feature vector $\mathbf{w}(u)$. Relations between neurons u and v are enforced by the neighborhood function $h(u, v) = \exp\left(-\frac{|u-v|^2}{2\delta^2}\right)$. The update of the feature vector of a neuron v can be written as

$$\Delta\mathbf{w}(v) = \alpha h(u_*, v)\left(\mathbf{z} - \mathbf{w}(u_*)\right). \tag{5.7}$$

Note that in each learning step the neuron u_*, showing maximal response to the current input $\mathbf{z}$, is updated in the direction of the input, weighted by a learning rate α. Depending on the *cortical* distance to the maximally activated neuron, the preferred features of the remaining neurons will be updated to a lesser extent in the same direction (mediated by the neighborhood function). In this rule we took the maximally activated neuron to be the nearest in feature space to the current input, $u_* = \text{argmin}|\mathbf{w}(u) - \mathbf{z}|$. We sampled the input features from uniform distribution (within the above bounds). k regulates the hierarchy between different features (Obermayer et al., 1992) and was set to $k = 5$. If one starts from a retinotopic initial condition, a high value for k ensures that cortical position corresponds to visual space in an approximate one-to-one map. The characteristic length scale δ was set to match the experimental observed correlation length in cortical orientation maps (corresponding to the distance of neighboring pinwheel center) of $d_{\text{pin}} = 660\mu\text{m}$ (Obermayer and Blasdel, 1993). We used the approximate formula $\delta = \sqrt{k}\, d_{\text{pin}}/D/8$, where D denotes the lateral extent of our V1 model.

5.2.2 Setup of the stimulus to the model

The stimulus, presented to the V1 model during simulation, resembled the one presented to the cat in the experiments. On the retina, the receptive field was centered at 5° eccentricity relative to the fovea. The stimulus presented to the retina model covered a visual field of 8° × 8° with letters covering about 5° × 5° of the visual field. The visual field covered by the V1 model is somewhat smaller (5° × 5°) than the stimulus to avoid boundary effects in the input connectivity. For analogous reasons, the LGN neurons were set to cover an intermediate area of 7.1° × 7.1° (50 × 50 grid). At 5° eccentricity a V1 model covering 5° × 5° has a lateral extent of $4.5 \times 4.5\,\text{mm}^2$ cortical surface. The neurons of the V1 model were positioned on a virtual grid of size 60 × 60 × 9 assuming a lateral grid spacing of 75 μm. Vertically, the grid spacing corresponds to 200 μm (see Method Section 5.2.1.1).

5.2.3 Simulation techniques

The model was implemented and simulated using the PCSIM simulation environment (Pecevski et al., 2009). All simulations were performed in a distributed fashion on a cluster of 30 quad core machines.

5.2.4 Classification tasks

The readout neurons performing a classification task were modeled as leaky integrate and fire (LIF) neurons. The recorded spike trains were convolved by an exponentially decaying kernel $\exp(-t/\tau)$ with a time constant $\tau = 20$ ms. This corresponds to a typical time course of excitatory postsynaptic potentials (EPSPs) in the readout neuron. The optimal firing threshold and weights of the readout neurons were learned by applying linear support vector machine (SVM). We used also SVMs with polynomial kernels (see (Schölkopf and Smola, 2002; Vapnik, 1998)) to evaluate the influence of higher order correlations on the readout performance. The parameter C of the SVMs was chosen to be 50. If not specified otherwise, we performed 150 simulation trials for each stimulus condition. All readout performance results represent the performance on a test set not shown during training. For most of the tasks different classifiers were trained for each time point t. We denote these classifiers as R_t. Classifiers trained on a time interval (300 ms) are denoted as R_{int}. These R_{int} classifiers use only one set of weights and a fixed threshold to classify the response over the time interval R_{int} classifiers typically perform worse than R_t classifiers because they have to generalize over the time interall, whereas the R_t classifiers can specialize on a given point in time. For all tasks the readout neuron had to classify two distinct classes, thus the chance level for a correct classification was 50 %.

5.3 Results

The model of a patch of cat V1 reproduces the experimental results reported in (Nikolic et al., 2009). In particular the model reproduces i) the non-linear temporal dynamics of information observed in the experiments, and ii) the performance drop caused by adding jitter to the recorded spike times. In addition the model suggest that iii) the variety of firing rate profiles reported in (Nikolic et al., 2009) may be partially explained by different positions (depths) of the electrodes, and iv) higher-order correlations contain little information about stimulus properties even if more data is available that would allow readouts to extract potential correlations. Finally, the model shows that the prolonged responses to single letter stimuli reported for the experiments can only be partially explained by effects of NMDA that only slightly increases the response durations.

A virtual 4×4 electrode array was positioned in the model to reproduce the setup of (Nikolic et al., 2009). We recorded the spiking activity of 64 neurons close to the 16 electrodes (4 neurons per electrode). The upper 12 electrodes were located in layer 2/3 and the remaining 4 lower electrodes in layer 4. Horizontally the electrodes were separated by 200 μm. The position relative to the retinotopic map of the cortex of the electrodes was chosen to be very similar as in the experiments (compare (Nikolic et al., 2009, Fig. 1 A) with Fig. 5.1 A). Spike trains recorded by the virtual electrode are shown in Fig. 5.2. The exact position of the virtual electrode in the model influences the temporal profile of the recorded mean firing rate. We analyzed this dependence in more detail in Section 5.3.4.

5.3.1 Similar neural spike responses

The spiking response recorded by the virtual electrode shows characteristic on and off responses for each letter of the stimulus sequence. The trial-to-trial variability of the spiking response of a single neuron (see Fig. 5.2 A, upper part) is comparable to the variability seen in experiment (compare with Fig. 1 A of (Nikolic et al., 2009)). However, in the experiment, neurons may be nearly silent for a full trial in a few cases. This effect cannot be seen in the simulation and may be an indication of some slow processes in the experiments, which are not present in the model. The Peri-stimulus time histogram (PSTH) plotted in the lower part of Fig. 5.2 A is very similar for two stimuli conditions (ABE, CBE). The variability in the responses of different neurons during a trial Fig. 5.2 B) resembles the variability observed in experimental recordings (compare with Fig. 1 B of (Nikolic et al., 2009)).

As in (Nikolic et al., 2009) we used two types of readout neurons, the R_t classifiers, trained for each time point and the R_{int} classifiers, trained to generalize on a time interval (see Method section for further details about the readout neurons).

Fig. 5.3 shows the performance of R_t classifiers for a single letter stimulus.

In the figures, the time interval at which letters are presented are indicated by gray boxes. The performance of the R_t classifiers correlates with the mean firing rate of the recorded neurons. The performance achieved very high (near perfect) values about 50 ms after the stimulus onset. Such a fast performance onset was also present for cat 4 in (Nikolic et al., 2009). This onset latency corresponds well to the 30-60 ms latencies after light stimulus reported by (Fries et al., 2001).

The performance stayed above chance level for about 220 ms with some additional small peaks significantly above chance level up to 420 ms after stimulus onset. The experiments showed a high variability across cats regarding the peak and duration of significant performance.

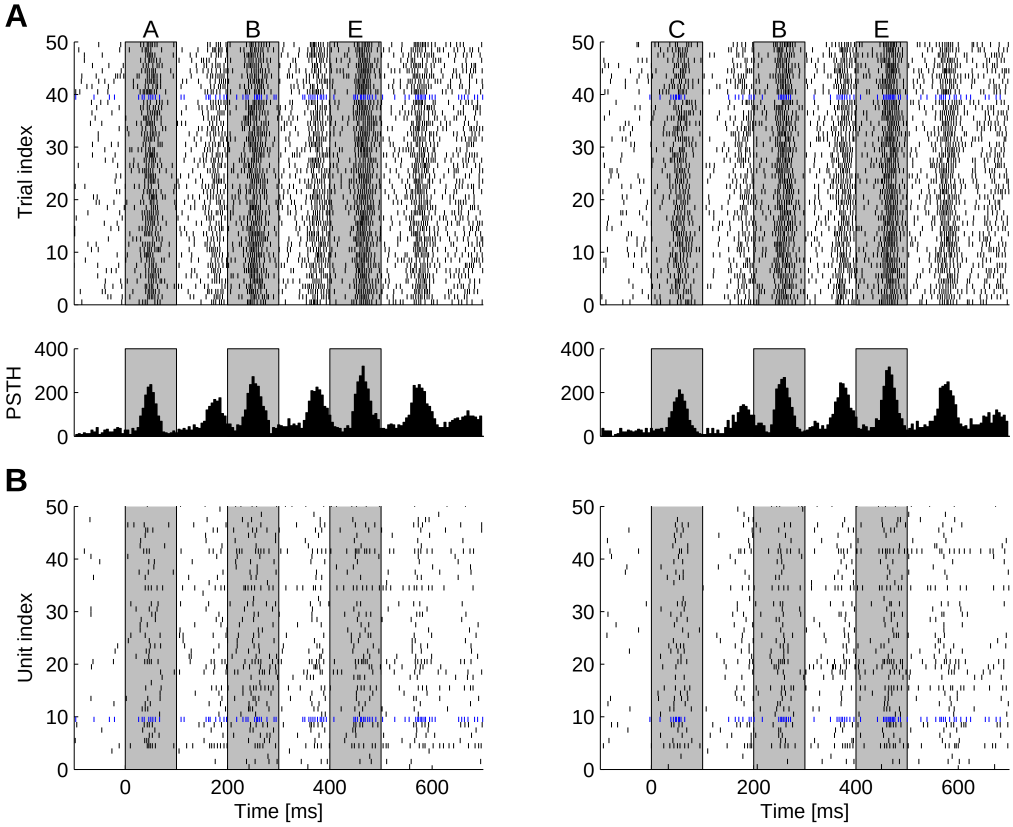

Figure 5.2: Circuit activity in response the Stimulus presentation
A: Upper part: Spike times of the recorded neuron with index 10 across 50 stimulation trials for two stimulus sequences (ABE and CBE). The gray boxes indicate the time intervals during which the letters of the stimulus were presented. Lower part: Peri-stimulus time histogram (PSTH) of the spikes of this neuron (5 ms bins). **B**: Spike trains of 50 recorded neurons for trial number 40 (compare with blue spike trains of **A**). The spikes of neuron with index 10 are plotted in blue. The spike trains of 50 trails of this neuron are shown in **A**.

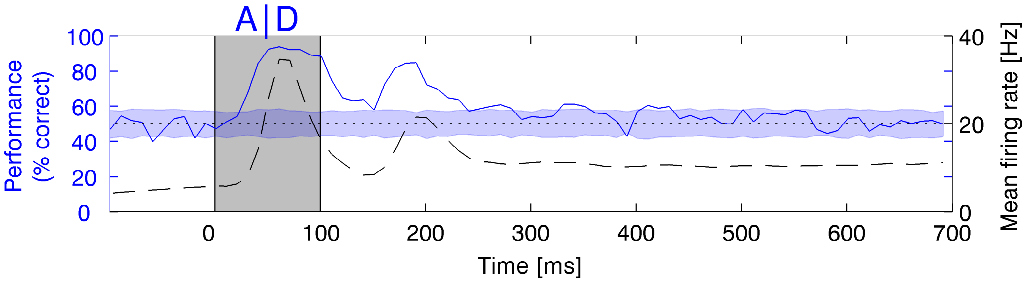

Figure 5.3: The ability of a readout to classify the identity of a letter used as stimulus.
The performance of R_t classifiers as a function of time is plotted as a solid line. The stimulus was presented for 100 ms starting at 0 ms. The presence of a stimulus is indicated by a gray box. The shaded region around the dotted line indicates the statistically non-significant deviations from the chance level ($p > 0.05$). The mean firing rate of the recorded neurons is plotted as a dash-dotted line.

5.3.2 Prolonged responses to single letters can not be explained by local synaptic connectivity

The mean firing rate responses in the experiments are typically longer than the responses of the model (compare Fig. 5.3 with Fig. 2 of (Nikolic et al., 2009)). The NMDA-to-AMPA ratio might have to be adjusted to create such a long response in the model. The NMDA-to-AMPA ratio shows a relatively high fluctuation in the experimental literature (Myme et al., 2003). Furthermore, the NMDA receptor function might be influenced by anesthesia (Narimatsu et al., 2002; Guntz et al., 2005). In order to adjust this parameter of the model to the experimental results, we increased the NMDA-to-AMPA ratio by a multiplicative factor of 2 and 3 in the model. Additionally, we simulated the model without NMDA receptors. The firing rate profile for the different setups is plotted in Fig. 5.4. The overall weight scale W_{scale} of the model was decreased in case of a higher NMDA-to-AMPA ratio and increased for the model without NMDA receptors, respectively, to result in about the same peak firing rate after the onset of the first letter of the stimulus. An increased NMDA-to-AMPA ratio resulted in an increased spontaneous firing rate but only weakly prolonged the response induced by the letter stimulus. The readout performance was only little affected by an increase of the NMDA-to-AMPA ratio (not shown).

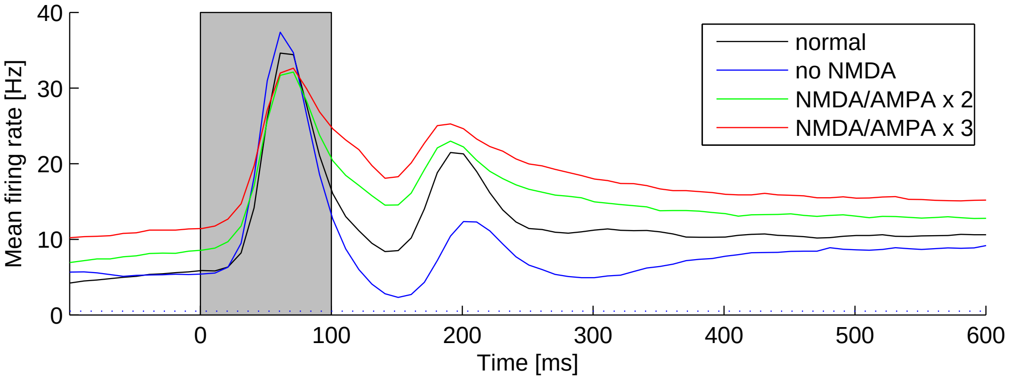

Figure 5.4: The impact of the NMDA-to-AMPA ratio on the firing rates of the recorded neurons.
The mean firing rate of the response is plotted for different NMDA-to-AMPA ratio. The spontaneous activity is enhanced by higher NMDA-to-AMPA ratios, but the response to the stimulus is only weakly prolonged.

5.3.3 Similar non-linear superposition of information

To investigated if a new stimulus erases the memory about the previous stimulus or if information about succeeding stimuli are superimposed in the circuit activity, we stimulated the model with sequences of letters.

The performance of the R_t classifiers stays above chance level for about 350 ms for the first letter (Fig. 5.5 A) and about 320 ms for the second letter (Fig. 5.5 B), respectively. These time intervals are even longer than for the single-letter stimulus condition. This behavior may be attributed to the firing rate increase during the presentation of the second letter and the third letter, respectively, because the readout performance is correlated to the firing rate of the recorded neurons. The time course of the R_t performance is comparable to the experimental results of cat 3 in (Nikolic et al., 2009) (compare with Fig. 4 C of (Nikolic et al., 2009)).

Additionally, we analyzed how the model combines two subsequent stimuli in a non-linear fashion (Fig. 5.5 C). The readout neurons were trained to do a XOR like classification, i.e. to be active if the first two letters of the sequence are AB or CD but to be inactive in case of CB or AD. For this task the linear spiking threshold was removed from the readout neuron to ensure that the non-linearity originates from the circuitry and not from the readout mechanism. Thus, we computed the so called point-biserial coefficient of correlation between the binary coded target value and the depolarization of the readout neuron to evaluate the performance for this task. For this task, the performance of the model peaks at about the same time as for cat 3 of (Nikolic et al., 2009), but the peak value is lower. The peak value of the model is comparable to the peak value during the experiments of cat 2 (compare with Supporting Fig. S21 of (Nikolic et al., 2009)).

5.3.4 The variety of firing rate profiles correlates with electrode locations

The time course of the mean firing rate for the three-letter experiments of (Nikolic et al., 2009) was different for each cat. This might be attributed to the actual depth, position and orientation (relative to the retinotopic map) of the electrode in the experiments. We checked if varying the depth, position and orientation of the virtual electrode in the model results in distinct temporal firing rate profiles. We computed the mean firing for all possible depths and positions of the electrode within the circuit for horizontally and vertically (with respect to the retinotopic map) oriented virtual electrodes. Fig. 5.6 shows three typical temporal firing rate profiles for vertical (Fig. 5.6**A**) and horizontally (Fig. 5.6**C**) oriented electrodes. These typical profiles were inferred by using a k-medoid clustering algorithm on the firing rate profiles. All the firing rate profiles resulted from simulations using the ABE letter sequence stimulus. For both orientations, the shape of the profiles are clustering mainly depending on the depth of the virtual electrode (compare with Fig. 5.6**B**,**D**). For virtual electrodes mainly located in layer 2/3, the first on response peak is typically smaller than the second on response peak. As the depth of the virtual electrode increases the ratio of the amplitude of the first to the second on response peak changes. For deeper electrode positions the first on response is stronger than the second on response. This observation holds for horizontally and vertically oriented virtual electrodes.

5.3.5 Similar impact of jitter on the computational performance

As in (Nikolic et al., 2009) we analyzed how strongly the performance of the readouts depends on the exact spike times of the recorded neurons. We tested this by applying a jitter to each spike time. Each spike time was moved by a random time drawn from a Gaussian distribution with zero mean and a predefined standard deviation (SD). The dependence of the R_t classifier performance on the amount of jitter added to each spike time is shown in Fig. 5.7 for three time points of the simulation.

For all three points in time, the classifier performance decreases as the SD of the jitter increases. During the the presentation of the first letter (Fig. 5.7 A), the performance of the model drops by about 5% for a jitter with SD = 20 ms. For the same case, the experimentally measured performance drop was in the range of about 10% to 17%. Particularly for the time point after the presentation of the first letter (Fig. 5.7 B), the results of the model correspond very well to the experimental results. For a jitter with SD = 20 ms the performance decreases by about 6% for the model and about 6%-10% in the experiment, respectivly. The maximal performance drop during the presentation of the second letter (Fig. 5.7 C) was about 8% for the model and about 9%-11% in the experiments.

The similar performance drop for the model and the experiment shows that this drop may not be attributed to a specific temporal coding of information in the spike times but is a common result for generic cortical microcircuits based on random connectivity pattern.

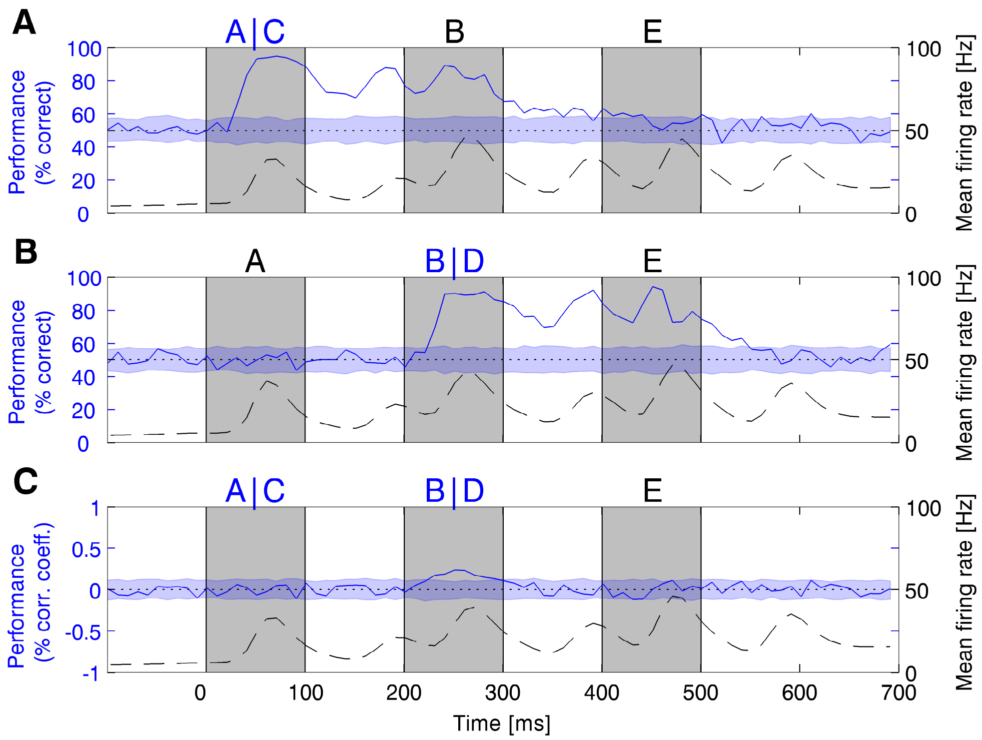

Figure 5.5: The information about different parts of the stimulus sequence is simultaneously available. **A**: Performance of readouts trained on individual time points to classify the the first letter of the three letter stimulus sequence. **B**: Performance as in **A** but for the second letter of the stimulus sequence. **C**: Performance in case of XOR classification. The performance for this task is measured by a point-biserial coefficient of correlation between the stimulus (binary code) and the depolarization of the readout neuron. Notations are the same as described in Fig. 5.3.

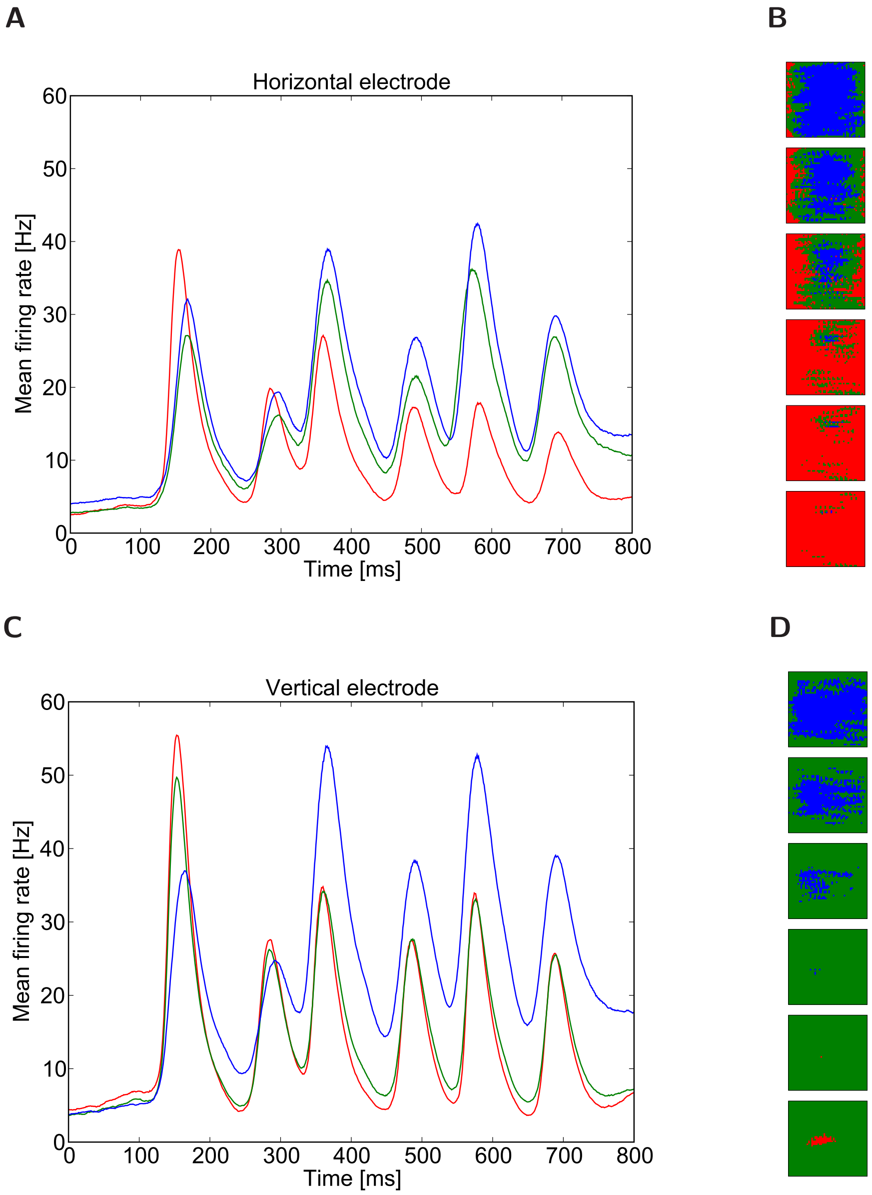

Figure 5.6: The temporal profile of the average firing rate of recorded neurons for different virtual electrode positions in response to a three letter sequence stimulus (ABE).
A,**C**: Three prototypical firing rate profiles inferred by a k-medoid clustering for vertically (**A**) or horizontally (**C**) positioned virtual electrodes.
B,**D**: The assignment of an electrode position to a cluster prototype for the different depths for vertically (**B**) and horizontally (**D**) positioned virtual electrodes. The depth of the virtual electrode increases by 200 μm per panel from top to bottom.

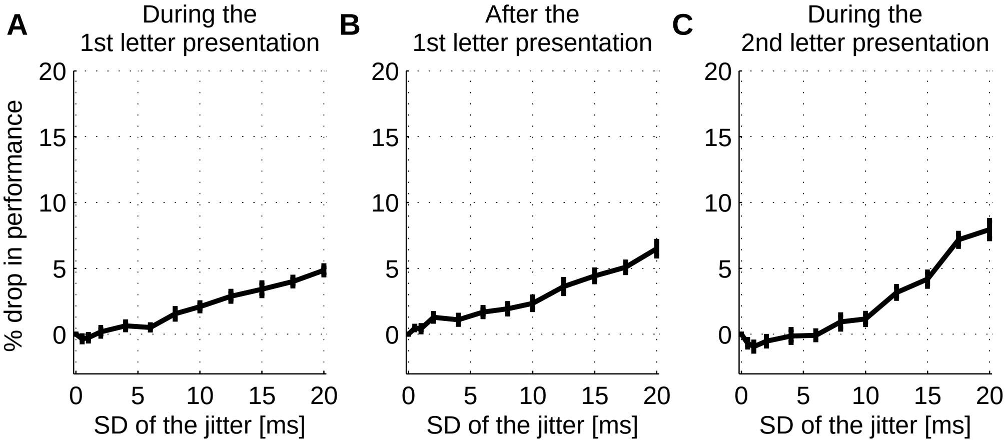

Figure 5.7: The classification performance of R_t readouts as a function of the amount of Gaussian jitter to spike times.
The drop in performance as a function of the SD of the jitter applied to the spike times during the presentation of the first letter (**A**), after the first letter (**B**), and during the presentation of the second letter (**C**).

Table 5.1: Average and maximum performance increase for polynomial kernels

Degree of polynomial kernel	2	3	4	5
Average increase (150 trials)	1.3 %	2.0 %	2.2 %	2.5 %
Average increase (2000 trials)	2.6 %	2.9 %	3.1 %	3.3 %
Maximum increase (150 trials)	9.3 %	13.3 %	17.0 %	15.7 %
Maximum increase (2000 trials)	13.2 %	14.6 %	14.7 %	13.5 %

5.3.6 Higher-order correlations contain little information about stimulus properties

(Nikolic et al., 2009) analyzed how the performance of R_{int} classifiers improves if one uses non-linear kernels. They used polynomial kernels of degree 2 to check if pairwise correlations can improve the readout performance. They found that for cat 1 and cat 3 a polynomial kernel of degree 2 increase the readout performance on average by 5.3 % and 2.0 %, respectively and the maximum performance increase was about 6 % and 10 %, respectively. As for the experiments with cat 3 of (Nikolic et al., 2009), we used the time interval $0 - 300$ ms relative to the onset of the stimulus. For polynomial kernels of degree 2, the average performance increase for the model for 150 trials per stimulus was 1.3 % and the the maximum increase was 9.3 % (for higher degrees see Table 5.1). These values agree well with the results of cat 3. The performance improvement for polynomial kernels peaks at about 250 ms after stimulus onset (Fig. 5.8**A**) when the second letter stimulus induces a change in the circuit activity. (Nikolic et al., 2009) argued that performance of classifiers with higher order kernels may improve if more trials would be available to the learning algorithm. By using the computer model of V1, we were able to generate many more trials (2000) than the 150 trials per stimulus in the experimental setup. Fig. 5.8**B** shows that the performance for higher order kernels increases only slightly for a dataset which is more than a magnitude bigger (2000 trials of the model vs. 150 experimental trials).

For other time intervals the dependence of the performance of the classifiers on the number of trials was very similar (not shown).

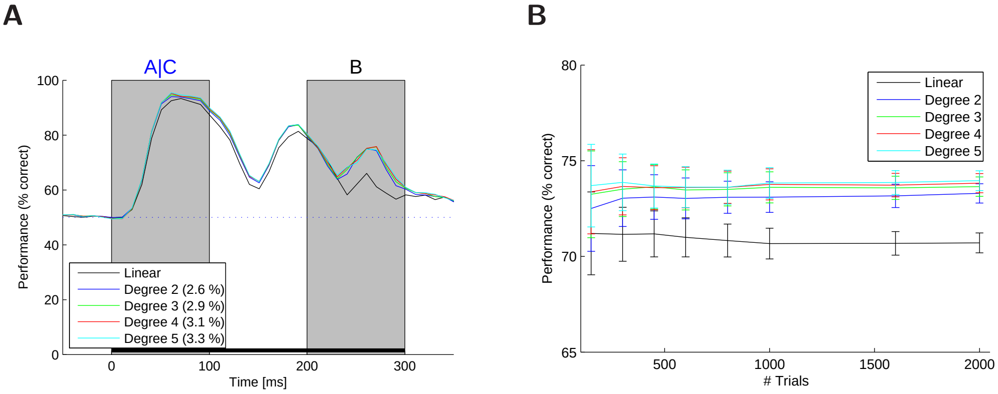

Figure 5.8: Performance of readouts using linear and polynomial kernels.
A: The performance of R_{int} readouts using linear and polynomial kernels of different degrees using 2000 trials per stimulus. The average performance improvement relative to the linear readout is given in the brackets of the legend. **B**: The dependence of the mean readout performance (average over the $0 - 300$ ms time interval) on the number of trials per stimulus for different different degrees of the polynomial kernel. The error bars denote the standard error (average over the $0 - 300$ ms time interval).

5.4 Discussion

We have shown that the temporal dynamics of information about sequences of visual stimuli contained in neural responses in cat primary visual cortex as reported in (Nikolic et al., 2009) can be replicated by a detailed cortical microcircuit model that implements stereotypical lamina-specific cortical synaptic connectivity.

The overall good agreement between model and experiment can be ascribed to certain model components. On one hand the results can be attributed to specific properties of the LGN model. First, the characteristic firing rate profiles of cortical neurons are similar for model and experiments. This is because the peaks in the mean firing rates of neurons within the cortical microcircuit model closely follow the corresponding peaks in the mean firing rates of cells in the LGN model. The latter result from characteristic on- and off-responses of LGN cells, which overlap for consecutive stimuli in a sequences presented at a frequency of 5 Hz. As long as the firing rates of LGN cells in response to a visual stimulus are high, information about this stimulus can be extracted from the circuit activity caused by the input. This was also observed experimentally.

Secondly, the information content in the fine temporal structure of neural responses is similar for model and experiments. In (Nikolic et al., 2009) it was shown that jitter with a standard deviation of a few millisecond causes a significant performance drop for readout neurons that were trained to extract information about previously presented letters from the spiking activity of cortical neurons. Furthermore, it was reported that this performance drop can be attributed to a neural code for which information is contained exclusively in the firing rates of neurons. These results are in agreement with the corresponding results for the model. In the model information about visual stimuli is explicitly coded in the instantaneous firing rates of cortical neurons because the models for the LGN and the retina transform visual stimuli into firing rates of LGN cells (inhomogeneous Poisson processes). However, the performance drop in the model after applying jitter with a standard deviation up to 5 ms is smaller than in the experiments. This indicates that the changes in the mean firing rates (i.e. the PSTHs) of cortical neurons in the model are slower compared to experiments.

On the other hand certain properties of the temporal dynamics of information about past input can be attributed to the recurrent synaptic connectivity within the cortical microcircuit model. First, the non-linear superposition of information about visual input presented at different moments in time as observed in experiments, occurs also in the model. This non-linear computation on visual input is carried out by the recurrent synaptic connectivity of the cortical microcircuit model because the model for the retina and the LGN perform only linear operations on visual input and therefore can not contribute to non-linear computations (as required for the XOR task).

Secondly, the prolonged neural activity in response to visual stimuli results from the recurrent synaptic connectivity as well. In particular NMDA synapses with large decay constants for their post-synaptic potentials with values of about 150 ms have a significant effect on the duration of the prolonged firing activity after stimulus offset. This is indicated by the positive correlation of the duration of the prolonged neural activity with the fraction of NMDA receptors within neurons. However, the fraction of NMDA receptors reported in the literature resulted in a significantly shorter duration of the neural responses in the model as observed in experiments.

Additionally, we apply the cortical microcircuit model to make predictions about three dynamical or computational properties of cat primary visual cortex that could not be answered in (Nikolic et al., 2009) based on experimental data. First, the firing rate profiles of neurons in response to sequences of visual stimuli depend on the location of the neurons within the neural tissue. The deeper the neurons are located within the neural tissue the more pronounced is their peak firing rate response to the first stimulus in a sequence. In contrast, neurons located in superficial layers have larger peak firing rates in response to subsequent stimuli in a sequence when compared to neurons located in deep layers. Apart from that the firing rate profiles show only little dependence on the horizontal location of the neurons. The variety of different firing rate profiles reported in (Nikolic et al., 2009) for different animals and recording sessions may therefore be attributed to different electrode positions, which were not identified in this study.

Secondly, all information about stimulus properties that could be extracted from the neural responses is contained in first and second order correlations between the discharge patterns of cortical neurons. The application of additional input samples did not provide significant additional information contained in higher-order correlations as was also observed in (Nikolic et al., 2009). This information could be missed in case of too small datasets due to the problem of overfitting. Our findings are consistent with results reported in (Hung et al., 2005) that show that an artificial readout system can extract stimulus-specific information from the spiking activity of neurons located in macaque inferotemporal (IT) cortex. Furthermore, previous studies examining synchronization among neuron populations (Schneidman et al., 2006; Yu et al., 2008) found that second-order correlations can account for almost all the correlation patterns observed in certain natural neuronal networks, and thus, higher-order correlations are unlikely to convey additional information.

Thirdly, the duration of the prolonged neural responses to single letter presentations can not only be attributed to local synaptic connectivity. Nevertheless the model exhibits prolonged neural activity on a time scale that is shorter compared to experiments. We have shown in the model that the duration of this responses depends on the fraction of NMDA receptors within a neuron, i.e. the more NMDA receptors the longer the response duration. The NMDA receptor function might be influenced by anesthesia (Guntz et al., 2005; Narimatsu et al., 2002). Previous reports (Narimatsu et al., 2002) indicate that ketamine anesthesia facilitates excitatory synaptic transmission by activating NMDA receptors. However, our results indicate that the prolonged neural responses reported in (Nikolic et al., 2009) can not only be attributed to an increased NMDA receptor activation. Tripling the ratio of NMDA to AMPA receptors in the model only results in an enhanced neural activity for up to 200 ms but not for up to 1 s as observed experimentally. These results support the hypothesis that the prolonged neural responses to single letter presentations reported in (Nikolic et al., 2009) may not simply be an artifact of the applied anesthetic but may be more likely related to the specific type of stimulus (white letters on black background were flashed for 100 ms).

All three predictions can be verified experimentally with i) precisely located recording electrodes, ii) larger datasets, and iii) *in vivo* recordings carried out in awake animals or under application of different anesthetics.

The results for the model and the experiments differ in two points. First, the changes in the mean firing rates (i.e. the PSTHs) of cortical neurons in the model in response to visual input are slower compared to experiments. This could be a result of shorter temporal integration time constant for LGN on- and off-cells in the cat brain. Alternatively, the peak in the firing rates of cortical neurons in response to LGN input in the cat brain could be sharpened due to more structured local cortical connectivity that implements for instance feedback mechanisms like gain control. Secondly, the prolonged neural responses to single letter stimuli are shorter for the model than for experimental data. A possible reason for this shorter responses are missing feedback loops with higher cortical areas that are not explicitly accounted for in the model. Input from higher cortical areas was modeled as conductance noise that is uncorrelated with visual input and therefore not capable to prolong stimulus induced activity. Further simulation studies that extend the current cortical microcircuit model by implementing higher cortical areas are required to clarify if a generic cortical model can account for the experimentally observed long lasting neural responses.

5.5 Acknowledgments

This chapter is based on the paper *Temporal dynamics of information in a model for cat primary visual cortex*, which was written by myself (KS), Stefan Häusler (SH), and Wolfgang Maass (WM). The experiments and analysis were performed by KS and SH, and the paper was written by SH and KS with additional input from WM.

Appendix A

List of Publications

1. D. Goswami, K. Schuch, Y. Zheng, T. De Marse, and J. C. Principe. *Towards the modeling of dissociated cortical tissue in the liquid state machine framework.* Proceedings of the International Joint Conference on Neural Networks (IJCNN), Montreal, Canada. 2005
2. K. Schuch. *Computational properties of threshold circuits and dissociated cortical tissue.* Master's thesis. Institute for Theoretical Computer Science, Graz University of Technology. 2006
3. S. Häusler, K. Schuch, and W. Maass. *Motif distribution and computational performance of two data-based cortical microcircuit templates.* 38th Annual Conference of the Society for Neuroscience. Program 220.9. 2008
4. S. Häusler, K. Schuch, and W. Maass. *Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates.* Journal of Physiology (Paris), 103(1-2). 2009.
5. D. Pecevski, T. Natschläger, and K. Schuch. *PCSIM: A Parallel Simulation Environment for Neural Circuits Fully Integrated with Python.* Frontiers in neuroinformatics 3. 2009
6. J. Bill, K. Schuch, D. Brüderle, J. Schemmel, K. Meier, and W. Maass. *Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity.* Submitted for publication. 2010
7. M. Rasch, K. Schuch, N. K. Logothetis, and W. Maass. *Statistical comparison of spike responses to natural stimuli in monkey area V1 with simulated responses of a detailed laminar network model for a patch of V1.* Submitted for publication. 2010
8. K. Schuch, S. Häusler, and W. Maass. *Temporal dynamics of information in a model for cat primary visual cortex.* In preparation. 2010

A.1 Comments and Contribution to Publications

The first two publications resulted from my master's thesis. These results are not contained in this thesis.

The paper *Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates* was a joint paper with Stefan Häusler (SH) and my supervisor Wolfgang Maass (WM). The experiments and analysis were performed by SH and KS, and the paper was written by SH and KS with additional input from WM. This paper is the basis for Chapter 2 of this thesis.

The paper *PCSIM: A Parallel Simulation Environment for Neural Circuits Fully Integrated with Python* is a joint paper with Thomas Natschläger (TN) of the Software Competence Center Hagenberg and Dejan Pecevski (DP). PCSIM was originally developed by TN and DP. I added some features to PCSIM that were necessary for the development of the V1 model. These results are not contained in this thesis.

The paper *Compensating inhomogeneities of neuromorphic VLSI devices via short-term synaptic plasticity* is a joint paper with Johannes Bill (JB), Daniel Brüderle (DB), Johannes Schemmel (JS), Karlheinz Meier (KM) from University of Heidelberg, and Wolfgang Maass (WM). The work for this paper was done within the FACETS project as a collaboration with the University of Heidelberg. The experiments were designed and planned by JB and KS during the visit of JB in our lab for his master thesis. The hardware experiments were performed by JB and DB. The software experiments were performed by KS and JB, and the paper was written by JB, KS, and DB with additional input from JS, KM, and WM. This paper is the basis for Chapter 3 of this thesis.

The paper *Statistical comparison of spike responses to natural stimuli in monkey area V1 with simulated responses of a detailed laminar network model for a patch of V1* is a joint paper with Malte Rasch (MR), Nikos Logothetis (NK), and Wolfgang Maass (WM). MR and KS equally contributed to this work. The in-vivo data was recorded in the lab of NL. The data analysis was performed by MR and KS. The circuit model was developed and implemented by MR and KS. The paper was written by MR, KS and WM with additional input from NL. This paper is the basis for Chapter 4 of this thesis.

The paper *Temporal dynamics of information in a model for cat primary visual cortex* is a joint paper with Stefan Häusler (SH) and Wolfgang Maass (WM). The experiments and analysis were performed by KS and SH, and the paper was written by SH and KS with additional input from WM. This paper is the basis for Chapter 5 of this thesis.

Bibliography

Abbott, L., Varela, J., Sen, K., and Nelson, S. (1997). Synaptic depression and cortical gain control. *Science*, 275:221–224.

Adorjan, P., Levitt, J. B., Lund, J. S., and Obermayer, K. (1999). A model for the intracortical origin of orientation preference and tuning in macaque striate cortex. *Vis Neurosci*, 16(2):303–18.

Alon, U. (2006). *An Introduction to Systems Biology: Design Principles of Biological Circuits*. Chapman & Hall/CRC, Boca Raton, US.

Baddeley, R., Abbott, L. F., Booth, M. C. A., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. (1997). Responses of neurons in primary and inferior temporal visual cortices to neural scenes. *Proc Biol Sci*, 264(1389):1775–1783.

Bair, W., Koch, C., Newsome, W., and Britten, K. (1994). Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J Neurosci*, 14(5 Pt 1):2870–92.

Bair, W., Zohary, W., and Newsome, W. (2001). Correlated firing in macaque visual area mt: Time scales and relationship to behavior. *J Neurosci*, 21(5):1676–1697.

Bartolozzi, C. and Indiveri, G. (2007). Synaptic dynamics in analog VLSI. *Neural Comput.*, 19(10):2581–2603.

Bartsch, A. P. and van Hemmen, J. L. (2001). Combined Hebbian development of geniculocortical and lateral connectivity in a model of primary visual cortex. *Biol Cybern*, 84(1):41–55.

Beaulieu, C., Kisvarday, Z., Somogyi, P., Cynader, M., and Cowey, A. (1992). Quantitative distribution of GABA-immunopositive and -immunonegative neurons and synapses in the monkey striate cortex (area 17). *Cereb Cortex*, 2(4):295–309.

Bedard, C., Kroger, H., and Destexhe, A. (2006). Does the 1/f frequency scaling of brain signals reflect self-organized critical states? *Phys Rev Lett*, 97(11):118102.

Beggs, J. M. and Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *J Neurosci*, 23(35):11167–11177.

Beggs, J. M. and Plenz, D. (2004). Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *J Neurosci*, 24(22):5216–5229.

Belitski, A., Gretton, A., Magri, C., Murayama, M., Montemurro, M., Logothetis, N., and Panzeri, S. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J Neurosci*, 28(22):5696–5709.

Bertschinger, N. and Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16:1413–1436.

Bi, G. and Poo, M. (1997). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Neural Computation*, 9:503–514.

Bill, J. (2008). Self-stabilizing network architectures on a neuromorphic hardware system. Diploma thesis (English), University of Heidelberg, HD-KIP-08-44.

Binzegger, T., Douglas, R. J., and Martin, K. A. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.*, 24(39):8441–8453.

Blasdel, G. G., Lund, J. S., and Fitzpatrick, D. (1985). Intrinsic connections of macaque striate cortex: axonal projections of cells outside lamina 4C. *J Neurosci*, 5(12):3350–69.

Blumenfeld, B., Bibitchkov, D., and Tsodyks, M. (2006). Neural network model of the primary visual cortex: from functional architecture to lateral connectivity and back. *J Comput Neurosci*, 20(2):219–41.

Boegershausen, M., Suter, P., and Liu, S.-C. (2003). Modeling short-term synaptic depression in silicon. *Neural Computation*, 15(2):331–348.

Borg-Graham, L., Monier, C., and Fregnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393:369.

Bredfeldt, C. E. and Ringach, D. L. (2002). Dynamics of spatial frequency tuning in macaque V1. *J Neurosci*, 22(5):1976–84.

Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J. M., Diesmann, M., Morrison, A., Goodman, P. H., Harris Jr, F. C., Zirpe, M., Natschlager, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A. P., Boustani, S. E., and Destexhe, A. (2007). Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience*, 23(3):349–398.

Brockmann, D., Bauer, H. U., Riesenhuber, M., and Geisel, T. (1997). SOM-model for the development of oriented receptive fields and orientation maps from non-oriented ON-center OFF-center inputs. In Gerstner, W., Germond, A., Hasler, M., and Nicoud, J. D., editors, *Artificial Neural Networks—ICANN '97. 7th International Conference Proceedings*, pages 207–12. Springer-Verlag, Berlin, Germany.

Brüderle, D. (2009). Neuroscientific modeling with a mixed-signal VLSI hardware system. Ph.D. thesis, University of Heidelberg, HD-KIP 09-30.

Brüderle, D., Müller, E., Davison, A., Muller, E., Schemmel, J., and Meier, K. (2009). Establishing a novel modeling tool: A python-based interface for a neuromorphic hardware system. *Front. Neuroinform.*, 3(17).

Brunel, N. (2000). Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. *J Physiol Paris*, 94(5–6):445–463. incl. erratum: vol. 96(1–6)2001.

Buesing, L., Schrauwen, B., and Legenstein, R. (2010). Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. *Neural Computation*, 22:1272–1311.

Buzas, P., Kovacs, K., Ferecsko, A. S., Budd, J. M. L., Eysel, U. T., and Kisvarday, Z. F. (2006). Model-based analysis of excitatory lateral connections in the visual cortex. *J Comp Neurol*, 499(6):861–81.

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci*, 21:47–74.

Callaway, E. M. (2004). Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Networks*, 17(5-6):625–632.

Cattaneo, A., Maffei, L., and Morrone, C. (1981). Two firing patterns in the discharge of complex cells encoding different attributes of the visual stimulus. *Exp Brain Res*, 43:115–118.

Chance, F. S., Nelson, S. B., and Abbott, L. F. (1999). Complex cells as cortically amplified simple cells. *Nat Neurosci*, 2(3):277–82.

Conway, B. R. and Livingstone, M. S. (2006). Spatial and temporal properties of cone signals in alert macaque primary visual cortex. *J Neurosci*, 26(42):10826–46.

Croner, L. J. and Kaplan, E. (1995). Receptive fields of P and M ganglion cells across the primate retina. *Vision Res*, 35(1):7–24.

Dally, W. J. and Poulton, J. W. (1998). *Digital systems engineering.* Cambridge University Press, New York, NY, USA.

Dantzker, J. L. and Callaway, E. M. (2000). Laminar sources of synaptic input to cortical inhibitory interneurons and pyramidal neurons. *Nature Neuroscience*, 3(7):701–707.

Davison, A. P., Brüderle, D., Eppler, J., Kremkow, J., Muller, E., Pecevski, D., Perrinet, L., and Yger, P. (2008). PyNN: a common interface for neuronal network simulators. *Front. Neuroinform.*, 2(11).

Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems.* The MIT Press.

Destexhe, A., Contreras, D., and Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J Neurosci*, 19(11):4595–608.

Destexhe, A., Mainen, Z. F., and Sejnowski, T. J. (1994). Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *J Comput Neurosci*, 1:195–230.

Destexhe, A., Rudolph, M., Fellous, J. M., and Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24.

Destexhe, A., Rudolph, M., and Pare, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*, 4:739–751.

Diesmann, M., Gewaltig, M.-O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529–533.

Dong, D. W. and Atick, J. J. (1995). Temporal decorrelation: A theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in Neural Systems*, 6(2):159–178.

Douglas, R., Mahowald, M., and Mead, C. (1995a). Neuromorphic analogue VLSI. *Annu. Rev. Neurosci.*, 18:255–281.

Douglas, R. J., Koch, C., Mahowald, M., Martin, K., and Suarez, H. (1995b). Recurrent excitation in neocortical circuits. *Science*, 269(5226):981–985.

Douglas, R. J. and Martin, K. A. (2004). Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.*, 27:419–451.

Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., and Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science*, 327:584–587.

Enroth-Cugell, C. and Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *J Physiol*, 187(3):517–52.

Ernst, U. A., Pawelzik, K. R., Sahar-Pikielny, C., and Tsodyks, M. (2001). Intracortical origin of visual maps. *Nat Neurosci*, 4(4):431–6.

Erwin, E., Obermayer, K., and Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput*, 7(3):425–68.

FACETS (2009). Fast Analog Computing with Emergent Transient States – project website. `http://www.facets-project.org`.

Felleman, D. J. and Essen, D. C. V. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1(1):1–47.

Fieres, J., Schemmel, J., and Meier, K. (2008). Realizing biological spiking network models in a configurable wafer-scale hardware system. In *Proceedings of the 2008 International Joint Conference on Neural Networks (IJCNN)*.

Fries, P., Neuenschwander, S., Engel, A. K., and Goebel, R. (2001). Rapid feature selective neuronal synchronization through correlated latency shifting. *Nature Neuroscience*, 4:194–200.

Ganguli, S., Huh, D., and Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proc Natl Acad Sci U S A*, 105(48):18970–5.

Gazeres, N., Borg-Graham, L. J., and Fregnac, Y. (1998). A phenomenological model of visually evoked spike trains in cat geniculate nonlagged X-cells. *Vis Neurosci*, 15(6):1157–74.

Gerstner, W. and Kistler, W. M. (2002). *Spiking neuron Models: Single Neurons, Populations, Plasticity.* Cambridge University Press.

Gilbert, C. D., Das, A., Ito, M., Kapadia, M., and Westheimer, G. (1996). Spatial integration and cortical dynamics. *Proc Natl Acad Sci U S A*, 93(2):615–22.

Girard, P., Hupe, J. M., and Bullier, J. (2001). Feedforward and feedback connections between areas V1 and V2 of the monkey have similar rapid conduction velocities. *J Neurophysiol*, 85(3):1328–31.

Grübl, A. (2007). *VLSI Implementation of a Spiking Neural Network*. PhD thesis, Ruprecht-Karls-University, Heidelberg. Document No. HD-KIP 07-10.

Guido, W., Lu, S. M., and Sherman, S. M. (1992). Relative contributions of burst and tonic responses to the receptive field properties of lateral geniculate neurones in the cat. *Journal of Neurophysiology*, 68:2199–2211.

Guido, W. and Sherman, S. M. (1998). Response latencies of cells in the cat's lateral geniculate neurones are less variable during burst than tonic firing. *Visual Neuroscience*, 15:231–237.

Guntz, E., Dumont, H., Roussel, C., Gall, D., Dufrasne, F., Cuvelier, L., Blum, D., Schiffmann, S., and Sosnowski, M. (2005). Effects of remifentanil on n-methyl-d-aspartate receptor: an electrophysiologic study in rat spinal chord. *Anesthesiology*, 102(6):1235–41.

Gupta, A., Wang, Y., and Markram, H. (2000). Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, 287:273–278.

Haeusler, S., Schuch, K., and Maass, W. (2009). Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates. *Journal of Physiology-Paris*, 103:73–87.

Häusler, S. and Maass, W. (2007). A statistical analysis of information processing properties of lamina-specific cortical microcircuit models. *Cerebral Cortex*, 17(1):149–62.

Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol*, 551(Pt 1):139–53.

Holt, G. R., Softky, W. R., and Douglas, R. J. (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *J Neurophysiol*, 75:1806–1814.

Hubel, D. H. and Wiesel, T. N. (1977). Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond B Biol Sci*, 198(1130):1–59.

Hung, C. P., Kreiman, G., Poggio, T., and DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, 310(5749):863–866.

Indiveri, G., Chicca, E., and Douglas, R. (2006). A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. *IEEE Transactions on Neural Networks*, 17(1):211–221.

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans Neural Netw*, 14(6):1569–72.

Izhikevich, E. M. (2006). *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting (Computational Neuroscience)*. The MIT Press.

Izhikevich, E. M. and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc Natl Acad Sci U S A*, 105(9):3593–8.

Izhikevich, E. M., Gally, J. A., and Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cerebral Cortex*, 14(8):933–44.

Jensen, H. J. (1998). *Self-Organized Criticality: Emergent Complex Behavior in Physical and Biological Systems (Cambridge Lecture Notes in Physics)*. Cambridge University Press.

Johansson, C. and Lansner, A. (2007). Towards cortex sized artificial neural systems. *Neural Netw*, 20(1):48–61.

Kaiser, M. and Hilgetag, C. (2004). Spatial growth of real-world networks. *Phys. Rev. E*, 69:036103.

Kalisman, N., Silberberg, G., and Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proc Natl Acad Sci*, 102(3):880–885.

Kampa, B. M., Letzkus, J. J., and Stuart, G. J. (2006). Cortical feed-forward networks for binding different streams of sensory information. *Nat Neurosci*, 9(12):1472–3.

Kannan, R., Tetali, P., and Vempala, S. (1999). Simple markov-chain algorithms for generating bipartite graphs and tournaments. *Random Structures and Algorithms*, 14(4):293–08.

Kaplan, B., Brüderle, D., Schemmel, J., and Meier, K. (2009). High-conductance states on a neuromorphic hardware system. In *Proceedings of the 2009 International Joint Conference on Neural Networks (IJCNN)*.

Koch, C., Poggio, T., and Torre, V. (1983). Nonlinear interactions in a dendritice tree: localization, timing, and role in information processing. *PNAS*, 80:2799–2802.

Kohn, A. and Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of macaque. *J Neurosci*, 25(14):3661–3673.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69.

Kremkow, J., Kumar, A., Rotter, S., and Aertsen, A. (2007). Emergence of population synchrony in a layered network of the cat visual cortex. *Neurocomput*, 70(10-12):2069–2073.

Legenstein, R. and Maass, W. (2007). Edge of chaos and prediction of computational performance for neural microcircuit models. *Neural Networks*, 20(3):323–334.

Lisman, J. (1997). Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci*, 20(1):38–43.

Logothetis, N., Merkle, H., Augath, M., Trinath, T., and Ugurbil, K. (2002). Ultra high-resolution fMRI in monkeys with implanted RF coils. *Neuron*, 35(2):227–42.

Logothetis, N. K., Guggenberger, H., Peled, S., and Pauls, J. (1999). Functional imaging of the monkey brain. *Nat. Neurosci.*, 2(6):555–62.

Lund, J. S., Angelucci, A., and Bressloff, P. C. (2003). Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cerebral Cortex*, 13(1):15–24.

Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560.

Maass, W., Natschläger, T., and Markram, H. (2004). Fading memory and kernel properties of generic cortical microcircuit models. *J Physiol Paris*, 98(4-6):315–30.

Markram, H. (1997). A network of tufted layer 5 pyramidal neurons. *Cerebral Cortex*, 7:523–33.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci*, 5(10):793–807.

Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *PNAS*, 95:5323–5328.

Maslov, S. and Sneppen, K. (2002). Specificity and stability in topology of protein networks. *Science*, 296(5569):910–3.

Maunsell, J. H., Ghose, G. M., Assad, J. A., McAdams, C. J., Boudreau, C. E., and Noerager, B. D. (1999). Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. *Vis Neurosci*, 16(1):1–14.

McGuire, B. A., Gilbert, C. D., Rivlin, P. K., and Wiesel, T. N. (1991). Targets of horizontal connections in macaque primary visual cortex. *J Comp Neurol*, 305(3):370–92.

Mead, C. A. (1989). *Analog VLSI and Neural Systems.* Addison Wesley, Reading, MA.

Merolla, P. A. and Boahen, K. (2006). Dynamic computation in a recurrent network of heterogeneous silicon neurons. In *Proceedings of the 2006 IEEE International Symposium on Circuits and Systems (ISCAS 2006)*.

Milo, R., Itzkovitz, S., Kashtan, N., Levitt, R., Shen-Orr, S., Ayzenshtat, I., Sheffer, M., and Alon, U. (2004). Superfamilies of evolved and designed networks. *Science*, 303(5663):1538–42.

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., and Alon, U. (2002). Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–7.

Mitra, S., Fusi, S., and Indiveri, G. (2009). Real-time classification of complex patterns using spike-based learning in neuromorphic VLSI. *IEEE Transactions on Biomedical Circuits and Systems*, 3:(1):32–42.

Montemurro, M., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr Biol*, 18:375–380.

Morrison, A., Diesmann, M., and Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics*, 98(6):459–478.

Morrison, A., Mehring, C., Geisel, T., Aertsen, A., and Diesmann, M. (2005). Advancing the boundaries of high connectivity network simulation with distributed computing. *Neural Comput.*, 17(8):1776–1801.

Mountcastle, V. B. (1998). *Perceptual Neuroscience: The Cerebral Cortex*. Harvard University Press, Cambridge.

Movshon, J. A., Thompson, I. D., and Tolhurst, D. J. (1978). Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *J Physiol*, 283:101–20.

Müller, E. (2008). Operation of an imperfect neuromorphic hardware device. Diploma thesis (English), University of Heidelberg, HD-KIP-08-43.

Myme, C. I., Sugino, K., Turrigiano, G. G., and Nelson, S. B. (2003). The NMDA-to-AMPA ratio at synapses onto layer 2/3 pyramidal neurons is conserved across prefrontal and visual cortices. *J Neurophysiol*, 90(2):771–779.

Narimatsu, E., Kawamata, Y., Kawamata, M., Fujimura, N., and Namiki, A. (2002). NMDA receptor-mediated mechanism of ketamine-induced facilitation of glutamatergic excitatory synaptic transmission. *Brain Res*, 953(1-2):272–5.

Natschläger, T. and Maass, W. (2005). Dynamics of information and emergent computation in generic neural microcircuit models. *Neural Networks*, 18(10):1301–1308.

Natschläger, T., Markram, H., and Maass, W. (2003). Computer models and analysis tools for neural microcircuits. In Kötter, R., editor, *Neuroscience Databases. A Practical Guide*, chapter 9, pages 123–138. Kluwer Academic Publishers (Boston).

Nelson, S. (2002). Cortical microcircuits: diverse or canonical? *Neuron*, 36(1):19–27.

Newman, M. E. (2003). The structure and function of complex networks. *SIAM Review*, 45:167–256.

Nikolic, D., Haeusler, S., Singer, W., and Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biology*, 7(12):1–19.

Obermayer, K. and Blasdel, G. G. (1993). Geometry of orientation and ocular dominance columns in monkey striate cortex. *J Neurosci*, 13(10):4114–29.

Obermayer, K., Blasdel, G. G., and Schulten, K. (1992). Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys Rev A*, 45(10):7568–7589.

Obermayer, K., Ritter, H., and Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proc Natl Acad Sci U S A*, 87(21):8345–9.

O'Kusky, J. and Colonnier, M. (1982). Postnatal changes in the number of astrocytes, oligodendrocytes, and microglia in the visual cortex (area 17) of the macaque monkey: a stereological analysis in normal and monocularly deprived animals. *J Comp Neurol*, 210(3):307–15.

Olshausen, B. A. and Field, D. J. (2005). How close are we to understanding v1? *Neural Comput*, 17(8):1665–99.

Olshausen, B. A. and Field, D. J. (2006). What is the other 85% of V1 doing? In Sejnowski, T. J. and van Hemmen, L., editors, *Problems in Systems Neuroscience*, pages 182–211. Oxford Univ. Press.

Pecevski, D., Natschläger, T., and Schuch, K. (2009). PCSIM: a parallel simulation environment for neural circuits fully integrated with python. *Front Neuroinformatics*, 3(11).

Perry, V. H., Oehler, R., and Cowey, A. (1984). Retinal ganglion cells that project to the dorsal lateral geniculate nucleus in the macaque monkey. *Neuroscience*, 12(4):1101–23.

Peters, A. and Payne, B. R. (1993). Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cereb Cortex*, 3(1):69–78.

Peters, A., Payne, B. R., and Budd, J. (1994). A numerical analysis of the geniculocortical input to striate cortex in the monkey. *Cerebral Cortex*, 4(3):215–29.

Potjans, T. C. and Diesmann, M. (2008). Integration of anatomical and physiological connectivity data sets for layered cortical network models. *BMC Neuroscience 2008*, 9(Suppl 1)(P60):Seventeenth Annual Computational Neuroscience Meeting: CNS*2008, July 19th – 24th 2008, Portland, Oregon, USA.

Quian Quiroga, R., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Computation*, 16:1661–1687.

Rasch, M. J., Gretton, A., Murayama, Y., Maass, W., and Logothetis, N. K. (2007). Inferring spike trains from local field potentials. *Journal of Neurophysiology*. in press.

Rasch, M. J., Logothetis, N. K., and Kreiman, G. (2009). From neurons to circuits: Linear estimation of local field potentials. *J Neurosci*, 29(44):13785–13796.

Reich, D. S., Mechler, F., and Victor, J. D. (2001). Independent and redundant information in nearby cortical neurons. *Science*, 294:2566–2568.

Reinagel, P., Godwin, D., Sherman, S. M., and Koch, C. (1999). The encoding of visual information by LGN bursts. *J Neurophysiol*, 81:2558–2569.

Rockland, K. S. and Virga, A. (1989). Terminal arbors of individual "feedback" axons projecting from area V2 to V1 in the macaque monkey: a study using immunohistochemistry of anterogradely transported Phaseolus vulgaris-leucoagglutinin. *J Comp Neurol*, 285(1):54–72.

Rodieck, R. W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Res*, 5(11):583–601.

Rossum, G. V. (2000). *Python Reference Manual: February 19, 1999, Release 1.5.2*. iUniverse, Incorporated.

Salinas, E. and Sejnowski, T. J. (2000). Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J Neurosci*, 20:6193–6209.

Schein, S. J. and de Monasterio, F. M. (1987). Mapping of retinal and geniculate neurons onto striate cortex of macaque. *J Neurosci*, 7(4):996–1009.

Schemmel, J., Brüderle, D., Meier, K., and Ostendorf, B. (2007). Modeling synaptic plasticity within networks of highly accelerated I&F neurons. In *Proceedings of the 2007 IEEE International Symposium on Circuits and Systems (ISCAS'07)*. IEEE Press.

Schemmel, J., Fieres, J., and Meier, K. (2008). Wafer-scale integration of analog neural networks. In *Proceedings of the 2008 International Joint Conference on Neural Networks (IJCNN)*.

Schemmel, J., Grübl, A., Meier, K., and Muller, E. (2006). Implementing synaptic plasticity in a VLSI spiking neural network model. In *Proceedings of the 2006 International Joint Conference on Neural Networks (IJCNN'06)*. IEEE Press.

Schneidman, E., 2nd Berry, M. J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–12.

Schölkopf, B. and Smola, A. J. (2002). *Learning with Kernels*. MIT Press, Cambridge, MA.

Schrauwen, B., Büsing, L., and Legenstein, R. (2009). On computational power and the order-chaos phase transition in reservoir computing. In *Proc. of NIPS 2008, Advances in Neural Information Processing Systems*, volume 21. MIT Press. in press.

Schrauwen, B., Wardermann, M., Verstraeten, D., Steil, J. J., and Stroobandt, D. (2008). Improving reservoirs using intrinsic plasticity. *Neurocomputing*, 71(7-9):1159–1171.

Serre, T., Kouh, M., Cadieu, C., Knoblich, U., Kreiman, G., and Poggio, T. (2005). A theory of object recognition: computations and circuits in the feedforward path of the ventral stream in primate visual cortex. *MIT, Cambridge, MA*, AI Memo 2005-036(CBCL Memo 259).

Shadlen, M. N. and Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci*, 18:3870–3896.

Shelley, M., McLaughlin, D., Shapley, R., and Wielaard, J. (2002). States of high conductance in a large-scale model of the visual cortex. *J. Comp. Neurosci.*, 13:93–109.

Shen-Orr, S. S., Milo, R., Mangan, S., and Alon, U. (2002). Network motifs in the transcriptional regulation network of escherichia coli. *Nature Genetics*, 31:64–8.

Sherman, S. M. (1996). Dual response modes in lateral geniculate nucleus: Mechanisms and function. *Visual Neuroscience*, 13:205–213.

Silberberg, G., Gupta, A., and Markram, H. (2002). Stereotypy in neocortical microcircuits. *Trends Neurosci.*, 25(5):227–230.

Sincich, L. C. and Horton, J. C. (2005). The circuitry of V1 and V2: integration of color, form, and motion. *Annu Rev Neurosci*, 28:303–26.

Smith, M. A. and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci*, 28(48):12591–12603.

Smith, M. A. and Kohn, A. (2009). Lamina dependence of neuronal correlation in macaque v1. *Program No. 166.8/V5. Neuroscience Meeting Planner. Chicago, IL: Society for Neuroscience, 2009. Online.*

Softky, W. R. and Koch, C. (1992). Cortical cells should fire regularly, but do not. *Neural Computation*, 4(5):643–46.

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci*, 13(1):334–50.

Song, S., Miller, K., and Abbott, L. (2000). Competitive hebbian learning through spiketiming-dependent synaptic plasticity. *Nat. Neurosci.*, 3:919–926.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*, 3(3):e68+.

Steriade, M. (2001). *The Intact and Sliced Brain.* MIT Press.

Steriade, M. (2003). *Neuronal Substrates of Sleep and Epilepsy.* Cambridge University Press, Cambridge, UK.

Steriade, M., McCormick, D. A., and Sejnowski, T. J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science*, 262(5134):679.

Steriade, M., Timofeev, I., and Grenier, F. (2001). Natural waking and sleep states: A view from inside neocortical neurons. *J Neurophysiol*, 85:1969 – 1985.

Stevens, C. F. and Zador, A. M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci*, 1(3):210–217.

Strogatz, S. H. (2001). Exploring complex networks. *Nature*, 410:268–76.

Sussillo, D., Toyoizumi, T., and Maass, W. (2007). Self-tuning of neural circuits through short-term synaptic plasticity. *J Neurophysiol*, 97(6):4079–4095.

Tadmor, Y. and Tolhurst, D. J. (2000). Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes. *Vision Res*, 40(22):3145–57.

Tao, L., Shelley, M., McLaughlin, D., and Shapley, R. (2004). An egalitarian network model for the emergence of simple and complex cells in visual cortex. *Proc Natl Acad Sci U S A*, 101(1):366–71.

Teich, A. F. and Qian, N. (2006). Comparison among some models of orientation selectivity. *J Neurophysiol*, 96(1):404–19.

Teich, M. C., Heneghan, C., Lowen, S. B., Ozaki, T., and Kaplan, E. (1997). Fractal character of the neural spike train in the visual system of the cat. *J Opt Soc Am A Opt Image Sci Vis*, 14(3):529–546.

Thomson, A. M. (2005). personal communication.

Thomson, A. M. and Destexhe, A. (1999). Dual intracellular recordings and computational models of slow inhibitory postsynaptic potentials in rat neocortical and hippocampal slices. *Neuroscience*, 92:1193–1215(23).

Thomson, A. M. and Lamy, C. (2007). Functional maps of neocortical local circuitry. *Frontiers in Neurosci*, 1(1):19–42.

Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2 - 5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. *Cerebral Cortex*, 12(9):936–953.

Tootell, R. B., Silverman, M. S., Switkes, E., and Valois, R. L. D. (1982). Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science*, 218(4575):902–4.

Triesch, J. (2007). Synergies between intrinsic and synaptic plasticitymechanisms. *Neural Comput*, 19(4):885–909.

Troyer, T. W., Krukowski, A. E., Priebe, N. J., and Miller, K. D. (1998). Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *J Neurosci*, 18(15):5908–27.

Tsubo, Y., Isomura, Y., and Fukai, T. (2009). Power-law distributions of interspike intervals in in vivo cortical neurons. In *Frontiers in Systems Neuroscience. Conference Abstract: Computational and systems neuroscience.*

Tyler, C. J., Dunlop, S. A., Lund, R. D., Harman, A. M., Dann, J. F., Beazley, L. D., and Lund, J. S. (1998). Anatomical comparison of the macaque and marsupial visual cortex: common features that may reflect retention of essential cortical elements. *J Comp Neurol*, 400(4):449–68.

Vapnik, V. N. (1998). *Statistical Learning Theory*. John Wiley (New York).

Vinje, W. E. and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287:1273–1276.

Vogels, T. P. and Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci*, 25(46):10786–10795.

Vogelstein, R. J., Mallik, U., Vogelstein, J. T., and Cauwenberghs, G. (2007). Dynamically reconfigurable silicon array of spiking neuron with conductance-based synapses. *IEEE Transactions on Neural Networks*, 18:253–265.

Wang, H., Spencer, D., Fellous, J., and Sejnowski, T. (2010). Synchrony of thalamocortical inputs maximizes cortical reliability. *Science*, 328(5974):106–109.

Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J., and Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat Neurosci*, 9(4):534–42.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393:440 – 442.

Wenisch, O. G., Noll, J., and van Hemmen, J. L. (2005). Spontaneously emerging direction selectivity maps in visual cortex through STDP. *Biol Cybern*, 93(4):239–47.

White, E. L. (1989). *Cortical Circuits*. Birkhaeuser, Boston.

Wielaard, D. J., Shelley, M., McLaughlin, D., and Shapley, R. (2001). How simple cells are made in a nonlinear network model of the visual cortex. *J Neurosci*, 21(14):5203–11.

Yen, S. C., Baker, J., and Gray, C. M. (2007). Heterogeneity in the responses of adjacent neurons to natural stimuli in cat striate cortex. *J Neurophysiol*, 97(2):1326–1341.

Yen, S. C., Baker, J., Lachaux, J. P., and Gray, C. M. (2008). Responses of primary visual cortical neurons to natural movies in anesthetized cat. In *Seventeenth Annual Computational Neuroscience Meeting*.

Yu, S., Huang, D., Singer, W., and Nikolic, D. (2008). A small world of neuronal synchrony. *Cereb Cortex*, 18(12):2891–901.

Zucker, R. S. and Regehr, W. G. (2002). Short-term synaptic plasticity. *Annu. Rev. Physiol.*, 64:355–405.